Algorithms for Two-sided Online Marketplaces

by

Reza Alijani

Department of Computer Science
Duke University

Date: ____________________
Approved:

____

Kamesh Munagala, Advisor

____
Siddhartha Banerjee

____
Vincent Conitzer

____
Debmalya Panigrahi

____

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Computer Science
in the Graduate School of
Duke University

2020

ABSTRACT

Algorithms for Two-sided Online Marketplaces

by

Reza Alijani

Department of Computer Science
Duke University

Date: ______________________

Approved:

____

Kamesh Munagala, Advisor

____

Siddhartha Banerjee

____

Vincent Conitzer

____

Debmalya Panigrahi

____

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Computer Science in the Graduate School of Duke University

2020

Abstract

The recent emergence of new successful two-sided online platforms has transformed the concept of a marketplace. Numerous two-sided mechanism design problems, including those studied in this thesis, are motivated by these platforms. Similar to any other mechanism design problem, we wish to optimize an objective in the presence of selfish agents. However, there are unique features to a two-sided market, such as supply uncertainty and the need for ensuring budget balance, which make these problems particularly challenging. In our problems, we design models to capture the two-sidedness, identify the challenges, and use algorithmic techniques to solve these problems.

We start with a Bayesian single item auction with n independent buyers. For this problem, we show how the existence of a trusted intermediary can result in a better outcome for buyers without hurting the seller's revenue. In this model, the intermediary gets to see the true valuations of buyers and will reveal some information after that in the form of a signal. The intermediary does not have any control over agents after sending this signal, and any agent only maximizes their utility. This essentially means that the seller will run an optimal auction conditioned on receiving any signal. For this problem, we design approximately optimal ways of revealing information.

The previous problem is an interesting model to show how a platform can mediate in the market and improve the outcome for every agent. That problem is a single item static problem. Next, we focus on pricing in two problems with multiple dynamic agents on the seller side. The first problem is an extension of the multiple-choice prophet inequality to the setting in which each item might disappear after an a

priori unknown amount of time. This can be viewed as a way to model the supply uncertainty arising because of the fact that different sellers might depart after some time. Considering the importance of prophet inequalities in online posted pricing mechanisms, we hope that incorporating features of two-sided markets in the model and finding new prophet inequalities will be useful and provide new insights.

In our last problem, we design a model to capture a general dynamic two-sided market. In our model, agents (buyers and sellers) with heterogeneous valuations/costs, service quality requirements, and patience levels (or deadlines) arrive over time. Both buyers and sellers arrive to the market following Poisson processes, and each buyer/seller has a private value/cost for service, as well as a private patience level for receiving/ providing service. In addition, each agent has a known location in an underlying metric space, where the metric distance between any buyer and seller captures the quality of the corresponding match. The platform knows the distribution over values/costs and patience levels and can post prices and wages at nodes in the metric space, as well as choose agents for matching. It uses these controls to maximize the social surplus subject to weak budget balance while guaranteeing a high match quality and a service time (with a high probability) smaller than their patience.

Contents

List of Figures

Acknowledgements

I would not have been able to finish my graduate studies without tremendous help from many people. First and foremost, I am extremely grateful to my advisor Kamesh Munagala for his support during my Ph.D. I have learned a lot from him both in academic and non-academic matters.

I would like to thank Siddhartha Banerjee for both working with me closely during these years and also serving on my committee. I am also grateful to the rest of my committee–Vincent Conitzer and Debmalya Panigrahi– who were also among many great teachers in graduate school that I learned a lot from.

Special thanks to my fellow student Kangning Wang. Kangning was a collaborator and a great friend during these years. I would also like to thank my other collaborators–Sreenivas Gollapudi and Kostas Kollias.

I appreciate the financial support from NSF and Duke University. Thanks also to all the professors, staff members, and students of the CS department at Duke, who made this experience very enjoyable.

Finally, I want to thank my family, who have constantly been supporting me in my life.

1 Introduction

While marketplaces as a physical place that producers and consumers gather to trade have existed for centuries, they have been drastically changed recently by the emergence of successful online platforms. These platforms have transformed the economic landscape of the modern world. Many of today's most important companies are platforms facilitating trade between agents: both for goods (Amazon, eBay), and increasingly, services (Lyft, Uber, Airbnb, etc). These platforms enable fine-grained monitoring, and greater control via pricing, recommendation, and directed search, etc. The challenge of harnessing this increase in data and control has led to a growing literature in online marketplace design.

The main theme of this thesis is to study different two-sided mechanism design problems motivated by these platforms. In all of these problems, the platform as an intermediary implements a mechanism to maximize an objective in the presence of strategic agents. Though this description has much in common with other one-sided mechanism design problems in settings such as sponsored search auction and monopolist pricing problems, there are a few features unique to two-sided markets which makes them particularly challenging.

One such feature is the supply uncertainty arising because of the two-sidedness of the market. This is particularly important in a dynamic marketplace in which agents can arrive and depart over time. This feature introduces new complexities to our problems and needs to be considered in our models and solutions. We provide some scenarios as examples to elaborate on how this feature can affect our problems. First, consider a setting in which there are some sellers available on our platform, but we believe that some of them might depart soon and no other seller substitute them in the near future. In this case, we might prefer to match buyers to those sellers even if they do not have a high value for the trade. Another interesting scenario is that we might want to accept an agent for a potential trade even if there is no match from the other side of the market at the moment, but we believe that it is likely that a match arrives soon after that. However, for this scenario, the agents do not want to wait for a long time, which might cause dissatisfaction with the platform. In other words, agents have limited patience. These complexities must be carefully considered when we design our models.

Another feature is the tension arising from the need for ensuring budget balance (i.e., non-negative profits). This introduces new challenges, even in a static problem. In particular, the problem of maximizing social welfare while respecting budget-balance is at the heart of the famous Myerson-Satterthwaite impossibility result and has received much attention in recent work in approximate mechanism-design. Next, we will explain this impossibility result and compare it to the one-sided problem to explain the challenges arising from this budget balance property.

In a one-sided market, the problem of designing a mechanism that maximizes welfare can be solved by the celebrated VCG mechanism, which is incentive compatible (IC) and individually rational (IR). The two-sided variant of this problem in the most basic form is called a double auction. In a double auction, there are unit-demand buyers and unit-supply sellers who want to trade, and they are indifferent between any two agents from the other side. Similar to the one-sided version of the problem, being IC and IR are two important desiderata of a solution. Though the VCG mechanism still outputs an efficient solution to this problem, the total amount of money paid to sellers can be larger than the total money paid by buyers. In other words, the mechanism is not budget balanced (BB). In addition, Myerson-Satterthwaite impossibility result shows that any Bayesian incentive compatible (BIC) and IR mechanism which is ex-post efficient, cannot be BB, i.e., it needs outside subsidies. This impossibility holds even for the simple setting of bilateral trade with a single buyer and a single seller.

We start with a Bayesian single-item auction with independent buyers in which both the seller and buyers seek to maximize their own utility. We show how a trusted intermediary who knows all the private values of buyers can share this information publicly such that the expected utility of buyers increases without decreasing the expected utility of the seller. This result can be seen as a justification for the existence of a platform or an intermediary, even in a static problem. In practice, it is common to have multiple sellers. Moreover, we need to add dynamicity to capture many applications. For this purpose, we add a specific type of supply uncertainty to the prophet inequalities. Given the importance of prophet inequalities and online posted pricing results in the one-sided setting, we hope that our results for this problem can be a first step towards developing the same set of results for two-sided variants of online posted pricing. These results are theoretically interesting and can lead to a better understanding of the challenges in a two-sided problem. However, they still do not capture a two-sided platform in its full generality. In the last chapter of this thesis, we present a model to capture different aspects of a real platform for dynamic marketplaces. We answer the question of optimal pricing for such a marketplace by studying different objectives.

1.1 Outline of Results

We explain how the dissertation is organized and mention the results presented in each chapter.

Chapter 2. In this chapter, we study how information intermediaries can increase the expected utility of buyers in Bayesian single-item auctions with n buyers and independent and discrete valuation distributions. The intermediary knows the true valuations of buyers and can provide additional information via sending a signal according to a signaling scheme. This can be considered as segmenting the prior distribution of the buyers so that consumer surplus is maximized in the presence of a revenue-maximizing seller. The work of Bergemann, *et al.* shows that a segmentation that raises the maximum possible consumer surplus always exists for one buyer. Our first result is an impossibility: Such a segmentation need not exist for $n = 2$ buyers. Indeed, no segmentation can achieve full social welfare while preserving any non-trivial of the consumer surplus without segmentation, and further, no segmentation can achieve consumer surplus that is more than a factor of 2 against a benchmark for maximum possible consumer surplus. Next, we consider approximation algorithms for consumer surplus against this benchmark. When the buyers' valuation distributions are identical, we present a $O(\min(\log n, \mathcal{K}))$ approximation where $\mathcal{K}$ is the support size of the valuations. For general independent distributions, we present a $O(\min(n \log n, \mathcal{K}^2))$ approximation.

Chapter 3. In this chapter, we consider the problem of selling multiple items to a stream of buyers whose values are drawn *i.i.d.* from a known distribution, in order to maximize social welfare. The items are identical from the perspective of the buyers. We consider the variant where the items do not last the entire stream but can disappear after an *a priori* unknown amount of time that we term the *horizon*. The mechanism knows the (possibly different) distribution of the horizon for each item, but not its realization till the item actually disappears. As with the classic prophet inequalities, the goal is to design an online pricing scheme that competes with the prophet that knows the horizon and extracts full social surplus (or welfare).

Our main results are for the setting with multiple items where the horizon distributions satisfy the monotone-hazard-rate (MHR) condition, and are independent across items, but not necessarily identical. For any number of items, we achieve a constant-competitive bound via a conceptually simple policy that balances the rate at which buyers are accepted with the rate at which items are removed from the system. We implement such a policy via a novel technique of matching via probabilistically simulating departures of the items at future times.

For a single item and MHR horizon distribution with mean μ, we show a tight result: There is a fixed pricing scheme that has a competitive ratio of at most $2-1/\mu$, and this is tight in the sense that there is a value distribution and an MHR horizon distribution for which any dynamic pricing scheme has a competitive ratio of at least $2 - 1/\mu$.

We further show that our results are best possible. First, we show that the competitive ratio is unbounded without the MHR assumption even for one item. Further, even when the horizon distributions are *i.i.d.* MHR and the number of items becomes large, the competitive ratio of any policy is lower bounded by a constant greater than 1, which is in sharp contrast to the setting with identical deterministic horizons.

Chapter 4. In this chapter, we model a dynamic two-sided market. In our model, agents arrive at nodes in an underlying metric space, where the metric distance between any buyer and seller captures the quality of the corresponding match. The platform posts prices and wages at the nodes, and opens a set of facilities to route the agents to. The agents at any facility are assumed to be matched. The platform ensures high match quality by imposing a distance constraint between a node and the facilities it is routed to. It ensures high service availability by ensuring flow to the facility is at least a pre-specified lower bound. Subject to these constraints, the goal of the platform is to maximize the social surplus subject to weak budget balance, i.e., profit being non-negative. We present an LP rounding based approximation algorithm for this problem that yields a $(1+\epsilon)$ approximation to surplus for any constant $\epsilon > 0$ while relaxing the match quality (i.e., maximum distance of any match) by a constant factor.

We also justify our models by considering a dynamic marketplace setting where agents arrive according to a stochastic process and have finite patience (or deadlines) for being matched. We perform queueing analysis to show that for policies that route agents to facilities and match them, ensuring a low abandonment probability of agents reduces to ensuring sufficient flow arrives at each facility.

Finally, we show how to extend our model and techniques to the case where the platform elicits deadlines truthfully by posting lotteries over different prices and wages for different deadlines.

2

Information Intermediaries in Single Item Auctions

In this chapter, we study the effects of an information intermediary on the utility of agents in a single item Bayesian auction with independent buyers. Without intermediary intervention, the seller runs an optimal auction to maximize her revenue. As a result of this auction, buyers might achieve positive utility in expectation since the seller only has distributional information about the valuations of buyers. The intermediary who knows the true values of the buyers can reveal additional information according to a signaling scheme. Upon receiving any signal, the seller can compute conditional value distributions and run a mechanism to maximize her revenue. This additional signal does not hurt the seller's revenue, and the goal of the intermediary is to increase the expected utility of the buyers via this process.

This problem can be considered as a two-sided market problem because of the fact that the seller is a strategic agent who wants to maximize her revenue, and the goals of the intermediary and the seller are not aligned. However, note that unlike many other two-sided problems such as bilateral trading, the seller does not have any private information. Besides, the platform is an information intermediary who only provides additional information. The seller is the one who runs the final mechanism.

The results of this chapter are from a working paper which is a joint work with Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang.

2.1 Introduction

Consider a seller selling an item to a buyer, whose private value V is drawn from some known distribution $\mathcal{D}$. The overall *social welfare* is maximized when the seller sells the item for \$0. In contrast, to maximize the (average) revenue, the seller's optimal strategy is to sell at a revenue-maximizing price, which may lead to welfare loss due to the item going unsold.

More generally, in a single-item Bayesian auction with n buyers with independent private valuations, a welfare-optimal mechanism is the second-price (or VCG) auction, which always gives the item to the highest-valued buyer. In contrast, even when the buyers have i.i.d. regular (continuous) valuations, the revenue-optimal mechanism was shown by Myerson [123] to be a second-price auction with a *reserve price*; this may lead to the item going unsold. The situation is more complex for non-i.i.d. buyers, where the revenue-optimal mechanism may in addition sell the item to a buyer with lower value than the highest, leading to more welfare loss. We visualize this via a *revenue-CS trade-off* diagram (Fig. 2.1b), where, for different mechanisms and value distributions, we plot expected consumer-surplus (i.e., value minus payment), denoted CS, versus expected seller-revenue, denoted by $\mathcal{R}$. Any welfare-maximizing mechanism including VCG (point V) lies on the 135° line with intercept $\mathcal{W}^*$, the maximum welfare. In contrast, Myerson's mechanism (point M) has revenue $\mathcal{R}^M$ greater than that under VCG, but also lies strictly below the maximum-welfare line.

Information Intermediary. Now consider the same setting, but with an additional *information intermediary*: a third-party who knows the *true* buyer values $\vec{V} = (V_1, V_2, \ldots, V_n)$ and can provide a "signal" or side-information to the seller and the buyers. Both the signal and signaling scheme are common knowledge to all agents (buyers and seller), who can thus use Bayes' rule to update the prior over valuations given the signal. The signal "re-shapes" the joint prior over the buyer valuations in a *Bayes-plausible* manner (i.e., such that the posterior averaged over signals equals the prior). Given the signal, a seller can then propose the revenue-maximizing mechanism, and buyers bid optimally, under the posterior distribution. We illustrate this in Fig. 2.1a.

Formally, consider a setting where n buyers have independent private valuations $\vec{V}$ drawn from a distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$. The valuations $\vec{V}$ are known to the intermediary, who maps them to a signal σ via a public signaling scheme $\mathcal{Z}$. Given σ, all agents compute the posterior $\mathcal{S}$ over buyer values; note these can now be correlated. The seller then proposes a mechanism $\mathcal{M}_{\mathcal{S}}$ (comprising allocation and payment rules) which maximizes the expected revenue assuming buyers act in a manner which is ex-interim incentive-compatible (IC) and individually-rational (IR) given $\mathcal{S}$. If σ is such that $\mathcal{S} = \mathcal{D}$, then $\mathcal{M}_{\mathcal{S}}$ is Myerson's auction (point M in Fig. 2.1b); on the other hand, if the signal fully reveals $\vec{V}$, then the seller can extract full surplus (i.e., get revenue $\mathcal{W}^*$, point A in Fig. 2.1b). Moreover, the seller gets revenue at least $\mathcal{R}^M$ under any signaling scheme, as she can always ignore the signal (see Section 2.2). Thus any signaling scheme $\mathcal{Z}$ must give a point in the shaded triangle with consumer surplus $\mathrm{CS}(\mathcal{Z})$ and revenue $\mathcal{R}(\mathcal{Z})$, and the maximum possible surplus OPT is achieved at point O in Fig. 2.1b. Now we can ask:

> What revenue-CS trade-offs can an information intermediary achieve via signaling? More specifically, what is the maximum possible consumer surplus that is achievable?

In the single-buyer case, Bergemann, Brooks and Morris [33] completely answer these questions by showing that the entire shaded region is always achievable. In particular, the point O is met by a simple signaling scheme where the revenue is *exactly* $\mathcal{R}^M$, and the item is *always* sold thus the mechanism is efficient.

In this chapter we study the effectiveness of an information intermediary in a multi-buyer (i.e., $n \geqslant 2$ buyers) Bayesian auction. In brief, we expose a sharp separation between the single and multi-buyer settings, as in the latter, no signaling scheme can guarantee more than a constant fraction of the optimal consumer surplus (OPT in Fig. 2.1b). On the positive side, we obtain a novel yet simple signaling scheme with strong approximation guarantees for a wide range of settings.

While our main focus is on theoretical results, our work has broader practical relevance. Consider for example an agency like the FCC with privileged information about bidders in a spectrum auction, or a bid optimizer working for multiple competing clients in an ad-exchange. These intermediaries have private information about the buyers, and can selectively release it to influence the auction. Our work also fits in a broader space of multi-criteria optimization where a third-party platform or government agency can release information about agents to a principal in charge of an activity such as admissions or hiring, so as to trade-off the principal's objective such as maximizing quality of hire, with a societal objective such as fairness or diversity.

2.1.1 Our Results

We consider a single-item auction with n buyers with discrete valuations. We assume the buyer valuations are independent, so $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$, where $\mathcal{D}_i$ has support size $\mathcal{K}_i$, and the size of the union of the supports is $\mathcal{K}$. We parametrize our results in terms of $n, \mathcal{K}_i$, and $\mathcal{K}$.

Our first set of results (Section 2.3) shows a sharp demarcation between the cases of $n = 1$ and $n \geqslant 2$ buyers. In contrast with the former case (where signaling achieves the entire shaded region in Fig. 2.1b), we show in the latter case, the entire segment BO is not achievable; indeed, the only achievable points on segment AO are arbitrarily close to A. Therefore, achieving full welfare needs sacrificing an arbitrarily

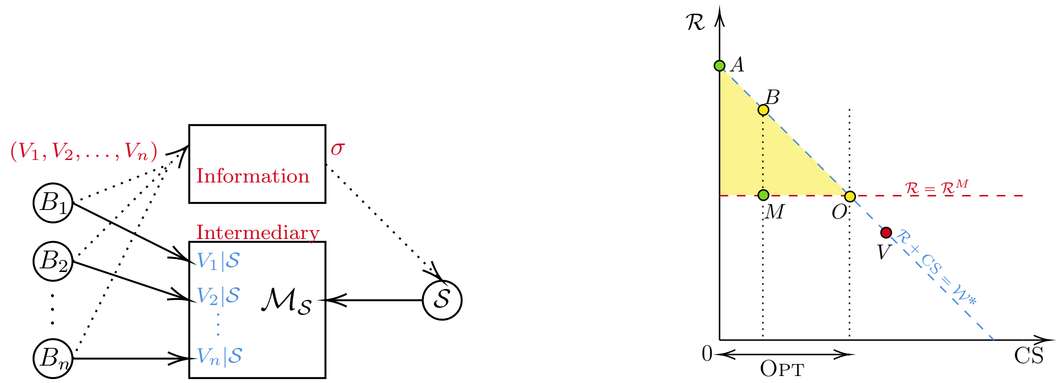

(a) Auction with Information Intermediary

(b) Revenue-CS Trade-off Diagram

Figure 2.1: *(a) The auction with information intermediary setting, where the intermediary has full knowledge of valuations $\vec{V}$, and can use this to provide a signal σ to the seller and the buyers. The seller then uses the revenue-optimal mechanism $\mathcal{M}_\mathcal{S}$ for the posterior distribution over valuations $\mathcal{S}$ given σ.*
(b) Two-dimensional space of seller revenue, $\mathcal{R}$, and consumer surplus, CS, of different signaling schemes. The points M and V correspond to the Myerson and VCG mechanisms, and the point A corresponds to selling to the highest-valued buyer at her value when the seller has full information.

large fraction of consumer surplus compared to the no-signaling baseline.

Theorem 1. *For $n = 2$ buyers each with $\mathcal{K}_i = 2$, for any given constant $\varepsilon > 0$, there are instances where any signaling scheme $\mathcal{Z}$ under which the revenue-optimal auction obtains full welfare (i.e., allocates to highest-value buyer), has $\mathrm{CS}(\mathcal{Z}) \leqslant \varepsilon \cdot \mathrm{CS}(\mathcal{D})$, where $\mathrm{CS}(\mathcal{D})$ is the consumer surplus of Myerson's auction without signaling.*

We next ask if we can sacrifice on welfare, but raise a consumer surplus arbitrarily close to OPT? We again answer in the negative, and show a lower bound of 2 on the approximation ratio.

Theorem 2. *For $n = 2$ buyers each with $\mathcal{K}_i = 2$, for any constant $\varepsilon > 0$, there are problem instances where any signaling scheme $\mathcal{Z}$ has $\mathrm{CS}(\mathcal{Z}) \leqslant \left(\frac{1}{2} + \varepsilon\right) \cdot \mathrm{OPT}$.*

We note that the above results are *existential impossibility results*, and do not depend on the complexity of the signaling scheme. Given these negative results, we focus on approximating the consumer surplus. To this end, we propose novel signaling schemes, which, informally, achieve a constant-approximation to OPT when either the number of buyers n is a constant (with arbitrary support size $\mathcal{K}$), or when $\mathcal{K}$ is a constant (for arbitrary n). Formally, our main result, presented in Sections 2.4 and 2.5, is the following:

Theorem 3 (Main)**.** *There are signaling schemes that achieve the following approximations to* OPT*:*

1. $O\big(\min(\log n, \mathcal{K})\big)$ *when $\mathcal{D}_i$'s are identical; and*
2. $O\big(\min(n \log n, \mathcal{K}^2)\big)$ *when $\mathcal{D}_i$'s are arbitrary.*

Further, this signaling scheme has computation time polynomial in n and $\mathcal{K}$.

2.1.2 Intuition and Techniques

For any n, the optimal signaling scheme for maximizing surplus can be obtained via an infinite-sized linear program (see Eq. (2.2) in Section 2.3) with variables for every possible signal, i.e., every possible joint distribution over buyer valuations. Further, for each such signal, the quantity of interest is the consumer surplus of the revenue-optimal auction given the signal. For $n = 1$ case, Bergemann et al. show this LP has a special structure in that it admits a basis comprising of "equal-revenue distributions" containing the revenue-maximizing price (see Section 2.2.2). Our work shows that this breaks down for optimal auctions with signaling involving $n \geqslant 2$ buyers.

To understand why things change dramatically from $n = 1$ to $n \geqslant 2$ buyers, in the former case, the optimal mechanism is a simple posted price scheme and its revenue is continuous in the distribution $\mathcal{D}$. However, with multiple buyers, the optimal auction does not have simple structure even for independent buyers (see Algorithm 1), and we need to analyze the consumer surplus of this auction, which can be a discontinuous function of the prior. Further, for correlated buyers, the revenue of the auction itself may not be continuous in the prior! Indeed, a celebrated result of Crémer and McLean [63] shows that slightly perturbing an independent prior to a correlated one can discontinuously increase the revenue to $\mathcal{W}^*$, hence decreasing consumer surplus to 0. (See Theorem 4 in Section 2.2.) This makes it tricky to reason about the optimal signaling scheme, leading to the gap between our upper and lower bounds.

In more detail: Our proofs of Theorems 1 and 2 use a special case of the Crémer-McLean characterization [63]: for $n = 2$ buyers each with $\mathcal{K}_i = 2$, under any non-independent prior the seller can extract full social surplus as revenue. This lets us focus on signaling schemes where buyers remain conditionally independent given each signal. Using Myerson's characterization of the optimal auction for discrete valuations [75], we show a structural characterization that reduces the space of optimal signals to a sufficiently simple form, yielding the desired counterexamples.

Theorem 3 is technically the most interesting result in this chapter. At a high level, our scheme balances the trade-off between revealing enough information about valuations so that the item is sold to a high-value buyer, and revealing too much information such that the seller extracts all the surplus. Balancing these is delicate; nevertheless, we show simple schemes, with polynomial computation time complexity.

Our schemes involve choosing a threshold value corresponding to that of the $(t + 1)^{\text{st}}$ buyer in descending order (for carefully chosen t), and then applying the single-buyer signaling scheme in [33] to the *excess value* (i.e., value minus threshold) of a randomly chosen buyer in the top t (while leaving the rest unchanged). Note that though the posterior conditioned on signal is a product distribution, it requires the intermediary to observe all the buyer valuations. We present this basic scheme in Section 2.4. To show our bounds that depend on n, we use an appropriate randomization over such schemes, while the bounds depending on $\mathcal{K}$ follow from a concentration lemma over independent Bernoulli trials (Lemma 6), which may be of independent interest.

2.1.3 Related Work

The general problem of information structure design considers how sharing additional information can influence the outcome of a mechanism. Different variants of this problem have been formulated and studied; we refer the reader to [35, 71] for surveys. Of particular importance to us is the *Bayesian persuasion problem* formulated by Kamenica and Gentzkow [105], where a receiver selects a utility-maximizing action based on incomplete information about the state of nature. A sender who knows the state of nature can signal side-information to the receiver so that the action taken by the receiver is utility-maximizing for the sender. This general problem has been widely studied in different domains such as bilateral trade and advertising [33, 54, 151]. For this problem, there is a distinction between existence and computational results, and the work of [73] studies the computational complexity of finding the optimal signaling scheme under different input models.

The restriction of our problem to one buyer is termed "bilateral trade". Here, the intermediary is the sender whose utility is consumer surplus, and the seller is the receiver whose action space is take-it-or-leave-it prices and whose utility is revenue. Beginning with the work of Bergemann et al. [33], several works [72, 138, 64, 49] have considered various extensions and modifications to this basic problem. Unlike bilateral trade where the buyer is perfectly informed, in our setting, not only the seller, but also all the buyers are receivers, in the sense that they have imperfect knowledge of the true valuations of other buyers, and modify their respective bidding strategies in response to the intermediary's signal to maximize their own utilities. Our setting is therefore a Bayesian persuasion problem with multiple receivers, and this aspect makes it significantly more complex.

There has been work on signaling in auctions that cannot be modeled as Bayesian persuasion, i.e., in which the common signal is not generated by an intermediary who knows all the true values of the buyers. For instance, in [36], the auctioneer has perfect information about buyer valuations and controls the precision to which buyers can learn it, and in [82], the seller's signal is drawn from a distribution that is correlated with the buyer's value, In both these works, the goal is to maximize seller revenue. Finally [139] studies equilibria of optimal auctions when each buyer commits to a signaling scheme with imperfect knowledge of other buyers' valuations, while [34] studies equilibria in first price auctions when buyers are provided correlated signals about other buyers' valuations. In contrast with the former, our work considers a richer space of signals via an information intermediary, while compared to the latter, in our setting the seller's mechanism is not fixed, but is instead also a function of the information structure.

2.2 Preliminaries

We consider Bayesian single-item auctions with n buyers, with independent private valuations $\vec{V} = (V_1, V_2, \ldots, V_n)$ drawn from a known product distribution $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$. Unless otherwise stated, we present our results for the setting in which each $\mathcal{D}_i$ is discrete. We denote by $\mathcal{K}_i$ the size of the support of $\mathcal{D}_i$, and by $\mathcal{K}$ the size of the union of these supports.

For distribution $\mathcal{D}_i$, we use f_{D_i} to denote its probability mass function, and define $S_{D_i}(x) = \mathbf{Pr}_{V \sim D_i}[V \geqslant x]$ and $F_{D_i}(x) = \mathbf{Pr}_{V \sim D_i}[V \leqslant x]$. For a joint distribution $\mathcal{D}$ and vector $\vec{v}$, we use $\mathbf{Pr}[\mathcal{D} = \vec{v}] = f_{\mathcal{D}}(\vec{v})$ as shorthand for denoting the probability of $\vec{v}$ drawn from $\mathcal{D}$.

2.2.1 Revenue-Maximizing Auctions

Given any shared prior $\mathcal{D}'$ on the valuations of the buyers, which in the case of signaling, can be different from $\mathcal{D}$ and arbitrarily correlated, the seller runs an optimal (revenue maximizing) auction that satisfies ex-interim incentive compatibility and individual rationality. Using the revelation principle [123], such an auction is specified by an allocation rule $x^*(\vec{v}) \geqslant 0$ and a payment rule $p^*(\vec{v})$ (that can be positive or negative) given any realized valuation profile $\vec{v}$. The quantity $x_i^*(\vec{v})$ is the probability buyer i gets the item given the valuation profile $\vec{v}$. Finally, for any prior $\mathcal{D}'$, let $(\mathcal{R}(\mathcal{D}'), \mathcal{W}(\mathcal{D}'), \mathrm{CS}(\mathcal{D}'))$ denote the expected revenue, welfare (or *total surplus*) and consumer surplus under the revenue-maximizing auction. Then we have $\mathrm{CS}(\mathcal{D}') = \mathcal{W}(\mathcal{D}') - \mathcal{R}(\mathcal{D}')$, and:

$$\mathcal{R}(\mathcal{D}') = \max \sum_{\vec{v}} \mathbf{Pr}[\mathcal{D}' = \vec{v}] \cdot \sum_i p_i^*(\vec{v}) \qquad \text{and} \qquad \mathcal{W}(\mathcal{D}') = \sum_{\vec{v}} \mathbf{Pr}[\mathcal{D}' = \vec{v}] \cdot \sum_i v_i x_i^*(\vec{v}).$$

First, we formalize the general revenue-maximizing auction. Our work builds on two special cases – independent valuations, and full surplus extraction. Next, we explain these cases.

General Setting

Give, $\mathcal{D}'$, the prior distribution on valuations (which can be arbitrarily correlated), let $\mathcal{D}_i'$ denote the marginal distribution of buyer i. The optimal auction is a mapping of each revealed valuation vector $\vec{v}$ to an allocation $\{x_i(\vec{v})\}$ and price $\{p_i(\vec{v})\}$. The allocation is always non-negative, while the price could be negative. The auction maximizes revenue:

$$\mathcal{R}(\mathcal{D}') = \max_{\vec{x}, \vec{p}} \sum_{\vec{v}} \mathbf{Pr}[\mathcal{D}' = \vec{v}] \cdot \sum_i p_i(\vec{v})$$

subject to the following ex-interim constraints. The incentive compatibility (IC) constraint states that the expected utility of any buyer does not increase by misreporting its valuation. Formally, for every agent i, for every value q, and for every valuation vector $\vec{v} = (q, \vec{v}_{-i})$ (where $\vec{v}_{-i}$ denotes the valuations of the other buyers), and every other possible report r, we have

$$\sum_{\vec{v}_{-i}} \frac{\mathbf{Pr}[\mathcal{D}' = (q, \vec{v}_{-i})]}{\mathbf{Pr}[\mathcal{D}'_i = q]} \cdot (x_i((q, \vec{v}_{-i})) - p_i((q, \vec{v}_{-i}))) \geqslant$$

$$\sum_{\vec{v}_{-i}} \frac{\mathbf{Pr}[\mathcal{D}' = (q, \vec{v}_{-i})]}{\mathbf{Pr}[\mathcal{D}'_i = q]} \cdot (x_i(r, \vec{v}_{-i}) - p_i(r, \vec{v}_{-i}))\,.$$

The individual rationality (IR) constraint says that for any buyer and any valuation, the expected utility under the mechanism is non-negative:

$$\sum_{\vec{v}_{-i}} \frac{\mathbf{Pr}[\mathcal{D}' = (q, \vec{v}_{-i})]}{\mathbf{Pr}[\mathcal{D}'_i = q]} \cdot (x_i((q, \vec{v}_{-i})) - p_i((q, \vec{v}_{-i}))) \geqslant 0, \qquad \forall i, q.$$

Finally, we have the constraint that the item is allocated probabilistically to at most one buyer:

$$\sum_i x_i(\vec{v}) \leqslant 1, \qquad \forall \vec{v}.$$

Optimal auction for independent valuations. When $\mathcal{D} = \mathcal{D}_1 \times \ldots \mathcal{D}_n$ is a product distribution, the optimal auction has a simple form given by Myerson [123]. For distribution $\mathcal{D}_i$ with support $z_1 < z_2 < \cdots < z_k$, its *virtual value* function $\varphi_{\mathcal{D}_i}$ is defined as:

$$\varphi_{\mathcal{D}_i}(z_k) = z_k \qquad \text{and} \qquad \varphi_{\mathcal{D}_i}(z_\ell) = z_\ell - (z_{\ell+1} - z_\ell)\frac{S_{\mathcal{D}_i}(z_{\ell+1})}{f_{\mathcal{D}_i}(z_\ell)}, \qquad \forall \ell < k. \tag{2.1}$$

If buyer i is the only buyer in the system, the optimal auction sets a fixed price, and the buyer buys the item when her valuation is at least this price. The *reserve price* of $\mathcal{D}_i$, denoted $r_{\mathcal{D}_i}$ is the smallest value r in the support of $\mathcal{D}_i$ that maximizes the corresponding revenue $rS_{\mathcal{D}_i}(r)$. It is easy to check that $\varphi_{\mathcal{D}_i}(r_{\mathcal{D}_i}) \geqslant 0$.

Throughout this chapter, for analytic simplicity, we assume the distributions $\mathcal{D}_i$ are *regular*, so that $\varphi_{\mathcal{D}_i}(z)$ is a non-decreasing function of z. Therefore, for all $v < r_{\mathcal{D}_i}$, we have $\varphi_{\mathcal{D}_i}(v) < 0$. Our proofs can be extended without this restriction.[1] For discrete regular distributions, Myerson's auction takes the form [75] in Algorithm 1. Note that this auction is also *ex-post* IC and IR.

Algorithm 1: Myerson's Auction with prior $\mathcal{D}$ and valuations $\vec{v}$.

---

1 Sort the buyers in decreasing order of $q_i = \varphi_{\mathcal{D}_i}(v_i)$. Assume no two values are identical (can be ensured by using a fixed tie-breaking rule).

2 Allocate to the bidder j with highest virtual value q_j, provided $q_j \geqslant 0$.

3 Let m be the bidder with second highest virtual value, and let $w = \max(0, q_m)$.

4 Charge j the smallest value z in the support of $\mathcal{D}_j$ such that $\varphi_{\mathcal{D}_j}(z) > w$.

Extracting full surplus as revenue. At the other extreme, a celebrated result of Crémer and McLean [63] shows that for distributions $\mathcal{D}'$ which are "sufficiently correlated", the optimal auction extracts full surplus (i.e., the revenue equals the maximum valuation in each valuation profile). Formally, the result requires that for each agent, their conditional distribution over others' values given their own value is full rank; for our purposes, we require a restriction of their result to $n = 2$ buyers, each with two possible valuations.

