How to Sell Online (Fast) via Pricing-Based Algorithms

DISSERTATION

zur

Erlangung des Doktorgrades (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Alexander Braun

aus

Fürth, Deutschland

Bonn 2024

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät
der Rheinischen Friedrich-Wilhelms-Universität Bonn

Gutachter/Betreuer: Prof. Dr. Thomas Kesselheim
Gutachter: Prof. Dr. Heiko Röglin

Tag der Promotion: 18. Juli 2024
Erscheinungsjahr: 2024

Abstract

Online resource allocation problems play a fundamental role in online decision making. In these problems, a sequence of agents arrives one-by-one, each with the goal of being assigned a subset of available items. An allocation decision is required to be made immediately and irrevocably. This introduces several challenges for the decision-maker: Future arrivals are unknown, agents can act strategically and the allocation decisions may be required within milliseconds.

The goal of this thesis is to design algorithms which address these challenges and for which we can theoretically quantify *the loss due to*

I. *limited information about the future*,

II. *limited computational power* and/or

III. *strategic behavior of agents*.

In order to bound the loss due to limited information about the future (I.), we work in the *Prophet Inequality* model with stochastic prior information about the arriving agents. Given this prior knowledge, the goal is to design online algorithms which perform reasonably well compared to the expected offline optimum. For this class of problems, the contribution in this thesis is twofold: First, we give simplified proofs for existing pricing-based Prophet Inequalities for combinatorial auctions and matroids. Second, we show that asymptotically optimal welfare can be achieved once we restrict the class of distributions from which the values are sampled.

Shifting the perspective to bound the loss due to limited computational power (II.), we provide a polynomial-time approximation algorithm for the optimal policy once buyers have multi-demand valuations. For these, we show how to use an LP-based rounding approach in order to beat existing results and derive state-of-the-art guarantees. As a side remark, all algorithmic approaches mentioned so far can also be applied when agents behave strategically. In particular, the algorithms are (or can be made) incentive compatible.

Beyond the assumption that items are always available for sale, we consider the related model of two-sided markets: Items are initially held by strategic sellers. In this setting, we are able to bound the loss due to strategic behavior of agents (III.). In particular, we show that for matroid, knapsack and combinatorial double auctions, a reasonable fraction of the expected welfare of the optimal allocation is achievable when restricting to incentive compatibility, individual rationality and budget balance constraints.

Complementing results in the Prophet model, we also argue about the loss due to limited information about the future (I.) in the *Secretary* model. Here, weights are chosen adversarially, but arrive in random order. We show that in this model, a single piece of information can help to beat prevalent bounds. In particular, we use a predicted additive gap as advice and are able to beat the guarantee of $1/e$, which was tight in the classical model.

Acknowledgments

There are several people which I would like to thank for their support during the last years.

First of all, I am extremely grateful to my PhD advisor Thomas Kesselheim. He supported me throughout the last couple of years, provided guidance whenever required, worked together with me on several research projects, but also gave me the freedom to pursue my own interests, ideas and plans. I could not imagine a better supervisor for my PhD studies.

In addition, I would like to thank all the colleagues at the CS Departments I and V at University of Bonn. Especially after the end of the remote period, you made the university such a nice place to work at. The same holds for the participants of the Data-Driven Decision Processes program at the Simons Institute at UC Berkeley during summer 2022. You made this semester in the United States a super exciting time. I will never forget our endless discussions on several research questions, Taco Tuesdays at Rayleigh's and trips on the weekends.

I am grateful to have worked with inspiring collaborators and coauthors: Matthias Buttkus, Thomas Kesselheim, Tristan Pollner, Amin Saberi and Sherry Sarkar. Thank you for a lot of discussions which lead to novel insights and helped me to gain a much deeper understanding of many topics in algorithm theory. It was so much fun working with you.

During my PhD, I was fortunate to spend one summer as an intern at Amazon Web Services in Berlin. Thanks to all colleagues there for supporting me and introducing me to the more practical side of computer science.

Parts of my PhD were funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Project No. 437739576.