Theorem 4 (Crémer-McLean [63]). *For $n = 2$ buyers, where each buyer i has $\mathcal{K}_i = 2$ and the joint distribution over the valuations is $\mathcal{D}'$, the seller can extract the entire social welfare (expected value of the maximum of the buyer's valuations) as her revenue when $\mathcal{D}'$ is a correlated distribution.*

[1]When the valuations are non-regular, we use the non-decreasing ironed virtual value function [123, 75] instead.

2.2.2 Auctions with an Information Intermediary

We next formalize the model of an information intermediary illustrated in Fig. 2.1a. Since the effect of the intermediary's signal is captured by the resulting posterior distribution over valuations, for ease of notation, we henceforth use "signal" to refer to a distribution $\mathcal{S}$ over valuations.

A signaling scheme $\mathcal{Z} = \{\gamma_q, \mathcal{S}_q\}_{q\in[m]}$ comprises a collection of *signals* (i.e., joint distributions over valuations) $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_m$ and corresponding non-negative weights $\gamma_1, \gamma_2, \ldots, \gamma_m$. The scheme $\mathcal{Z}$ is feasible (or "Bayes plausible" [105]) if it satisfies $\sum_q \gamma_q = 1$ and $\sum_q \gamma_q \mathcal{S}_q = \mathcal{D}$. The intermediary commits to scheme $\mathcal{Z}$ before the auction, and it is known to the seller and all buyers.

The intermediary maps observed valuation profile $\vec{v} \sim \mathcal{D}$ to signal $\mathcal{S}_q$ with probability $\frac{\gamma_q \mathbf{Pr}[\mathcal{S}_q = \vec{v}]}{\mathbf{Pr}[\mathcal{D} = \vec{v}]}$. The seller uses $\mathcal{S}_q$ as the shared prior and runs an optimal auction on the buyers. Note that though $\mathcal{D}$ is a product distribution, the $\{\mathcal{S}_q\}$ can be correlated. Abusing the notation introduced in Eq. (2.1), we denote the revenue generated by signaling scheme $\mathcal{Z}$ as $\mathcal{R}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{R}(\mathcal{S}_q)$, its consumer surplus by $\mathrm{CS}(\mathcal{Z}) = \sum_q \gamma_q \mathrm{CS}(\mathcal{S}_q)$, and its welfare by $\mathcal{W}(\mathcal{Z}) = \sum_q \gamma_q \mathcal{W}(\mathcal{S}_q)$.

When $\mathcal{D}$ is a product distribution, the revenue from any signaling scheme must be at least the optimal revenue of Myerson's auction without signaling, $\mathcal{R}(\mathcal{D})$. To see this, we note that Myerson's auction on $\mathcal{D}$ is ex-post IC and IR. This means that this allocation and payment rule is still a feasible (ex-interim IC and IR) mechanism conditioned on receiving any signal, completing the argument. Therefore, the consumer surplus $\mathrm{CS}(\mathcal{Z})$ under any signaling scheme $\mathcal{Z}$ is bounded by the difference of the maximum possible welfare $\mathcal{W}^* = \mathbf{E}_{\vec{V}\sim\mathcal{D}}[\max_i V_i]$ and the maximum revenue without signaling $\mathcal{R}(\mathcal{D})$. We henceforth denote this bound as $\textsc{Opt}$, which is defined as follows:

$$\textsc{Opt} = \mathcal{W}^* - \mathcal{R}(\mathcal{D}).$$

We say that $\mathcal{Z}$ is a θ-approximation signaling scheme if $\text{CS}(\mathcal{Z}) \geqslant \frac{\text{OPT}}{\theta}$. Our goal is to find the best approximation factor θ via a signaling scheme whose computation time is polynomial in n and $\mathcal{K}$. In the rest of the chapter, we omit the dependence on $\mathcal{D}$ when clear from context.

Optimal signaling for a single buyer. For $n = 1$ buyer, Bergemann et al. [33] present a signaling scheme with consumer surplus exactly equal to OPT (i.e., implementing the point O in Fig. 2.1b. We present this in more detail in this section, and henceforth use $\text{BBM}(v, D)$ to refer to this scheme when the buyer has valuation distribution D and the realized value is $v \sim D$ but first we state some critical properties of the BBM scheme which we use in our results.

Lemma 1 (Implicit in Bergemann et al. [33])**.** *For a single buyer with value distribution $\mathcal{D}$ (with reserve price $r_{\mathcal{D}}$), the* BBM *mechanism satisfies the following properties:*

1. *For any signal $\mathcal{S}_q$, $\varphi_{\mathcal{S}_q}(v) \geqslant 0$ for all v in the support of $\mathcal{S}_q$.*
2. $\text{CS}(\text{BBM}) = \text{OPT} \geqslant \mathbf{Pr}_{V\sim\mathcal{D}}[V < r_{\mathcal{D}}] \cdot \mathbf{E}_{V\sim\mathcal{D}}[V \mid V < r_{\mathcal{D}}] = \sum_{v<r_{\mathcal{D}}} v f_{\mathcal{D}}(v)$.

The Single-Buyer Signaling Scheme of Bergemann et al. [33]

We describe one of the signaling schemes in [33] that achieves optimal consumer surplus OPT via a simple greedy algorithm. This scheme constructs distributions (signals) $\mathcal{S}_1, \mathcal{S}_2, \ldots \mathcal{S}_m$ and assigns weights $\gamma_1, \gamma_2, \ldots \gamma_m$ to them such that $\sum_q \gamma_q \mathcal{S}_q = \mathcal{D}$.

Let prior $\mathcal{D}$ takes value v_i with probability p_i, where $0 < v_1 < \ldots < v_k$. Let $\vec{p} = (p_1, p_2, \ldots, p_k)$. In each iteration ℓ, the algorithm constructs an *equal revenue* distribution $\mathcal{S}_\ell$ and subtracts it from the prior $\mathcal{D}$. This equal revenue distribution assigns positive probability $p_{i\ell}$ to v_i if $p_i > 0$ and assigns $p_{i\ell} = 0$ if $p_i = 0$. In $\mathcal{S}_\ell$, the seller raises equal revenue by setting the price to be any of the values v_i with $p_i > 0$.

It is easy to see that the equal revenue condition specifies a unique distribution $\mathcal{S}_\ell$. Note that since this signal is equal revenue, the seller sets the lowest value as price, so that the item always sells and the consumer surplus is maximum possible.

Let $\vec{p_\ell}$ be the probability vector of $\mathcal{S}_\ell$. We set the largest weight γ_ℓ such that $\vec{p} - \gamma_\ell \vec{p_\ell} \geqslant 0$. We update $\mathcal{D}$ by setting $\vec{p}$ to $\vec{p} - \gamma_\ell \vec{p_\ell}$, normalize it so that $\sum_i p_i = 1$, and increase ℓ by one. We repeat this till the support of $\mathcal{D}$ becomes empty. The $\{\gamma_\ell, \mathcal{S}_\ell\}$ specifies the signaling scheme.

We illustrate this procedure by an example below.

Example 1. *Suppose the type space is* $\{1, 2, 3\}$ *and* $\mathcal{D} = \langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \rangle$ *are the probabilities of these types. The monopoly price is* $p = 2$ *with revenue* $\mathcal{R}(\mathcal{D}) = \frac{4}{3}$*, while the point* A *in Fig. 2.1b has social welfare* $\mathcal{R}(A) = \mathcal{W}^* = \mathbf{E}[\mathcal{D}] = 2$*. Suppose* $\mathcal{S}_1 = \langle \frac{1}{2}, \frac{1}{6}, \frac{1}{3} \rangle$ *with* $\gamma_1 = \frac{2}{3}$*;* $\mathcal{S}_2 = \langle 0, \frac{1}{3}, \frac{2}{3} \rangle$ *with* $\gamma_2 = \frac{1}{6}$*; and* $\mathcal{S}_3 = \langle 0, 1, 0 \rangle$ *with* $\gamma_3 = \frac{1}{6}$*. It is easy to check that the monopoly price for each signal is the lowest price in its support so that the item always sells, and* $\sum \gamma_i \mathcal{R}(\mathcal{S}_i) = \frac{4}{3}$*. Therefore,* $\sum \gamma_i \mathrm{CS}(\mathcal{S}_i) = 2 - \frac{4}{3} = \frac{2}{3} = \textsc{Opt}$*, which corresponds to point* O *in Fig. 2.1b.*

2.3 Lower Bounds

In this section, we prove Theorems 1 and 2. Our lower bounds are based on a 2-buyer instance illustrated in Fig. 2.2: given values $a > b > c > d$, buyer 1 has value $V_1 \in \{a, b\}$ with probabilities α and $1 - \alpha$ respectively, while buyer 2 has value $V_2 \in \{c, d\}$ with probabilities β and $1 - \beta$ respectively. We choose $\alpha a = b$ and $\beta c = d$; thus, the virtual values satisfy: $\varphi_1(a) = a$, $\varphi_2(c) = c$, and $\varphi_1(b) = \varphi_2(d) = 0$. We call this distribution $\mathcal{D}$.

Characterization of optimal signaling. By Theorem 4, we know any signal that correlates the buyers raises zero consumer surplus. Therefore, the only signals $\mathcal{S}$ of

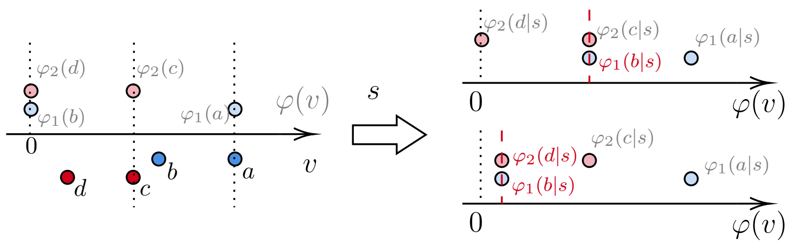

Figure 2.2: *Illustrating the setting for Theorems 1 and 2: On the left (below the axis) we show the setting without signaling, where buyer 1 (blue) has values (a, b) and buyer 2 (red) has values (c, d); we also show the corresponding virtual values (above the axis). On the right, we show the two settings characterized by Theorem 5 under which a signal s has non-zero consumer surplus (the changed virtual values are highlighted).*

interest are those under which buyer values are independent. Abusing notation we denote such a signal as $s = (\alpha', \beta')$, where $\mathbf{Pr}[v_1 = a] = \alpha'$ and $\mathbf{Pr}[v_2 = c] = \beta'$. Note that in this instance, for a signal to get maximum welfare the resulting optimal mechanism must always award buyer 1, and for non-zero consumer surplus it must award the item to buyer 1 at price b, or buyer 2 at price d.

Let $\mathrm{CS}(s)$ denote the consumer surplus under any such a signal s, and $\varphi_1(b|s)$ and $\varphi_2(d|s)$ denote the new virtual values (note that by definition, $\varphi_1(a|s) = a$ and $\varphi_2(c|s) = c$ under any signal s with $\alpha', \beta' > 0$). We can use Myerson's characterization (Section 2.2.1) to exhaustively characterize the resulting optimal mechanisms as a function of $(\varphi_1(b|s), \varphi_2(d|s))$ – we do so in Proposition 1 and our main insight is that the setting can be further simplified to get the following structural property for the optimal signaling scheme.

Theorem 5 (Structural Theorem)**.** *In an optimal signaling scheme, the only signals $s = (\alpha', \beta')$ that raise non-zero consumer surplus have the following form:*

(1') Under signal s, $\varphi_1(b|s) = \varphi_2(c|s) = c$ and $\mathrm{CS}(s) = \alpha'(a - b)$.

(2') Under signal s, $\varphi_2(d|s) = \varphi_1(b|s) \geqslant 0$ and $\mathrm{CS}(s) = \alpha'(1 - \beta')(a - b)$.

Proof. Recall that we restrict ourselves to signals $\mathcal{S}$ under which the buyer valuations remain independent. Any such signal can be alternately written as $s = (\alpha', \beta')$ where

$\alpha' = \mathbf{Pr}[v_1 = a]$ and $\beta' = \mathbf{Pr}[v_2 = c]$. For ease of notation, we henceforth drop the conditioning of virtual valuations on signal s (i.e., write $\varphi(\cdot)$ for $\varphi(\cdot|s)$) when clear from context.

The following proposition uses Myerson's characterization [123, 75] to exhaustively list out the different forms of auctions that can be realized under such signals, and the corresponding revenue and consumer surplus.

Proposition 1. *Conditioned on receiving a signal s, we have the following cases:*

1. *If $\varphi_1(b) \geqslant c$, then the optimal mechanism is to sell to Buyer 1 at price b.* $\mathrm{CS}(s) = (a-b)\alpha'$.
2. *If $\varphi_2(d) \geqslant \max(0, \varphi_1(b))$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price d.* $\mathrm{CS}(s) = (1-\alpha')\beta'(c-d)$.
3. *If $\varphi_1(b) \leqslant 0$ and $\varphi_2(d) \leqslant 0$, then the optimal mechanism is to try selling to Buyer 1 at price a then to Buyer 2 at price c.* $\mathrm{CS}(s) = 0$.
4. *If $0 \leqslant \varphi_1(b) \leqslant c$ and $\varphi_2(d) \leqslant \varphi_1(b)$, then the optimal mechanism is to sell to Buyer 1 at price b if Buyer 2 has valuation d; otherwise, it tries selling to Buyer 1 at price a then to Buyer 2 at price c.* $\mathrm{CS}(s) = \alpha'(1-\beta')(a-b)$.

Next, let γ_s denote the weight of any signal $s = (\alpha', \beta')$. The signaling scheme that maximizes consumer surplus is the solution to the following linear program written over signals $s = (\alpha', \beta')$:

$$
\begin{array}{lrcll}
\text{Maximize} & \sum_s \gamma_s \mathrm{CS}(s) & & & \\
\text{Subject to} & \sum_{s=(\alpha',\beta')} \gamma_s \alpha' \beta' & \leqslant & \alpha\beta & \\
 & \sum_{s=(\alpha',\beta')} \gamma_s \alpha'(1-\beta') & \leqslant & \alpha(1-\beta) & \\
 & \sum_{s=(\alpha',\beta')} \gamma_s (1-\alpha')\beta' & \leqslant & (1-\alpha)\beta & \\
 & \sum_{s=(\alpha',\beta')} \gamma_s (1-\alpha')(1-\beta') & \leqslant & (1-\alpha)(1-\beta) & \\
 & \gamma_s & \geqslant & 0 & \forall s
\end{array}
\tag{2.2}
$$

We examine the cases in Proposition 1 with positive consumer surplus, and characterize the optimal solution:

- In Case (1), we have $\varphi_1(b) = c$. To see this, consider any signal s with $\varphi_1(b) > c$. Suppose we increase α' and decrease γ_s while preserving the product $\alpha'\gamma_s$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s \alpha'(a-b)$, this is preserved by the change. Therefore, the objective of LP (2.2) is preserved, and so are the first two constraints. Further, since $(1-\alpha')$ decreases, this only makes the third and fourth constraints more feasible. This transformation decreases $\varphi_1(b)$.
- In Case (2) and (4), we have $\varphi_2(d) = \varphi_1(b)$. It does not help to make them unequal by a similar argument as above: In case (2), if $\varphi_2(d) > \varphi_1(b)$, we can increase β' while preserving $\gamma_s\beta'$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s(1-\alpha')\beta'(c-d)$, this does not change the contribution to the objective of LP (2.2), and preserves all constraints. This transformation decreases $\varphi_2(d)$ In case (4), if $\varphi_2(d) < \varphi_1(b)$, we can increase α' while preserving $\gamma_s\alpha'$. Since $\gamma_s \mathrm{CS}(s) = \gamma_s\alpha'(1-\beta')(a-b)$, this does not change the contribution to the objective of LP (2.2), and preserves all constraints. This transformation decreases $\varphi_1(b)$

Therefore, the only two types of signals s that give positive CS are

(1') If $\varphi_1(b) = c$, then $\mathrm{CS}(s) = \alpha'(a-b)$.

(2') If $\varphi_2(d) = \varphi_1(b) \geqslant 0$, then $\mathrm{CS}(s) = \max((1-\alpha')\beta'(c-d), \alpha'(1-\beta')(a-b))$.

As $(1-\beta')(b-d) \geqslant 0$, we have

$$\left(b - d + \frac{\beta'}{1-\beta'}(c-d)\right)(1-\beta') \geqslant \beta'(c-d).$$

Notice that in Case (2'), we have $\varphi_1(b) = b - \frac{\alpha'}{1-\alpha'}(a-b) = d - \frac{\beta'}{1-\beta'}(c-d) = \varphi_2(d)$. This gives

$$\alpha'(1-\beta')(a-b) \geqslant (1-\alpha')\beta'(c-d).$$

Thus, the two types of signals s that give positive CS become

(1') If $\varphi_1(b) = c$, then $\text{CS}(s) = \alpha'(a - b)$.

(2') If $\varphi_2(d) = \varphi_1(b) \geqslant 0$, then $\text{CS}(s) = \alpha'(1 - \beta')(a - b)$. □

Using the above structural theorem, the proofs of Theorems 1 and 2 follow by different choices of the parameters (a, b, c, d). For Theorem 1, we set $d = \left(1 - \frac{\varepsilon}{2}\right) b$, with $\varphi_1(b)$ and $\varphi_2(d)$ slightly above zero; for Theorem 2, we set $\alpha = \beta = 1 - \delta$; $a = \frac{1}{(1-\delta)^2}$, $b = \frac{1}{1-\delta}$, $c = 1$, $d = 1 - \delta$ with $\delta \to 0^+$. We provide complete proofs of Theorem 1 in Section 2.3.1 and Theorem 2 in Section 2.3.2.

2.3.1 First Lower Bound

In the instance presented in the beginning of Section 2.3, suppose the virtual values of b and d are slightly above zero with $\varphi_1(b) > \varphi_2(d)$ so that Case (4) in Proposition 1 is uniquely optimal for the seller. The optimal auction generates consumer surplus $\text{CS}(\mathcal{D}) = \alpha(1 - \beta)(a - b) = \frac{b}{c} \cdot \frac{c-d}{a} \cdot (a - b)$ according to Proposition 1.

To prove Theorem 1, we set $b \to c^+$.[2] Now in Proposition 1, in Case (1), we must have $\alpha' \to 0^+$ since $b \to c^+$, so that $\text{CS} \to 0$. Also if $\alpha' = 1$ in a signal then $\text{CS} = 0$ here. The only other signal where the item is allocated to the higher bidder is in Case (4) when $\beta' = 0$. Let γ denote the probability of the signal of this type $s = (\alpha', 0)$. (Having multiple signals of this form gives the same CS as having a single signal as their average.) Since $\varphi_1(b) \geqslant \varphi_2(d)$, we have $\alpha' \leqslant \frac{b-d}{a-d}$.

By the constraints of LP (2.2), we have:

$$\mathbf{Pr}[v_1 = b \wedge v_2 = d] = \gamma(1 - \alpha') \leqslant (1 - \alpha)(1 - \beta),$$

[2] $\varphi_1(b) - \varphi_2(d)$ and $\varphi_2(d)$ can be arbitrarily small as long as positive, so we take the limits for them first, i.e., we are calculating $\lim_{b \to c^+} \lim_{\varphi_2(d) \to 0^+, \varphi_1(b) \to \varphi_2(d)^+} \text{CS}$ in the following part of the proof. This allows us to treat $\alpha = \frac{b}{a}$ and $\beta = \frac{d}{c}$ in calculating $(1 - \alpha)(1 - \beta)$, as α and β are not infinitesimally close to 1 for any fixed ε.

which simplifies to $\gamma \leqslant \frac{(a-d)\cdot(c-d)}{ac}$. The consumer surplus in this case is therefore:

$$\mathrm{CS} = \gamma \mathrm{CS}(s) = \gamma\alpha'(a-b) \leqslant \frac{(b-d)\cdot(c-d)}{ac}\cdot(a-b) \leqslant \frac{b-d}{b}\cdot \mathrm{CS}(\mathcal{D}).$$

Setting $d = \left(1-\frac{\varepsilon}{2}\right)b$ and combining with the fact that CS $\to 0$ in Case (1), we have the consumer surplus of any efficient signaling, CS $\to \frac{\varepsilon}{2}\cdot \mathrm{CS}(\mathcal{D})$ so CS $< \varepsilon \cdot \mathrm{CS}(\mathcal{D})$.

2.3.2 Second Lower Bound

Without signaling, $\mathbf{E}[\max v_i] = \alpha a + (1-\alpha)b$ and $\mathcal{R}(\mathcal{D}) = \alpha a + (1-\alpha)\beta c$. (Case (2), (3) and (4) in Proposition 1 give the same revenue $\mathcal{R}(\mathcal{D})$.) Therefore

$$\textsc{Opt} = \mathbf{E}[\max v_i] - \mathcal{R}(\mathcal{D}) = (1-\alpha)(b-\beta c).$$

Now we assign the values as: $\alpha = \beta = 1-\delta$; $a = \frac{1}{(1-\delta)^2}$, $b = \frac{1}{1-\delta}$, $c = 1$, $d = 1-\delta$ with $\delta \to 0^+$. Then we plug in the values and the two possible types of signals s in Theorem 5 become

(1') If $\alpha' = \frac{1-\delta}{2-\delta} < \frac{1}{2}$, then $\mathrm{CS}(s) \leqslant \frac{1}{2}\delta(1+o(1))$.

(2') If $\alpha' \leqslant 1-\delta$; $\beta' = \frac{1-3(1-\alpha')+3\delta(1-\alpha')-\delta^2(1-\alpha')}{1-2(1-\alpha')+\delta(1-\alpha')} > \frac{1-3(1-\alpha')}{1-2(1-\alpha')}$, then $\mathrm{CS}(s) \leqslant \alpha'(1-\beta')\delta(1+o(1))$.

The consumer surplus maximizing signaling scheme should use t signals $S_{2,i}$ of type (2'), with $\alpha'_{2,i}$, $\beta'_{2,i}$ and weight $w(S_{2,i})$. There is an additional signal S_1 (with weight $w(S_1)$) of type (1') with α'_1 and β'_1. (Having multiple signals of type (1') gives the same CS as having a single signal as their average.) Denoting the valuation of the first buyer by v_1 and the second buyer by v_2, the constraints in LP (2.2) imply the two constraints:

$$\mathbf{Pr}[v_1 = b] = (1-\alpha'_1)w(S_1) + \sum_{i=1}^{t}(1-\alpha'_{2,i})\cdot w(S_{2,i}) \leqslant 1-\alpha = \delta, \quad \text{(Constraint (I))}$$

$$\mathbf{Pr}[v_1 = b \wedge v_2 = d] = \sum_{i=1}^{t}(1-\alpha'_{2,i})(1-\beta'_{2,i}) \cdot w(S_{2,i}) \leqslant (1-\alpha)(1-\beta) = \delta^2. \qquad \text{(Constraint (II))}$$

Note that $\textsc{Opt} = (1-\alpha)(b-\beta c) = 2\delta^2(1+o(1))$.

The total consumer surplus therefore is:

$$\begin{aligned}
\text{CS} \leqslant & \frac{1}{2}\delta(1+o(1)) \cdot w(S_1) + \sum_{i=1}^{t} \alpha'_{2,i}(1-\beta'_{2,i})\delta(1+o(1)) \cdot w(S_{2,i}) \\
\leqslant & \frac{1}{2}\delta(1+o(1)) \cdot 2\left(\delta - \sum_{i=1}^{t}(1-\alpha'_{2,i}) \cdot w(S_{2,i})\right) + \delta(1+o(1))\sum_{i=1}^{t}\alpha'_{2,i}(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
= & \delta^2(1+o(1)) + \delta(1+o(1))\sum_{i=1}^{t}(\alpha'_{2,i}(1-\beta'_{2,i}) - (1-\alpha'_{2,i})) \cdot w(S_{2,i}) \\
\leqslant & \delta^2(1+o(1)) + \delta(1+o(1))\sum_{i=1}^{t}(1-\alpha'_{2,i})(1-\beta'_{2,i}) \cdot w(S_{2,i}) \\
\leqslant & \delta^2(1+o(1)) + \delta(1+o(1)) \cdot \delta^2 \\
= & \delta^2(1+o(1)).
\end{aligned}$$

Here the second inequality follows from Constraint (I), and $\alpha'_1 < \frac{1}{2}$. The third inequality uses the implication of $\varphi_2(d) = \varphi_1(b)$ that $\beta'_{2,i} > \frac{1-3(1-\alpha'_{2,i})}{1-2(1-\alpha'_{2,i})}$. The fourth inequality uses Constraint (II). This establishes a lower bound of 2, since $\textsc{Opt} = 2\delta^2(1+o(1))$.

2.4 Ranking-Based Multi-Buyer Signaling Scheme

In this section, we introduce a family of signaling schemes, which forms the basic subroutine for obtaining our upper bounds. We first need one additional definition. For agent i with $V_i \sim \mathcal{D}_i$, we use $\mathcal{D}_{i|>a}$ to denote the conditional distribution of V_i given $V_i > a$, and $\mathcal{D}_{i|<a}$ to denote the conditional distribution of V_i given $V_i < a$. Moreover, we use $\mathcal{D}_{i|>a} - b$ to denote the distribution of $V_i - b$ given $V_i > a$; we refer to it as the distribution of V_i truncated at a and reduced by b. We now state a simple result relating the reserve price of the truncated and the original distributions.

Lemma 2. *Let $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$. Then we have $r_{\mathcal{D}'_i} \geqslant r_{\mathcal{D}_i} - v^\circ$, and moreover, for any $v > v^\circ$, we have $\varphi_{\mathcal{D}'_i}(v - v^\circ) = \varphi_{\mathcal{D}_i}(v) - v^\circ$.*

Proof. We first prove the result about reserve prices. If $r \leqslant v^\circ$ the inequality is trivial. Otherwise, suppose for contradiction that for some $r' < r - v^\circ$, $r' \cdot S_{\mathcal{D}'_i}(r') \geqslant (r - v^\circ) \cdot S_{\mathcal{D}'_i}(r - v^\circ)$. Then we would have: $(r' + v^\circ) \cdot S_{\mathcal{D}'_i}(r') \geqslant r \cdot S_{\mathcal{D}'_i}(r - v^\circ)$, since $S_{\mathcal{D}'_i}(r') \geqslant S_{\mathcal{D}'_i}(r - v^\circ)$. Thus, $(r' + v^\circ) \cdot S_{\mathcal{D}_i}(r' + v^\circ) \geqslant r \cdot S_{\mathcal{D}_i}(r)$, a contradiction to the assumption that r is the smallest optimal reserve price of $\mathcal{D}_i$.

To see the second part, if we condition the distribution on $v > v^\circ$, this does not change the virtual value function for values $v_j > v^\circ$, since both the numerator and denominator in Eq (2.1) scale by the same amount. If we now subtract v° from the support of the distribution, it reduces the virtual value by the same amount. This completes the proof.

□

Note that $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ represents buyer i's *excess value* compared to v°. Lemma 2 shows that for any threshold v° and any buyer i, given the side-information that $V_i > v^\circ$, her new reserve price is greater than her original reserve price.

The $\textsc{Rank}_t$ signalling scheme. We now introduce a family of signaling schemes $\textsc{Rank}_t$ parameterized by $t \in \{1, \ldots n\}$, For any realized joint valuation profile $\vec{v} = (v_1, v_2, \ldots, v_n)$, the signal sent by $\textsc{Rank}_t$ consists of two parts. First, $\textsc{Rank}_t$ observes $\vec{v}$ and outputs (v°, T), where v° is the value of $(t+1)^{\text{st}}$ largest realized value (or 0 when $t = n$), and T is the subset of buyers with realized value strictly greater than v°. For the second part of the signal, $\textsc{Rank}_t$ chooses a buyer j uniformly at random from T, and computes her excess distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. It then reveals both the identity of j, as well as the signal $\text{BBM}(v_j - v^\circ, \mathcal{D}_{j|>v^\circ} - v^\circ)$ generated by the single-buyer BBM scheme on a buyer with value distribution $\mathcal{D}_{j|>v^\circ} - v^\circ$. The scheme is formalized in Algorithm 2.

Algorithm 2: $\text{RANK}_t(\vec{v}, \mathcal{D})$

```
v° ← (t+1)^st largest value in v⃗
T ← {i : v_i > v°}
if T ≠ ∅ then
    j ← Buyer chosen uniformly at random from T
    s ← BBM(v_j − v°, D_{j|>v°} − v°)
    return v°, T, j, and s
else
    return v°, and T = ∅
```

Optimal mechanism under RANK_t. Conditioned on receiving the signal generated by RANK_t, the seller is guaranteed a revenue of v° from the $(t+1)^{\text{st}}$ largest buyer, and knows that only buyers in T can pay more than v°. The seller can now charge at least v° to any buyer in T, and can further run an auction over the excess value of buyers in T, where for buyer $i \in T$, her excess value has distribution $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$. Note that for any buyer $i \in T$ except the randomly chosen buyer j, a value drawn from $\mathcal{D}'_i$ represents how much more than v° she is willing to pay. Moreover, distributions $\mathcal{D}'_i$ are independent, and also, since the identity of j is chosen uniformly at random, the BBM scheme modifies the distribution of buyer j in a fashion that is independent of $\mathcal{D}'_i$.

By Lemma 1, we know that the BBM scheme ensures the virtual value of buyer j is always non-negative. From the characterization of the optimal auction [123, 75], since the item is always allocated to the highest virtual value bidder as long as this value is non-negative, the item will always be allocated to buyer j if all other buyers

$i \in T, i \neq j$ have excess values $v_i - v^\circ < r_{\mathcal{D}'_i}$ (and hence, negative virtual values).

Consumer surplus under Rank_t**.** Let $p_i := 1 - S_{\mathcal{D}_i}(r_{\mathcal{D}_i}) = \mathbf{Pr}_{v_i \sim \mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$ for any buyer i, where $r_{\mathcal{D}_i}$ is the reserve price of $\mathcal{D}_i$. Let $Y_i := \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ denote the distribution of $\mathcal{D}_i$ conditioned on being smaller than $r_{\mathcal{D}_i}$. Suppose we draw a sample independently from each distribution Y_i. Let Z_ℓ denote the distribution for the ℓ^{th} largest value among these n draws. The following key lemma gives a lower bound for the consumer surplus generated under Rank_t.

Lemma 3. *For $1 \leqslant t \leqslant n$, the consumer surplus of Rank_t satisfies:*

$$\text{CS}(\text{Rank}_t) \geqslant \left(\prod_{i=1}^{n} p_i\right) \cdot \left(\left(\frac{1}{t} \cdot \sum_{l=1}^{t} \mathbf{E}[Z_l]\right) - \mathbf{E}[Z_{t+1}]\right).$$

Proof. Fix a buyer b, and any valuation profile $\vec{v}_{-b} = \{v_i, i \neq b\}$ such that $v_i < r_{\mathcal{D}_i} \,\forall\, i \neq b$. Define v_b^t to denote the t^{th} largest value in $\{v_i, i \neq b\}$. Now consider the event

$$Q(\vec{v}_{-b}, b, t) = \{V_b > v_b^t \text{ AND } b \text{ selected for BBM signaling}\}.$$

Conditioned on $Q(\vec{v}_{-b}, b, t)$, we have that the Rank_t scheme (Algorithm 2) with parameter t sets threshold value as $v^\circ = v_b^t$. By Lemma 2, we have that for every $i \in T, i \neq b$, their value v_i is smaller than their new reserve price $v^\circ + r_{\mathcal{D}'_i}$, since $v_i < r_{\mathcal{D}_i}$, and modifying $\mathcal{D}_i$ to $\mathcal{D}'_i = \mathcal{D}_{i|>v^\circ} - v^\circ$ does not decrease the reserve price. Therefore, conditioned on $Q(\vec{v}_{-b}, b, t)$, the auction behaves like the single item mechanism $\text{BBM}(v'_b, \mathcal{D}'_b)$, where $v'_b = v_b - v_b^t$, $\mathcal{D}'_b = \mathcal{D}_{b|>v_b^t} - v_b^t$. Let r'_b denote the reserve price of $\mathcal{D}'_b$; again using Lemma 2 we have $r'_b \geqslant r_b - v_b^t$. Now, using Lemma 1, we get that the expected consumer surplus generated by Rank_t under $Q(\vec{v}_{-b}, b, t)$ is at least:

$$\mathbf{E}[\text{CS}(\text{Rank}_t) \mid Q(\vec{v}_{-b}, b, t)] \geqslant \sum_{v'_b < r'_b} v'_b \,\mathbf{Pr}[\mathcal{D}'_b = v'_b] \geqslant \sum_{v_b^t < v_b < r_b} (v_b - v_b^t) \frac{f_{\mathcal{D}_b}(v_b)}{S_{\mathcal{D}_b}(v_b^t)}.$$

Note also that $\mathbf{Pr}[Q(\vec{v}_{-b}, b, t)] = \frac{1}{t} \cdot \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i) \mathbb{1}_{\{v_i < r_{\mathcal{D}_i}\}}\right) \cdot S_{\mathcal{D}_b}(v_b^t)$. Thus for any buyer b, and any valuation profile $\vec{v}_{-b}$ with $v_i < r_i$ for all $i \neq b$, we have

$$\begin{aligned}
\mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b},b,t)}] &= \mathbf{Pr}[Q(\vec{v}_{-b}, b, t)] \cdot \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \mid Q(\vec{v}_{-b}, b, t)] \\
&\geqslant \frac{1}{t} \left(\prod_{i \neq b} f_{\mathcal{D}_i}(v_i)\right) \sum_{v_b^t < v_b < r_b} (v_b - v_b^t) f_{\mathcal{D}_b}(v_b) \\
&= \frac{1}{t} \sum_{v_b < r_b} \left(\prod_i f_{\mathcal{D}_i}(v_i)\right) \max\{(v_b - v_b^t), 0\}.
\end{aligned}$$

For any $\vec{v}$ let $v^{(t)}$ denote the t^{th} largest value, and $I_t(\vec{v})$ to be the indices corresponding to the top t values in $\vec{v}$. Summing up over all b, and all $\vec{v}_{-b}$ such that $v_i < r_i \, \forall \, i \neq b$, we have

$$\begin{aligned}
\sum_b \sum_{\vec{v}_{-b}} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b},b,t)}] &\geqslant \sum_{\vec{v} | v_i < r_i} \frac{1}{t} \cdot \left(\prod_i f_{\mathcal{D}_i}(v_i)\right) \cdot \left(\sum_b \max\{(v_b - v_b^t), 0\}\right) \\
&= \sum_{\vec{v} | v_i < r_i} \left(\prod_i f_{\mathcal{D}_i}(v_i)\right) \left(\sum_{i \in I_t(\vec{v})} \frac{1}{t} \left(v_i - v^{(t+1)}\right)\right) \\
&= \sum_{\vec{v} | v_i < r_i} \left(\prod_i f_{\mathcal{D}_i}(v_i)\right) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t}\right) - v^{(t+1)}\right).
\end{aligned} \tag{2.3}$$

Let $\mathcal{D}'_i = \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ be the distribution of buyer i's value conditioned on $V_i < r_{\mathcal{D}_i}$. Recall we define $p_i = \mathbf{Pr}_{v_i \sim \mathcal{D}_i}[v_i < r_{\mathcal{D}_i}]$; thus $f_{\mathcal{D}'_i}(v) = f_{\mathcal{D}_i}(v)/p_i$. Suppose we independently sample $Y_i \sim \mathcal{D}_{i|<r_{\mathcal{D}_i}}$ for each i, and define Z_ℓ as the ℓ^{th} largest value in $\{Y_i\}$. Then Eq. (2.3) can be written as

$$\begin{aligned}
\sum_b \sum_{\vec{v}_{-b} | v_i < r_i \, \forall \, i \neq b} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b},b,t)}] &\geqslant \sum_{\vec{v} | v_i < r_i} \left(\prod_i p_i\right) f_{\mathcal{D}'}(\vec{v}) \left(\left(\sum_{i \in I_t(\vec{v})} \frac{v_i}{t}\right) - v_{(t+1)}\right) \\
&= \left(\prod_i p_i\right) \left(\left(\sum_{\ell=1}^{t} \frac{\mathbf{E}[Z_\ell]}{t}\right) - \mathbf{E}[Z_{t+1}]\right).
\end{aligned}$$

Finally, noting that the $Q(\vec{v}_{-b}, b, t)$ events are all non-overlapping, we can write

$$\mathrm{CS}(\textsc{Rank}_t) \geqslant \sum_{b} \sum_{\vec{v}_{-b} | v_i < r_i \, \forall \, i \neq b} \mathbf{E}[\mathrm{CS}(\textsc{Rank}_t) \cdot \mathbb{1}_{Q(\vec{v}_{-b}, b, t)}],$$

thereby completing the proof.

□

Lemma 3 forms the crux of our subsequent analysis, helping us quantify how our selective BBM signal recovers much of the surplus lost by the Myerson auction. The difficulty in proving it arises from the correlation between the probability the chosen buyer j wins the auction, and the surplus raised by the BBM scheme, and we get around it by *coupling* the surplus generated when buyer values are above the reserve with the order statistics of buyer valuations below the reserve.

2.5 Approximating Consumer Surplus

We now prove our main theorem. For this, we first decompose our benchmark consumer surplus $\textsc{Opt}$, and then use Lemma 3 to quantify how we recover each term via signaling. Recall we start with a product distribution $\mathcal{D} = \mathcal{D}_1 \times \mathcal{D}_2 \times \cdots \times \mathcal{D}_n$, where $\mathcal{D}_i$ has reserve price r_i. Let $x_1 < x_2 < \ldots < x_{\mathcal{K}}$ be the union of the supports of value distributions. $\mathcal{R}(\mathcal{D})$ denotes the revenue of the optimal auction on $\mathcal{D}$, and $\textsc{Opt} = \mathbf{E}_{\vec{v} \sim \mathcal{D}}[\max_i v_i] - \mathcal{R}(\mathcal{D})$. We can now decompose $\textsc{Opt}$ into three components:

Myerson's surplus. The consumer surplus generated by Myerson's auction, denoted by $\mathrm{CS}(\mathcal{D})$.

Non-allocation surplus. The loss in consumer surplus due to Myerson's auction not allocating the item, denoted by CS_0. In Lemma 4, we show that

$$\mathrm{CS}_0 = \prod_i \mathop{\mathbf{Pr}}_{v_i \sim \mathcal{D}_i} [v_i < r_i] \cdot \mathop{\mathbf{E}}_{\vec{v} \sim \mathcal{D}} \left[\max_i v_i \;\middle|\; v_i < r_i \quad \forall i \right].$$

Mis-allocation surplus. The loss in surplus because the highest value and highest virtual-value buyers are different. We denote the expected loss due to mis-allocating the item to buyer i as CS_i, and the expected loss due to mis-allocating to a buyer with value x_k as $\widehat{\mathrm{CS}}_k$.

For i.i.d. values, we can write $\textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0$, since the highest virtual-value and highest value buyers coincide. For non-i.i.d. settings, we have two ways of decomposing $\textsc{Opt}$:

$$\textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0 + \sum_{i=1}^{n} \mathrm{CS}_i \qquad \text{and} \qquad \textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0 + \sum_{k=1}^{\mathcal{K}} \widehat{\mathrm{CS}}_k. \tag{2.4}$$

In the remainder, when seeking a θ approximation to $\textsc{Opt}$, we assume $\mathrm{CS}(\mathcal{D})$ is not already a θ-approximation. We first present an $O(\min(\log n, \mathcal{K}))$-approximation for the non-allocation surplus CS_0, and then use it to approximate the remaining terms CS_i and $\widehat{\mathrm{CS}}_k$ in the above expressions.

2.5.1 Approximating the Non-allocation Surplus

The scheme we use to approximate CS_0 is simple: We choose the parameter $t \in \{1, 2, \ldots, n\}$ that maximizes $\mathrm{CS}(\textsc{Rank}_t)$ and run $\textsc{Rank}_t$. This choice of t only depends on the distribution $\mathcal{D}$. Formally, we denote this scheme as $\mathcal{S}_0(\vec{v}, \mathcal{D})$, and present it in Algorithm 3.

Algorithm 3: $\mathcal{S}_0(\vec{v}, \mathcal{D})$

Choose $t = \mathrm{argmax}_{t'=1}^{n} \mathrm{CS}(\textsc{Rank}_{t'})$.

return $\textsc{Rank}_t(\vec{v}, \mathcal{D})$

In the rest of this section, we show the following theorem which is our main technical contribution:

Theorem 6. *The consumer surplus of the signaling scheme* $\mathcal{S}_0(\vec{v}, \mathcal{D})$ *is an* $O\big(\min(\log n, \mathcal{K})\big)$*-approximation to the non-allocation surplus,* CS_0*.*

The proof of the $O(\log n)$-approximation involves analyzing an appropriate randomization of $\textsc{Rank}_t$ over different t (Algorithm 4), while the proof of the $O(\mathcal{K})$-approximation involves showing a stronger version of Markov's inequality for independent Bernoulli random variables (Lemma 6), which may be of independent interest.

We first derive an expression for CS_0. We denote a realization from $\mathcal{D}$ by $\vec{v} = \{v_i\}$. Let p_i, r_i, Y_i, Z_t be as defined in Section 2.4. We have the following lemma:

Lemma 4. *Let* $P = \prod_{i=1}^{n} p_i$. *Then,* $\mathrm{CS}_0 = P \cdot \mathbf{E}[Z_1]$.

Proof. Note that CS_0 is the expected surplus lost due to not allocating the item in Myerson's mechanism. This happens only when all realized values are below their corresponding reserve price. In this case, the value lost is the maximum valuation, since this value contributes to the welfare, and the revenue raised is zero. Therefore, we have:

$$\mathrm{CS}_0 = \left(\prod_{i=1}^{n} p_i\right) \cdot \mathbf{E}\left[\max_{i=1,2,\ldots,n} v_i \;\middle|\; \forall i, v_i < r_i\right]$$

where the expectation is over $\vec{v} \sim \mathcal{D}$. This is equal to $P \cdot \mathbf{E}[Z_1]$. □

$O(\log n)$-Approximation of CS_0

For this purpose, we construct a signaling scheme denoted by $\mathcal{S}_0^2$ such that $CS(\mathcal{S}_0^2) \leqslant \mathrm{CS}(\mathcal{S}_0)$. We will then prove that consumer surplus of $\mathcal{S}_0^2$ is an $O(\log n)$-approximation to CS_0.

This scheme can be constructed by randomizing over different values for parameter t in $\textsc{Rank}_t$. We assign a weight w_j to each rank $j \in \{1, 2, \ldots, n\}$ and pick a rank t with probability proportional to these weights. Note that this choice of the rank t does not depend on $\vec{v}$. Subsequently, we run $\textsc{Rank}_t$. Formally, the signaling scheme is as follows.

Algorithm 4: $\mathcal{S}_0^2(\vec{v}, \mathcal{D})$

$w_j \leftarrow \frac{1}{j+1}$ for $j \in \{1, 2, \ldots, n-2\}$

$w_{n-1} \leftarrow 1$; $w_n \leftarrow 1$

Choose rank $t \in \{1, 2, \ldots, n\}$ where rank j is chosen with probability proportional to w_j.

return $\textsc{Rank}_t(\vec{v}, \mathcal{D})$

Lemma 5. *The consumer surplus of $\mathcal{S}_0^2$ is an $O(\log n)$-approximation to* CS_0.

Proof. We have:

$$\mathrm{CS}(\mathcal{S}_0^2) = \frac{\sum_{t=1}^n w_t \cdot \mathrm{CS}(\textsc{Rank}_t)}{\sum_{t=1}^n w_t}.$$

Recall the definition of P from Lemma 4. We now use Lemma 3 to lower bound the numerator of the above formula:

$$\begin{aligned}
&\sum_{t=1}^{n-1} w_t \mathrm{CS}(\textsc{Rank}_t) + w_n \mathrm{CS}(\textsc{Rank}_n) \\
\geqslant &P\left(\left(\sum_{t=1}^{n-2} \mathbf{E}[Z_t] \cdot \sum_{j=t}^{n-2} \frac{1}{j(j+1)}\right) - \left(\sum_{t=2}^{n-1} \mathbf{E}[Z_t] \cdot \frac{1}{t}\right) + \left(\sum_{t=1}^{n-1} \mathbf{E}[Z_t] \cdot \frac{1}{n-1}\right) - \mathbf{E}[Z_n]\right) + \\
&\mathrm{CS}(\textsc{Rank}_n) \\
= &P \cdot (\mathbf{E}[Z_1] - \mathbf{E}[Z_n]) + \mathrm{CS}(\textsc{Rank}_n).
\end{aligned}$$

It is easy to see that $\mathrm{CS}(\textsc{Rank}_n) \geqslant P \cdot \mathbf{E}[Z_n]$. Using Lemma 4, we have:

$$\sum_{i=1}^n w_i \mathrm{CS}(\textsc{Rank}_i) \geqslant P \cdot \mathbf{E}[Z_1] = \mathrm{CS}_0.$$

Now, we compute the denominator:

$$\sum_{t=1}^n w_t = 1 + \sum_{t=1}^{n-1} \frac{1}{t} \leqslant 1 + \ln n.$$

Therefore, we have:

$$\mathrm{CS}(\mathcal{S}_0^2) = \frac{\sum_{t=1}^n w_t \cdot \mathrm{CS}(\textsc{Rank}_t)}{\sum_{t=1}^n w_t} \geqslant \frac{\mathrm{CS}_0}{1 + \ln n}.$$ □

Since the scheme $\mathcal{S}_0$ chooses the Rank_t with largest value, we now have the following corollary:

Corollary 1. *Consumer surplus of $\mathcal{S}_0$ is an $O(\log n)$-approximation to* CS_0.

$O(\mathcal{K})$-Approximation of CS_0

For proving this part, we will assume $n > 5$. For smaller values of n, the analysis in the previous section already yields a constant-approximation to CS_0. Using Lemma 4, we know $\text{CS}_0 = P \cdot \mathbf{E}[Z_1]$. Setting $x_0 = 0$, we therefore have:

$$\mathbf{E}[Z_1] = \sum_{k=1}^{\mathcal{K}} \mathbf{Pr}[Z_1 \geqslant x_k] \cdot (x_k - x_{k-1}).$$

Therefore, there is a $k^* \in \{1, 2, \ldots, \mathcal{K}\}$ such that $\mathbf{E}[Z_1] \leqslant \mathcal{K} \cdot \mathbf{Pr}[Z_1 \geqslant x_{k^*}] \cdot (x_{k^*} - x_{k^*-1})$. We fix this k^*, and deduce:

$$\text{CS}_0 \leqslant \mathcal{K} \cdot P \cdot \mathbf{Pr}[Z_1 \geqslant x_{k^*}] \cdot (x_{k^*} - x_{k^*-1}).$$

Therefore, if we show that $\text{CS}(\mathcal{S}_0) \geqslant \Omega(1) P \cdot \mathbf{Pr}[Z_1 \geqslant x_k^*] \cdot (x_k^* - x_{k^*-1})$, then we have an $O(\mathcal{K})$-approximation to CS_0.

In order to show the former statement, we need the following probability lemma:

Lemma 6. *Given n independent Bernoulli random variables, $X_i \in \{0, 1\}$ for $i \in \{1, 2, \ldots, n\}$, let $N = \sum_i X_i$. Then there exists a value $j \in \{1, 2, \ldots, n\}$ such that:*

$$\min\left(1, \frac{\mathbf{E}[N]}{j}\right) - \mathbf{Pr}[N \geqslant j+1] \geqslant \frac{\mathbf{Pr}[N \geqslant 1]}{5}.$$

Proof. Let α_q be the probability that $N = q$. At least one of the following holds:

- If $\mathbf{E}[N] < 5$, then for $j = 5$, we have:

$$\begin{aligned}\min\left(1, \frac{\mathbf{E}[N]}{5}\right) - \mathbf{Pr}[N \geqslant 6] &= \frac{1}{5}\mathbf{E}[N] - \mathbf{Pr}[N \geqslant 6] \\ &\geqslant \frac{1}{5}\left(\sum_{q=1}^{n} q\alpha_q\right) - \alpha_6 \geqslant \frac{1}{5} \cdot \left(\sum_{q=1}^{n} \alpha_q\right) = \frac{\mathbf{Pr}[N \geqslant 1]}{5}.\end{aligned}$$

- If $5 \leqslant \mathbf{E}[N] \leqslant \frac{n}{3}$, then we set $j = 2 \cdot \lceil \mathbf{E}[N] \rceil$. It is easy to see that for this setting, we have $\min\left(1, \frac{\mathbf{E}[N]}{j}\right) \geqslant 0.4$. By a standard application of Chernoff bounds we have:

$$\mathbf{Pr}[N \geqslant j+1] < e^{\frac{-5}{3}} < 0.19.$$

  Therefore we have:

$$\min\left(1, \frac{\mathbf{E}[N]}{j}\right) - \mathbf{Pr}[N \geqslant j+1] \geqslant 0.2 \geqslant \frac{\mathbf{Pr}[N \geqslant 1]}{5}.$$

- If $\mathbf{E}[N] \geqslant \frac{n}{3}$, then we set $j = n$, so that

$$\min\left(1, \frac{\mathbf{E}[N]}{n}\right) - \mathbf{Pr}[N \geqslant n+1] \geqslant \frac{1}{3} \geqslant \frac{\mathbf{Pr}[N \geqslant 1]}{5}. \qquad \square$$

We now complete the proof of the $O(\mathcal{K})$-approximation in the lemma below.

Lemma 7. *There exists a value t' such that* $\mathrm{CS}(\textsc{Rank}_{t'}) \geqslant \Omega(1) P \cdot \mathbf{Pr}[Z_1 \geqslant x_{k^*}] \cdot (x_{k^*} - x_{k^*-1})$.

Proof. Recall that Y_i is the distribution of $\mathcal{D}_i$ conditioned on being strictly below the reserve price r_i. We draw one random variable independently from each Y_i. We define G_k to be the random variable corresponds to the number of draws with value at least x_{k^*} among those n draws. Using Lemma 3, we have:

$$\begin{aligned}\mathrm{CS}(\textsc{Rank}_t) &\geqslant P \cdot \mathbf{E}\left[\frac{\sum_{i=1}^{t} Z_i}{t} - Z_{t+1}\right] \\ &\geqslant P \cdot (x_{k^*} - x_{k^*-1}) \cdot \left(\min\left(1, \frac{\mathbf{E}[G_{k^*}]}{t}\right) - \mathbf{Pr}[G_{k^*} \geqslant t+1]\right).\end{aligned}$$

We define a Bernoulli random variable X_i that is 1 when $Y_i \geqslant x_{k^*}$ and zero otherwise. Applying Lemma 6, there exists a value of t' between 1 and n such that the following holds:

$$\left(\min\left(1, \frac{\mathbf{E}[G_k^*]}{t'}\right) - \mathbf{Pr}[G_k^* \geqslant t'+1]\right) = \Omega(1) \cdot \mathbf{Pr}[G_k^* \geqslant 1] = \Omega(1) \cdot \mathbf{Pr}[Z_1 \geqslant x_{k^*}].$$

This when combined with the previous inequality, completes the proof. $\square$

According to the previous lemma, there is a t' such that consumer surplus of $\textsc{Rank}_{t'}$ is an $O(\mathcal{K})$-approximation and $\mathrm{CS}(\mathcal{S}_0) \geqslant \mathrm{CS}(\textsc{Rank}_{t'})$ for any t'. Therefore, we have the following corollary:

Corollary 2. *Consumer surplus of $\mathcal{S}_0$ is an $O(\mathcal{K})$-approximation to CS_0.*

Combining Corollaries 1 and 2 completes the proof of Theorem 6.

Identically distributed buyers. For i.i.d. values, we use the better of two schemes: Using $\mathcal{S}_0(\vec{v}, \mathcal{D})$ or sending no signal. This immediately implies an $O\big(\min(\log n, \mathcal{K})\big)$-approximation to $\textsc{Opt}$, as in this setting we have: $\textsc{Opt} = \mathrm{CS}(\mathcal{D}) + \mathrm{CS}_0$. This completes the first part of Theorem 3.

2.5.2 Approximating the Mis-allocation Surplus

Finally, we prove Theorem 3 when $\mathcal{D}_i$ are not all identical. From Eq (2.4), this requires approximating CS_i and $\widehat{\mathrm{CS}}_k$ for any i and k. Recall CS_i is the surplus lost when a non highest-value bidder i wins the auction. Assuming we break ties in favor of the higher valued buyer, we have:

$$\mathrm{CS}_i = \sum_{\vec{v}: i = \operatorname{argmax}_j(\varphi_{\mathcal{D}_j}(v_j))} \mathbf{Pr}[\mathcal{D} = \vec{v}] \cdot \left(\max_j(v_j) - v_i \right). \tag{2.5}$$

Similarly, $\widehat{\mathrm{CS}}_k$ is the surplus lost when the item is allocated to a buyer with value x_k. Again assuming ties are broken in favor of higher-valued buyers, we have:

$$\widehat{\mathrm{CS}}_k = \sum_{\vec{v}: x_k = \max_j(\varphi_{\mathcal{D}_j}(v_j))} \mathbf{Pr}[\mathcal{D} = \vec{v}] \cdot \left(\max_j(v_j) - x_k \right). \tag{2.6}$$

To approximate these quantities, we run the signaling scheme for approximating CS_0 on a *modified* product distribution. In this new scheme, we fix a cut-off value c, and reveal the identity of all the buyers with realized value strictly greater than

c. Let a denote the largest realized value that is at most c, and T denote the set of buyers with values strictly bigger than a. We first modify the distribution $\mathcal{D}_i$ for $i \in T$ as follows: Recall that $\mathcal{D}_{i|>c}$ denotes the distribution of V_i conditioned on $V_i > c$; let X_i denote the corresponding random variable. We change the distribution to be that of $X_i - a$, that we denote $\mathcal{D}_{j|>c} - a$. We now subtract a from all the valuations $v_i, i \in T$, and run the signaling scheme $\mathcal{S}_0$ in this instance. The details of the scheme can be found in Algorithm 5, where c is the cutoff parameter, and where we denote by $\vec{v}_T$ the $|T|$ dimensional vector made by choosing the indices in T from $\vec{v}$. Note that c could be different from a when there is no buyer whose value coincides with c.

Algorithm 5: $\textsc{Trunc}(\vec{v}, \mathcal{D}, c)$

$a \leftarrow \max_{v_i \leq c}(v_i)$; and $T \leftarrow \{j : v_j > a\}$

$\mathcal{D}_{T,a} \leftarrow \prod_{j \in T_i}(\mathcal{D}_{j|>c} - a)$ // Modified distributions for $j \in T$.

$\vec{v'} \leftarrow \vec{v}_T - \mathbf{a}$ // $\mathbf{a}$ is a vector with all elements equal to a.

$s \leftarrow \mathcal{S}_0(\vec{v'}, \mathcal{D}_{T,a})$ // s is the signal returned by $\mathcal{S}_0$.

return (a, T_i, s) as final signal

Approximating Opt **for small** n. We choose $i^* \in \{1, 2, \ldots, n\}$ that maximizes CS_i. By Eq. (2.4), to get an $O(n \log n)$-approximation to Opt, it suffices to demonstrate an $O(\log n)$-approximation to CS_{i^*}. The scheme $\mathcal{S}_1$ chooses $i^* \in \{1, 2, \ldots n\}$ that maximizes CS_i and returns $\textsc{Trunc}(\vec{v}, \mathcal{D}, v_{i^*})$.

Theorem 7. *The consumer surplus of $\mathcal{S}_1$ is an $O(\log n)$-approximation to CS_{i^*}.*

Proof. Recall that the final scheme $\mathcal{S}_1$ chooses $i^* \in \{1, 2, \ldots n\}$ that maximizes CS_i as defined in Eq (2.5) and returns $\textsc{Trunc}(\vec{v}, \mathcal{D}, v_{i^*})$ as defined in Algorithm 5. Applying Corollary 1 to the modified distribution $\mathcal{D}_{T,a}$ defined in Algorithm 5, the consumer surplus of $\mathcal{S}_1$ is at least:

$$\frac{1}{\ln n + 1} \sum_{T,a} \mathbf{Pr}(T, a) \text{CS}_0(\mathcal{D}_{T,a})$$

where $\mathbf{Pr}(T, a)$ is the probability of the event that $v_{i*} = a$ *and* the set of all the buyers with value strictly larger than a is T.

In order to prove the lemma, we need to show the following:

$$\text{CS}_{i*} \leqslant \sum_{T,a} \mathbf{Pr}(T, a)\text{CS}_0(\mathcal{D}_{T,a}).$$

From Eq (2.5), we have:

$$\text{CS}_{i*} = \sum_{\vec{v}: i* = \text{argmax}_j(\varphi_{\mathcal{D}_j}(v_j))} \mathbf{Pr}[\mathcal{D} = \vec{v}] \cdot \left(\max_j(v_j) - v_{i*}\right).$$

Let $\mathcal{D}'_j = \mathcal{D}_{j|>a} - a$ and let its reserve price be r'_j. Then,

$$\mathbf{Pr}(T, a)\text{CS}_0(\mathcal{D}_{T,a}) = \sum_{\vec{v'}: v'_j < r'_j \forall j \in T} \mathbf{Pr}(T, a) \cdot \mathbf{Pr}[\mathcal{D}_{T,a} = \vec{v'}] \cdot \max_{j \in T}(v'_j).$$

Note that for every vector $\vec{v} \in \mathcal{D}$ there is a corresponding set T, value a, and vector $\vec{v'} \in D_{T,a}$ as defined in Algorithm 5. We show that if a positive value is added to CS_{i*} in the above formula, at least the same amount will be added to the corresponding $\mathbf{Pr}(T, a)\text{CS}_0(\mathcal{D}_{T,a})$ by the corresponding vector v'. First, since $\mathcal{D}_{T,a}$ is the product of conditional distributions of values of buyers in T being larger than a, we have:

$$\mathbf{Pr}[\mathcal{D} = \vec{v}] = \mathbf{Pr}(T, a) \cdot \mathbf{Pr}[\mathcal{D}_{T,a} = \vec{v'}].$$

Now, there is a positive contribution to CS_{i*} in the event where $v_j = \max_i v_i$ and $\varphi_{\mathcal{D}_j}(v_j) < \varphi_{\mathcal{D}_{i*}}(v_{i*}) \leqslant v_{i*} = a$. This assumes ties are broken by allocating to the buyer with higher valuation. The contribution to CS_{i*} is $v_j - v_{i*}$.

Consider the distribution $\mathcal{D}_j$. By Lemma 2, $\varphi_{\mathcal{D}'_j}(v_j - a) = \varphi_{\mathcal{D}_j}(v_j)$. Since $\varphi_{\mathcal{D}_j}(v_j) < a = v_{i*}$, we have $\varphi_{\mathcal{D}'_j}(v'_j) < 0$ so that $v'_j < r'_j$. Therefore, the contribution to $\text{CS}_0(\mathcal{D}_{T,a})$ is precisely $v'_j = v_j - v_{i*}$. This completes the proof.

□

Since there are n possible choices of $i*$, this directly implies:

Corollary 3. *Any n-dimensional product distribution $\mathcal{D}$ admits $O(n \log n)$-approximate signaling.*

Approximating $\textsc{Opt}$ **for small** $\mathcal{K}$**.** Similar to the previous section, consider $k^* \in \{1, 2, \ldots, \mathcal{K}\}$ that maximizes $\widehat{\mathrm{CS}}_k$. The scheme, denoted by $\mathcal{S}_2$ executes $\textsc{Trunc}(\vec{v}, \mathcal{D}, x_{k^*})$. Note that there may be no bidder with valuation x_{k^*}, which motivates the way Algorithm 5 is presented.

Theorem 8. *The consumer surplus of* $\mathcal{S}_2$ *is an* $O(\mathcal{K})$*-approximation to* $\widehat{\mathrm{CS}}_{k^*}$*.*

Proof. The proof is similar to the proof of Theorem 7 and we omit details. First, by applying Corollary 2 to the modified distribution $\mathcal{D}_{T,x_{k^*}}$, the consumer surplus of $\mathcal{S}_2$ defined in Section 2.5.2 is at least:

$$\Omega\left(\frac{1}{\mathcal{K}}\right)\sum_T \mathbf{Pr}(T)\mathrm{CS}_0(\mathcal{D}_{T,x_{k^*}})$$

where $\mathbf{Pr}(T)$ is the probability that the set of all the buyers with value strictly larger than x_{k^*} is T. Furthermore using Eq (2.6) and by replacing v_{i^*} by x_{k^*} in the proof in Theorem 7, we have

$$\widehat{\mathrm{CS}}_{k^*} \leqslant \sum_T \mathbf{Pr}(T)\mathrm{CS}_0(\mathcal{D}_{T,x_{k^*}}).$$

Combining the above two bounds proves the theorem.

□

Since there are $\mathcal{K}$ possible values of k^*, the previous theorem directly implies:

Corollary 4. *Any* n*-dimensional product distribution* $\mathcal{D}$ *admits* $O(\mathcal{K}^2)$*-approximate signaling.*

By combining Corollaries 3 and 4, the proof of Theorem 3 is completed for the non-i.i.d. case.

2.6 Conclusion

The immediate and challenging open questions involve tightening our bounds. First, we do not have any lower bound for the i.i.d. case, since the bounds in Theorems 1 and 2 both require non-i.i.d. distributions. Indeed, for two-valued i.i.d. distributions, the upper bound on approximation ratio is also one. Improving our lower bounds will require a better characterization of the consumer surplus generated by signals that correlate buyer valuations, beyond the Crémer-McLean characterization [63]. Next, the space of signals we have considered for the upper bounds builds on the $n = 1$ case [33], and it is conceivable that signals that modify the valuation distribution of multiple buyers at once may lead to improved bounds. Such signals are challenging to analyze for reasons we describe in Section 2.1.2. Finally, it is quite likely there is a separation between existence results and computational results, i.e., there could be an existence result showing a better approximation via an exponential complexity signaling scheme. Such separations are known for the Bayesian persuasion problem [73], and it would be interesting to derive such results for our multi-agent setting.

3

Prophet Inequalities with Uncertain Supply

The results of this chapter are published in [13], which is a joint work with Siddhartha Banerjee, Sreenivas Gollapudi, Kamesh Munagala, and Kangning Wang.

3.1 Introduction

Online posted pricing problems are one of the canonical examples in online decision-making and optimal control. The basic model comprises of a fixed supply of non-replenishable items; buyers (demand) arrive in an online fashion over a fixed time interval, and the platform sets prices to maximize some objective such as social surplus (welfare) or revenue. Another variant of this setting is found in internet advertising, where the number of advertisements (supply) is assumed to be fixed (for example, based on contracts between the publisher and advertisers), while keywords/impressions (demand) arrive online, and are matched to ads via some policy. The demand is typically assumed to obey some underlying random process, which allows the problem to be cast as a Markov Decision Process (MDP); however, in many settings, such a formulation suffers from a "curse of dimensionality", making it infeasible to solve optimally.

An important idea for circumventing the computational intractability of optimal pricing is that of *prophet inequalities* — heuristics with performance guarantees with respect to the optimal policy in hindsight (*i.e.*, the performance of a *prophet* with full information of future arrivals). The simplest prophet inequality has its origins in the statistics community [115] — given a single item and T arriving buyers with values drawn from known distributions, there is a pricing scheme using only a single price that extracts at least half the social surplus earned by the prophet (moreover, this is tight). More recently, there has been a long line of work generalizing this setting to incorporate multiple (possibly non-identical) items, as well as combinatorial buyer valuations [95, 57, 113, 80, 69, 133, 61, 3, 74].

The aim of this chapter is to develop a theory of prophet inequalities for settings with *uncertainty in future supply*. This is a natural extension of the basic posted-price setting, and indeed special cases of our framework have been considered before [135, 95] (in the context of optimal secretary problems with a random "freeze" on hiring). What makes these problems of greater relevance today is the rise of online 'sharing economy' marketplaces, such as those for transportation (Lyft, Uber), labor (Taskrabbit, Upwork), lodging (Airbnb), medical services (PlushCare), *etc.* The novelty in such marketplaces arises because of their *two-sided nature*: in addition to buyers who arrive online, the supply is now controlled by "sellers" who can arrive and depart in an online fashion. For example, in the case of ridesharing/lodging platforms, the units of supply (empty vehicles/vacant listings) arrive over time, and have some *patience* interval after which they abandon the system (get matched to rides on other platforms/remove their listings). Supply uncertainty also arises in other settings, for instance, if items are *perishable* and last for *a priori* random amounts of time. Our work aims to understand the design of pricing policies for such settings, and characterize how the resulting prophet inequalities depend on the characteristics of the supply uncertainty.

3.1.1 Model

We introduce "supply uncertainty" into the basic prophet inequality setting as follows: There are m items present initially, but these do not last till the end of the buyer arrivals, but instead, depart after an *a priori* unknown amount of time. Formally, we assume each item i samples a *horizon* from a distribution H_i, at which time it departs. We assume the horizon lengths for items are mutually independent, and also independent of the valuation distribution of the buyers. Note though that the items can have different horizon distributions. We denote the maximum possible horizon length for any item as n.

On the demand side, we assume there is an infinite stream of unit-demand buyers arriving online, where the valuation of the h-th arriving buyer is a random variable X_h drawn *i.i.d.* from a distribution V. From the perspective of a buyer, all items are interchangeable, and hence being matched to any item that has not yet departed yields value X_h. Note that assuming an infinite stream of buyers is without loss of generality, because we can encode any upper bound on the number of buyers in the horizon distributions.

The algorithm designer knows the horizon distribution H_i for each item, and the buyer value distribution V, but not the realized horizons for each item (until the item actually departs), or the value for any buyer. The goal is to design an online pricing scheme that competes with a prophet that knows the realized horizons of each item and the valuation sequence of buyers, and extracts full social surplus (or welfare).

The main outcome of the standard prophet inequality is that there are constant-competitive algorithms for maximizing welfare, even when buyers are heterogeneous and arrive in arbitrary order. This however turns out to be impossible in the presence of item horizons without additional assumptions. First, even with *i.i.d.* horizons, achieving a constant factor turns out to be impossible for general horizon distributions (*cf.* Theorem 12); thus to make progress, we need more structure on the horizons. One natural assumption is that each item is more and more likely to depart as time goes on, which can be formalized as follows.

Definition 1. A horizon distribution H satisfies the *monotone-hazard-rate (MHR) condition* if:

$$\Pr_{h \sim H}[h \geqslant h^* + 2 \mid h \geqslant h^* + 1] \leqslant \Pr_{h \sim H}[h \geqslant h^* + 1 \mid h \geqslant h^*], \quad \forall\, h^* \geqslant 1.$$

Several distributions satisfy the MHR condition, including uniform, geometric, deterministic, and Poisson; note also that truncating an MHR distribution preserves the condition.

Finally, even with MHR horizons, buyer heterogeneity is a barrier for obtaining a constant-competitive algorithm, as demonstrated by the following example, with *deterministic valuations and known order of arrivals.*

Example 2. *Given $m = 1$ item with horizon following a geometric distribution with parameter* 0.5*, consider a sequence of n buyers with $v_h = 2^h$ for $h = 1, 2, \ldots, n$. The expected value of the prophet is $\Theta(n)$ while any algorithm can only achieve a constant value in expectation.*

3.1.2 Our Results

The above discussion motivates us to study settings with *i.i.d.* buyers, and items with MHR horizons. Our main result is that these two assumptions are *sufficient to obtain a constant-competitive approximation to the prophet welfare.* In particular, our main technical result is the following theorem, which we prove in Section 3.2.

Theorem 9. *There is a constant-competitive online policy for social surplus for any $m \geqslant 1$ items with independent and* possibly non-identical *MHR horizon distributions, and unit-demand buyers arriving with* i.i.d. *valuations.*

Though the complete algorithm is somewhat involved, at a high level, it is based on a simple underlying idea: to be constant-competitive against the prophet, we need to choose prices so as to balance the rate of matches and departures. Achieving this in the general case is non-trivial, and requires some new technical ideas. However, for the special case of a single item, balancing can be achieved via a simple fixed pricing scheme. In Section 3.3, we use this to obtain the following tight result for the $m = 1$ setting (this also serves as a primitive for our overall algorithm):

Theorem 10. *There is a fixed pricing scheme for a single item with an MHR horizon distribution with mean μ that has competitive ratio $2 - 1/\mu$. Further, this bound is tight for the geometric horizon distribution with mean μ.*

Intuitively, the factor of two in the above theorem corresponds to the prophet considering matching and departures as the same, which an algorithm cannot do. The surprising aspect is that this simple policy is worst-case optimal within the class of instances with MHR horizons — this is in contrast to deterministic horizons, where fixed pricing is known to be suboptimal for the special case of one item with known (deterministic) horizon and *i.i.d.* buyers [100, 74].

We complement our positive results by showing several lower bounds that establish their tightness. As mentioned above, in Section 3.3, we show a (tight) lower bound of $2 - 1/\mu$ for $m = 1$ items with MHR horizons. Our main lower bounds in Section 3.4 generalizes this to $m \geqslant 1$ items.

Theorem 11. *For the multi-item setting with* i.i.d. *geometric horizons:*

- *For any number of items, there is a lower bound of* 1.57 *on the competitive ratio of any dynamic pricing scheme; in the limit when the number of items goes to infinity, this improves to* 2.
- *No fixed pricing scheme can be $o(\log \log m)$-competitive where m is the number of items.*

The above theorem implies that the MHR horizon setting, even with *i.i.d.* horizons, is significantly different from the setting with multiple items and a single deterministic horizon (where fixed pricing extracts $\left(1 - O\left(\frac{1}{\sqrt{m}}\right)\right)$-fraction of surplus [10]). Put differently, the lower bound emphasizes that even with *i.i.d.* horizons, to obtain a constant-competitive algorithm, it is *not* sufficient to replace the horizon distributions by their expectations and use standard prophet inequalities — the stochastic nature of the horizons allows for significant deviations in the order of departures of the items, and a policy that knows this ordering can potentially extract much more welfare. Given this, it is quite surprising that a simple dynamic pricing scheme achieves a constant approximation.