Finally, I would like to thank my family and friends who helped me in various stages of life and provided unconditional and constant support during the last decades. Thank you so much for always having an open ear for me. All of this would not have been possible without you.

Contents

Chapter 1

Introduction

After a concert in the local theater, several hundreds to thousands of people would like rides to destinations including their homes, the railway station, and local restaurants. Initially, the number of potential passengers exceeds the number of available cabs. Over time, cabs become available for the remaining concert-goers. How should we assign passengers to cabs?

Or consider the problem of bringing aid supplies to a region in which a natural disaster took place. As a transportation plane becomes available, we can pick a subset of goods which are needed in order to help the people on-site. Still, some goods such as water or medical equipment might be more urgent than others. How should we proceed in order to maximize social welfare?

Both these problems can be modeled as *online resource allocation problems*. Generally speaking, in this thesis, we will talk about a set of *items* which is available offline (in the examples above: the passengers or the aid goods) and a sequence of *buyers* who arrive online one-by-one (the cab drivers or the transportation planes). In every time step, a new buyer arrives and we get to know her preference for being allocated different subsets of items. With this knowledge, we are required to make an immediate and irrevocable allocation decision. In other words, without knowing the realizations of future arrivals, we need to decide which subset of items we allocate to the current buyer. Our goal is to maximize the social welfare: the sum of buyers' values for their assigned bundles.

Examples for this kind of problems are widely spread, also beyond the two introductory examples mentioned above:

- In internet advertising, upon the arrival of an online search query, an immediate and irrevocable decision for which ads to display must be made. Advertisers may have different values for their ad appearing below different queries. How should we decide which ads to display?

- Every airplane has a capacity of seats for a flight. Customers can purchase subsets of available tickets which they desire the most. Still, some passengers might not show up at the end. How should we set prices in order to ensure a good allocation of tickets?

- Another example is the problem faced by an over-demanded hospital: patients are waiting to be admitted, but all beds are occupied. As soon as beds become available, we can assign some of the waiting patients. How should we do so to maximize social welfare?

One of the goals of this thesis is to design algorithms which are able to deal with (at least parts of) the complexity of the introductory examples. Notably, buyers may behave strategically (as e.g. in ticket purchasing), future arrivals are not necessarily known (as e.g. in healthcare) and decisions may be required within milliseconds (as e.g. in ad auctions). To this end, we want to design algorithms for which we can provably *bound the loss due to*

I. *limited information about the future*, in Chapter 3, Chapter 6 and Chapter 7

II. *limited computational power*, in Chapter 4, or

III. *strategic behavior of agents*, in Chapter 5.

Before we dive deep into tackling these questions, we start with a vanilla problem in order to gain a better understanding from a mathematical perspective. In particular, some of the following online single-selection problems will turn out to be special cases of many settings which are considered in this thesis.

1.1 Online Single-Selection Problems

Let us start with a highly simplified scenario of the introductory examples: the case of a single item, which is also called the *online selection problem*. Here, we assume that we are given numbers $w_1, \ldots, w_n \in \mathbb{R}_{\geq 0}$ one-by-one. The number w_i corresponds to the value which buyer i has for being allocated the item. Our goal is to find a *policy* or an *algorithm* which is likely to pick a "reasonably large" number. We determine the quality of an algorithm by comparing the selected value to the largest value in the sequence. Since values are not necessarily bounded, we measure the quality by the ratio of the value selected by the algorithm and $\max_i w_i$. This quantity is called the *competitive ratio*, the value achieved by an online algorithm divided by the optimal offline solution.

It is fairly easy to see that if we perform a worst-case analysis, we will not learn much about the structure of favorable algorithms. To see this, consider the following example. For the first couple of steps until some arrival i^*, the weights are defined as $w_i := W^i$ for some very large number W. After this time i^*, we suddenly only observe zero values, i.e. $w_i := 0$ for all $i > i^*$. Obtaining a "reasonably good" competitive ratio is hence only possible by picking the largest number. In other words, we need to find the index i^*. Still, by an application of Yao's lemma [Yao, 1977], no algorithm can ever achieve a competitive ratio better than $1/n$ in this setting. On the other hand, it is trivial to obtain this competitive ratio: Simply sample one index uniformly at random upfront and select the element with this index. Chances are $1/n$ that we pick the largest number in the sequence.