Finally, we consider the general case where there is no restriction on the horizon distribution. In this setting, the presence of supply uncertainty severely limits the performance of any non-anticipatory dynamic pricing scheme in comparison to the omniscient prophet. In particular, we show that for any number of items and *i.i.d.* buyer valuations, the ratio between the welfare of any algorithm and the prophet grows with the horizon, *even if the algorithm knows the realized valuations.*

Theorem 12. *For any $m \geqslant 1$ items, there exists a family of instances such that the prophet has welfare $\Omega\left(\frac{\log n}{\log\log n}\right)$-factor larger than any online policy, even if the policy knows all the realized values, but not the realized horizons. Here, $n = \max_i\{\text{supp}(H_i)\}$.*

This generalizes similar lower bounds for settings where the horizon is unknown [101, 95]. The proof of this result is provided in Section 3.5.

3.1.3 Technical Highlights

At a high level, we achieve our results via a conceptually simple and natural class of *balancing policies* that generalizes policies for the deterministic-horizon case:

Balancing Policy. Balance the rate at which buyers are accepted to the rate at which items depart the system because their horizon is reached.

Converting this high-level description of balancing into a concrete policy requires new technical ideas. We first note the technical challenges we encounter. In the setting with deterministic identical horizons [115, 80], we can achieve constant-competitive algorithms (or even better) via a *global* expected value relaxation that yields a fixed pricing scheme. Indeed, such an argument can safely assume buyers are non-identical with adversarial arrival order. However, the setting with stochastic horizons is very different. First, as Example 2 shows, even for $m = 1$ item with geometric horizon, there is an $\Omega(n)$ lower bound when buyer valuations are not identically distributed. Secondly, for $m > 1$ items, we need dynamic pricing even in the simplest settings — when horizons are *i.i.d.* geometric (see Theorem 11), or when they are deterministic. This precludes the use of a global one-shot analysis.

At this point, we could try using techniques from stochastic optimization, particularly stochastic matchings [58, 31] and multi-armed bandits [89, 91]. Here, the idea is to come up with a *weakly coupled* relaxation, say one policy per item, and devise a feasible policy by combining these. However, these algorithms crucially require the state of the system to only change via policy actions, and our problem more is similar to a *restless bandit* problem [90] where item departures cause the state of the system can change *regardless* of policy actions taken. Indeed, the actual departure process itself may significantly deviate from its expected values, making it non-trivial to use a global relaxation.

Simulating Departures. This brings up our technical highlight: *Instead of encoding the departure process in a fine-grained way into a relaxation, we simulate its behavior in our final policy.* In more detail, we first write a weak relaxation of the prophet's welfare separately in a sequence of stages with geometrically decreasing number of items. This only uses the expected number of items that survive in the stage, and not the identity of these items. The advantage of such a weak relaxation is that it yields a solution with nice structure: this policy non-adaptively sets a fixed price in each stage to balance the departure rate with the rate of matches. However, it is non-trivial to construct a feasible policy from this relaxation, since the relaxation decouples the allocations of the prophet across different stages, while any feasible algorithm's allocations are clearly coupled. Indeed, the optimal feasible policy is the solution to a dynamic program with state space exponential in m, and the prophet is further advantaged by knowing which items depart earlier in the future.

Surprisingly, we show that our simple relaxation is still enough to achieve a constant-competitive algorithm. We do so by simulating the departure process, that is, by choosing items for matching with the same probability that they would have departed at a future point in time. This couples the stochastic process that dictates the number of items available in the policy with that in the prophet's upper bound, albeit with a constant-factor speedup in time. This yields a *non-adaptive* policy that makes its pricing decisions for the entire horizon, as well as the (randomized) sequence in which to sell the items, in advance. We believe such a policy construction that simulates the evolution of state of the system may find further applications in the analysis of restless MDPs.

Lower Bounds from Time-Reversal. Our lower bounds are all based on demonstrating particular bad settings as in Example 2. From a technical perspective, the most interesting construction is that in Theorem 11 — here, we first consider a canonical, asymptotic regime where the horizon distribution is geometric with mean approaching infinity, and show that we can closely approximate the behavior of the prophet and the algorithm via an appropriate Markov chain. We then define and analyze a novel time-reversed Markov chain encoding the prophet's behavior, that captures matching a departing item to the optimal buyer that arrived previously.

3.1.4 Related Work

The first prophet inequalities are due to Krengel and Sucheston [115, 116]. It was subsequently shown [134], there is a 2-competitive fixed pricing scheme that is oblivious to the order in which the buyers arrive, and this ratio is tight in the worst case over the arrival order. Motivated by applications to online auctions, since then there have been several extensions to multiple items [109, 95, 10], matching setting [11, 143], matroid constraints [113] and general combinatorial valuation functions [80, 133].

Our work is a generalization of the single-item setting where buyer valuations are *i.i.d.* and the horizon is known, to the case where the horizon is stochastic and there are multiple items. The setting with known horizons was first considered in Hill and Kertz [100]. In this case, the optimal pricing scheme can be computed by a dynamic program, and a sequence of results [110, 3, 61] show a tight competitive ratio of 1.342 for this dynamic program against the prophet. In contrast, we show that when the horizon is MHR, a simple fixed pricing scheme has optimal competitive ratio of 2.

A generalization of the *i.i.d.* setting is the recently-introduced *prophet secretary* problem where the buyers are not identical, but the order of arrival is a random permutation. In this case, fixed pricing is a tight $\frac{e}{e-1}$-approximation [77, 74]; and a dynamic pricing scheme can beat this bound [20, 62] by a slight amount. Though our results extend to this setting, it is not the focus of our work since the *i.i.d.*-valuations case is sufficient to bring out our conceptual message.

The random horizon setting has been extensively studied in the context of the classic secretary problem. When the horizon is unknown (that is, no distributional information at all), no constant-competitive algorithm is possible [101]. In the context of prophet inequalities, the unknown-horizon setting was considered by Hajiaghayi *et al.* [95], who show again that no constant-competitive algorithm is possible. We use a similar example to extend this lower bound to the case where the horizon is stochastic from a known distribution.

3.2 Prophet Inequality for Heterogeneous Items with MHR Horizons

In this section, we present the proof of Theorem 9. We first give an overview of our algorithm. At a high level, this scheme attempts to balance the rate that items are assigned to buyers and the rate that items naturally depart. In Section 3.2.1, we first introduce a way to divide the entire time horizon into disjoint stages in a way such that during the k-th stage, $\frac{m}{2^k}$ items depart in expectation. We then bound the prophet's welfare separately for each stage (Section 3.2.2) — we do so via a relaxation that ignores the identity of the items, and only captures the constraint that the expected number of matches in a stage is at most the expected number of items present at the beginning of that stage.

The key technical hurdle at this point is that when we make a matching, we do so without knowing exactly when items depart in the future. This changes the distribution of the items available in subsequent stages. To get around this, in each stage, we first *simulate* the future departure of items, and use this to select items available for matching in the current stage. In more detail, in Section 3.2.3, we split the stages alternately into even and odd stages, and develop an algorithm whose welfare approximates the welfare of the relaxed prophet from the odd stages (and by symmetry, another algorithm that approximates the welfare from the even stages).

For approximating the welfare from the odd stages, the algorithm re-divides time into a new set of stages corresponding to the odd stages under the old division (See Figure 3.1). We then use each new stage to approximate the welfare generated in the corresponding odd stage in the old division; to do so, we sample candidate items for matching in the current stage with the probability they would leave in the subsequent even stage under the old division. Consequently, for every item, the probability of departure during an even stage under the old division is the same as of being selected for matching in the current stage. We show that this process couples the behavior of the algorithm and the benchmark, assuming the departure processes are MHR. Using concentration bounds, we show that this approach yields a constant approximation.

In addition to the above process, our algorithm needs to separately handle any stage of length 1 (*i.e.*, any single time period where the expected number of available items reduces by at least half), as well as a final stage where the expected number of available items is constant. We show that the welfare in the length 1 phases is approximated by a blind matching algorithm which matches all incoming buyers (Section 3.2.4), while the welfare of the final period is approximated by an algorithm that randomly selects only one item for matching at the beginning, and discards the rest (Section 3.2.5). For the latter setting (*i.e.*, for a single item setting), we present a tight 2-competitive fixed pricing scheme for the $m = 1$ setting in Section 3.3. Finally, the overall algorithm is based on randomly choosing one of the four candidate algorithms (*i.e.*, for approximating the prophet welfare in odd stages, even stages, short stages, and the final stage), with an appropriately chosen distribution.

3.2.1 Splitting Time into Stages

As a first step, we divide the time horizon into $s+1$ stages. The k-th stage corresponds to an interval $[\ell_k, r_k)$. For $k = 1, 2, \dots, s$, we define r_k by

$$r_k := \min\left\{t+1 : \mathbf{E}[\text{number of remaining items after time } t] \leqslant \frac{m}{2^k}\right\}.$$

Also $\ell_{k+1} := r_k$ for $k = 1, 2, \dots, s$; $\ell_1 = 1$ and $r_{s+1} = \infty$.

We set s to be the smallest non-negative integer so that $\frac{m}{2^s} \leqslant 10$, *i.e.*, $s := \max\left(0, \left\lceil \log_2 \frac{m}{10} \right\rceil\right)$. Within the first s stages, we separate stages of length $r_k - \ell_k = 1$ from the rest. We term the stages of length at least 2 as LONG stages, and those of length 1 as SHORT stages. We term the stage $s+1$ as the FINAL stage. Note that based on our choice of s, the expected number of items which remain in the final stage is at most 10, and unless $s = 0$, at least 5 items in expectation survive at one time step earlier into the final stage.

3.2.2 Upper Bound on Prophet's Welfare

In this section, we develop a tractable upper bound for the prophet. Let PRO denote the optimal welfare obtainable by the prophet. We term the total welfare of PRO in the LONG stages as PROLONG, the total welfare in the SHORT stages as PROSHORT, and the welfare in the FINAL stage as PROFINAL. Clearly, we have:

Lemma 8. PRO = PROLONG + PROSHORT + PROFINAL.

We bound PROLONG and PROSHORT separately for each stage. Let PRO_k denote the welfare from stage k, so that $\text{PROLONG} + \text{PROSHORT} = \sum_{k=1}^{s} \text{PRO}_k$.

Lemma 9. *For $1 \leqslant k \leqslant s$, we have:*

$$\text{PRO}_k \leqslant \min\left(r_k - \ell_k, \frac{m}{2^{k-1}}\right) \cdot \mathop{\mathbf{E}}_{v \sim V}[v \mid v \geqslant p_k],$$

where p_k satisfies $\mathbf{Pr}_{v \sim V}[v \geqslant p_k] = \min\left(1, \frac{m/2^{k-1}}{r_k - \ell_k}\right)$.[1]

Proof. Fix a stage k. Let W_i be the expected welfare that the prophet gets from buyer i, and let y_i be the probability that buyer i is matched by the prophet ($\ell_k \leqslant i < r_k$).

[1] The existence of such p is without loss of generality: Let $t = \min\left(1, \frac{m/2^{k-1}}{r_k - \ell_k}\right)$. When there exists some p^* such that $\mathbf{Pr}[v \geqslant p^*] > t$ and $\mathbf{Pr}[v > p^*] < t$, we could accept all values greater than p^* and accept p^* with probability $\frac{t - \mathbf{Pr}[v > p^*]}{\mathbf{Pr}[v = p^*]}$.

Notice that in expectation, at most $\frac{m}{2^{k-1}}$ items have horizons of at least ℓ_k by the definition of stages. Therefore, $\sum_{i=\ell_k}^{r_k} y_i \leqslant \frac{m}{2^{k-1}}$.

Let F_V be the CDF of the distribution V. We have $W_i \leqslant y_i \cdot \mathbf{E}_{v\sim V}\left[v \mid v \geqslant F_V^{-1}(1-y_i)\right]$, since when buyer i is matched with probability y_i, the prophet cannot do better than getting the top y_i-percentile of the distribution V from the buyer. With these constraints, we write a relaxation for the welfare of the prophet during stage k:

$$
\begin{aligned}
\max \quad & \sum_{i=\ell_k}^{r_k-1} W_i \\
\text{s.t.} \quad & W_i \leqslant y_i \cdot \mathop{\mathbf{E}}_{v\sim V}\left[v \mid v \geqslant F_V^{-1}(1-y_i)\right], \quad \forall i=\ell_k, \ell_k+1, \ldots, r_k-1, \\
& \sum_{i=\ell_k}^{r_k-1} y_i \leqslant \frac{m}{2^{k-1}}, \\
& y_i \in [0,1], \qquad \forall i=\ell_k, \ell_k+1, \ldots, r_k-1.
\end{aligned}
$$

Clearly y_i's should be equal in the optimal solution. Therefore,

$$
\sum_{i=\ell_k}^{r_k-1} W_i \leqslant (r_k-\ell_k)\cdot \min\left(1, \frac{m/2^{k-1}}{r_k-\ell_k}\right)\cdot \mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant p_k],
$$

where $\mathbf{Pr}_{v\sim V}[v \geqslant p_k] = \min\left(1, \frac{m/2^{k-1}}{r_k-\ell_k}\right)$. Summing over the s stages finishes the proof. □

Notice that in our upper bound for $\sum_{k=1}^{s} \textsc{Pro}_k$, if an item departs during stage k, we allow it to be matched once in stage 1, once in stage 2, ..., and once in stage k. However, since the expected number of departures in each stage exponentially decreases, only a constant factor is lost comparing with the finer relaxation where we enforce the constraint that each item is only matched once across the stages. Our coarser relaxation enables a cleaner benchmark to work on.

We next bound $\textsc{ProFinal}$. Let $\textsc{ProSingle}_i$ be the optimal welfare of the prophet (from all stages) if item i is the only item available in the system, *i.e.*, the single-item setting. We consider this setting in detail in Section 3.3.

Lemma 10. $\textsc{ProFinal} \leqslant \sum_{i=1}^{m} \mathbf{Pr}_{h_i\sim H_i}[h_i \geqslant \ell_{s+1}] \cdot \textsc{ProSingle}_i$.

Proof. Let W_i be the welfare that the prophet can get from item i during the final stage. We have

$$\begin{aligned} W_i \leqslant & \mathop{\mathbf{Pr}}_{h_i\sim H_i}[h_i \text{ reaches the final stage}] \cdot \\ & \mathop{\mathbf{E}}_{h_i\sim H_i}[\text{welfare from item } i \text{ in the final stage} \mid h_i \text{ reaches the final stage}] \\ & \leqslant \mathop{\mathbf{Pr}}_{h_i\sim H_i}[h_i \text{ reaches the final stage}] \cdot \mathbf{E}[\text{welfare from item } i] \\ & = \mathop{\mathbf{Pr}}_{h_i\sim H_i}[h_i \geqslant \ell_{s+1}] \cdot \textsc{ProSingle}_i, \end{aligned}$$

where the second inequality comes from the MHR condition of H_i: $\mathbf{Pr}_{h_i\sim H_i}[h_i \geqslant \ell_{s+1} + k \mid h_i \geqslant \ell_{s+1} - 1 + k] \leqslant \mathbf{Pr}_{h_i\sim H_i}[h_i \geqslant 1 + k \mid h_i \geqslant k]$ — item i would depart faster if it started at time ℓ_{s+1}.

Summing up the items, we have:

$$\textsc{ProFinal} \leqslant \sum_{i=1}^{m} W_i \leqslant \sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i\sim H_i}[h_i \geqslant \ell_{s+1}] \cdot \textsc{ProSingle}_i. \qquad \square$$

Lemmas 8, 9 and 10 together give an upper bound for our benchmark as:

$$\begin{aligned} \textsc{Pro} \leqslant & \left[\sum_{k\leqslant s,\, r_k-\ell_k>1} \textsc{Pro}_k\right] + \left[\sum_{k\leqslant s,\, r_k-\ell_k=1} \textsc{Pro}_k\right] + \\ & \left[\sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i\sim H_i}[h_i \geqslant \ell_{s+1}] \cdot \textsc{ProSingle}_i\right] \end{aligned}$$

where the three term correspond to an upper bound on the prophet's welfare in the Long, Short and Final stages respectively (*i.e.*, ProLong, ProShort, and ProFinal). In the next three sections, we describe three separate algorithms, each one of which, if run independently, provides an approximation to one of the terms. Our overall algorithm is then based on randomly choosing between the three algorithms with appropriately chosen distribution.

3.2.3 Approximating Prophet's Welfare in Long Stages

We first approximate upper bound given in Lemma 9. Within this, we approximate ProLong and ProShort separately. We first focus on ProLong, since this is technically the most interesting, and postpone approximating ProShort to Section 3.2.4.

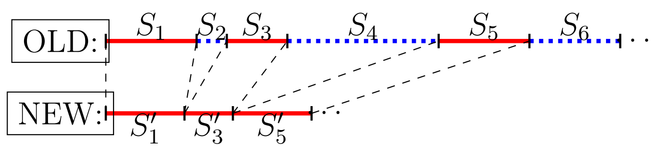

Figure 3.1: Redivision of the Time Horizon

We approximate ProLong by Algorithm 6. We divide all the s stages into alternate odd and even stages. We focus on illustrating the approximation for odd stages, and that for even stages is identical. We then re-divide time into stages corresponding to the original odd stages, as illustrated in Figure 3.1, where S_k stands for the old stage k and S'_k stands for the new stage k. At each odd stage, we sample items according to their departure rates during the next (fictitious) even stage. During the new process when items become unavailable by being sampled, each item is as least as likely to survive a stage as before, since the sampling is only as frequent as the natural departures during the original even stages.

Note that we set each S'_k to be 1 time step shorter than the corresponding S_k and make each fictitious even stage 1 time step longer (unless the length of S_k is 0). We do this to ensure enough items will be sampled: Because of integrality constraints, an even stage may be too short (*e.g.*, of length 0) and if so, little (or nothing if the stage has length 0) can be sampled there. This is also the reason why Short stages are separately considered.

Note that Algorithm 6 can be easily modified to work with even stages instead of

Algorithm 6: DEPARTURESIMULATION: Odd Stages Version

$A \leftarrow \{1, 2, \ldots, m\}$ `// A = Set of available items`

for *each odd stage* $k = 1, 3, \ldots,$ *till stage* s **do**

- $C_k \leftarrow \varnothing$ `// C_k = Set of items considered in this stage`
- For each $i \in A$, with probability $\mathbf{Pr}_{h_i \sim H_i}[h_i < r_{k+1} \mid h_i \geqslant \ell_{k+1} - 1]$, place in C_k
- $A \leftarrow A \backslash C_k$
- **if** $r_k - \ell_k \geqslant 2$ **then**
  - $p_k \leftarrow F_V^{-1}\left(\max\left(0, 1 - \frac{m/2^{k-1}}{r_k - \ell_k}\right)\right)$
  - For each of the next $r_k - \ell_k - 1$ arriving buyers, if this buyer has valuation $\geqslant p_k$, match to any item in C_k and remove this item from C_k
  - If any item departs, remove it from A and C_k

odd stages, and will yield the corresponding version of the theorem below with "odd" replaced by "even". In order to show Theorem 9, we will use either the odd stages or even stages algorithm depending on which yields larger expected welfare. Note that it is entirely possible that one of these stages yields very low welfare compared to the other.

Theorem 13. *Algorithm 6 is a* 15.1*-approximation to the sum of* $\textsc{Pro}_k$ *over odd stages* $k \leqslant s$ *with* $r_k - \ell_k \geqslant 2$.

Proof. We use y^+ to denote $\max(0, y)$. For any odd k with $r_k - \ell_k \geqslant 2$, let the random variable M_k be the number of items in the set C_k that has horizon of at least $\sum_{k'=1}^{(k+1)/2}(r_{2k'-1} - \ell_{2k'-1} - 1)^+$, *i.e.*, the end of (new) stage k. We denote $\sum_{k'=1}^{j}(r_{2k'-1} - \ell_{2k'-1} - 1)^+$ by $\mathcal{S}_j$ in the rest of the proof.

M_k is the sum of m independent Bernoulli random variables, where the i-th one denotes whether item i is in C_k and has horizon of at least $\mathcal{S}_{(k+1)/2}$. We have

$$\begin{aligned}
\mathbf{E}[M_k] &= \sum_{i=1}^{m} \mathbf{Pr}\left[\text{item } i \text{ is in } C_k \text{ and has horizon of at least } \mathcal{S}_{(k+1)/2}\right] \\
&= \sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i \sim H_i}\left[h_i \geqslant \mathcal{S}_{(k+1)/2}\right] \cdot \\
&\qquad \left(\prod_{j=1}^{(k-1)/2} \left(1 - \mathop{\mathbf{Pr}}_{h_i \sim H_i}[h_i < r_{2j} \mid h_i \geqslant \ell_{2j} - 1]\right)\right) \cdot \\
&\qquad \mathop{\mathbf{Pr}}_{h_i \sim H_i}[h_i < r_{k+1} \mid h_i \geqslant \ell_{k+1} - 1],
\end{aligned}$$

where we calculate the probability that item i has horizon of at least $\sum_{k'=1}^{(k+1)/2}(r_{2k'-1} - \ell_{2k'-1} - 1)^+$, was never selected into C_{2j-1}'s during previous stages $2j - 1 < k$, and was selected into C_k. Further simplifying it, we have

$$\mathbf{E}[M_k] = \sum_{i=1}^{m} \left(\prod_{j=1}^{(k+1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant \mathcal{S}_j \mid h_i \geqslant \mathcal{S}_{j-1}] \right) \cdot \left(\prod_{j=1}^{(k-1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant r_{2j} \mid h_i \geqslant \ell_{2j} - 1] \right) \cdot \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i < r_{k+1} \mid h_i \geqslant \ell_{k+1} - 1]$$

Since the MHR condition implies the item is more likely to survive in earlier time steps, we have:

$$\begin{aligned}
\mathbf{E}[M_k] &\geqslant \sum_{i=1}^{m} \left(\prod_{j=1}^{(k+1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant r_{2j-1} - 1 \mid h_i \geqslant \ell_{2j-1}] \right) \cdot \left(\prod_{j=1}^{(k-1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant r_{2j} \mid h_i \geqslant \ell_{2j} - 1] \right) \cdot \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i < r_{k+1} \mid h_i \geqslant \ell_{k+1} - 1] \\
&= \sum_{i=1}^{m} \left(\prod_{j=1}^{(k+1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant \ell_{2j} - 1 \mid h_i \geqslant \ell_{2j-1}] \right) \cdot \left(\prod_{j=1}^{(k-1)/2} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant l_{2j+1} \mid h_i \geqslant \ell_{2j} - 1] \right) \cdot \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i < r_{k+1} \mid h_i \geqslant \ell_{k+1} - 1] \\
&= \sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [\ell_{k+1} - 1 \leqslant h_i < r_{k+1}]
\end{aligned}$$

Now, $\sum_{i=1}^{m} \mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{k+1} - 1] \geqslant \frac{m}{2^k}$ and $\sum_{i=1}^{m} \mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant r_{k+1}] \leqslant \frac{m}{2^{k+1}}$. Thus,

$$\mathbf{E}[M_k] \geqslant \left(\sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant \ell_{k+1} - 1] \right) - \left(\sum_{i=1}^{m} \mathop{\mathbf{Pr}}_{h_i \sim H_i} [h_i \geqslant r_{k+1}] \right) \geqslant \frac{m}{2^{k+1}}$$

Note that $\frac{m}{2^{k+1}} \geqslant \frac{m}{2^s}$ for $k < s$. For $k = s$, $\mathbf{E}[M_k] \geqslant \sum_{i=1}^{m} \mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{k+1} - 1] \geqslant \frac{m}{2^k} = \frac{m}{2^s}$. Thus, $\mathbf{E}[M_k] \geqslant \frac{m}{2^s} > 5$ for any $k \leqslant s$. By Chernoff bound,

$$\mathbf{Pr}\left[M_k \geqslant \frac{1}{4} \cdot \frac{m}{2^k} \right] \geqslant 1 - \left(\frac{e^{-0.5}}{0.5^{0.5}} \right)^5 > 0.535.$$

Now let p_k be $F_V^{-1}\left(\max\left(0, 1-\frac{m/2^{k-1}}{r_k-\ell_k}\right)\right)$ where F_V is the CDF of distribution V, just as in Algorithm 6. Let the random variable N_k denote the number of buyers with valuation of at least p_k among the next $r_k-\ell_k-1$ buyers. We have

$$\mathbf{E}[N_k] = (r_k-\ell_k-1)\cdot\min\left(1, \frac{m/2^{k-1}}{r_k-\ell_k}\right).$$

If $\frac{m/2^{k-1}}{r_k-\ell_k} \geqslant 1$, then $p_k=-\infty$ and $N_k = r_k-\ell_k-1$ with probability 1. In this case, Algorithm 6 gets at least $\min(M_k, r_k-\ell_k-1)\,\mathbf{E}[V] \geqslant \min\left(M_k, \frac{1}{2}(r_k-\ell_k)\right)\mathbf{E}[V]$ in this stage. Since $M_k \geqslant \frac{1}{4}\cdot\frac{m}{2^k} \geqslant \frac{1}{8}\cdot(r_k-\ell_k)$ with probability at least 0.535, we know Algorithm 6 gets at least $\frac{0.535}{4}\cdot(r_k-\ell_k)\cdot\mathbf{E}[V]$ and thus is an $\frac{4}{0.535} < 8$-approximation during the stage.

If $\frac{m/2^{k-1}}{r_k-\ell_k} < 1$, then $r_k-\ell_k > 10$ and $\mathbf{E}[N_k] = (r_k-\ell_k-1)\cdot\frac{m/2^{k-1}}{r_k-\ell_k} > 0.9\cdot\frac{m}{2^{k-1}} = 1.8\cdot\frac{m}{2^k} > 9$. By Chernoff bound,

$$\mathbf{Pr}\left[N_k \geqslant \frac{1}{4}\cdot\frac{m}{2^k}\right] \geqslant 1-\left(\frac{e^{-\left(1-\frac{1}{4\times 1.8}\right)}}{\left(\frac{1}{4\times 1.8}\right)^{\frac{1}{4\times 1.8}}}\right)^9 > 0.994.$$

When $\min(N_k, M_k) \geqslant \frac{1}{4}\cdot\frac{m}{2^k}$, Algorithm 6 gets at least $\frac{1}{8}$ the benchmark during the stage. Therefore, it is an $\frac{8}{0.535\cdot 0.994} < 15.1$-approximation. □

3.2.4 Approximating Prophet's Welfare in Short Stages

In this section, we deal with length-1 stages using Algorithm BlindMatch, that simply matches each arriving buyer i to any available item.

Theorem 14. *Algorithm* BlindMatch *is a* 2.3*-approximation to* $\sum_{k=1}^{s}\textsc{Pro}_k\cdot\mathbb{1}_{(r_k-\ell_k=1)} = \mathbf{E}[V]\cdot|\{k\in\{1,2,\ldots,s\} \mid r_k-\ell_k=1\}|$.

Proof. Let $z = |\{k\in\{1,2,\ldots,s\} \mid r_k-\ell_k=1\}|$, the number of length-1 stages. Consider the time $t=\left\lceil\frac{z}{2}\right\rceil$. Since there are still at least $\left\lfloor\frac{z}{2}\right\rfloor$ length-1 stages after time t, at least $5\cdot 2^{\left\lfloor\frac{z}{2}\right\rfloor} \geqslant 5\cdot\left\lceil\frac{z}{2}\right\rceil$ items in expectation have horizons of at least t, by the definition of the stages. Using Chernoff bound, the probability that at least $\lceil \frac{z}{2} \rceil$ items with horizons of at least t is greater than $1 - \left(\frac{e^{-0.8}}{0.2^{0.2}}\right)^5 > 0.9$. If this happens, the first t items will be matched. Therefore, Algorithm BLINDMATCH is a $\frac{2}{0.9} < 2.3$-approximation to $\mathbf{E}[V] \cdot z$, completing the proof. □

3.2.5 Approximating Prophet's Welfare in Final Stage

We now approximate PROFINAL from Lemma 10. $\sum_{i=1}^{m} \mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{s+1}] \leqslant 10$ by the definition of the stages. We run Algorithm 7. We randomly sample an item and focus on the item in our algorithm. The probability that item i is sampled is proportional to $\mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{s+1}]$. If item i is sampled, we run an algorithm for the single-item setting (lines 3 and 4 in Algorithm 7). The single-item policy is analyzed in Section 3.3 where it is shown to achieve welfare at least $\frac{1}{2} \cdot \text{PROSINGLE}_i$.

Algorithm 7: SINGLEITEM

1 For $i = 1, 2, \ldots, m$, set $q_i \leftarrow \mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{s+1}]$

2 $i^* \leftarrow$ item $i \in \{1, 2, \ldots, m\}$ with probability $\frac{q_i}{\sum_{i=1}^{m} q_i}$

3 Set the reserve price p so that $\mathbf{Pr}_{v \sim V}[v \geqslant p] = \frac{1}{\mathbf{E}[H_{i^*}]}$

4 For each arriving buyer, try selling item i^* with reserve price p

Theorem 15. *Algorithm 7 is a* 20*-approximation of* PROFINAL *in expectation.*

Proof. By Theorem 16 and Theorem 17, if item $i^* = i$, the algorithm gets $\frac{1}{2} \cdot \text{PRO}_i$ in expectation. Thus, the expected welfare achieved by algorithm 7 is at least

$$\sum_{i=1}^{m} \frac{\mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{s+1}]}{\sum_j \mathbf{Pr}_{h_j \sim H_j}[h_j \geqslant \ell_{s+1}]} \cdot \frac{1}{2} \cdot \text{PROSINGLE}_i$$
$$\geqslant \sum_{i=1}^{m} \frac{\mathbf{Pr}_{h_i \sim H_i}[h_i \geqslant \ell_{s+1}]}{10} \cdot \frac{1}{2} \cdot \text{PROSINGLE}_i \geqslant \frac{1}{20} \cdot \text{PROFINAL}.$$
□

3.2.6 Approximating Prophet's Welfare

Now we are ready to prove our main theorem.

Proof of Theorem 9. To summarize our previous discussion:

(1) Theorem 13 yields a 15.1-approximation to $\sum_k \textsc{Pro}_k$, where the sum is over odd stages $k \leqslant s$ with $r_k - \ell_k \geqslant 2$.

(2) If we replace "odd" with "even" in Theorem 13 and the corresponding algorithm, we have a 15.1-approximation $\sum_k \textsc{Pro}_k$ over even stages k with $r_k - \ell_k \geqslant 2$.

(3) Theorem 14 is a 2.3-approximation to $\sum_k \textsc{Pro}_k$ over stages $k \leqslant s$ with $r_k - \ell_k = 1$.

(4) Theorem 15 yields a 20-approximation to ProFinal.

An algorithm can do one of (1) to (4) with probability $\frac{15.1}{52.5}, \frac{15.1}{52.5}, \frac{2.3}{52.5}$ and $\frac{20}{52.5}$ respectively, yielding a 52.5-approximation to Pro. □

3.3 Prophet Inequality for Single Item with MHR Horizon

In this section, we consider the case where there is $m = 1$ item, and present a proof of Theorem 10. The algorithm also serves as our approximation for ProSingle, which we use for the overall algorithm with multiple items

We show that the following fixed-price *balancing* scheme is a 2-approximation, and this bound is tight for geometric distributions:

> Pretend the item departs uniformly over time at rate $1/\mu$, where $\mu = \mathbf{E}[H]$. Choose a price p s.t. the rate of acceptance of buyers matches the rate of departure of the item.

We bound the performance of this policy by using a simple linear programming upper bound on PRO that only uses expected values. Though the relaxation is simple, just as in Section 3.2.2, it brings out the key insight that the upper bound also behaves like a balancing scheme, except it assumes the item lasts forever when performing the matching. Surprisingly, such a simple relaxation yields the worst-case optimal bound over all MHR distributions.

Theorem 16. *Let* $\alpha = 1 - \mathbf{E}_{h\sim H}[(1-\mu^{-1})^h]$. *Then for* $m = 1$ *items, there is a fixed pricing policy that is* $\frac{1}{\alpha}$*-competitive. This policy sets the price* p *such that* $\mathbf{Pr}_{X\sim V}[X \geqslant p] = \frac{1}{\mu}$ *where* $\mu = \mathbf{E}[H]$.

Proof. First we find an upper bound for PRO. Let X be a random variable with distribution V. Consider the following LP:

$$\begin{aligned} \text{maximize} \quad & \sum_v y(v)\cdot v \\ \text{subject to} \quad & \sum_v y(v) \leqslant 1, \\ & y(v) \leqslant \mu \cdot \mathop{\mathbf{Pr}}_{X\sim V}[X = v], \forall v. \end{aligned}$$

Variable $y(v)$ is the probability that a buyer with realized value v is chosen by prophet. The first constraint requires the item to be sold at most once in expectation. The second constraint says each value can be chosen only when it appears. Both of the constraints are relaxations as they should hold for any realization while the constraints are in expectation. The optimal objective is thus an upper bound for the expected value of the prophet.

Let λ be the Lagrange multiplier associated with the first constraint. The partial Lagrangian of the LP is:

$$\begin{aligned} \mathcal{L}(\lambda) = \lambda + \sum_v y(v)\cdot(v-\lambda), \\ y(v) \leqslant \mu \cdot \mathop{\mathbf{Pr}}_{X\sim V}[X = v], \quad \forall v. \end{aligned}$$

The partial Lagrangian is decoupled for each value v and is maximized when $y(v) = \mu \cdot \mathbf{Pr}_{X \sim V}[X = v]$ for any $v \geqslant \lambda$ and $y(v) = 0$ otherwise. For any λ, this gives us an upper bound on the prophet's welfare. Let p be the value such that $\mathbf{Pr}_{X \sim V}[X \geqslant p] = \frac{1}{\mu}$. If we set $\lambda = p$, we get the following upper bound for the prophet's value:

$$\textsc{Pro} \leqslant \sum_{v \geqslant p} \mu \cdot v \cdot \mathbf{Pr}[X = v] = \mathop{\mathbf{E}}_{X \sim V}[X \mid X \geqslant p].$$

Essentially, the prophet pretends that the horizon is infinite and it can always find a buyer with value at least p. Now we look at $\textsc{Alg}$ which is an algorithm with a single price p. The algorithm has to also consider the event that the horizon ends before the item is matched.

$$\begin{aligned} \textsc{Alg} = & \mathop{\mathbf{E}}_{X \sim V}[X \mid X \geqslant p] \cdot \\ & \quad \mathbf{Pr}[\text{a value at least } p \text{ was seen during the time horizon}] \\ = & \mathop{\mathbf{E}}_{X \sim V}[X \mid X \geqslant p] \cdot \mathop{\mathbf{E}}_{h \sim H}[1 - (1 - \mu^{-1})^h]. \end{aligned}$$

Therefore,

$$\frac{\textsc{Pro}}{\textsc{Alg}} \leqslant \mathop{\mathbf{E}}_{h \sim H}[1 - (1 - \mu^{-1})^h]^{-1}. \qquad \square$$

Now, we show that for MHR horizons, this algorithm is $(2 - \mu^{-1})$-competitive. The key idea is to use second order stochastic dominance to show that the upper bound is maximized for geometric distributions with the same mean. Somewhat surprisingly, we also show in Theorem 18 that this result is tight in the sense that for geometric distributions, no online policy can do better.

Theorem 17. *For any MHR distribution with mean μ, $\mathbf{E}_{h \sim H}[1 - (1 - \mu^{-1})^h]^{-1} \leqslant 2 - \mu^{-1}$.*

In order to prove the above theorem, we use *second-order stochastic dominance.*

Definition 2. Let A and B be two probability distributions on $\mathbb{R}$. Let F_A be the cumulative distribution function of A and F_B be the CDF of B. We say A is *second-order stochastically dominant* over B if for all $x \in \mathbb{R}$,

$$\int_{-\infty}^{x} (F_B(t) - F_A(t))\mathrm{d}t \geqslant 0.$$

Proposition 1. *If distribution A is second-order stochastically dominant over B, and A and B have the same mean, then for any convex function $f : \mathbb{R} \to \mathbb{R}$, $\mathbf{E}_{x\sim B}[f(x)] \geqslant \mathbf{E}_{x\sim A}[f(x)]$.*

We now use second order stochastic dominance to show the following.

Lemma 11. *Geometric distribution with mean μ is second-order stochastically dominated by any other MHR horizon distribution with the same mean.*

Proof. Let $\phi_c(x) : \mathbb{N}^+ \to \mathbb{R}$ be the following convex function:

$$\phi_c(x) = \begin{cases} c - x & \text{if } x \leqslant c \\ 0 & \text{if } x > c \end{cases}$$

where c is a positive integer. Let G be the geometric distribution with mean μ. From Definition 2, the lemma holds if and only if $\mathbf{E}_{x\sim D}[\phi_c(x)] \leqslant \mathbf{E}_{x\sim G}[\phi_c(x)]$ for any c and any MHR distribution D with the same mean μ.

We prove this by contradiction. Let D be an MHR distribution with mean μ which satisfies $\mathbf{E}_{x\sim D}[\phi_c(x)] > \mathbf{E}_{x\sim G}[\phi_c(x)]$ for some c. The set of MHR distributions with the same tail after c (the same $\mathbf{Pr}_{x\sim D}[x = x^* \mid x > c]$ for any $x^* > c$) is homeomorphic to a closed and bounded set in $\mathbb{R}^c$, which means it's compact. The function $\mathbf{E}_{x\sim D}[\phi_c(x)]$ is continuous in D under L^1-norm, so there is a $D = D^*$ maximizing $\mathbf{E}_{x\sim D}[\phi_c(x)]$ among MHR distributions with the same tail after c. This

D^* differs from G at some $x \leqslant c$. Define $q_i = \mathbf{Pr}_{x\sim D^*}[x \geqslant i+1 \mid x \geqslant i]$ and $q = \mathbf{Pr}_{x\sim G}[x \geqslant i+1 \mid x \geqslant i] = 1-\mu^{-1}$. Because D^* is MHR, q_i's are decreasing. Also $q_1 > q$ as otherwise the mean cannot be μ, and $q_c < q$ as otherwise $\mathbf{E}_{x\sim D^*}[\phi_c(t)] > \mathbf{E}_{x\sim G}[\phi_c(x)]$ cannot hold. Thus there is some $i^* < c$ such that $q_{i^*} > q$ and $q_{i^*+1} \leqslant q$.

We are going to show for a pair of small enough ε and ε', decreasing q_{i^*} by ε and increasing q_{i^*+1} by ε' such that the mean is preserved will increase $\mathbf{E}_{t\sim D^*}[\phi_c(x)]$. Let $r = 1 + q_{i^*+2} + q_{i^*+2}q_{i^*+3} + q_{i^*+2}q_{i^*+3}q_{i^*+4} + \cdots$. When $\varepsilon \to 0$, we have $\varepsilon(1+q_{i+1}r) = \varepsilon' q_i r$. This implies $\varepsilon' q_i - \varepsilon q_{i+1} > 0$, which means $\mathbf{E}_{x\sim D^*}[\phi_c(x)]$ is increased. It contradicts with the fact that D^* maximizes $\mathbf{E}_{x\sim D^*}[\phi_c(x)]$. □

Proof. (of Theorem 17) From Theorem 16, we know $\frac{\textsc{Pro}}{\textsc{Alg}} \leqslant 1/\mathbf{E}_{h\sim H}[\phi(h)]$ where $\phi(h) = 1-(1-\mu^{-1})^h$ is a concave function. From Lemma 11 and Proposition 1, among all MHR distributions H with mean μ, $\mathbf{E}_{h\sim H}[\phi(h)]$ is minimized by a geometric one. For geometric departure with mean μ, $\mathbf{E}_{h\sim H}[\phi(h)] = 2 - \mu^{-1}$. □

Theorem 18. *No online algorithm is better than* $(2-\mu^{-1})$*-competitive for* $m = 1$ *items when the horizon distribution* H *is geometric with mean* μ.

Proof. Let $q \in [0, 1)$ be the probability that the process continues after each step. We have $q = 1 - \mu^{-1}$.

Define $\textsc{Alg}^*$ as the expected value of the optimal algorithm and $\textsc{Pro}$ as that of the prophet. Let the valuation distribution be: v_{L} with probability $1-p$ and v_{H} with probability p, $v_{\mathrm{L}} < v_{\mathrm{H}}$. At each step, $\textsc{Alg}^*$ will set the price to v_{H} if it expects to get more than v_{L} afterwards. Otherwise it will set the price to v_{L}. Randomizing over v_{L} and v_{H} cannot help $\textsc{Alg}^*$. Also, because the geometric distribution is memoryless, $\textsc{Alg}^*$ will make the same decision every time, *i.e.*, the optimal algorithm is single-threshold. We have

$$\textsc{Alg}^* = \max\left\{v_{\mathrm{L}} \cdot (1-p) + v_{\mathrm{H}} \cdot p,\ v_{\mathrm{H}} \cdot \frac{p}{1-q(1-p)}\right\}$$

and

$$\textsc{Pro} = v_{\mathrm{H}} \cdot \frac{p}{1-q(1-p)} + v_{\mathrm{L}} \cdot \left(1 - \frac{p}{1-q(1-p)}\right).$$

When $\mu = 1$ and $q = 0$, the theorem holds because $2 - \mu^{-1} = 1$. Otherwise, we set v_{H} so that $\textsc{Alg}^*$ is indifferent between its two options. In that case,

$$\begin{aligned}
\lim_{p\to 0} \frac{\textsc{Pro}}{\textsc{Alg}^*} &= \lim_{p\to 0} \left(1 + \frac{v_{\mathrm{L}}}{v_{\mathrm{H}}} \cdot \frac{1-q(1-p)}{p}\right) \\
&= \lim_{p\to 0} \left(1 + \left(\frac{p}{1-q(1-p)} - p\right) \cdot \frac{1-q(1-p)}{p}\right) \\
&= 1 + q = 2 - \mu^{-1}.
\end{aligned}$$

□

3.4 Lower Bounds for MHR Horizons

Next we provide a proof of Theorem 11. For this, we first show a lower bound of 2 for any dynamic pricing scheme in the limit when m becomes large, and 1.57 for any finite m. We will subsequently show that no fixed pricing scheme can extract constant fraction of the welfare for $m > 1$ items. For showing these results, we consider a special family of *i.i.d.* MHR horizon distributions, which we call *low-rate geometric*: Let H be a geometric distribution with mean μ, so the probability of survival at each step is $q = 1 - \mu^{-1}$. We call H low-rate geometric when $q \to 1^-$. Let $\lambda = 1 - q$ be the rate of departure for each item. This goes to 0^+ when H is low-rate geometric.

Low-rate geometric distributions correspond to the canonical setting where items are long-lasting, yet their departures are memoryless. In addition to being canonical, the reason we consider this setting is its analytic tractability: It allows us to ignore events where multiple items depart simultaneously, leading to tractable Markov chains for both the prophet and the algorithm. The proof of lower bound of 2 for large m involves analyzing an interesting time-reversed Markov chain for the prophet's welfare.

3.4.1 Tractable Approximation

Denote by $\textsc{Alg}^*_m(\lambda)$ the optimal online policy when there are m items and the rate of departures is λ. Similarly, we define $\textsc{Pro}_m(\lambda)$ to denote the prophet. Since we are considering the limit as $\lambda \to 0^+$, we will assume throughout that $\lambda < \frac{1}{m}$.

Define the *state* of the system to be k if there are k items in the system. Note that since departures are geometric, any online policy will use a fixed price in each state. The state of the system therefore *decreases* over time. For both of the processes (corresponding to prophet and the optimal algorithm) given the current state is k, there is a positive probability that the next state will be k' for any $k' \leqslant k$. However, the probability that multiple items depart together (or a match and departures happen together for the algorithm) is extremely small when $\lambda \to 0^+$. In light of this, we introduce alternative processes for the ease of analysis.

In an alternative process, we will assume two events (departures, matches) do not simultaneously happen. In other words, for the prophet, given state k, the state transitions to $k-1$ with probability $k\lambda$ per time step. We do not consider state changes due to matching. Instead and equivalently, we will assume that in hindsight, the prophet can optimally match arriving buyers to items that had not departed by that time. Call this prophet $\textsc{Pro}'_m(\lambda)$. For the algorithm, we assume that if the state is k, the price is set so that the rate at which a buyer is matched is $\pi_k = \beta_k \lambda k$. Since items also depart at rate λk, we will assume the state transitions from k to $k-1$ at rate $(1+\beta_k)\lambda k$. Denote the optimal such algorithm as $\textsc{Alg}'_m(\lambda)$.

Lemma 12. *For any $m \geqslant 1$: $\frac{\textsc{Pro}_m(\lambda)}{\textsc{Pro}'_m(\lambda)} \to 1$ and $\frac{\textsc{Alg}^*_m(\lambda)}{\textsc{Alg}'_m(\lambda)} \to 1$ as $\lambda \to 0$*

Proof. We only show that $\frac{\textsc{Pro}_m(\lambda)}{\textsc{Pro}'_m(\lambda)} \to 1$. The proof of the second part that $\frac{\textsc{Alg}^*_m(\lambda)}{\textsc{Alg}'_m(\lambda)} \to 1$ uses a similar argument. We consider the following two processes: the main process based on the actual departure of items in which two departures might happen simultaneously and the alternative process in which at each time step at most one item can depart. The alternative process might modify the number of items in the system at some point during the process with a very small probability. In that case, states of the two processes differ at some point and the two corresponding prophets might achieve different values. Otherwise, they are always at the same state during the process and their values are exactly the same.

There exist two sources of differences (only consider the first time step that they are not at the same state during the process). The first one which we call type 1 is as follows: If two departures happen at the same time, alternative process will only consider one of them. In other words, if the main process goes from state k to k' such that $k' < k-1$, the alternative process will go from k to $k-1$ and will assume there are still $k-1$ items in the system at the next time step. The probability of such a difference for a state k is not more than $\frac{2^k(1-q)^2}{kq^{k-1}(1-q)}$ which goes to 0 as q approaches 1. Therefore, using the union bound and the fact that m is finite, the probability of such a difference during the process at some state k denoted by p_1 also approaches 0.

The second source of differences (type 2) is: If the current state of the main process is k and it remains unchanged after a time step (no departures happens) with a very small probability ($\frac{q^k-1+k(1-q)}{q^k}$), the state of the alternative process will change to state $k-1$ at this time step. The probability of such a difference at state k is $\frac{q^k-1+k(1-q)}{k(1-q)q^{k-1}}$. We can see that this probability goes to 0 as q approaches 1 and since m is finite, using the union bound, the probability of such a difference during the process denoted by p_2 goes to 0 as q approaches 1.

Note that the value of $\textsc{Pro}'$ (alternative prophet) can only be greater than $\textsc{Pro}$ (main prophet) if two departures happen at the same time step during the actual departure process (type 1 difference). However, note that the conditional expectation of $\textsc{Pro}'$ given that such a difference exists is not greater than $\textsc{Pro}'$ (the expected welfare of $\textsc{Pro}'$). Therefore, we have:

$$\textsc{Pro} \geqslant (1 - p_1)\textsc{Pro}'.$$

In addition, $\textsc{Pro}$ can be only greater then $\textsc{Pro}'$ if a type 2 difference exists. Similarly, the conditional expectation of $\textsc{Pro}$ given that a type 2 difference exists is not greater than $\textsc{Pro}$. Therefore, we also have:

$$\textsc{Pro}' \geqslant (1 - p_2)\textsc{Pro}.$$

Using the last two inequalities,

$$1 - p_1 \leqslant \frac{\textsc{Pro}}{\textsc{Pro}'} \leqslant \frac{1}{1 - p_2}.$$

Using that p_1 and p_2 both go to 0, we have $\frac{\textsc{Pro}}{\textsc{Pro}'} \to 1$.

□

Therefore, we will analyze the quantity $c_m(\lambda) = \frac{\textsc{Pro}'_m(\lambda)}{\textsc{Alg}'_m(\lambda)}$ as the competitive ratio of the algorithm against prophet for any m, λ and subsequently take the limit as $\lambda \to 0^+$. In the remainder of this section, without creating ambiguity we omit the m and λ in notation and use $\textsc{Alg}'$ and $\textsc{Pro}'$ instead.

3.4.2 Lower Bound Construction for Dynamic Pricing

To show the lower bounds, we consider the valuation distribution V such that for any $x \in [1, \infty)$, $\mathbf{Pr}_{v \sim V}[v \geqslant x] = x^{-\alpha}$ where $\alpha \in (1, +\infty)$ is a constant that will be determined later. Note that $\mathbf{E}_{v \sim V}[v]$ is finite. We first give an upper bound for $\textsc{Alg}'$ for this valuation distribution:

$$\textsc{Alg}' \leqslant \sum_{k=1}^{m} \max_{\beta_k} \frac{\beta_k k \lambda}{(1 + \beta_k) k \lambda} \cdot \mathbf{E}[v \mid v \geqslant F_V^{-1}(1 - \beta_k k \lambda)]$$

where F_V is the cumulative distribution function for V. The probability of accepting a buyer in state k is at most $\frac{\beta_k k\lambda}{(1+\beta_k)k\lambda}$ (because acceptance and departure are disjoint events in the alternative process).

Simplifying it, we have:

$$\begin{aligned}
\textsc{Alg}' &\leqslant \sum_{k=1}^{m} \max_{\beta_k} \frac{\beta_k k\lambda}{(1+\beta_k)k\lambda} \cdot \mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant (\beta_k k\lambda)^{-\frac{1}{\alpha}}] \\
&= \sum_{k=1}^{m} \max_{\beta_k} \frac{\beta_k k\lambda}{(1+\beta_k)k\lambda} \cdot \frac{\alpha}{\alpha-1} \cdot (\beta_k k\lambda)^{-\frac{1}{\alpha}} \\
&= \sum_{k=1}^{m} (k\lambda)^{-\frac{1}{\alpha}} \cdot \frac{\alpha}{\alpha-1} \cdot \max_{\beta_k} \frac{\beta_k^{1-\frac{1}{\alpha}}}{1+\beta_k}
\end{aligned}$$

Optimizing over β_k, we have:

$$\textsc{Alg}' = \sum_{k=1}^{m} (k\lambda)^{-\frac{1}{\alpha}} \cdot (\alpha-1)^{-\frac{1}{\alpha}}. \tag{3.1}$$

Now we solve for $\textsc{Pro}'$. Note that for the prophet, we assume the state only changes due to departure of items. Let $p_k(v)$ denote the probability that the item departing in state k is matched to a buyer with valuation at least v by the prophet. In the rest of this section, we call the item departing at state k to be item k. We have:

$$\textsc{Pro}' = \sum_{k=1}^{m} \int_0^{+\infty} p_k(v)\mathrm{d}v \tag{3.2}$$

We now present different bounds for the above quantity depending on whether m is finite, or we are considering the limit $m \to \infty$.

Lower Bound for Dynamic Pricing: Finite m

This bound is simpler. Clearly, if a buyer with value at least v arrives at state k, $\textsc{Pro}'$ always can assign the item k to a buyer with value at least v. Therefore,

$$\begin{aligned}\textsc{Pro}' \geqslant & \\ \sum_{k=1}^{m}\int_{0}^{+\infty} & \mathbf{Pr}[\text{some buyer with valuation at least } v \text{ arrives in state } k]\mathrm{d}v \\ &= \sum_{k=1}^{m}\left(1+\int_{1}^{+\infty}\frac{(1-k\lambda)v^{-\alpha}}{v^{-\alpha}+k\lambda-k\lambda v^{-\alpha}}\mathrm{d}v\right).\end{aligned}$$

Therefore, we have:

$$\begin{aligned}\liminf_{\lambda\to 0^+}\frac{\textsc{Pro}'}{\textsc{Alg}'} \geqslant \lim_{\lambda\to 0^+}\frac{\int_1^{+\infty}\frac{1}{1+k\lambda v^{\alpha}}\mathrm{d}v}{(k\lambda)^{-\frac{1}{\alpha}}(\alpha-1)^{-1/\alpha}} &= \frac{\int_0^{+\infty}\frac{1}{1+u^{\alpha}}\mathrm{d}u}{(\alpha-1)^{-1/\alpha}} \\ = \frac{\frac{1}{\alpha}\cdot B(\frac{1}{\alpha},1-\frac{1}{\alpha})}{(\alpha-1)^{-1/\alpha}} &= \frac{\frac{\pi}{\alpha}/\sin(\frac{\pi}{\alpha})}{(\alpha-1)^{-1/\alpha}},\end{aligned}$$

where $B(\cdot,\cdot)$ is the beta function. Comparing to Equation (3.1), we have that $\liminf_{\lambda\to 0^+}\frac{\textsc{Pro}'}{\textsc{Alg}'}$ is maximized at $\alpha = 2$ and in that case,

$$\frac{\textsc{Pro}'}{\textsc{Alg}'} \geqslant \pi/2 \approx 1.5708$$

The bound holds for any $m \in \mathbb{N}^+$.

Lower Bound for Dynamic Pricing: Large m

We now consider the more interesting case when $m \to \infty$. We present a tighter lower bound for Equation (3.2). To achieve this goal, we need to to analyze $p_k(v)$ more carefully. Previously, we used the fact that if a buyer with value at least v arrives during state k, then a buyer with value at least v will be assigned to the item k by the prophet. However, the prophet might assign a buyer with value at least v to item k even if no such buyer arrives in state k.

It is easy to see that the optimal policy for the prophet is the following: It considers the items in increasing order of realized horizon, and matches each item to the highest valued unmatched buyer arriving no later than the horizon of the item. A buyer with value at least v is matched to the item k if and only if there is an $i \geqslant 0$ such that between beginning of the state $k+i$ and end of state k, at least $i+1$ buyers with value at least v arrive. Note that in the previous section, we only considered the case of $i = 0$ to give a lower bound for $p_k(v)$.

Time-reversed Markov Chain. In order to analyze the new process, we start from the end of state k and go back in time. There are two possible types of events:

- An item departs, so that the state increases by 1 (note we are going back in time); or
- A buyer with valuation at least v arrives.

We maintain a counter q initially set to 1. Each time an item departs, we increase q by 1, and each time a buyer with valuation at least v arrives, we decrease q by 1. It is easy to see that the item k is matched to a buyer with valuation at least v by the prophet if and only if q reaches 0, *i.e.*, $p_k(v) = \mathbf{Pr}[q = 0 \text{ at some time}]$.

Note that as we are going back in time, when the state is $k+j-1$, the probability an item departs is $(k+j)\lambda$. Similarly, the probability a buyer with valuation at least v arrives is $v^{-\alpha}$. This yields a Markov chain in which when the state is the $(k+j, q)$, the former event causes the state to become $(k+j+1, q+1)$ and the latter causes the state to become $(k+j, q-1)$.

As $p_k(v)$'s themselves are hard to analyze, we approximate them by a sequence of functions $\{f_j(x)\}_{j=0}^{\infty}$. Each $f_j(x)$ is defined on $[0, 1]$, and it represents the probability that the following random walk ever reaches 0 in j steps: A point starts at 1 on the number line. Independently in each step, it goes left by 1 with probability x, and goes right by 1 otherwise. Note that as $j \to \infty$, $f_j(x) \to \min\left(1, \frac{x}{1-x}\right)$. That is, the point-wise limit of $\{f_j(x)\}_{j=1}^{\infty}$ as $j \to \infty$ is $f(x) = \min(1, x/(1-x))$. (We slightly abuse notation at $x = 1$ and $f(1) = 1$.)

Lemma 13. *$p_k(v) \geqslant f_j(v^{-\alpha}/(v^{-\alpha} + (k+j)\lambda))$ for any integer $j \in [1, m-k]$.*

Proof. From state k to state $k+j$, exactly j departures happen so the process for $p_k(v)$ has at least j moves in this period. For each move, the probability that the counter q decreases is at least $v^{-\alpha}/(v^{-\alpha} + (k+j)\lambda)$. Therefore, we can couple these two processes so that if the random walk ever reaches 0, the counter must have visited 0 too. □

We now show that these functions uniformly converge.

Lemma 14. *$\{\log f_j(x)\}_{j=1}^{\infty}$ uniformly converges to $\log f(x)$ on $(0, 1]$. This implies $\forall \varepsilon > 0, \exists k, \forall j > k, \forall x, f_j(x) > (1-\varepsilon)f(x)$.*

Proof. Notice $f_j(x)$ is continuous on x and increasing in j. For any $c > 0$, on the compact set $[c, 1]$, each $\log f_j(x)$ is continuous in x, and their limit $\log f(x)$ is continuous too. Further, $\log f_j(x)$ is increasing in j. By Dini's theorem, the convergence on $[c, 1]$ is uniform.

For any $\varepsilon > 0$, for any $x \in (0, \varepsilon)$ and any $j \geqslant 1$, $f_j(x) \geqslant x > \frac{x}{1-x} \cdot (1-\varepsilon) = (1-\varepsilon)f(x)$. Because $\{\log f_j(x)\}_{j=1}^{\infty}$ uniformly converges on $[\varepsilon, 1]$, there is a k so that for any $j > k$ and any $x \in [\varepsilon, 1]$, $f_j(x) > (1-\varepsilon)f(x)$. This completes the proof. □

Now we are ready to explicitly compute a lower bound for Pro′ as $m \to \infty$. We start with Equation (3.2).

$$
\begin{aligned}
\textsc{Pro}' &= \sum_{k=1}^{m} \int_0^{+\infty} p_k(v)\mathrm{d}v \\
&\geqslant \sum_{k=1}^{m-\sqrt{m}} \int_0^{+\infty} p_k(v)\mathrm{d}v \\
&\geqslant \sum_{k=1}^{m-\sqrt{m}} \int_0^{+\infty} f_{\sqrt{m}}(v^{-\alpha}/(v^{-\alpha} + (k+\sqrt{m})\lambda))\mathrm{d}v
\end{aligned}
$$

where the final inequality follows from Lemma 13.

Let $c_k = \inf_{x\in(0,1]} f_k(x)/f(x)$. Then we have:

$$
\begin{aligned}
\textsc{Pro}' &\geqslant c_{\sqrt{m}} \sum_{k=1}^{m-\sqrt{m}} \int_0^{+\infty} f(v^{-\alpha}/(v^{-\alpha} + (k+\sqrt{m})\lambda))\mathrm{d}v \\
&= c_{\sqrt{m}} \sum_{k=1}^{m-\sqrt{m}} \int_0^{+\infty} \min(1, v^{-\alpha}/((k+\sqrt{m})\lambda))\mathrm{d}v \\
&= c_{\sqrt{m}} \sum_{k=\sqrt{m}}^{m} \int_0^{+\infty} \min(1, v^{-\alpha}/(k\lambda))\mathrm{d}v \\
&= c_{\sqrt{m}} \sum_{k=\sqrt{m}}^{m} \left((k\lambda)^{-\frac{1}{\alpha}} + \int_{(k\lambda)^{-\frac{1}{\alpha}}}^{+\infty} v^{-\alpha}/(k\lambda)\mathrm{d}v \right) \\
&= c_{\sqrt{m}} \sum_{k=\sqrt{m}}^{m} \frac{\alpha}{\alpha-1} \cdot (k\lambda)^{-\frac{1}{\alpha}}.
\end{aligned}
$$

When $m \to \infty$, $c_{\sqrt{m}}$ goes to 1 by Lemma 14, and $\frac{\sum_{k=\sqrt{m}}^{m} k^{-\frac{1}{\alpha}}}{\sum_{k=1}^{m} k^{-\frac{1}{\alpha}}}$ goes to 1 too. Thus,

$$
\liminf_{m\to\infty} \left(\liminf_{\lambda\to 0^+} \frac{\textsc{Pro}'}{\sum_{k=1}^{m} \frac{\alpha}{\alpha-1} \cdot (k\lambda)^{-\frac{1}{\alpha}}} \right) \geqslant 1.
$$

Together with the bound for $\textsc{Alg}'$ from Equation (3.1), this gives us:

$$
\liminf_{m\to\infty} \left(\liminf_{\lambda\to 0^+} \frac{\textsc{Pro}'}{\textsc{Alg}'} \right) \geqslant \frac{\frac{\alpha}{\alpha-1}}{(\alpha-1)^{-\frac{1}{\alpha}}},
$$

which reaches its maximum of 2 at $\alpha = 2$. This completes the proof of Theorem 11.

3.4.3 Lower Bound for Fixed Pricing Schemes

A natural question is whether there is a single-threshold algorithm that is a constant approximation. Note that this is indeed the case when the horizons H_i's are identical and deterministic; in fact, in this case, the competitive ratio approaches 1 as $m \to \infty$. In contrast, when the horizons are not deterministic — even if they are *i.i.d* geometric, we show that no fixed pricing scheme can be constant-competitive. This shows the second part of Theorem 11.

Theorem 19. *There exists a family of instances with* i.i.d *geometric horizons, such that any fixed pricing algorithm is* $\Omega(\log\log m)$*-competitive, where* m *is the number of items.*

Proof. For any $m \geqslant 2^5$ such that $\log_2 m$ is an integer, consider a geometric horizon distribution H whose mean is m: Let q_m be the probability that the horizon is greater than the mean, *i.e.* $q_m = \mathbf{Pr}_{h\sim H}[h > m]$. It is easy to verify $\frac{1}{4} \leqslant q_m \leqslant \frac{1}{e}$ since H is geometric. Let the value distribution V satisfy: $\mathrm{supp}(V) = \{1/(q_m^t t^2) \mid t = 3, 4, \dots, \log_2 m\}$ and $\mathbf{Pr}_{v\sim V}[v \geqslant 1/(q_m^t t^2)] = q_m^t$ for $t = 3, 4, \dots, \log_2 m$. Straightforward calculation shows

$$\begin{aligned}\mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant 1/(q_m^t t^2)] &= \Theta(1) \cdot \frac{1}{q_m^t} \cdot \sum_{k=t}^{\log_2 m} q_m^k \cdot \frac{1}{q_m^k k^2} \\ &= \Theta(1) \cdot \frac{1}{q_m^t} \cdot \left(\frac{1}{t} - \frac{1}{(\log_2 m) + 1}\right).\end{aligned}$$

Without loss of generality, for any single-threshold algorithm Sing, assume the threshold is $1/(q_m^t t^2)$. We know in time interval $[jm + 1, (j + 1)m]$, the expected number of transactions is at most the minimum of the expected number of buyers with valuations at least $1/(q_m^t t^2)$, and the expected number of items alive at the start of the interval. Therefore,

$$\begin{aligned}
\textsc{Sing} &\leqslant \mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant 1/(q_m^t t^2)] \cdot \sum_{j=0}^{\infty} \min(mq_m^t, mq_m^j) \\
&\leqslant m \cdot \mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant 1/(q_m^t t^2)] \cdot (t+1)q_m^t/(1-q_m) = O(m).
\end{aligned}$$

We know from previous discussion that the upper bound from Lemma 9 is at most $53 \cdot \textsc{Pro}$. Previously, we set the stages so that about $\frac{1}{2}$ of items depart in each stage. The factor of $\frac{1}{2}$ is not essential and we can change it to any constant strictly between 0 and 1, *e.g.* q_m. Doing this only costs us a constant.

If we set the reserve price in the interval $[jm+1,(j+1)m]$ to be $q_m^j j^2$, we have:

$$\begin{aligned}
\textsc{Pro} &= \Omega(1) \cdot m \cdot \sum_{j=3}^{(\log_2 m)-5} q_m^j \cdot \mathop{\mathbf{E}}_{v\sim V}[v \mid v \geqslant 1/(q_m^j j^2)] \\
&= \Omega(1) \cdot m \cdot \sum_{j=3}^{(\log_2 m)-5} 1/j = \Omega(m \log\log m).
\end{aligned}$$

Therefore, $\textsc{Pro} = \Omega(\log\log m) \cdot \textsc{Sing}$ for the constructed family of instances. □

3.5 Lower Bound for General Horizons

In this section, we prove Theorem 12. We consider the general case without the MHR restriction on the horizon distributions and provide lower bound for this case.

We assume $m = 1$ in this proof. The same ideas apply to any $m \geqslant 1$. Without loss of generality, assume $n = 2^{ck}$ for c that will be fixed later. The horizon is 2^{ci} with probability 2^{-i-1} for $i = 0, 1, 2, \ldots, k-1$, and is n with probability 2^{-k}. Intuitively, there are $k+1$ possible horizons, where each one is exponentially longer, yet exponentially less probable than the previous one. Denote the valuation distribution by: a_1 with probability p_1, a_2 with probability p_2, …, a_m with probability p_m where $a_1 < a_2 < \cdots < a_m$. Here we set $m = ck$, $a_i = 2^{i/c}$ and $p_i = 2^{-i}$ except $p_{ck} = 2^{-ck+1}$.

Let VPRO be any policy that knows realized valuations but not realized horizon, and PRO be the omniscient prophet. The only information VPRO does not know beforehand is the realized horizon, and during execution it cannot do anything once the horizon ends. Therefore it should aim for a specific buyer in advance:

$$\begin{aligned}
\text{VPRO} &= \mathop{\mathbf{E}}_{v_1,\ldots,v_n}\left[\max_i \left(\sum_{j\geqslant i}\pi_j\right) M_i(v_1,\ldots,v_n)\right] \\
&\leqslant 2 \mathop{\mathbf{E}}_{v_1,\ldots,v_n}\left[\max_i \pi_i M_i(v_1,\ldots,v_n)\right],
\end{aligned}$$

where π_i is the probability for the horizon to be 2^{ci} and $M_i(v_1,\ldots,v_n)$ is the maximum of the first 2^{ci} values. Then we have

$$\begin{aligned}
\text{VPRO} &\leqslant 4 \cdot \sum_{i=0}^{k} 2^i \mathop{\mathbf{Pr}}_{v_1,\ldots,v_n}[\exists j, \pi_j M_j \geqslant 2^i] \\
&\leqslant 4 \sum_{i=0}^{k} 2^i \min\left(1, \sum_j \mathop{\mathbf{Pr}}_{v_1,\ldots,v_n}[\pi_j M_j \geqslant 2^i]\right) \\
&\leqslant 4 \sum_{i=0}^{k} 2^i \min\left(1, \sum_j 2^{cj} \mathop{\mathbf{Pr}}_{v_1,\ldots,v_n}[2^{-j} v_1 \geqslant 2^i]\right) \\
&\leqslant 4 \sum_{i=0}^{k} 2^i \min\left(1, 2\sum_j 2^{cj} 2^{-ci-cj}\right) \\
&\leqslant 4 \sum_{i=0}^{k} 2^i \min\left(1, 2(k+1)2^{-ci}\right) \quad = O(1)
\end{aligned}$$

when $k = 2^c$. Here the first inequality is an approximation of the Lebesgue integral of VPRO. The second and third inequalities are union bounds.

On the other hand, we have

$$
\begin{aligned}
\textsc{Pro} &\geqslant \frac{1}{2} \cdot \sum_{i=0}^{k} 2^i \sum_{j=0}^{k} \pi_j \cdot \mathop{\mathbf{Pr}}_{v_1,\ldots,v_n} [M_j \geqslant 2^i] \\
&\geqslant \frac{1}{2} \cdot \sum_{i=0}^{k} 2^i \sum_{j=0}^{k} 2^{-j} \cdot \min \left((1-e^{-1}), (1-e^{-1}) 2^{cj} \cdot \mathop{\mathbf{Pr}}_{v_1,\ldots,v_n} [v_1 \geqslant 2^i] \right) \\
&\geqslant \frac{1}{2} \cdot \sum_{i=0}^{k} 2^i \sum_{j=0}^{k} 2^{-j} \cdot \min \left((1-e^{-1}), (1-e^{-1}) 2^{cj-ci} \right) \\
&\geqslant \frac{1}{2} \cdot \sum_{i=0}^{k} 2^i 2^{-i} \cdot \min \left((1-e^{-1}), (1-e^{-1}) 2^{ci-ci} \right) \quad = \Omega(k).
\end{aligned}
$$

Here the first inequality is an approximation of the Lebesgue integral of Pro. The second inequality uses the fact that: if sum of the probabilities of several independent events is $p \leqslant 1$, then the union of them happens with probability at least $(1-e^{-1}) \cdot p$. As $n = 2^{ck} = 2^{k \log_2 k}$, we know $k = \Theta\left(\frac{\log n}{\log \log n}\right)$.

4

Pricing in Dynamic Two-sided Markets

The results of this chapter are published in [12], which is a joint work with Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, and Kamesh Munagala.

4.1 Introduction

The basic algorithmic challenge facing a marketplace platform can be summarized as follows: it must decide *which buyer should match to which seller, at what time, and for what price and wage*, in order to maximize some desired objective, such as the social surplus. In a sense, this combines the challenges of online bipartite matching, job scheduling, and pricing and mechanism design.

We start with the classic two-sided marketplace setting as in [125, 48, 43], where buyers and sellers are characterized by their *type* and there is a single service they are interested in trading. Buyers have values, and sellers have costs, and these are drawn from known type-dependent distributions. For instance, on a room-sharing platform, the type of a room can represent some combination of geographic location and quality of room; on a freelancing platform, the type of a consultant can capture their skill set and the type of a task can capture its requirements; and so on.

Suppose we now incorporate dynamics as follows. Buyers and sellers belong to discrete types, and agents of any type arrive at a steady `Poisson` rate. Buyers accept prices that are lower than their value, and sellers accept wages that are higher than their costs, and the prices and wages set for different types determine the rates of arrival. We further assume buyer and seller types are located in a metric space. Typically, the buyer would be happy as long as the seller a reasonable match with her requirements, so that there is an upper bound on the metric distance of any feasible match. We will use "type" and "node" interchangeably.

A policy for the platform is to (i) set a per-type price (for buyers) and wage (for sellers); and (ii) schedule feasible matches between the arriving buyers and sellers who accept the price/wage. The goal of this process is to maximize social surplus. Finding the optimal policy is not difficult – it will be a convex flow program with prices/wages set in a fashion to balance the rates of flow from buyers to sellers and maximize social surplus.

Scheduling Impatient Agents via Thick Markets. Now consider what happens when buyers and sellers have a *patience level*. They observe the price/wage, and accept/reject it. If they accept, the platform may not immediately be able to match them because a feasible match may not exist at that very moment. This will cause the buyers and sellers to wait. However, if they are not matched by the time their patience expires, they will abandon the system. The pricing and scheduling policy of the platform needs to make sure the abandonment probability is very small, since this causes dissatisfaction with the platform. Patience is a real constraint in two-sided platforms; for instance, Uber has recently introduced pricing for buyers with different patience levels [1].