The example above shows that with *adversarial weights* arriving in *adversarial order*, there is not much hope to learn anything about how to distinguish promising algorithms from unfavorable ones for online selection problems. This motivates the study of relaxations of these two assumptions which will be the two cornerstones for all problems in this thesis: The Prophet Inequality model and the Secretary model.

- *Prophet Inequality model*: We keep the assumption that the arrival order is adversarial, but assume that values are drawn from probability distributions. These distributions are given to the decision-maker upfront.

- *Secretary model*: Weights are chosen adversarially, but they arrive in an order sampled uniformly at random from all permutations.

Interestingly, one can achieve constant-factor guarantees in both, the Prophet as well as the Secretary model, which are independent of the length of the sequence n. We first give a brief introduction to Prophet Inequalities before discussing the Secretary problem afterwards.

1.1.1 Prophet Inequalities

In the standard Prophet Inequality model, we assume that in every round, the value[1] v_i is sampled independently from a publicly known distribution $\mathcal{D}_i$ over non-negative real numbers. The distributions are known upfront, but the realizations are only observed one at a time. After a new realization is revealed, the decision-maker is required to make an immediate and irrevocable accept or reject decision.

This problem corresponds to a Markov Decision Process with the two actions "accept" and "reject". The value achieved by the optimal policy for this problem can be easily computed via backwards induction (see e.g. Chow et al. [1971]). Denote by

$$\omega^{(n)} = \mathbf{E}\left[v_n\right] \quad \text{and} \quad \omega^{(i)} = \mathbf{E}\left[\max\{v_i, \omega^{(i+1)}\}\right] \quad \text{for } i < n \ .$$

It is not too hard to see that the optimal policy will accept a value v_i if and only if it exceeds $\omega^{(i+1)}$. In other words, the optimal policy accepts v_i if this value is larger than what the optimal policy can achieve in the remaining sequence. As a consequence, $\omega^{(1)}$ corresponds to the value which is achievable by the optimal policy starting from the first element in the sequence. This directly introduces thresholds (or dynamic prices) for every element via

$$p^{(n)} = 0 \quad \text{and} \quad p^{(i)} = \mathbf{E}\left[\max\{v_{i+1}, p^{(i+1)}\}\right] \quad \text{for } i < n \ ,$$

and the optimal policy accepts v_i if and only if $v_i \geq p^{(i)}$.

Having this, a natural question to ask is how much we lose due to not knowing future arrivals upfront but only having distributional information. In other words, one might wonder how well this policy performs compared to the expected offline optimum $\mathbf{E}\left[\max_i v_i\right]$, also called *prophet*. That is, the offline optimum (or the prophet) can see all realizations upfront and simply pick the largest number. Before we dive into any algorithmic approaches, let us consider a very simple and well-known example.

Example 1.1.1. (Folklore) *The first buyer has a value $v_1 = 1$ deterministically, the second buyer has a value of zero with probability $1 - \varepsilon$ and a high value of $1/\varepsilon$ with tiny probability ε. While the expected offline optimum in this case is $2 - \varepsilon$, any online algorithm (which does only observe the realization for the second buyer* after *making the decision for the first buyer) can only achieve a value of one in expectation.*

This example limits our expectations. In particular, Example 1.1.1 shows that the best competitive ratio we can hope for is $1/2$. Luckily, there are positive results with reasonable guarantees, originally introduced by Krengel and Sucheston [1977, 1978]. As a matter of fact, instead of analyzing the optimal policy and comparing its expected value

[1] In order to distinguish weights from distributions and adversarial weights, we call weights which are sampled from distributions "values" and denote them by v_i instead of w_i.

to the offline optimum, we will see that a cleaner and simpler algorithm can already achieve a $1/2$-fraction compared to the prophet.

To this end, consider Algorithm 1 which already dates back to the 80s by Samuel-Cahn [1984]: Set a static price p for selling the item and the first buyer who is willing to pay the price, i.e. $v_i \geq p$, gets the item.

Algorithm 1: Sequential Posted Pricing

```
Compute a price p based on D_1, ..., D_n before the first arrival
Accept the first i with v_i ≥ p
```

For this *posted pricing* algorithm, there are several different proofs for the following theorem in the literature.

Theorem 1.1.2 (e.g. in Samuel-Cahn [1984], Kleinberg and Weinberg [2019])**.** *Denote by* ALG *the value selected by Algorithm 1. Then, for* OPT $:= \max_i v_i$*, we have*

$$\mathbf{E}[\text{ALG}] \geq 1/2 \cdot \mathbf{E}[\text{OPT}] \quad .$$

Theorems of this kind are usually referred to as *Prophet Inequalities* as we compare the performance of an online algorithm without knowledge of future arrivals to the above-mentioned *prophet* who has foresight about all values in the sequence. At the same time, from the perspective of competitive analysis, Theorem 1.1.2 derives a bound on the competitive ratio.

The proof techniques for this theorem became a standard tool over the last couple of years [Feldman et al., 2015, Dütting and Kleinberg, 2015, Gravin and Wang, 2019, Kleinberg and Weinberg, 2019, Dütting et al., 2020, Dütting et al., 2020, Correa et al., 2022, Braun and Kesselheim, 2023a] with different choices of the price p. For example, one can set $p = 1/2 \cdot \mathbf{E}[\max_i v_i]$, introduced by Kleinberg and Weinberg [2012][2]. The first proof by Samuel-Cahn [1984] in the 1980s uses a price p with the property that $\mathbf{Pr}[\exists i : v_i \geq p] = 1/2$. Also, there are many other prices which lead to the desired competitive ratio. For example, setting p such that $p = \mathbf{E}[\sum_i (v_i - p)^+]$ [Correa et al., 2022] or a combination of the different options presented so far [Samuel-Cahn, 1984] also works.

For the sake of completeness, a proof of Theorem 1.1.2 can be found in Appendix A.1, which the experienced reader may skip. As an important remark, recall that by Example 1.1.1, the guarantee of $1/2$ is tight.

1.1.2 The Secretary Problem

In contrast to the Prophet Inequality model, in the Secretary problem, an adversary fixes non-negative, real-valued weights $w_1 \geq w_2 \geq \cdots \geq w_n$ which are revealed online in *random order*. We model this by drawing an *arrival time* $t_i \overset{\text{iid}}{\sim} \text{Unif}[0, 1]$ for every weight. Weight w_i is revealed at time t_i and we immediately and irrevocably need to decide if we want to accept or reject this element. As before, we are allowed to pick at most one element with the goal of maximizing the selected weight.

A tight competitive ratio of $1/e$ is known for this problem since the seminal work of Lindley [1961] and Dynkin [1963] (also see Ferguson [1989] or Freeman [1983]) and can

[2] Kleinberg and Weinberg [2012] is the conference version of Kleinberg and Weinberg [2019].

be achieved with a very simple threshold policy: Wait until some time $\tau \in [0,1]$ and only observe weights without accepting. At time τ, take the weight which is *best-so-far* as a threshold, i.e. $\mathrm{BSF}(\tau) = \max_{i:t_i \le \tau} w_i$ (define $\mathrm{BSF}(\tau) = 0$ if no arrival took place before time τ). After time τ, accept the first element whose weight exceeds $\mathrm{BSF}(\tau)$.

Algorithm 2: Secretary Algorithm

Input: Time $\tau \in [0,1]$
Before time τ: Observe weights w_i
At time τ: Compute $\mathrm{BSF}(\tau) = \max_{i:t_i \le \tau} w_i$
After time τ: Accept first element with $w_i \ge \mathrm{BSF}(\tau)$

The following seminal result dates back to the work from the 1960s.