Finding the optimal policy for scheduling matches under an abandonment constraint is a difficult stochastic control problem (see Section 4.5 for details). We therefore consider a simpler class of policies that allow us to analytically model basic tradeoffs the platform faces. These policies "pool" the arriving demand and supply at close-by nodes, which we call "facilities". Demand and supply is probabilistically routed to facilities, and agents routed to a single facility are matched optimally. For such policies to have abandonment probability below a threshold, we show in Section 4.5 that two natural conditions are sufficient: (i) the rate of arrival of supply and demand to a facility are equal; and (ii) this rate is at least a certain threshold, so that agents find feasible matches with high probability, *i.e.*, the facility is sufficiently *thick*. In essence, the pricing policy will ensure the time for finding a match is on average much smaller than the patience of any buyer or seller.

Remarks on the Model. We assume the set of different types and their embedding into a metric space is constructed by the platform based on features of the agents in a way that metric distance captures the quality of a match. However, these are not necessarily known to the agents. Hence we assume the value (or cost) of each agent is decoupled from the metric distance. For the same reason, agents cannot choose to arrive as a different type. Finally, note that the facilities are not physical locations, but rather canonical types. An agent can be assigned this canonical type if (s)he has a type that is within a distance threshold and our policy will be constrained to match agents such that the corresponding canonical types are the same. For instance, these could be canonical room types on a renting portal, or canonical skill sets on a freelancing platform. Hence we assume there are no facility costs.

4.1.1 High Level Problem Statement

One of our contributions (Section 4.5) is the reduction of a natural class of stochastic control policies for the *dynamic* pricing and scheduling problem described above to a *static* problem of locating facilities in the metric space so that these facilities are sufficiently thick. In the resulting facility location problem, buyers and sellers flow into nodes located in a metric space. The price (resp. wage) set at any buyer (resp. seller) node determines the volume of flow at that node. The platform needs to (i) open facilities in the metric space; (ii) assign prices and wages to each node; and (iii) probabilistically route the resulting flow to the open facilities so that the following service guarantees are satisfied:

1. **Quality of service guarantees**: The flow assigned to a facility is from supply and demand nodes within distance R. (This ensures any matched demand/supply is within distance $2R$.)

2. **Service availability guarantees**: Each facility needs to have *flow balance*, that is, equal amount of supply and demand is routed there. Further, there is a *lower bound* $\mathbf{L}$ on the flow routed to each facility, capturing the thickness constraint. As mentioned before, in Section 4.5, we show that these constraints arise from natural stochastic matching policies for dynamically arriving *impatient* agents.

The platform's objective is to maximize the *gains from trade* or *social surplus*, which is total value of buyers *minus* total cost of sellers, subject to weak *budget balance*, meaning the platform has non-negative total profit. We term this optimization problem as *Two-sided Facility Location.*

Such a facility location model enables us to abstract out the stochastic dynamics of scheduling with deadlines, and focus on studying the interplay between pricing and the service guarantee of finding good quality matches with high probability. Note that if we ignore the flow lower bound constraint, there is indeed a convex program for maximizing surplus. However, this may set prices and wages in such a fashion that there may not be enough nearby supply (resp. demand) to ensure thickness and prevent abandonment. On the other hand, if we do take the lower bounds into account, this may cause the prices (resp wages) to be set so low (resp. high) that the platform loses money. We illustrate this trade-off in an example in Section 4.3, which shows that with a flow lower bound constraint, the platform can obtain much larger surplus by losing money on one facility, and making up for this loss at other facilities.

In summary, there is a three-way tension between the goals of maximizing surplus, platform profitability, and the service guarantees described above. Our problem formulation of *Two-sided Facility Location* captures this trade-off, and understanding the computational complexity of this problem is the main focus of this chapter. Our model is sufficiently flexible to accommodate more complex extensions, and we present an extension to pricing patience in Section 4.6.

4.1.2 Our Results

Our main contribution in Sections 4.2–4.4 is to show an approximation algorithm for *Two Sided Facility Location.* We present a new LP rounding framework that for any constant $\epsilon > 0$, achieves a $(1+\epsilon)$ approximation to the social surplus objective. It relaxes the distance bound constraint by a factor of 4, while preserving the budget balance constraint, as well as the flow balance and lower bound constraints at each facility. If we allow a tiny additive error Δ in the surplus objective, our algorithm requires solving $O\left(\frac{n^{1/\epsilon}}{\epsilon}\log\frac{nW_{\max}}{\epsilon\Delta}\right)$ LPs, where n is the number of nodes, and $W_{\max}$ is the maximum possible surplus.

We show in Section 4.4.1 that the surplus objective is NP-Hard to approximate to a factor $o(\mathbf{L}^c)$ for some constant $c > 0$, unless the distance bound is relaxed by at least a factor of 2.

Techniques. Our facility location variants mirror the *profit earning facility location* problem in [120]. Just like that setting, we have lower bounds on demand served at each facility and an upper bound on how far the facility can be from an assigned demand. However, there are key differences that preclude the application of existing techniques from lower-balanced facility location [81, 88, 106, 120, 140]: First, the demand or supply at each node is a variable that can be adjusted using pricing. This means the demand/supply can be zero at some "outlier" nodes, so that they do not need to be served by any facility. Secondly, each facility needs to satisfy flow balance between supply and demand, and finally, both surplus and profit involve *differences*, so the platform can potentially lose money at some facilities, but recover it at others.

The above differences make formulating an LP relaxation tricky. Note that even in [120], the version with outliers and profits that can become negative has unbounded integrality gap, because the optimal profit can be zero while the LP achieves positive profit. Unlike [120], since we can control demand/supply by pricing, we have greater flexibility in modifying the LP variables. Despite this, the integrality gap of the straightforward LP formulation for our problem is large, because there could be a facility that generates a bulk of the surplus, but has large negative profit that is compensated by other facilities. (See the integrality gap example in Section 4.4.4).

This brings up our main technical contributions: We first observe that if we focus on the LP variables corresponding to a facility, we can scale these up or down by changing the fraction to which this node is an outlier. This enables us to use techniques reminiscent of improved greedy algorithms for budgeted coverage problems [111, 65]: In particular, we *strengthen the LP formulation* via guessing a few of the facilities that are opened in the optimal solution. Next, we use the guesses to develop a structural characterization for this stronger LP based on modifying variables for *pairs* of facilities. In effect this shows that there is some integrality in the neighborhood of any partially open facility, which helps us consolidate these facilities while preserving all constraints.

4.1.3 Related Work

Two-sided Markets. Our objective maximizes social surplus subject to budget balance (and individual rationality). This is a classic objective in two-sided market mechanisms, and originates in the celebrated work of Myerson and Satterthwaite [125], where it is termed *gains of trade*. They considered the case of a single buyer and seller. This has inspired a recent line of work on truthful mechanisms for approximate surplus maximization in markets of multiple buyers and sellers [119, 48, 43], ultimately resulting in a 2-approximation to gains of trade. This line of work assumes buyers and sellers are matched in one shot. The novelty in our work is in modeling a *dynamic setting* and incorporating *service availability* guarantees while preserving the same objectives. We therefore consider the more natural class of mechanisms that post prices and wages. Posted price mechanisms have been extensively studied in two-sided marketplaces [119, 130, 16, 147], and the main idea we borrow from this literature is the notion of *insulating tariffs* [147], which posits that market design is easier if the prices seen by buyers is disconnected from the wages seen by service providers.

Another recent line of work shows approximately optimal mechanisms for maximizing welfare in two sided markets with *goods* [59, 60]; however, theirs is a sum objective defined in terms of the final sets of items allocated to each buyer and seller, which is different from the gains of trade.

Dynamic Marketplaces. Our work on dynamic marketplaces is related to several recent works on online scheduling under stochastic arrivals of tasks on machines with limited resources [96, 55, 56]. Tasks have (private) types comprising their value, arrival time, and deadline; the platform's goal is to maximize welfare while truthfully eliciting the type. While similar to our work on pricing resources or tasks, they allow agents to choose assignments based on posted prices (*envy-freeness*). Another difference is the markets considered in their studies are one-sided.

Dynamic two-sided markets also serve as motivation for recent work on "online matching with delays" [76, 19]. Here, buyers and sellers arrive online in a metric space, and can be matched at any time subsequently. The goal is to minimize the total distance cost plus waiting cost, and the authors present a log-competitive algorithm. These models do not incorporate pricing. Further, our dynamic marketplace models are more closely related to dynamic matchings with *stochastic* arrivals, and we review this literature in Section 4.5.

4.2 Problem Statement

There is a metric space $G(V, E)$ with an associated distance function c. The two sided facility location problem is parametrized as TWO-SIDED FAC-LOC($\mathbf{L}, R$). Here, R is the distance threshold for assignment to a facility and $\mathbf{L}$ is the lower bound on flow routed to each open facility.

Each node $j \in V$ is associated with a demand function F_j and a supply function H_j. When offered price p, we assume the demand (*i.e.*, buyers) at node j is $d_j F_j(p)$, where $F_j(p)$ is a non-increasing function of p corresponding to the survival function of a continuous density function f_j on valuations; formally $F_j(p) = \int_{v=p}^{\infty} f_j(v)dv$. In other words, the volume of buyers is d_j, and when quoted a price p, only buyers with valuations at least p choose to participate. We assume there is a finite price $p_{\max}$ so that $F_j(p_{\max}) = 0$ for all $j \in V$.

Similarly, when offered wage w, the supply of sellers at node j is $s_j H_j(w)$, where $H_j(w)$ is a non-decreasing function of w, corresponding to the CDF of a continuous density function h_j on costs; formally $H_j(w) = \int_{c=0}^{w} h_j(c)dc$. When offered wage w, all sellers with cost at most w participate, resulting in supply $s_j H_j(w)$. We assume that $H_j(0) = 0$, *i.e.*, sellers accrue 0 utility by not participating in the platform.

Let $\mathcal{F} \subseteq V$ the set of all candidate facilities; we set $\mathcal{F} = V$. For each node j, $B_R(j) \subseteq \mathcal{F}$ denotes the set of all *compatible* facilities, i.e., $i \in \mathcal{F}$ such that $c(i, j) \leqslant R$. Similarly, for each facility i, we define $B_R(i)$ as the set of all *compatible* nodes, i.e., $j \in V$ such that $c(i, j) \leqslant R$. A solution to TWO-SIDED FAC-LOC($\mathbf{L}, R$) is specified by the following:

- An assignment of price p_j and wage w_j to each node $j \in V$. If the price (resp. wage) at node j is $p_{\max}$ (resp. 0), we assume this node generates no demand (resp. supply).
- A set of locations $S \subseteq \mathcal{F}$ for opening the facilities; and
- A routing scheme $\vec{x_j^d}$ (resp. $\vec{x_j^s}$) for each demand (resp. supply) node $j \in V$ that generates non-zero demand (resp. supply). For facility $i \in S$, if $x_{ij}^d > 0$ then $i \in B_R(j)$ (i.e., i is a compatible node for facility j). Further, $\sum_{i \in B_R(j)} x_{ij}^d = 1$ for all nodes $j \in V$ that generate non-zero demand; similarly, $\sum_{i \in B_R(j)} x_{ij}^s = 1$ for each $j \in V$ with non-zero supply.

Note that the flow of demand (resp. supply) from node $j \in V$ to facility $i \in S$ is $d_j F_j(p_j) x_{ij}^d$ (resp. $s_j H_j(w_j) x_{ij}^s$). As motivated in Section 4.1.1, the flows need to satisfy the *flow balance* and *flow lower bound* conditions at each facility $i \in S$:

Flow Balance. The total amount of supply and demand are equal,

Flow Lower Bound. The total amount of supply (resp. demand) routed there is at least $\mathbf{L}$.

Both these desiderata can be formalized via the following constraint:

$$\sum_j d_j F_j(p_j) x_{ij}^d = \sum_j s_j H_j(w_j) x_{ij}^s \geqslant \mathbf{L}, \text{ for all } i \in S.$$

Surplus (Gains from Trade) Objective. We first define the following quantities:

$$\mathcal{V}_j(p) = d_j \int_{v=p}^{\infty} v f_j(v) dv \qquad \text{and} \qquad \mathcal{C}_j(w) = s_j \int_{c=0}^{w} c h_j(c) dc$$

respectively denote the total value of buyers generated by node j when the price there is p and the cost of sellers at node j when the wage there is w. The surplus objective can then be written as:

$$\text{Social Surplus} = \sum_{j \in V} (\mathcal{V}_j(p_j) - \mathcal{C}_j(w_j))$$

The goal of the platform is to maximize social surplus subject to the platform profit being non-negative. This is termed weak budget balance, and is written as:

$$\text{Profit} = \sum_{j \in V} (d_j p_j F_j(p_j) - s_j w_j H_j(w_j)) \geqslant 0$$

Note that simply maximizing surplus may not guarantee non-negative profit because of the flow lower bound constraint. We illustrate this in an example in Section 4.3.

In this section, we make the standard regularity assumptions (*à la* Myerson-Satterthwaite [125]) on the density functions f_j and h_j. In particular, we assume $xF_j^{-1}(x)$ is concave in x and $yH_j^{-1}(y)$ is convex in y. This is true for instance, for all log-concave densities f_j and h_j, which includes Normal, Exponential, and Uniform distributions.

4.3 Tradeoff Between Surplus, Profit, and Thickness

We show that the surplus optimum solution can open a facility with negative profit. To be more specific, for any given constant $c < 1$ we present a simple example in which c fraction of the total surplus is generated by a facility with negative profit.

Let $\{v, v'\}$ be two nodes infinitely far apart. Let L be the lower bound for the total amount of demand (supply) at each open facility. For node v, assume that the volume of demand and supply are $d_v = s_v = L$, the valuation of buyers is uniformly distributed over the interval $[2, 3]$, and the cost of sellers is uniformly distributed over the interval $[0, 1]$. For node v', assume the volume of demand and supply are $d_{v'} = s_{v'} = L$, the valuation of buyers is uniformly distributed over the interval $[c' - 1, 2c' + 1]$, and the cost of sellers is uniformly distributed over the interval $[0, c']$ where $c' = \frac{2c}{1-c}$. We claim that the optimum solution for this example is to open a facility at each of the nodes and set the price and wage at node v to 2 and 1 respectively, and set the price and wage at node v' to $c' - 1$ and c' respectively.

First, we show that this solution is feasible. At each node the price is not more than the valuation of any arriving buyer. Therefore, all the buyers choose to participate. Similarly, since the wage is not less than the cost of any arriving seller, all the sellers choose to participate. This solution satisfies flow balance for each of the facilities because the volume of sellers and buyers are equal at the corresponding node, and all of them choose to participate. In addition, flow lower bound is also satisfied, since this volume is at least L. Finally, the profit of the facility at v is d_v and the loss of the facility at node v' is $d_{v'}$. Therefore, the total profit is 0 and profit of the facility at node v compensates for the loss at the other facility.

Now, we show that the welfare of the facility with negative profit is a fraction c of the total welfare. The welfare at node v is $d_v \times (2.5 - 0.5) = 2L$ and the welfare at node v' is $d_{v'} \times (3c'/2 - c'/2) = c'L$. Therefore, $c'/(c' + 2) = c$ fraction of the welfare is generated at node v'.

Finally, we need to show that this solution is optimum. The nodes are far from each other and we cannot send the buyers and sellers from different nodes to a common facility. The only option for opening a facility at each of the nodes is to set the price and wage at each node in a way that all the arriving buyers and sellers choose to participate (otherwise, the flow lower bound cannot be satisfied). Therefore, this problem has three feasible integral solutions: no facility is opened, a facility at node v

is opened, and a facility at each of the nodes is opened. Note that the solution which only opens a facility at v' is not feasible because it does not satisfy budget balance – in order to generate L volume of demand (supply), the platform must lose money here. The welfare of those solutions are 0, $2L$, and $(2 + c')L$ respectively. Therefore, the third solution is optimum.

4.4 Approximation Algorithm

We characterize the approximation ratio of any algorithm for Two-Sided FacLoc($\mathbf{L}, R$) as (α, γ), if the resulting solution relaxes the distance bound of an assignment to a facility to αR, ensures lower bound $\mathbf{L}$, and has surplus OPT/γ, where OPT is the optimal surplus. First, we show that it is NP-Hard to obtain $\gamma = o(\mathbf{L}^{\xi})$ for some constant $\xi > 0$, unless $\alpha \geqslant 2$. Subsequently, we present a $(4, 1 + \epsilon)$ approximation. For the algorithm to have polynomial running time, they also need lose a small additive amount in the objective; as we show later, this quantity can be exponentially small.

4.4.1 Hardness of Approximation

Theorem 20. *It is* NP-Hard *to find a* (α, γ) *approximation for* Two-Sided FacLoc($\mathbf{L}, R$) *unless* $\alpha \geqslant 2$ *or* $\gamma \geqslant \mathbf{L}^{\xi}$ *for some constant* $\xi > 0$*.*

Proof. We reduce from Maximum Independent Set in k-regular graphs (k-MIS). Given a k-MIS instance with n vertices and $m = kn/2$ edges, construct a metric space where each edge in the k-regular graph $G(V, E)$ has length $2R$. Place a demand node at the mid-point of each edge, and a supply node at each vertex. We set $\mathbf{L} = k$. Each supply node has $s_j = k$, and supply function $H^{-1}(r) = 1 - \delta$ for $r \in [0, 1]$. Similarly, each demand node has $d_j = 1$, and demand function $F^{-1}(q) = 1$ for $q \in [0, 1]$. Since the distance threshold is R, the facilities are opened at vertices of the graph. Each such facility must see k units of supply and demand, which means all neighboring demand is routed there, leading to welfare (resp. profit) $k\delta$ at that facility. Since two open facilities cannot share a demand, this means the open facilities form an independent set. Therefore, the surplus of Two-Sided Fac-Loc(k, R) is δ times the size of the maximum independent set in G. This is NP-hard to approximate to within a factor of k^{ξ} for some constant $\xi > 0$; see [14, 99]. Therefore, we need to relax the distance bound by at least a factor of 2. □

4.4.2 Linear Programming Relaxation

We now formulate Two-Sided Fac-Loc$(\mathbf{L}, R)$ as an integer linear program. For ease of exposition, we compare against an optimal solution that is restricted to using prices from a fixed set $\mathcal{P}$ and wages from a fixed set $\mathcal{W}$. Our solution is not restricted to using prices and wages from this set. In Section 4.4.3, we show that our LP admits to a polynomial time solution of arbitrary additive accuracy when this assumption is relaxed, and demand/supply distributions are continuous.

Note that we assume $p_{\max} \in \mathcal{P}$ and $0 \in \mathcal{W}$, and at this price (resp. wage) the demand (resp. supply) is identically zero. This is the price (resp. wage) where this node becomes an outlier and the solution is not required to open a facility nearby.

Instead of writing our LP using prices and wages, we use the associated demand/supply values. Let $\mathcal{Q}_j = \{q \mid q = F_j(p), p \in \mathcal{P}\}$ and $\mathcal{R}_j = \{r \mid r = H_j(w), w \in \mathcal{W}\}$. The case where the node is an outlier now corresponds to setting $q = 0$ (resp. $r = 0$). Moreover, the valuations/costs can also be redefined using supply/demand values as follows:

$$\mathcal{V}_j(q) = d_j \int_{v=F_j^{-1}(q)}^{\infty} v f_j(v) dv \qquad \text{and} \qquad \mathcal{C}_j(r) = s_j \int_{c=0}^{H_j^{-1}(r)} c h_j(c) dc \tag{4.1}$$

respectively denote the total value of buyers generated by node j when the price there is $F_j^{-1}(q)$, and the cost of sellers at node j when the wage there is $H_j^{-1}(r)$.

Variables. For each candidate facility $i \in \mathcal{F}$, let $y_i \in \{0, 1\}$ be the indicator variable that a facility is opened at that location in the metric space. Let $\alpha_{jq} = 1$ if the price at node $j \in V$ corresponds to $q \in \mathcal{Q}_j$. Similarly define β_{jr} for $r \in \mathcal{R}_j$. The variable z_{ijq} is non-zero only if $\alpha_{jq} = 1$ and $i \in B_R(j)$. In this case, it is the fraction of j's demand that is routed to i. We define z_{ijr} similarly for supply. Note that the actual flow from j to i is $d_j q z_{ijq}$; similarly for sellers.

Objective and Weak Budget Balance. The objective of social surplus and the profit being non-negative can be captured by:

- Surplus Objective:

$$\max \sum_{j \in V} \left(\sum_{q \in \mathcal{Q}_j} \alpha_{jq} \mathcal{V}_j(q) - \sum_{r \in \mathcal{R}_j} \beta_{jr} \mathcal{C}_j(r) \right) \tag{4.2}$$

- Weak Budget Balance:

$$\sum_{j \in V} \left(\sum_{q \in \mathcal{Q}_j} \alpha_{jq} d_j q F_j^{-1}(q) - \sum_{r \in \mathcal{R}_j} \beta_{jr} s_j r H_j^{-1}(r) \right) \geqslant 0 \tag{4.3}$$

Feasibility. The following constraints connect the variables together. We present these constraints only for buyers (that is, $q \in \mathcal{Q}_j$); the constraints for sellers is obtained by replacing q with $r \in \mathcal{R}_j$. First, for each $q \in \mathcal{Q}_j$, we need to choose one price for buyers (resp. sellers).

$$\sum_{q \in \mathcal{Q}_j} \alpha_{jq} = 1 \quad \forall j \in V \tag{4.4}$$

$$\sum_{i \in B_R(j)} z_{ijq} = \alpha_{jq} \quad \forall j \in V, q \in \mathcal{Q}_j \tag{4.5}$$

Next, if demand is fractionally routed from j to i, then i should be open and within distance R. Note that we need to ignore the case where $q = 0$ (resp. $r = 0$) since in this case, the demand (resp. supply) routed is zero, so that there is no need for a nearby facility.

$$\sum_{q \in \mathcal{Q}_j, q > 0} z_{ijq} \leqslant y_i \quad \forall j \in V, i \in B_R(j) \tag{4.6}$$

Service Availability. We finally encode *flow balance* and *flow lower bound* at each facility:

$$\sum_{j \in B_R(i)} d_j \sum_{q \in \mathcal{Q}_j} q z_{ijq} = \sum_{j \in B_R(i)} s_j \sum_{r \in \mathcal{R}_j} r z_{ijr} \quad \forall i \in \mathcal{F} \tag{4.7}$$

$$\sum_{j \in B_R(i)} d_j \sum_{q \in \mathcal{Q}_j} q z_{ijq} \geqslant \mathbf{L} y_i \quad \forall i \in \mathcal{F} \tag{4.8}$$

If we replace the integrality constraints on $\{y_i\}$ and the $\{\alpha_{jq}, \beta_{jr}\}$ with $y_i, \alpha_{jq}, \beta_{jr} \in [0,1]$, the above is a linear programming relaxation of the problem.

4.4.3 Solving the LP Formulation

We now show how to use the Ellipsoid algorithm to efficiently solve the LP formulation in the previous section to arbitrary additive accuracy even when the demand and supply distributions are continuous, so that the sets $\mathcal{Q}_j$ (resp. $\mathcal{R}_j$) are continuous. First we get rid of weak budget balance by take a Lagrangian of surplus and the profit. For any parameter $\lambda \geqslant 0$, define:

$$\mathcal{V}_j^\lambda(q) = \mathcal{V}_j(q) + \lambda d_j q F_j^{-1}(q)$$

and

$$\mathcal{C}_j^\lambda(r) = \mathcal{C}_j(r) + \lambda s_j r H_j^{-1}(r)$$

Since we assumed regular supply and demand distributions, it is easy to show that $\mathcal{V}_j^\lambda(q)$ is concave in q and $\mathcal{C}_j^\lambda(r)$ is convex in r. The Lagrangian objective is then:

$$\text{Maximize} \sum_{j \in V} \left(\sum_{q \in \mathcal{Q}_j} \sum_{i \in B_R(j)} z_{ijq} \mathcal{V}_j^\lambda(q) - \sum_{r \in \mathcal{R}_j} \sum_{i \in B_R(j)} z_{ijr} \mathcal{C}_j^\lambda(r) \right)$$

$$
\begin{array}{rcll}
\sum_{q \in \mathcal{Q}_j} \sum_{i \in \mathcal{F}} z_{ijq} & \leqslant & 1 & \forall j \in V \\
\sum_{r \in \mathcal{R}_j} \sum_{i \in \mathcal{F}} z_{ijr} & \leqslant & 1 & \forall j \in V \\
\sum_{q \in \mathcal{Q}_j} z_{ijq} & \leqslant & y_i & \forall j \in V, i \in B_R(j) \\
\sum_{q \in \mathcal{R}_j} z_{ijr} & \leqslant & y_i & \forall j \in V, i \in B_R(j) \\
\sum_{j \in B_R(i)} d_j \sum_{q \in \mathcal{Q}_j} q z_{ijq} & = & \sum_{j \in B_R(i)} s_j \sum_{r \in \mathcal{R}_j} r z_{ijr} & \forall i \in \mathcal{F} \\
\sum_{j \in B_R(i)} d_j \sum_{q \in \mathcal{Q}_j} q z_{ijq} & \geqslant & \mathbf{L} y_i & \forall i \in \mathcal{F} \\
z_{ijq}, z_{ijr}, y_i & \geqslant & 0 & \forall i, j, q, r
\end{array}
$$

The dual is the following:

$$
\text{Minimize} \sum_{j \in V} (a_j + b_j)
$$

$$
\begin{array}{rcll}
a_j + \eta_{ij} + d_j q(\zeta_i - \rho_i) & \geqslant & \mathcal{V}_j^{\lambda}(q) & \forall j \in V, i \in B_R(j), q \in \mathcal{Q}_j \\
b_j + \theta_{ij} - s_j r \zeta_i + \mathcal{C}_j^{\lambda}(r) & \geqslant & 0 & \forall j \in V, i \in B_R(j), r \in \mathcal{R}_j \\
\mathbf{L} \rho_i & \geqslant & \eta_{ij} + \theta_{ij} & \forall j \in V, i \in B_R(j) \\
\eta_{ij}, \theta_{ij}, \rho_i & \geqslant & 0 & \forall j \in V, i \in B_R(j)
\end{array}
$$

For fixed dual variables, since $\mathcal{V}_j^{\lambda}(q)$ is concave in q and $\mathcal{C}_j^{\lambda}(r)$ is convex in r, it is easy to check that for each i, j, the separation oracle either involves maximizing a concave function in q (for the first set of constraints) or minimizing a convex function in r (for the second set of constraints). In either case, finding the separating hyperplane involves one-dimensional convex optimization. This implies the LP admits to an efficient additive approximation even for continuous distributions over a bounded domain. We omit the standard details.

4.4.4 Integrality Gap and Stronger LP Formulation

The main technical hurdle arises because of the flow lower bound constraint: The LP optimum (and even an integer optimum) can now open facilities i which have positive surplus but negative profit, and compensate for the loss in profit by other facilities with positive profit. (See Section 4.3 for an example.) Note that Constraints (4.6) and (4.4) together imply:

$$\sum_{i \in B_R(j)} y_i \geqslant 1 - \alpha_{j0} \quad \forall \text{ Demand nodes } j$$

and similarly for supply nodes. We call the quantities α_{j0} (resp. β_{j0}) the *outlier fraction* of node j, and correspond to the case where the node is priced in such a way that it does not generate flow. In this case, there is no need to open a facility to satisfy j. Therefore, if $\alpha_{j0}\beta_{j0} > 0$, then the above constraints could imply $\sum_{i \in B_R(j)} y_i < 1$. This means there could only be a small fractional facility open in the vicinity of j, which can account for a lot of the surplus. This makes the LP have super-polynomial integrality gap and we present an example to show this.

Integrality Gap Example. Now we slightly modify the example in Section 4.3 to show that the LP has unbounded integrality gap. We only change the distribution of the valuation of the buyers at node v'. The valuation of the buyers is now uniformly distributed over the interval $[c' - 1 - \epsilon, 2c' + 1 + \epsilon]$ for a small positive constant ϵ. After this change, the integral solution which opens a facility at each node is not feasible anymore because it violates weak budget balance constraint. Therefore, the optimum integral solution has $2L$ welfare.

On the other hand we claim that there is a fractional solution which has $(\frac{1}{1+\epsilon} \times \frac{2c}{1-c} + 2)L$ welfare. Set the price and wage at node v to 2 and 1 and open the facility at that node ($y_v = 1$). For the node v' we can only open the facility partially. Set $y_{v'} = \frac{1}{1+\epsilon}$ and the price and wage at node v' to $c' - 1 - \epsilon$ and c' with probability $\frac{1}{1+\epsilon}$ and to p_{max} and 0 with probability $\frac{\epsilon}{1+\epsilon}$. In other words, set $\alpha_{v'1} = \beta_{v'1} = \frac{1}{1+\epsilon}$ and $\alpha_{v'0} = \beta_{v'0} = \frac{\epsilon}{1+\epsilon}$. This solution is feasible and generates $(\frac{1}{1+\epsilon} \times \frac{2c}{1-c} + 2)L$ welfare, while the optimum integer solution generates only $2L$ welfare. Note that c can be arbitrarily close to 1 and therefore the integrality gap is unbounded.

Our first technical contribution involves adding constraints to the above LP formulation to bound its integrality gap. Before showing how to strengthen the LP, we present the following easy claim, which implies that once we round $\{y_i\}$, the remaining solution can easily be made integral.

Lemma 15. *Given any feasible LP solution, there is an equivalent solution that assigns only one price (resp. wage) per demand (resp. supply) node, that preserves all constraints and does not decrease the objective.*

Proof. The rounding of α_{jq}, β_{jr} is simple. Let $\hat{q}_j = \sum_{q \in \mathcal{Q}_j} q\alpha_{jq}$ and $\hat{r}_j = \sum_{r \in \mathcal{R}_j} r\beta_{jr}$. Set the price of location j to be $F_j^{-1}(\hat{q}_j)$ and the wage at j to be $H_j^{-1}(\hat{r}_j)$. In other words, set $\hat{\alpha}_{j\hat{q}_j} \leftarrow 1$ and $\hat{\beta}_{j\hat{r}_j} \leftarrow 1$. Further set $\hat{z}_{ij\hat{q}_j} \leftarrow \sum_{q \in \mathcal{Q}'_j} z_{ijq} \frac{q}{\hat{q}_j}$ and $\hat{z}_{ij\hat{r}_j} \leftarrow \sum_{r \in \mathcal{R}'_j} z_{ijr} \frac{r}{\hat{r}_j}$.

Note that this process preserves the demand and supply from node j to facility i, which preserves all the constraints in the LP formulation. Note further that the function $qF_j^{-1}(q)$ is concave in q by the regularity of the demand function. Therefore,

$$\sum_q \alpha_{jq} q F_j^{-1}(q) \leqslant \left(\sum_q \alpha_{jq} q \right) F_j^{-1} \left(\sum_q \alpha_{jq} q \right) = \hat{q}_j F_j^{-1}(\hat{q}_j)$$

Similarly, since we assumed $rH_j^{-1}(r)$ is convex in r (by regularity of supply), we have $\sum_r \beta_{jr} r H_j^{-1}(r) \geqslant \hat{r}_j H_j^{-1}(\hat{r}_j)$. Therefore, this transformation preserves weak budget balance. Next, we note that $\mathcal{V}_j(q)$ is always a concave function of q and $\mathcal{C}_j(r)$ is always a convex function. Therefore, the above argument also implies the welfare (social surplus) does not decrease in the above transformation. $\square$

Define a variable for the surplus W_i and profit R_i of facility i respectively as:

$$W_i = \sum_{j \in B_R(i)} \left(\sum_{q \in \mathcal{Q}'_j} \mathcal{V}_j(q) z_{ijq} - \sum_{r \in \mathcal{R}'_j} \mathcal{C}_j(r) z_{ijr} \right) \tag{4.9}$$

$$R_i = \sum_{j \in B_R(i)} \left(\sum_{q \in \mathcal{Q}'_j} d_j q F_j^{-1}(q) z_{ijq} - \sum_{r \in \mathcal{R}'_j} s_j r H_j^{-1}(r) z_{ijr} \right) \tag{4.10}$$

Then the objective can be rewritten as: Maximize $\sum_i W_i$, and weak budget balance is $\sum_i R_i \geqslant 0$. Further note that $W_i \geqslant R_i$ since for any q, r, we have $\mathcal{V}_j(q) \geqslant d_j q F_j^{-1}(q)$, and $\mathcal{C}_j(r) \leqslant s_j r H_j^{-1}(r)$ if we integrate the expressions in Equation (4.1) by parts.

Stronger LP Formulation. Let $\epsilon > 0$ be any constant, and let $\theta = \frac{1}{\epsilon}$. We guess the θ facilities in the optimum solution that have the most surplus. There are two cases. First, if the optimum solution opens fewer than θ facilities, we can perform a brute force search over all integer solutions that open at most θ facilities. This can be done in $O(n^\theta)$ time, where $n = |V|$. For each selection of facilities, Lemma 15 implies that solving the LP formulation with the corresponding y_i set to 1 and the rest to zero yields the optimal surplus (or results in declaring infeasibility). We can therefore find the surplus maximizing solution among these in polynomial time; call this surplus W_1.

In the other case, the optimum solution opens more than θ facilities. In this case, for every choice of parameter $\mathcal{W} \geqslant 0$ scaled to powers of $(1+\epsilon)$, and every subset $S \subseteq \mathcal{F}$ with $|S| = \theta$, define $LP(\mathcal{W}, S)$ as having all of Constraints (4.3) – (4.8), *plus* the following new ones:

$$W_i \leqslant \mathcal{W} y_i \qquad \forall i \in \mathcal{F} \backslash S \tag{4.11}$$

$$y_i = 1 \qquad \forall i \in S \tag{4.12}$$

$$\sum_{i \in S} W_i \geqslant \mathcal{W} \theta (1 - \epsilon) \tag{4.13}$$

Let OPT denote the optimum surplus, and also define $W_2 = \max\{LP(\mathcal{W}, S) \mid \mathcal{W} \geqslant 0,\ S \text{ s.t. } |S| = \theta\}$. Then, it is easy to see that $OPT \leqslant \max(W_1, W_2)$: If OPT opens fewer than θ facilities, then clearly $W_1 \geqslant OPT$, since W_1 opens all possible choices of at most θ facilities. Otherwise, let W_i^* denote the surplus generated by open facility i in OPT. Let W^* denote the θ^{th} largest value of W_i^*. Choose $\mathcal{W} \in [W^*, W^*(1+\epsilon)]$, and S as the set of θ facilities in OPT with $W_i^* \geqslant W^*$. This induces a feasible solution to the above constraints, so that the LP optimum is at least OPT.

4.4.5 Structural Characterization of LP Optimum

Our second technical contribution is a new structural characterization about the LP optimum. This is crucial for the rounding that we present subsequently, since it allows sufficient mass of facility to be located in roughly the same neighborhood.

Recall α_{j0}, β_{j0} are the fractions to which node j is an outlier, *i.e.* has zero flow. These variables are the reason the simpler LP had large integrality gap, since they allow facilities in $B_R(j)$ to be open to small fractions. Our main observation is the following:

Lemma 16 (Structural Characterization)**.** *There is a* $(1+\epsilon)$ *approximation to the objective of* $LP(\mathcal{W}, S)$ *that satisfies:*

$$\forall i \in \mathcal{F}, \qquad y_i \in (0,1) \qquad \Rightarrow \qquad \exists j \in B_R(i) \ s.t. \ \alpha_{j0}\beta_{j0} = 0$$

High level Idea. Before presenting the proof, we present the high level idea. Consider a facility that violates the statement. If it has $R_i > 0$, then consider all LP variables $\{y_i, z_{ijq}, z_{ijr}\}$ corresponding to some such facility i and uniformly increase them. This increases both profit and surplus. We can decrease the fractions $\{\alpha_{j0}, \beta_{j0}\}$ to which any node j connected to i is assigned as outlier to compensate the fraction to which it is assigned i. Note that Constraints (4.7) and (4.8) are local to a single facility. Since we scale up *all* variables corresponding to a facility, we preserve these constraints. If we keep up this process, then either the facility is completely open ($y_i = 1$); or some demand/supply node assigned to it has $\alpha_{j0} = 0$ or $\beta_{j0} = 0$. (This must hold in the LP optimum.)

On the other hand, if $W_i > 0$ but $R_i < 0$, then increasing its LP variables would hurt profit, which may violate the budget balance constraint; while reducing the variables would increase profit but hurt the surplus. The idea now is the following: Take any pair of such facilities; increase the variables for one facility while decrease them for the other. There is always a way of doing this so that *both* the total profit and surplus do not decrease – this is essentially a fractional knapsack argument. Again, since we uniformly scale all variables corresponding to a facility, we preserve all constraints. Note that the process can also stop when a facility closes ($y_i = 0$). Eventually, we run out of pairs, so that for all but one facility, the above characterization holds.

At this point, the strengthened LP kicks in. The singleton facility violating the above lemma was fractionally open and had $R_i < 0$. It has surplus at most $\mathcal{W}$ by Constraint (4.11). But we have integrally open facilities that generate surplus at least $\mathcal{W}\frac{(1-\epsilon)}{\epsilon}$ by Constraint (4.13), which means closing the singleton facility reduces surplus by at most $(1-\epsilon)$, and preserves budget balance.

Proof of Lemma 16. We first simplify the LP. Let $\mathcal{Q}'_j = \mathcal{Q}_j \backslash \{0\}$ and $\mathcal{R}'_j = \mathcal{R}_j \backslash \{0\}$. Let $\eta_j = \sum_{q \in \mathcal{Q}'_j} \alpha_{jq}$ and $\phi_j = \sum_{r \in \mathcal{R}'_j} \beta_{jr}$ respectively denote the fractions to which j is assigned prices (wages) that correspond to non-zero demand (supply). We can rewrite the constraints (4.4) and (4.5) as:

$$\eta_j = \sum_{q \in \mathcal{Q}'_j} \sum_{i \in B_R(j)} z_{ijq} \leqslant 1 \qquad \text{and} \qquad \phi_j = \sum_{r \in \mathcal{R}'_j} \sum_{i \in B_R(j)} z_{ijr} \leqslant 1 \qquad \forall j \in V \qquad (4.14)$$

and set $\alpha_{j0} = 1 - \eta_j$, and $\beta_{j0} = 1 - \phi_j$. Recall from Equations (4.9) and (4.10) that W_i and R_i are respectively the surplus and profit of facility i in the LP optimum.

We call node j *fully demand-utilized* if $\eta_j = 1$, and *fully supply-utilized* if $\phi_j = 1$. We say that node j is *partially demand-connected* to facility $i \in \mathcal{F}$ if $\sum_{q \in \mathcal{Q}'_j} z_{ijq} > 0$, and *partially supply-connected* if $\sum_{r \in \mathcal{R}'_j} z_{ijr} > 0$. Let $J_D(i)$ denote the set of nodes that are partially demand-connected to $i \in \mathcal{F}$, and $J_S(i)$ be the set that is partially supply-connected.

Lemma 17. *In the LP optimum, for any $i \in \mathcal{F}$ with $y_i > 0$, we have $W_i > 0$. Furthermore, all except one facility with $y_i > 0$ satisfy the following condition: either $y_i = 1$; or there exists $j \in J_D(i)$, such that j is* fully demand-utilized*; or there exists $j \in J_S(i)$ such that j is* fully supply-utilized.

Proof. First, note that $W_i \geqslant R_i$. Suppose an open facility has $W_i \leqslant 0$. This implies $R_i \leqslant 0$. Consider a different solution that sets $y_i = 0$ and $z_{ijq} = z_{ijr} = 0$ for all $j \in V, q \in \mathcal{Q}'_j, r \in \mathcal{R}'_j$. We adjust η_j and ϕ_j for each $j \in V$ to preserve constraint (4.14). This new solution has $W_i = R_i = 0$ and has at least as large surplus and profit. Since we set all LP variables corresponding to i to zero, this satisfies constraints (4.6), (4.7), and (4.8), and is therefore feasible for the LP.

We therefore only focus on facilities whose $W_i > 0$. Consider the set of these facilities and split them into two groups. Let

$$S_1 = \{i \in \mathcal{F} \mid y_i \in (0,1) \text{ and } R_i < 0\} \qquad \text{and} \qquad S_2 = \{i \in \mathcal{F} \mid y_i \in (0,1) \text{ and } R_i \geqslant 0\}$$

Assume that for all of these facilities, there is no $j \in J_D(i)$, such that j is *fully demand-utilized* and no $j \in J_S(i)$ such that j is *fully supply-utilized*

First consider the facilities in set S_2, we can increase the LP variables till the condition of the lemma is satisfied; this process only increases both profit and surplus, preserving all constraints. We do this as follows: Suppose no $j \in J_D(i)$ is fully demand-utilized and no $j \in J_S(i)$ is fully supply-utilized. In this case, let

$$\theta = \min\left(\frac{1}{y_i}, \min_{j \in J_D(i)}\left(\frac{1 - \sum\limits_{q \in \mathcal{Q}'_j} \sum_{i' \neq i} z_{i'jq}}{\sum\limits_{q \in \mathcal{Q}'_j} z_{ijq}}\right), \min_{j \in J_S(i)}\left(\frac{1 - \sum\limits_{r \in \mathcal{R}'_j} \sum\limits_{i' \neq i} z_{i'jr}}{\sum\limits_{r \in \mathcal{R}'_j} z_{ijr}}\right)\right)$$

Since $\eta_j < 1$ for all $j \in J_D(i)$ and $\phi_j < 1$ for all $j \in J_S(i)$, we have $\theta > 1$. Suppose we increase y_i, z_{ijq} for all $j \in J_D(i), q \in \mathcal{Q}'_j$, and z_{ijr} for all $j \in J_S(i)$, $r \in \mathcal{R}'_j$ by a factor of θ. We will still have $\eta_j \leqslant 1$ for all $j \in J_D(i)$ and $\phi_j \leqslant 1$ for all $j \in J_S(i)$. However, either y_i or one of these values will become exactly 1. Note that since we scaled all LP variables corresponding to i by the same factor, this preserves constraints (4.6), (4.7), and (4.8). The surplus and profit of this facility increase by a factor $\theta > 1$, which contradicts the optimality of the LP solution. Therefore, the facilities in S_2 all have a neighboring j that is either fully demand-utilized or fully supply-utilized.

Next consider the facilities in set S_1. Suppose the condition in the lemma is not satisfied, so that there are two facilities i and i' with $y_i, y_{i'} \in (0, 1)$, and with no neighboring j that is either fully demand-utilized or fully supply-utilized. Suppose $W_i/|R_i| = a$ and $W_{i'}/|R_{i'}| = b$ with $a \geqslant b$. We multiply each LP variable corresponding to i by a factor of $(1 + \delta)$, and multiply each LP variable corresponding to i' by a factor of $\left(1 - \frac{W_i}{W_{i'}}\delta\right)$. Using the same argument as above, this process preserves the constraints that are specific to a facility, since all variables are changed by the same factor. The increase in surplus of facility i is δW_i, and the decrease in surplus of facility i' is δW_i, so the overall surplus is preserved. The decrease in profit of facility i is $|R_i|\delta$, and the increase in profit of facility i' is $|R_{i'}|\frac{W_i}{W_{i'}}\delta \geqslant |R_i|\delta$ by our assumption that $a \geqslant b$. Therefore, this process cannot decrease profit, hence all constraints are preserved. We choose δ as the smallest value that either makes facility i have $y_i = 1$ or one neighboring j either fully supply or demand utilized, or that sets the variables of facility i' to zero. In all cases, the size of set S_1 reduces by one. We repeat this process till there is only one facility in S_1, completing the proof. □

The following corollary now restates Lemma 16, completing its proof.

Corollary 5. *There is a $(1+\epsilon)$-approximation to the LP optimum where any facility with $y_i > 0$ satisfies the following condition: either $y_i = 1$; or there exists $j \in J_D(i)$, such that j is* fully demand-utilized*; or there exists $j \in J_S(i)$ such that j is* fully supply-utilized*.*

Proof. By Constraint (4.12), there is a set of facilities S that are fully open (*i.e.*,

$y_i = 1$) and $\sum_{i \in S} W_i \geqslant \mathcal{W}\frac{1-\epsilon}{\epsilon}$ by Constraint (4.13). The rounding in Lemma 17 does not touch these facilities, since we only increase/decrease variables corresponding to partially open facilities (*i.e.*, those with $y_i \in (0,1)$). Lemma 17 implies there is only facility i that violates the condition of the corollary. This facility must have surplus $W_i \leqslant \mathcal{W} y_i \leqslant \mathcal{W}$ by Constraint (4.11). This means closing this facility (setting all its associated variables to zero) reduces the LP optimum by at most a factor of $(1-\epsilon)$. Since the previous lemma implies this facility had $R_i < 0$, this means closing it only increases profit, preserving weak budget balance. □

4.4.6 Rounding the LP Relaxation

The rounding now follows approaches similar to those in [81, 149]. We first present the high-level idea. Note that if a node j has $\alpha_{j0} = 0$ or $\beta_{j0} = 0$, then Constraint (4.6) implies $\sum_{i \in B_R(j)} y_i \geqslant 1$. Consider an *independent set* of such nodes, such that no two are fractionally assigned to the same facility. For any j in this set, move all partially open facilities in $B_R(j)$ to j itself, so that there is a facility integrally opened at j. Since we move an entire facility, we preserve all flows, so that flow balance and lower bound are preserved, and so is profit. Now a demand/supply can be assigned a distance $2R$ away, and the opened facilities are integral.

At this point, consider any fractionally open facility i. It must have a node j adjacent to it that satisfies the condition in Lemma 16. If j has a facility completely open at its location, then move i to location j. Otherwise, j was not part of the independent set in the previous step, which means j and j' shared a fractionally open facility, and the previous step opened a facility completely at j'. In this case, we move i to j', again preserving all flows. This means any demand/supply moves distance at most $4R$, preserving all the LP constraints.

Rounding Facilities

We now present the rounding algorithm in detail. Initially, all facilities $i \in \mathcal{F}$ with $y_i > 0$ are *partially open*. Node $j \in V$ is *untouched* if for all i such that $j \in J_D(i) \cup J_S(i)$, the facility i is partially open. Let U be the set of *untouched* nodes, and let Z be the set that is either fully demand-utilized or fully supply-utilized. Let $U_f = U \cap Z$.

Phase 1. Consider any $j \in U_f$. W.l.o.g., assume $\eta_j = 1$; the case where $\phi_j = 1$ is symmetric. Let $N(j) = \{i | j \in J_D(i)\}$. For every $i \in N(j)$, we "move" i to location j; call the new facility at location j as i^*. This means we set

- $\bar{y}_{i^*} \leftarrow \sum_{i \in N(j)} y_i$ and $\bar{y}_i \leftarrow 0 \qquad \forall i \in N(j)$;
- $\bar{z}_{i^*j'q} \leftarrow \sum_{i \in N(j)} z_{ij'q}$ and $\bar{z}_{ij'q} \leftarrow 0 \qquad \forall j' \in V, q \in \mathcal{Q}_{j'}, i \in N(j)$
- $\bar{z}_{i^*j'r} \leftarrow \sum_{i \in N(j)} z_{ij'r}$ and $\bar{z}_{ij'r} \leftarrow 0 \qquad \forall j' \in V, r \in \mathcal{R}_{j'}, i \in N(j)$

From constraint (4.6), and the fact that $\eta_j = 1$, we have:

$$\bar{y}_{i^*} = \sum_{i \in N(j)} y_i \geqslant \sum_{i \in B_R(j)} \sum_{q \in \mathcal{Q}'_j} z_{ijq} = \eta_j = 1$$

Subsequently, we mark every $i \in N(j)$ as *closed*, and mark i^* as *completely open*. Furthermore, we mark every j' that was reassigned in the above steps as *touched*.

Note that in the last three steps, any agents at a node j' that was initially assigned to $i \in N(j)$ is now assigned to i^*. Since each of the distances $j' \to i$ and $i \to j$ is at most R, the distance from j' to i^* is at most $2R$. Therefore, this step relaxes the distance of a feasible assignment to a facility from R to $2R$.

This process trivially preserves the objective and weak budget balance, as well as constraints (4.14). Moreover, constraints (4.6) and (4.7) are satisfied since we add both sides of the constraints corresponding to $i \in N(j)$ to obtain the constraint for i^*. Finally, to see that (4.8) is satisfied for i^*, note that

$$\sum_{j' \in J_D(i*)} \sum_{q \in \mathcal{Q}'_{j'}} d_{j'} q \bar{z}_{i*j'q} = \sum_{i \in N(j)} \sum_{j' \in J_D(i)} \sum_{q \in \mathcal{Q}'_{j'}} d_{j'} q z_{ij'q} \geqslant \mathbf{L} \sum_{i \in N(j)} y_i \geqslant \mathbf{L}$$

We continue this process, finding a node $j \in U_f$, and merging all facilities in $N(j)$ to one location. At the end of this process, the set U_f is empty.

Phase 2. At the end of Phase 1, each node j which is *touched* (including all *fully utilized* nodes) route some fraction of their demand (or supply) to at least one facility that is completely open. However, there could still be partially open facilities with $y_i \in (0,1)$ to which demand and supply are assigned. Consider these partially open facilities in arbitrary order. Suppose we are considering facility i and there exists *touched* and *fully utilized* node j such that $\sum_{q \in \mathcal{Q}'_j} z_{ijq} > 0$ (resp. $\sum_{r \in \mathcal{R}'_j} z_{ijr} > 0$). Consider the completely open facility i^* such that $\sum_{q \in \mathcal{Q}'_j} \bar{z}_{i*jq} > 0$ or $\sum_{r \in \mathcal{R}'_j} \bar{z}_{i*jr} > 0$. We move the facility i to location i^*, updating the variables just as in Phase 1; *i.e.*, we set

- $\bar{y}_{i*} \leftarrow \bar{y}_{i*} + y_i$ and $\bar{y}_i \leftarrow 0$;
- $\bar{z}_{i*j'q} \leftarrow \bar{z}_{i*j'q} + z_{ij'q}$ and $\bar{z}_{ij'q} \leftarrow 0 \qquad \forall j' \in V, q \in \mathcal{Q}_{j'}$
- $\bar{z}_{i*j'r} \leftarrow \bar{z}_{i*j'r} + z_{ij'r}$ and $\bar{z}_{ij'r} \leftarrow 0 \qquad \forall j' \in V, r \in \mathcal{R}_{j'}$

The argument that all constraints are preserved follows just as before. For any j' that was partially assigned to i, the new assignment is to i^*. This distance is at most

$$c(j', i^*) \leqslant c(j', i) + c(i, j) + c(j, i^*) \leqslant R + R + 2R = 4R$$

where we note that the distance $j \to i^*$ was at most $2R$ because j was potentially reassigned to i^* in Phase 1. We mark all nodes j' that are reassigned in this process as *touched*.

At the end of this process, suppose there are still partially open facilities with $y_i \in (0,1)$. By Corollary 5, each of these facilities i must have some $j \in Z$ partially assigned to it. At the end of Phase 1, we have the invariant that $j \notin U_f$, since U_f is empty. This means j was touched on Phase 1. But in that case, i must have been reassigned in Phase 2, which is a contradiction. Therefore, at this point, all facilities are either closed ($\bar{y}_i = 0$) or completely open ($\bar{y}_i \geqslant 1$). Furthermore, for any variable $\bar{z}_{ijq} > 0$ (resp. $\bar{z}_{ijr} > 0$), the facility i is completely open; the distance from j to i is at most $4R$; for each completely open facility, the rate of supply equals the rate of demand (Constraint (4.7)), and finally, the total flow is at least $\mathbf{L}$ (Constraint (4.8)).

Final Steps and Running Time

At this point, the facilities are opened integrally. Lemma 15 now implies that we can choose one price/wage per node preserving all constraints and the objective. This completes the proof of the following theorem.

Theorem 21. *For any constant $\epsilon > 0$, there is a $(4, 1+\epsilon)$ approximation algorithm for* TWO-SIDED FAC-LOC$(\mathbf{L}, R)$.

Note that the surplus can become arbitrarily close to zero. Therefore, for parameter $\Delta > 0$, we will allow additive error Δ in the surplus objective. Note that the maximum possible surplus is $W_{\max} = (\sum_j d_j)p_{\max}$, which is an upper bound on $\mathcal{W}$. If we assume the surplus is at least Δ, then $\max_i W_i^* \geqslant \Delta/n$. Since the top $1/\epsilon$ facilities on OPT have surplus $\mathcal{W}(1-\epsilon)/\epsilon$, this means we can set $\mathcal{W} \geqslant \frac{\epsilon\Delta}{2n}$ Therefore, the number of choices of $\mathcal{W}$ is $O\left(\frac{1}{\epsilon}\log\frac{nW_{\max}}{\epsilon\Delta}\right)$. For each choice of $\mathcal{W}$, we need to solve $O(n^{1/\epsilon})$ LPs, so that the overall number of LPs is $O\left(\frac{n^{1/\epsilon}}{\epsilon}\log\frac{nW_{\max}}{\epsilon\Delta}\right)$. Note that Δ can be exponentially small, and our algorithm for solving the LP in Section 4.4.3 will lose such an additive factor in the objective anyway. Omitting details, we note that if the objective is profit, we can achieve optimal objective by directly rounding the *single LP* in Section 4.4.2 (details similar to Section 4.6).

4.5 Queueing-theoretic Justification: Dynamic Marketplaces

In the facility location model discussed above, we imposed a lower bound $\mathbf{L}$ on the flow routed to any facility. We now present a *dynamic marketplace* model that provides queueing-theoretic justification for these constraints. This model has the following features:

- Buyer and seller types are located in a metric space just as before.
- Buyers at node j arrive as a `Poisson` process with rate $d_j F_j(p)$ when quoted price p, and when quoted a wage w; similarly, sellers follow a `Poisson` process with rate $s_j H_j(w)$.
- We assume each buyer and seller has a private patience level or deadline; if not matched within their deadline, they abandon the system. The platform knows the patience distribution.

The stochastic control problem that we term *dynamic marketplace problem* can be summarized by two control decisions:

1. *Pricing decision.* Choose static prices p_j and wages w_j at each node $j \in V$; and
2. *Scheduling decision.* This matches feasible buyer-seller pairs and removes them from the system. This decision is dynamic, depending on the entire state of the system as captured by the number of unmatched buyers and sellers at different nodes at any point of time.

The goal is to design a stochastic control policy that maximizes the long-term average surplus subject to long-term budget balance. We insist all scheduled matches must involve a current buyer and seller with metric distance at most R. The key difference is in the *service availability guarantee*: Given the stochastic nature of our arrivals, there is always some probability that an incoming buyer or seller exhausts her patience before being matched. A more realistic goal is to design policies that guarantee a minimum level of service availability. We quantify this via the long-term average probability of abandonment of agents. Formally, given a parameter $\epsilon > 0$ as input, the goal of the platform is to make the abandonment probability at most ϵ.

Scheduling Policies. Constructing the optimal policy for the dynamic marketplace problem is closely related to several lines of work in *dynamic matchings* over a compatibility graph – in kidney exchanges [9] where patients abandon the exchange if their health fails; in control of matching queues for housing allocation [51]; and more generally in service system design [94, 7, 6], wherein customers and servers arrive stochastically and are matched according to a compatibility graph. In all these models, the choice of whom to match an arriving agent to depends on the entire set of agents waiting at different nodes, leading to the "curse of dimensionality".

Given this curse of dimensionality, we consider the restricted sub-class of matching policies where the platform creates facilities in the metric space, and uses each facility to cater to a different set of mutually-compatible agent types. Arriving agents are randomly routed to a compatible facility, where they are queued up to be matched to agents on the other side. The probabilistic routing is fixed over time, and does not depend on the state of the facilities.

Each facility maintains a queue of active buyers and sellers that have been assigned there, ignores what location they came from, and matches them up using an optimal scheduling policy for minimizing abandonment rate using only the current state of that particular queue. We will enforce the constraint that for any facility, the long-term abandonment probability is at most ϵ, which in turn will ensure the overall abandonment probability is at most ϵ.

Now assume buyer deadlines are distributed as $\texttt{Exponential}(\kappa)$, and seller deadlines are distributed as $\texttt{Exponential}(\gamma)$. Though these distributions are known, the scheduling decisions at any facility are made without knowing the patience level of any individual agent. Then, any work-conserving policy (including FIFO) is optimal. If an agent's deadline expires and there is no agent to match it with in the queue, this agent is considered abandoned. We build on results from queueing theory [146] to bound this abandonment rate tightly as follows:

Theorem 22. *Suppose λ and μ be the (Poisson) arrival rates of buyers and sellers into a facility. Assume buyer deadlines are distributed as $\texttt{Exponential}(\kappa)$, and seller deadlines are distributed as $\texttt{Exponential}(\gamma)$. Then the FIFO policy has abandonment rate at most ϵ when:*

1. *There is* flow balance*, that is, $\lambda = \mu$; and*
2. *There is a* flow lower bound*, that is, $\lambda \geqslant \frac{3}{2}\left(\frac{\min(\gamma,\kappa)}{\epsilon^2}\right)$.*

Proof. The behavior of a facility is captured via the following birth-death Markov chain:

consider the state-space $\{\ldots, s(2), s(1), 0, b(1), b(2), \ldots\}$, where $0 = s(0) = b(0)$ denotes the state that the facility is empty, while for any $n \geqslant 1$, the state $b(n)$ denotes that there are n buyers queued up, and state $s(n)$ denote that there are n sellers queued up. For any $n \geqslant 1$, the transition rate from $b(n)$ to $b(n+1)$ is λ, and for $s(n)$ to $s(n+1)$ is μ; on the other hand, the rate of transition from $s(n)$ to $s(n-1)$ is $n\gamma + \lambda$, while from $b(n)$ to $b(n-1)$ is $n\kappa + \mu$. Here, the term $n\gamma$ corresponds to the rate of abandonment of sellers as their deadlines expires, and similarly $n\kappa$ is the rate of abandonment of buyers.

Assume that $\lambda = \mu$. Let q_0 denote the steady state probability of the queue being empty. Let

$$\mathbf{Pr}[\text{State } = s(n)] = \alpha_n \qquad \mathbf{Pr}[\text{State } = b(n)] = \beta_n$$

where $\alpha_0 = \beta_0 = q_0$. We have the following balance equations:

$$\alpha_n n\gamma = \lambda(\alpha_{n-1} - \alpha_n) \qquad \text{and} \qquad \beta_n n\kappa = \lambda(\beta_{n-1} - \beta_n)$$

Adding these equations, we have: $\sum_{n=1}^{\infty} (\alpha_n n\gamma + \beta_n n\kappa) = 2\lambda q_0$. Note that the LHS here is the total abandonment rate, and since the total arrival rate of agents (buyers and sellers) is 2λ, this means the abandonment probability is exactly q_0.

Since $\alpha_n = \frac{\lambda}{\lambda+n\gamma}\alpha_{n-1}$ and $\beta_n = \frac{\lambda}{\lambda+n\kappa}\beta_{n-1}$, we have by telescoping:

$$\alpha_n = q_0 \frac{\lambda^n}{\prod_{j=1}^{n}(\lambda + j\gamma)} = q_0 \prod_{j=1}^{n} \frac{1}{\left(1 + j\frac{\gamma}{\lambda}\right)}$$

$$\beta_n = q_0 \frac{\lambda^n}{\prod_{j=1}^{n}(\lambda + j\kappa)} = q_0 \prod_{j=1}^{n} \frac{1}{\left(1 + j\frac{\kappa}{\lambda}\right)}$$

Since these probability values sum to one, this implies

$$\frac{1}{q_0} = 1 + \sum_{n=1}^{\infty} \left(\prod_{j=1}^{n} \frac{1}{\left(1 + j\frac{\gamma}{\lambda}\right)} + \prod_{j=1}^{n} \frac{1}{\left(1 + j\frac{\kappa}{\lambda}\right)} \right) \tag{4.15}$$

For given κ and γ, this is an increasing function of λ. Therefore, $q_0 \leqslant \epsilon$ translates to a bound of the form $\lambda \geqslant \mathbf{L}$. An upper bound $\mathbf{L}_\epsilon$ on $\mathbf{L}$ can be computed as follows. Let $c = \frac{\min(\gamma,\kappa)}{\lambda}$. Then,

$$\begin{aligned}
\frac{1}{q_0} &\geqslant 1 + \sum_{n=1}^{\infty} \left(\prod_{j=1}^{n} \frac{1}{\left(1 + j\frac{\min(\gamma,\kappa)}{\lambda}\right)} \right) \\
&\geqslant \sum_{n=0}^{\infty} e^{-cn^2/2} - e^{-c} \geqslant \int_0^{\infty} e^{-cx^2/2} dx - e^{-c} \\
&= \sqrt{\frac{\pi}{2c}} - e^{-c} \geqslant \sqrt{\frac{2}{3c}}
\end{aligned}$$

where the second inequality uses $1 + x \leqslant e^x$ for all $x \geqslant 0$. Therefore, if we insist $\sqrt{\frac{2\lambda}{3\min(\gamma,\kappa)}} \geqslant \frac{1}{\epsilon}$, this ensures the abandonment probability is at most ϵ. This translates to the following lower bound:

$$\lambda \geqslant \mathbf{L}_\epsilon = \frac{3}{2} \cdot \frac{\min(\gamma,\kappa)}{\epsilon^2}$$

□

We have therefore shown that a sufficient condition for bounding the abandonment probability at any facility by ϵ reduces to saying the flow to the facility is balanced, and facility is *thick* – there is a lower bound $\mathbf{L} = \frac{3}{2}\left(\frac{\min(\gamma,\kappa)}{\epsilon^2}\right)$ on how much demand or supply needs to be routed there. This reduces the *dynamic* control problem with atomic agents to a *static* problem where demand/supply are fluid – exactly TWO-SIDED FAC-LOC($\mathbf{L}, R$) for suitable $\mathbf{L}$ that depends on ϵ.

4.6 Envy-Free Pricing and Profit Maximization

The idea from Section 4.2 of independently scaling up/down LP variables corresponding to individual facilities is fairly general, and leads naturally to approximation algorithms for more complex variants that are motivated by different scheduling policies for the dynamic marketplace problem. In this section, we present one such formulation that generalizes the model discussed in Section 4.2. In Section 4.6.2, we show that this model corresponds to the dynamic marketplace setting when the platform uses prices to elicit patience of agents, and uses Earliest Deadline First (EDF) scheduling in each virtual market.

We assume each node (type) j of buyer/seller has a collection of subtypes $\mathcal{S}_j$. There is a DAG $G_j(\mathcal{S}_j, E_j)$ on $\mathcal{S}_j$ that captures *envy.* If there is an edge $(k_1, k_2) \in E_j$, then sub-type k_1 envies sub-type k_2. The platform announces a price (resp. wage) p_{jk} (resp. w_{jk}) for each sub-type $k \in \mathcal{S}_j$. In order to preserve incentive compatibility, we require that if $(k_1, k_2) \in E_j$, then $p_{jk_1} \leqslant p_{jk_2}$; resp. $w_{jk_1} \geqslant w_{jk_2}$. This prevents an agent of sub-type k_1 from reporting its type to be k_2. Note that since the graph G_j is a DAG, such a price (resp. wage) assignment is feasible. We term such an assignment of prices (resp. wages) at each j as a *price (resp. wage) ladder.*

As before, there is a non-increasing demand function $d_{jk}F_{jk}(p_{jk})$ for each buyer sub-type $k \in \mathcal{S}_j$, and a non-decreasing supply function $s_{jk}H_{jk}(w_{jk})$ for each seller sub-type $k \in \mathcal{S}_j$. Each sub-type $k \in \mathcal{S}_j$ is also associated with a weight $\mathcal{G}_{jk}$. The platform learns which sub-type any agent chooses.

Lottery Pricing and Assignment. The platform opens a set of "virtual markets". For each node j and $k \in \mathcal{S}_j$, buyers (resp. sellers) arriving at the node and choosing that type are probabilistically routed to virtual markets which are within distance R from the node. We assume the platform shows a lottery over price (resp. wage) ladders as follows: For each node $j \in V$ the platform maintains a distribution $\mathcal{Z}_j$ of virtual markets within distance bound R, and for each virtual market in this set, it maintains a distribution $\mathcal{L}_{ij}$ of price (resp. wage) ladders. Given an agent arriving at this node, the platform first chooses a market i from $\mathcal{Z}_j$, and then a ladder from $\mathcal{L}_{ij}$ and shows it to the agent. After the agent chooses the price or wage (hence revealing its sub-type), she is routed to market i. We note that the routing policy makes the market the agent is routed to be independent of the sub-type elicited. Though this assumption is somewhat restrictive, it prevents the agent from choosing a sub-type to optimize for the virtual market they get assigned to.

Service Availability Guarantee. As before, we capture service availability by ensuring that each virtual market i has balanced supply and demand, and is also sufficiently *thick*. However, we now capture thickness by a lower bound $\mathbf{L}$ on the *total weight* of the sub-types assigned there. Formally, let x_{ijk} denote the expected flow of sub-type $k \in \mathcal{S}_j$ to virtual market i.

Flow Balance. The expected amount of supply and demand assigned there are the same.

Weight Lower Bound. The expected weight assigned there is large: $\sum_{j \in V, k \in \mathcal{S}_j} \mathcal{G}_{jk} x_{ijk} \geqslant \mathbf{L}$.

The objective is to maximize the expected profit of the solution. We term this problem Envy-Free FL($\mathbf{L}, R$). We note that similar ideas can be used to maximize other objectives; we present the profit objective for simplicity.

In the dynamic marketplace setting presented in Section 4.6.2, the sub-types correspond to different deadlines, and the weight of a sub-type is precisely the deadline value. We show there that the weight lower bound corresponds to the condition for the EDF scheduling policy to have low abandonment rate.

4.6.1 Approximation Algorithm

Our LP formulation and rounding are similar to the one for Two-Sided FacLoc($\mathbf{L}, R$), and we highlight the differences. As before, we assume there is a candidate set $\mathcal{P}$ and $\mathcal{W}$ of prices and wages for each node, respectively. The set of all candidate virtual markets in the metric space is denoted by $\mathcal{F}$; since we assume the metric space is explicitly specified as input, we set $\mathcal{F} = V$. For each node j, $B_R(j) \subseteq \mathcal{F}$ denotes the set of all the virtual markets $i \in \mathcal{F}$ such that $c(i,j) \leqslant R$. For each virtual market i, define $B_R(i)$ as the set of all the nodes $j \in V$ such that $c(i,j) \leqslant R$.

Linear Programming Relaxation

Variables. For each candidate virtual market $i \in \mathcal{F}$, let $y_i \in \{0,1\}$ be the indicator variable that a virtual market exists at that location in the metric space. These are the only integer variables in our formulation. Variables x^d_{ij} and x^s_{ij} are non-zero only if $y_i = 1$ and $i \in B_R(j)$. In this case, those are respectively the probability that buyers and sellers at node j are routed to virtual market i. Note that there is some probability that all prices at node j are set to $p_{\max}$, which corresponds to not routing node j anywhere. Let z_{ijkp} be the probability that buyers at node j with sub-type $k \in \mathcal{S}_j$ are assigned to virtual market i and offered price p. Similarly, z_{ijkw} denotes the probability that sellers at node j with sub-type k are assigned to virtual market

i and offered wage w.

Objective and Constraints. The objective is to maximize the profit.

$$\max \sum_{j\in V}\sum_{k\in\mathcal{S}_j}\sum_{i\in B_R(j)}\left(\sum_{p\in\mathcal{P}} p d_{jk} F_{jk}(p) z_{ijkp} - \sum_{w\in\mathcal{W}} w s_{jk} H_{jk}(w) z_{ijkw}\right) \tag{4.16}$$

The following constraints connect the variables together. We present these constraints only for buyers (that is, $p \in \mathcal{P}$); the constraints for sellers is obtained by replacing p and x^d with $w \in \mathcal{W}$ and x^s. Since we route the buyers at node j probabilistically to one of the virtual markets, or to no market by offering all deadlines a price $p_{\max}$:

$$\sum_{i\in B_R(j)} x^d_{ij} \leqslant 1 \quad \forall j \in V \tag{4.17}$$

Next, a price should be offered to each buyer with sub-type k at node j assigned to market i:

$$\sum_{p\in\mathcal{P}} z_{ijkp} = x^d_{ij} \quad \forall j \in V, i \in B_R(j), k \in \mathcal{S}_j \tag{4.18}$$

Next, if demand is fractionally routed from j to i, then i should be open and within distance R:

$$x^d_{ij} \leqslant y_i \quad \forall j \in V, i \in B_R(j) \tag{4.19}$$

We next enforce that the prices and wages form a distribution over ladders. Note that the policy first chooses the virtual market to route to, and then chooses from a distribution over ladders. This reduces to a stochastic dominance condition for the distributions corresponding to z:

$$\sum_{p'\leqslant p, p'\in\mathcal{P}} z_{ijkp'} \leqslant \sum_{p'\leqslant p, p'\in\mathcal{P}} z_{ijk'p'} \quad \forall p \in \mathcal{P}, (k,k') \in E_j, \forall j \in V, \forall i \in B_R(j) \tag{4.20}$$

$$\sum_{w'\leqslant w, w'\in\mathcal{W}} z_{ijkw'} \geqslant \sum_{w'\leqslant w, w'\in\mathcal{W}} z_{ijk'w'} \quad \forall w \in \mathcal{W}, (k,k') \in E_j, \forall j \in V, \forall i \in B_R(j) \tag{4.21}$$

Finally, we encode the service availability constraints. We first capture *flow balance* at each virtual market: the rate of arrival of buyers and sellers are equal.

$$\sum_{j \in B_R(i)} \sum_{k \in \mathcal{S}_j, p \in \mathcal{P}} d_{jk} F_{jk}(p) z_{ijkp} = \sum_{j \in B_R(i)} \sum_{k \in \mathcal{S}_j, w \in \mathcal{W}} s_{jk} H_{jk}(w) z_{ijkw} \quad \forall i \in \mathcal{F} \tag{4.22}$$

We finally encode *weighted flow lower bound* on the total deadline of buyers and sellers at the market:

$$\sum_{j \in B_R(i)} \sum_{k \in \mathcal{S}_j} \mathcal{G}_{jk} \left(\sum_{p \in \mathcal{P}} d_{jk} F_{jk}(p) z_{ijkp} + \sum_{w \in \mathcal{W}} s_{jk} H_{jk}(w) z_{ijkw} \right) \geqslant \mathbf{L} y_i \quad \forall i \in \mathcal{F} \tag{4.23}$$

Rounding

If we ignore the integrality constraints on y_i, the above is a linear programming relaxation of the problem. We will now show how to round the resulting solution.

We generalize Lemma 17 using the following definitions of *fully utilized.* We say that $j \in V$ is *fully demand utilized* if $\sum_{i \in B_R(j)} x_{ij}^d = 1$; similarly, it is *fully supply-utilized* if $\sum_{i \in B_R(j)} x_{ij}^s = 1$. We say j is *partially demand-connected* to market $i \in \mathcal{F}$ if $x_{ij}^d > 0$, and *partially supply-connected* if $x_{ij}^s > 0$. Let $J_D(i)$ denote the set of nodes that are partially demand-connected to $i \in \mathcal{F}$, and $J_S(i)$ be the set that is partially supply-connected. As before, we define the profit of a virtual market $i \in \mathcal{F}$ as:

$$R_i = \sum_{j \in B_R(i)} \sum_{k \in \mathcal{S}_j} \left(\sum_{p \in \mathcal{P}} p d_{jk} F_{jk}(p) z_{ijkp} - \sum_{w \in \mathcal{W}} w s_{jk} H_{jk}(w) z_{ijkw} \right)$$

Lemma 18. *In the LP optimum, for any $i \in \mathcal{F}$, $R_i \geqslant 0$. Further, if $R_i > 0$, either $y_i = 1$; or there exists $j \in J_D(i)$, s.t. j is* fully demand-utilized*; or there exists $j \in J_S(i)$, s.t. j is* fully supply-utilized*.*

The proof of the above lemma follows the same argument as Lemma 17: If a market has negative R_i, we can set all its variables to zero without violating any constraints. If the condition in the Lemma is violated for $i \in \mathcal{F}$, then we can increase all variables corresponding to i by the same factor till the condition is satisfied. Since all constraints involve single markets, this process preserves them while increasing the objective. For this transformation to work, it is crucial Constraints (4.20) are defined separately for each (i, j) pair; in other words, we crucially need to assume the policy chooses a market first and then chooses a distribution over ladders for that market.

The rounding now proceeds in the same way as in Section 4.2: In Phase 1, we identify untouched and *fully utilized* j and merge all i to which it is partially connected to one market. Note that the total y_i of these markets is at least 1 by the LP constraints. At the end of this phase, we move the remaining partially open i as in Phase 2 of the rounding scheme. This preserves the profit, and satisfies the flow balance and lower bound constraints (B_R is replaced by B_{4R} in the constraints), yielding the following theorem:

Theorem 23. *There is a feasible solution $\{\bar{x}, \bar{y}, \bar{z}\}$ to the above linear program, whose objective is optimal, and all of whose constraints are satisfied. For each $i \in \mathcal{F}$, either $\bar{y}_i = 0$ or $\bar{y}_i \geqslant 1$.*

Final Policy. The final choice of prices and wages, and the routing policy is the following. We present it only for buyers; the policy for sellers is symmetric.

- At node j, choose a market i with probability $\bar{x}_{ij}$. If no market is chosen, the price is set to $p_{\max}$.
- If market i is chosen, choose α uniformly at random in $[0, 1]$. For each $k \in \mathcal{S}_j$, find $p_k \in \mathcal{P}$ such that $\sum_{p' < p_k, p' \in \mathcal{P}} \frac{\bar{z}_{jkp'}}{\bar{x}_{ij}} \leqslant \alpha$ and $\sum_{p' \leqslant p_k, p' \in \mathcal{P}} \frac{\bar{z}_{jkp'}}{\bar{x}_{ij}} > \alpha$. Post prices $\{p_1, p_2, \ldots, p_K\}$.

- If the buyer accepts price p_k, route her to virtual market i.

Constraints (4.20) imply that regardless of the choice of α, the prices $\{p_1, p_2, \dots, p_K\}$ in the second step form a ladder, so that $p_1 \geqslant p_2 \geqslant \cdots \geqslant p_K$. A similar statement holds for wages. Therefore, the second step produces a lottery over ladders. Further, if Z_{ijkp} denote the event that the price for sub-type $k \in \mathcal{S}_j$ is p and market i is chosen, then it is an easy exercise to check that $\mathbf{E}[Z_{ijkp}] = \bar{z}_{jkp}$. Therefore, the randomized policy exactly implements the solution found in Theorem 23, so that it maximizes profit. We omit the details and state the final theorem.

Theorem 24. *There is a polynomial time* $(4, 1)$ *approximation for* ENVY-FREE FL$(\mathbf{L}, R)$. *That is, we obtain the optimal expected profit by relaxing the distance constraint by a factor of* 4.

4.6.2 Justification via Dynamic Marketplace Model

In this section, we present a dynamic marketplace model that justifies the problem statement of ENVY-FREE FL. As in Section 4.5, we assume buyers and sellers have an inherent patience level or deadline. If they are not matched within their deadline, they drop out of the system. We assume every agent m is associated with a patience level ν_m; unlike Section 4.5, we do not assume these are Exponentially distributed. The platform advertises a fixed set of patience levels, or deadlines, denoted by $\mathcal{S}_j = \{\nu_{j1}, \nu_{j2}, \dots, \nu_{jK}\}$, which is a guarantee on the time by which a buyer or seller choosing that deadline is guaranteed to be matched. We assume $\nu_{j1} \leqslant \nu_{j2} \leqslant \cdots \leqslant \nu_{jK}$. For simplicity, we use $k \in \mathcal{S}_j$ and $\nu_k \in \mathcal{S}_j$ interchangeably.

Incentive-compatibility. We assume the platform sets a lottery of prices and wages at each node j, that are independent of time. Consider the issue of eliciting deadlines truthfully. Consider buyers first. At node j, suppose the platform offers price p_{jk} for deadline ν_{jk}. Every buyer can choose one deadline in $\mathcal{S}_j$, in which case he pays price p_{jk}, and is guaranteed to be matched within time ν_{jk} from his arrival. We assume any buyer m has very large negative utility for being matched after his patience level ν_m, therefore he will choose a k such that $\nu_{jk} \leqslant \nu_m$. Subject to this, he will choose k with smallest p_{jk}, since this maximizes his valuation minus price. A symmetric model can be posited for sellers, where we replace price with wage, and the seller chooses the largest wage such that the corresponding deadline is smaller than his own patience level.

Since the goal of the platform is to elicit patience levels truthfully, the platform chooses a *price ladder* $p_{j1} \geqslant p_{j2} \geqslant \cdots \geqslant p_{jK}$ and *wage ladder* $w_{j1} \leqslant w_{j2} \leqslant \cdots \leqslant w_{jK}$ at each node j. This ensures that agents with $\nu_m \in [\nu_{jk}, \nu_{jk+1}]$ report deadline ν_{jk}.

Each deadline level $\nu_{jk} \in \mathcal{S}_j$ gets associated with non-increasing demand function $d_{jk}F_{jk}(p_{jk})$, which is the Poisson rate at which buyers m with patience $\nu_m \in [\nu_{jk}, \nu_{jk+1}]$ arrive when the price of deadline ν_{jk} is p_{jk}. Similarly, deadline level $k \in \mathcal{S}_j$ is associated with a non-decreasing supply function $s_{jk}H_{jk}(w_{jk})$, which is the Poisson rate at which sellers m with patience $\nu_m \in [\nu_{jk}, \nu_{jk+1}]$ arrive when the wage for deadline ν_{jk} is w_{jk}. These deadline levels correspond to the sub-types described before.

Scheduling Policy. As in Section 4.5, the platform opens a set of "virtual markets". For each node j and deadline level k, buyers (resp. sellers) arriving at the node and choosing that deadline are probabilistically routed to virtual markets which are within distance R from the node. Buyers and sellers arriving at the virtual market are queued up, and optimally matched to minimize abandonment. Since the platform knows which deadline was chosen by the agent, the optimal matching policy is now a variant of Earliest Deadline First (EDF): When the deadline of some buyer (resp. seller) expires, it is matched to that seller (resp. buyer) in the queue whose deadline will expire earliest in the future. If an agent's deadline expires and there is no agent to match it with in the queue, this agent is abandoned. It is an easy exercise to show that this policy maximizes the number of matches made in any virtual market.

As in Section 4.5, the goal of the platform is to design a joint pricing and scheduling policy to maximize profit, while ensuring bounded match distance and bounded abandonment probability.

Bounding Abandonment Rate

We will now show that the weight lower bound can be interpreted as a sufficient condition for the abandonment rate of the EDF policy to be at most ϵ, where the weight of a sub-type is simply its deadline value.

The main technical assumption we require in this part is that the desired abandonment probability, ϵ is small, in particular that $\epsilon \ll \frac{\nu_{\min}}{\nu_{\max}}$. As noted above, the scheduling policy within a virtual market is a variant of EDF. Unlike the PATIENCE-OBLIVIOUS model where the behavior of a virtual market could be modeled as a variant of a $M/M/1$ queue, the optimal abandonment probability in a two-sided EDF queue clearly depends on the entire distribution of deadlines of buyers and sellers, which in turn depends on the pricing scheme and assignment policy. However, we crucially need a closed-form bound on this probability in order to plug into an LP relaxation for the overall problem. We use recent results from queueing due to Kruk *et al.* [117] to construct such a closed-form bound, whose very existence we find non-trivial and surprising!

Kruk *et al.* [117] present an approximation to the abandonment probability of a one-sided queue $M/M/1$ queue with EDF scheduling. They approximate the queueing process via a reflected Brownian motion. We adapt their result to our setting, and rephrase it below. Consider the queue associated with a virtual market. Let $\bar{S}$

denote the average deadline of a seller arriving to this queue, and $\bar{D}$ denote the average deadline of a buyer arriving to the queue. Note that the distribution of deadlines as well as arrival rate depends on the overall pricing and assignment policies.

Recall that we assumed ϵ is small, in particular that $\epsilon \ll \frac{\nu_{\min}}{\nu_{\max}}$. We first enforce that supply and demand arrive to the queue at the same rate; call this rate λ. Next suppose w.l.o.g. that $\bar{D} > \bar{S}$. Consider the policy that instantaneously matches arriving sellers to the queued buyer with earliest deadline; if the queue is empty, the seller is abandoned. This exactly mimics a one-sided $M/M/1$ queue with EDF scheduling. We quote the following result informally from [117]:

> Consider a one sided $M/M/1$ queue with arrival rate and service rate equal to λ. Suppose deadlines of jobs are independently distributed with mean $\bar{D}$, and the scheduler uses the EDF policy. Then holding λ and $\frac{\nu_{\max}}{\nu_{\min}}$ fixed, in the regime where $\nu_{\min}$ becomes very large, the abandonment probability approaches $\frac{1}{\lambda\bar{D}}$.

Though part of their argument is heuristic, they perform simulations to show that this approximation is indeed accurate. Since we need abandonment probability of $\frac{1}{\lambda\bar{D}}$ to be at most $\epsilon \ll 1$, and since we assumed $\frac{\nu_{\min}}{\nu_{\max}} \gg \epsilon$, this automatically enforces that all deadlines are much larger than the mean inter-arrival time, satisfying their precondition for our setting.

Since the optimal policy for a two-sided queue only has lower abandonment probability, we use $\frac{1}{\lambda\bar{D}}$ as an upper bound on this quantity. Since we assumed $\bar{D} \geqslant \bar{S}$, we will instead use $\frac{2}{\lambda(\bar{D}+\bar{S})}$ as the upper bound, which we will set to be at most ϵ.

We now show that this is the best possible upper bound that only depends on $\bar{D}$ and $\bar{S}$. Suppose buyers deadlines are deterministic with value $\bar{D}$, and seller deadlines are deterministic with value $\bar{S}$. Then the optimal policy matches without waiting in a FIFO fashion. This means the loss probability assuming the queue has buyers is the same as that of a $M/M/1$ queue with deadlines $\bar{D}$, which from [44] is exactly

$$P_1 = \frac{1}{1 + \lambda \bar{D}}$$

Similarly, when there are sellers in the queue, the loss probability is

$$P_2 = \frac{1}{1 + \lambda \bar{S}}$$

Conditioned on the queue being empty and a buyer arriving, the expected time after which the queue next becomes empty is $T_b = \frac{1}{P_1} = 1 + \lambda \bar{D}$, in which period the loss probability is P_1. Similarly, if a seller arrives when the queue is empty, the expected time after which the queue again becomes empty is $T_s = \frac{1}{P_2} = 1 + \lambda \bar{D}$, in which period the loss probability is P_2. Since a buyer or seller arrives with equal probability when a queue is empty, the expected loss probability is

$$P = \frac{T_b P_1 + T_s P_2}{T_b + T_s} = \frac{2}{2 + \lambda \bar{D} + \lambda \bar{S}} \approx \frac{2}{\lambda \left(\bar{D} + \bar{S}\right)}$$

assuming $\lambda(\bar{D} + \bar{S}) \gg 1$.

In summary, each virtual market needs to satisfy the following two sufficient conditions for its abandonment probability to be at most ϵ:

1. The rate of arrival of supply and demand should be the same; call this rate λ.

2. If $\bar{S}$ denote the average deadline of a seller, and $\bar{D}$ denote the average deadline of a buyer, then $\lambda(\bar{D} + \bar{S}) \geqslant \frac{2}{\epsilon}$.

Therefore, to reduce the scheduling policy to an instance of Envy-Free FL$(R, \mathbf{L})$, we set $\mathcal{G}_{jk} = \nu_{jk}$ and $\mathbf{L} = \frac{2}{\epsilon}$, so that the second condition above translates to the weight lower bound. This justifies the Envy-Free FL$(R, \mathbf{L})$ problem as capturing the optimal scheduling policy for the dynamic marketplace problem presented above.

Note that the resulting lower bound on λ derived by the above condition is a significant improvement over the patience-oblivious case, since the lower bound now depends on $\frac{1}{\epsilon}$ instead of $\frac{1}{\epsilon^2}$. This intuitively means that in order to achieve comparable profit and abandonment probability, we can aim for a higher quality of match by reducing the radius R. A similar observation that even partial information about deadlines significantly reduces abandonment is made in [9], albeit for a different model.

4.7 Conclusion

Our work is a first step in understanding the problem of jointly pricing and scheduling in dynamic matching facilities. We now mention several open questions that arise. For the facility location problems, there are other variants that we do not yet have good algorithms for, for instance, general policies for the envy-free version where the routing can be correlated with the sub-type elicited. Further, our model imposes a uniform distance bound the match of any agent. Extending it to average match distance will require new techniques; the basic filtering step in facility location rounding fails in our case since the demand value itself is a variable. Finally, it is an interesting question to extend our techniques to when facilities can be priced, and agents choose facilities to optimize their utility, in particular, extending techniques for stochastic scheduling in one-sided facilities [55, 56].

5

Conclusion

We have studied different optimization problems that a platform as an intermediary might need to solve. The main objective of our problems is to maximize welfare. However, the goals of other agents (buyers and sellers) are not necessarily aligned with that objective, and each agent only wants to maximize its utility. Therefore, in all of these problems, we considered the strategic behavior of these agents. In this section, we briefly summarize the problems that we studied and some future directions.

Information intermediary in auctions. First, we considered a Bayesian single item auction with independent buyers and an additional intermediary who knows the true values of the buyers. We studied the problem of designing a signaling scheme by the intermediary to maximize the consumer surplus in the presence of a revenue-maximizing seller. We considered a benchmark for maximum consumer surplus and presented schemes with provable grantees against this benchmark. We also provided lower bounds for any signaling scheme. In this problem, we assumed the seller always runs an ex-interim IC and IR mechanism that maximizes her revenue. One future direction is to consider the setting in which the seller needs to run ex-post IC and IR mechanisms. It would be interesting to see if we can still find lower bounds for this setting.

Going beyond our specific setting, it would be interesting to explore the equilibria in optimal auctions when the intermediary can send different signals to the seller and to the buyers, much like in [34, 139]. At an even higher level, our work can be considered a special case of a larger problem of information intermediaries for multi-agent mechanisms. As mentioned before, in our case, the optimal auction is the mechanism, and the intermediary can change the information to this mechanism in order to achieve "fairness" between producer and consumer surplus. It would be interesting to explore the question of achieving fairness by selectively regulating information to a black-box optimizer or mechanism in more general settings.

Prophet inequalities with uncertain supply. In the second problem, we added supply uncertainty to the multiple-choice prophet inequality. In our setting, each item might disappear after an a priori unknown amount of time that we called horizon. The mechanism only knows the horizon distribution for each item. We showed that if the horizon distributions satisfy MHR condition, there is a constant competitive pricing scheme. Unlike the classic multi-choice prophet inequalities where the approximation ratio goes to 1 when the number of items becomes large, we showed a 1.57 (improves to 2 when the number of items becomes large) approximation lower bound even when the horizons are *i.i.d.* geometric. Our constant is tight for the single-item setting.

We now list several open questions. First, our constant factor for the upper bound in the multi-item setting does not match the lower bound, and the gap is pretty large. Closing the gap would be an interesting open question. Next, is it possible to have stochastic horizons in more general prophet-inequality settings such as [69]?

An important future direction is to capture a more general form of supply uncertainty. In other words, it would be interesting to extend our work to the case where items arrive and depart in a stochastic fashion.

Pricing in dynamic two-sided markets. Finally, we considered a dynamic marketplace setting where agents arrive according to a stochastic process and have finite patience (or deadlines) for being matched. The goal here is to maximize social surplus subject to budget balance while guaranteeing a high match quality and low abandonment probability. We reduce a special type of scheduling policy that uses facilities to our two-sided facility location problems, and we design approximation algorithms for the two-sided facility location problems. The main future direction is to approximate the overall optimal policy?

Our model here also assumed Poisson arrivals whose rate is constant over time. A different approach is to use online algorithms. In particular, it would be interesting to incorporate pricing and wages into the “online matching with delays" models considered in [19, 76].

Bibliography

[1] Uber is experimenting with letting riders wait longer in exchange for cheaper fares. *Quartz Magazine*, June 2018.

[2] Melika Abolhassani, Mohammad Hossein Bateni, MohammadTaghi Hajiaghayi, Hamid Mahini, and Anshul Sawant. Network cournot competition. In *Web and Internet Economics (WINE)*, pages 15–29, 2014.

[3] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi HajiAghayi, Robert Kleinberg, and Brendan Lucier. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 61–71. ACM, 2017.

[4] I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 3–16, 2012.

[5] Zoe Abrams, Adam Meyerson, Kamesh Munagala, and Serge Plotkin. The integrality gap of capacitated facility location. *CMU CS Technical Report CMU-CS-02-199*, 2002.

[6] Ivo Adan, Ana Busic, Jean Mairesse, and Gideon Weiss. Reversibility and further properties of fcfs infinite bipartite matching. *arXiv preprint arXiv:1507.05939*, 2015.

[7] Ivo Adan and Gideon Weiss. Exact FCFS matching rates for two infinite multitype sequences. *Operations research*, 60(2):475–489, 2012.

[8] A.A. Ageev and M.I. Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, Sep 2004.

[9] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Dynamic matching market design. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, page 355, 2014.

[10] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[11] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35. ACM, 2012.

[12] Reza Alijani, Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, and Kamesh Munagala. The segmentation-thickness tradeoff in online marketplaces. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(1):1–26, 2019.

[13] Reza Alijani, Siddhartha Banerjee, Sreenivas Gollapudi, Kamesh Munagala, and Kangning Wang. Predict and match: Prophet inequalities with uncertain supply. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–23, 2020.

[14] Noga Alon, Uriel Feige, Avi Wigderson, and David Zuckerman. Derandomized graph products. *Comput. Complex.*, 5(1):60–75, January 1995.

[15] Hyung-Chan An, Mohit Singh, and Ola Svensson. LP-based algorithms for capacitated facility location. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 256–265, 2014.

[16] Mark Armstrong. Competition in two-ï£¡ï£¡sided markets. *RAND Journal of Economics*, 37(3):668–691, 2006.

[17] Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in decentralized matching markets. 2014. Draft available at SSRN 2427960.

[18] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: A meta algorithm and applications. *Theory of Computing*, 8:121–164, 2005.

[19] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Polylogarithmic bounds on the competitiveness of min-cost perfect matching with delays. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1051–1061, 2017.

[20] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the 1-1/e barrier. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 303–318. ACM, 2018.

[21] Yossi Azar, Arun Ganesh, Rong Ge, and Debmalya Panigrahi. Online service with delay. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 551–563, 2017.

[22] E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Available at SSRN 2260567*, 2014.

[23] Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandrs Slivkins. Dynamic pricing with limited supply. *ACM Transactions on Economics and Computation*, 3(1):4, 2015.

[24] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 2005.

[25] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Multi-objective pricing for shared vehicle systems. *arXiv preprint arXiv:1608.06819*, 2016.

[26] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 517–517. ACM, 2017.

[27] Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, and Kamesh Munagala. Segmenting two-sided markets. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, pages 63–72, 2017.

[28] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 639–639, 2015.

[29] Siddhartha Banerjee, Ramesh Johari, and Zhengyuan Zhou. The importance of exploration in online marketplaces. *IEEE Internet Computing*, 20(1):20–26, 2016.

[30] Siddhartha Banerjee, Zhengyuan Zhou, and Ramesh Johari. The importance of exploration in online marketplaces. In *53rd IEEE Conference on Decision and Control*, pages 3499–3504. IEEE, 2014.

[31] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012.

[32] Richard E Barlow and Albert W Marshall. Tables of bounds for distributions with monotone hazard rate. *Journal of the American Statistical Association*, 60(311):872–890, 1965.

[33] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

[34] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1):107–143, 2017.

[35] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[36] Dirk Bergemann and Martin Pesendorfer. Information structures in optimal auctions. *Journal of economic theory*, 137(1):580–609, 2007.

[37] Dirk Bergemann and Karl H Schlag. Pricing without priors. *Journal of the European Economic Association*, 6(2-3):560–569, 2008.

[38] Omar Besbes and Ilan Lobel. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, 61(1):92–110, 2015.

[39] Omar Besbes and Ilan Lobel. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, 61(1):92–110, 2015.

[40] Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.

[41] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilkiliç. Cournot competition in networked markets. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 733–733, 2014.

[42] Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '02, pages 971–980, 2002.

[43] Liad Blumrosen and Yehonatan Mizrahi. Approximating gains-from-trade in bilateral trading. In *Proceedings of the 12th International Conference on Web and Internet Economics*, WINE 2016, pages 400–413, 2016.

[44] Nam Kyoo Boots and Henk Tijms. A multiserver queueing system with impatient customers. *Management Science*, 45(3):444–448, 1999.

[45] Christian Borgs, Ozan Candogan, Jennifer Chayes, Ilan Lobel, and Hamid Nazerzadeh. Optimal multiperiod pricing with service guarantees. *Manage. Sci.*, 60(7), 2014.

[46] S. Bose, D. W. H. Cai, S. H. Low, and A. Wierman. The role of a market maker in networked cournot competition. *CoRR*, abs/1403.7286, 2014.

[47] Shelby L Brumelle and Jeffrey I McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41(1):127–137, 1993.

[48] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. *CoRR*, abs/1706.04637, 2017.

[49] Yang Cai, Federico Echenique, Hu Fu, Katrina Ligett, Adam Wierman, and Juba Ziani. Third-party data providers ruin simple mechanisms. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–31, 2020.

[50] Bernard Caillaud and Bruno Jullien. Chicken & Egg: Competition among Intermediation Service Providers. *RAND Journal of Economics*, 34(2):309–28, 2003.

[51] René Caldentey, Edward H Kaplan, and Gideon Weiss. Fcfs infinite bipartite matching of servers and customers. *Advances in Applied Probability*, 41(3):695–730, 2009.

[52] René Caldentey, Ying Liu, and Ilan Lobel. Intertemporal pricing under minimax regret. *Available at SSRN 2357083*, 2015.

[53] Deeparnab Chakrabarty and Gagan Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM Journal on Computing*, 39(6):2189–2211, 2010.

[54] Archishman Chakraborty and Rick Harbaugh. Persuasive puffery. *Marketing Science*, 33(3):382–400, 2014.

[55] Shuchi Chawla, Nikhil R. Devanur, Alexander E. Holroyd, Anna R. Karlin, James Martin, and Balasubramanian Sivan. Stability of service under time-of-use pricing. In *ACM STOC*, 2017.

[56] Shuchi Chawla, Nikhil R. Devanur, Janardhan Kulkarni, and Rad Niazadeh. Truth and regret in online scheduling. In *ACM EC*, 2017.

[57] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 311–320. ACM, 2010.

[58] Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra. Approximating matches made in heaven. In *36th International Colloquium on Automata, Languages, and Programming*, pages 266–278, 2009.

[59] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1424–1443, 2016.

[60] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 591–608, 2017.

[61] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 169–186. ACM, 2017.

[62] Jose Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1946–1961. SIAM, 2019.

[63] Jacques Crémer and Richard P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

[64] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. Algorithmic price discrimination. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2432–2451. SIAM, 2020.

[65] Brian C. Dean, Michel X. Goemans, and Jan Vondák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.

[66] Nikhil R Devanur and Jason D Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 41–50. ACM, 2009.

[67] Peter A Diamond. Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy*, 90(5):881–94, 1982.

[68] Peter A Diamond. Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy*, 90(5):881–94, 1982.

[69] Paul Düetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 540–551. IEEE, 2017.

[70] Shaddin Dughmi. On the hardness of signaling. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 354–363. IEEE, 2014.

[71] Shaddin Dughmi. Algorithmic information structure design: a survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.

[72] Shaddin Dughmi, David Kempe, and Ruixin Qiang. Persuasion with limited communication. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 663–680, 2016.

[73] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. *SIAM Journal on Computing*, (0):STOC16–68, 2019.

[74] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 700–714. SIAM, 2018.

[75] Edith Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pages 736–745. SIAM, 2007.

[76] Yuval Emek, Shay Kutten, and Roger Wattenhofer. Online matching: haste makes waste! In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 333–344, 2016.

[77] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.

[78] U. Feige, M. Feldman, N. Immorlica, R. Izsak, B. Lucier, and V. Syrgkanis. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 872–878, 2015.

[79] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM Journal on Computing*, 39(1):122–142, 2009.

[80] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 123–135. SIAM, 2014.

[81] Zachary Friggstad, Mohsen Rezapour, and Mohammad R. Salavatipour. Approximating connected facility location with lower and upper bounds via LP rounding. In *15th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2016, June 22-24, 2016, Reykjavik, Iceland*, pages 1:1–1:14, 2016.

[82] Hu Fu, Christopher Liaw, Pinyan Lu, and Zhihao Gavin Tang. The value of information concealment. In *Proceedings of the 29^{th} Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, page 2533âĂŞ2544, 2018.

[83] Guillermo Gallego and Garrett Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science*, 40(8):999–1020, 1994.

[84] Alex Gershkov, Benny Moldovanu, and Philipp Strack. Revenue maximizing mechanisms with strategic customers and unknown demand: name-your-own-price. *Available at SSRN 2527653*, 2014.

[85] Atiyeh Ashari Ghomi, Allan Borodin, and Omer Lev. Seasonal goods and spoiled milk: Pricing for a limited shelf-life. *arXiv preprint arXiv:1801.02263*, 2018.

[86] Yiannis Giannakopoulos, Elias Koutsoupias, and Philip Lazos. Online market intermediation. *CoRR*, abs/1703.09279, 2017.

[87] Renato Gomes and Alessandro Pavan. Price discrimination in many-to-many matching markets. Technical report, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 2011.

[88] Sudipto Guha, Adam Meyerson, and Kamesh Munagala. Hierarchical placement and network design problems. In *41st Annual Symposium on Foundations of Computer Science, FOCS*, pages 603–612, 2000.

[89] Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 104–113, 2007.

[90] Sudipto Guha, Kamesh Munagala, and Peng Shi. Approximation algorithms for restless bandit problems. *J. ACM*, 58(1):3:1–3:50, 2010.

[91] Anupam Gupta, Ravishankar Krishnaswamy, Marco Molinaro, and R. Ravi. Approximation algorithms for correlated knapsacks and non-martingale bandits. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 827–836, 2011.

[92] Varun Gupta and Ana Radovanovic. Online stochastic bin packing. *arXiv preprint arXiv:1211.2687*, 2012.

[93] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1164–1173. Society for Industrial and Applied Mathematics, 2005.

[94] Itai Gurvich, Amy Ward, et al. On the dynamic control of matching queues. *Stochastic Systems*, 4(2):479–523, 2014.

[95] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1*, AAAI'07, pages 58–65, 2007.

[96] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, Mohammad Mahdian, and David C. Parkes. Online auctions with re-usable goods. In *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005), Vancouver, BC, Canada, June 5-8, 2005*, pages 165–174, 2005.

[97] Hanna Halaburda, Mikolaj Jan Piskorski, and Pinar Yildirim. Competing by restricting choice: the case of search platforms. *Harvard Business School Strategy Unit Working Paper*, 10-098, 2015.

[98] Jason D Hartline. *Mechanism design and approximation.*

[99] Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k-set packing. *Comput. Complex.*, 15(1):20–39, May 2006.

[100] Theodore P Hill and Robert P Kertz. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability*, 10(2):336–345, 1982.

[101] Theodore P Hill and Ulrich Krengel. Minimax-optimal stop rules and distributions in secretary problems. *The Annals of Probability*, 19(1):342–353, 1991.

[102] Shan-Yuan Ho and Abijith Krishnan. A secretary problem with a sliding window for recalling applicants. *arXiv preprint arXiv:1508.07931*, 2015.

[103] Ramesh Johari, Vijay Kamble, and Yash Kanoria. Know your customer: Multi-armed bandits with capacity constraints. *CoRR*, abs/1603.04549, 2016.

[104] Ramesh Johari, Vijay Kamble, and Yash Kanoria. Know your customer: Multi-armed bandits with capacity constraints. *arXiv preprint arXiv:1603.04549*, 2016.

[105] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[106] David R. Karger and Maria Minkoff. Building steiner trees with incomplete global knowledge. In *41st Annual Symposium on Foundations of Computer Science, FOCS*, pages 613–623, 2000.

[107] Frank P Kelly. *Reversibility and stochastic networks.* Cambridge University Press, 2011.

[108] Frank P Kelly, Aman K Maulloo, and David KH Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, pages 237–252, 1998.

[109] DP Kennedy. Prophet-type inequalities for multi-choice optimal stopping. *Stochastic Processes and their applications*, 24(1):77–88, 1987.

[110] Robert P Kertz. Stop rule and supremum expectations of iid random variables: a complete comparison by conjugate duality. *Journal of multivariate analysis*, 19(1):88–112, 1986.

[111] Samir Khuller, Anna Moss, and Joseph (Seffi) Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39 – 45, 1999.

[112] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 594–605. IEEE, 2003.

[113] Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136. ACM, 2012.

[114] Elias Koutsoupias and Philip Lazos. Online trading as a secretary problem. In Xiaotie Deng, editor, *Algorithmic Game Theory*, pages 201–212, Cham, 2018. Springer International Publishing.

[115] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745–747, 1977.

[116] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

[117] Aukasz Kruk, John Lehoczky, Kavita Ramanan, and Steven Shreve. Heavy traffic analysis for EDF queues with reneging. *Ann. Appl. Probab.*, 21(2):484–545, 2011.

[118] Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46(1):259–271, 1990.

[119] R Preston McAfee. The gains from trade under fixed price mechanisms. *Applied Economics Research Bulletin*, 1(1):1–10, 2008.

[120] Adam Meyerson. Profit-earning facility location. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 30–36, 2001.

[121] Dale T Mortensen and Christopher A Pissarides. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.

[122] Dale T Mortensen and Christopher A Pissarides. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.

[123] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[124] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[125] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265 – 281, 1983.

[126] Michael J Neely. Stochastic network optimization with application to communication and queueing systems. *Synthesis Lectures on Communication Networks*, 3(1):1–211, 2010.

[127] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.

[128] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press Cambridge.

[129] Ernst L Presman and Isaac Mikhailovich Sonin. The best choice problem for a random number of objects. *Theory of Probability & Its Applications*, 17(4):657–668, 1973.

[130] Jean-Charles Rochet and Jean Tirole. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

[131] Richard Rogerson, Robert Shimer, and Randall Wright. Search-theoretic models of the labor market: A survey. *Journal of Economic Literature*, 43(4):959–988, 2005.

[132] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

[133] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687. SIAM, 2017.

[134] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984.

[135] Ester Samuel-Cahn. Optimal stopping with random horizon with application to the full-information best-choice problem with random freeze. *Journal of the American Statistical Association*, 91(433):357–364, 1996.

[136] L. S. Shapley and M. Shubik. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111–130, 1971.

[137] Amit Sharma, Jake M Hofman, and Duncan J Watts. Estimating the causal impact of recommendation systems from observational data. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 453–470. ACM, 2015.

[138] Weiran Shen, Pingzhong Tang, and Yulong Zeng. A closed-form characterization of buyer signaling schemes in monopoly pricing. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1531–1539, 2018.

[139] Weiran Shen, Pingzhong Tang, and Yulong Zeng. Buyer signaling games in auctions. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1591–1599, 2019.

[140] Zoya Svitkina. Lower-bounded facility location. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 1154–1163, 2008.

[141] Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback.

[142] Kalyan T Talluri and Garrett J Van Ryzin. *The theory and practice of revenue management*, volume 68. Springer Science & Business Media, 2006.

[143] Van-Anh Truong and Xinshang Wang. Prophet inequality with correlated arrival probabilities, with application to two sided matchings. *arXiv preprint arXiv:1901.02552*, 2019.

[144] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 416–425. IEEE, 2002.

[145] Jan Vondrak. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 67–74, 2008.

[146] Amy R. Ward and Peter W. Glynn. A diffusion approximation for a markovian queue with reneging. *Queueing Systems*, 43(1):103–128, 2003.

[147] E Glen Weyl. A price theory of multi-sided platforms. *The American Economic Review*, pages 1642–1672, 2010.

[148] E. Glen Weyl. A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72, September 2010.

[149] David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.

[150] Haifeng Xu. On the tractability of public persuasion with no externalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2708–2727. SIAM, 2020.

[151] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *29th AAAI Conference on Artificial Intelligence, AAAI*, 2015.

[152] Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. Society for Industrial and Applied Mathematics, 2011.

Biography

Reza Alijani earned a bachelor's degree in software engineering from Sharif University of Technology (Iran) in 2015. He obtained his master's and Ph.D. in computer science from Duke University in 2017 and 2020, respectively. He will join Google as a software engineer in 2020.