Theorem 1.1.3 (implied by e.g. Lindley [1961], Dynkin [1963])**.** *Denote by* ALG *the weight selected by Algorithm 2. For time $\tau = 1/e$, we have*

$$\mathbf{E}\left[\mathrm{ALG}\right] \ge 1/e \cdot \mathrm{OPT} \ ,$$

where the expectation is over the random arrival times and $\mathrm{OPT} := \max_i w_i = w_1$.

To be precise, Lindley [1961] and Dynkin [1963] provide a stronger guarantee. They derive a lower bound on the probability of accepting the largest weight which implies Theorem 1.1.3. Also, they did not model the random order via arrival times but rather considered one permutation drawn uniformly at random from the set of all permutations. Still, their ideas are directly applied in order to prove the above guarantee in Theorem 1.1.3 with arrival times. Even further, the arrival time and random permutation models are equivalent for this problem, meaning that an algorithm can be used for an input from the other model and vice versa.

As before, we give a proof for Theorem 1.1.3 in Appendix A.2, which the experienced reader may skip. Note that the guarantee of $1/e$ is also tight [Dynkin, 1963, Buchbinder et al., 2014].

1.2 Beyond Single-Selection: Structure of this Thesis

As a matter of fact, the single-selection problems from Section 1.1 will not be sufficient to properly model our introductory examples on resource allocation, ad auctions and purchase processes. To this end, we consider generalizations in several directions in this thesis. We give a very brief, high-level overview with a connection to the single-selection problem now and discuss a much more fine-grained picture of the results and techniques with a comparison to previous work afterwards in Section 1.3 and Section 1.4.

Combinatorial Auctions with Bayesian Priors. In Section 3.2 and Section 3.3 in Chapter 3, we consider the setting of online combinatorial auctions. That is, instead of a single item, there is a set of multiple, heterogeneous items M available for sale and buyers have valuation functions $v_i : 2^M \to \mathbb{R}_{\ge 0}$ for item bundles which are drawn from distributions. For these problems, we give a simplified proof for the existence of static item prices via LP duality. This extends the guarantee of $1/2$ from Theorem 1.1.2 to this more general model.

Complementing this, in Chapter 4, we will see that there are classes of valuation functions for which we can even improve the guarantee of $1/2$ when changing the benchmark: Instead of comparing to the expected offline optimum, we now measure the quality of a polynomial time algorithm with respect to the expected welfare of the optimal online algorithm with unbounded computational power.

In the extension in Section 5.5 in Chapter 5, we consider the setting in which items are not always available but are sold by sellers. Here, sellers may act strategically to maximize their own utility and have a private value for the item. This implies that items only become available once we offer sellers a price exceeding their value. Also, the guarantee of $1/2$ from Theorem 1.1.2 can be generalized to this model.

In a slightly different model in Chapter 6, we make a stronger assumption on the distributions from which the valuations are drawn. This allows to obtain much better guarantees than $1/2$.

Resource Allocation with Feasibility Constraints. Complementing the work on combinatorial auctions, we also work in extensions of the single-selection case towards settings with constraints on the set of buyers which are allocated an item. Usually, one assumes that items are homogeneous, so buyers are simply interested in being allocated one copy. The crux is that the set of buyers with items is required to fulfill a feasibility constraint. In Section 3.4 in Chapter 3, we give an LP-based proof for a Prophet Inequality when the set of buyers with items is required to be independent in a matroid. As an interesting remark, the same guarantee of $1/2$ also appears in this context again.

Feasibility constraints also play a role in Chapter 5 when items are held by sellers, but the set of buyers which are allocated an item is restricted by a constraint family. In particular, we consider matroid constraints in Section 5.3 and Section 5.4 and knapsack constraints in Section 5.6 and Section 5.7.

Online Selection in the Secretary Model. In Chapter 7, we turn our perspective towards random arrival orders for adversarial weights. In contrast to Theorem 1.1.3, we assume to get a single additional piece of information upfront. One can think of this information as a prediction from some machine-learning model which the algorithm may or may not incorporate in order to improve its performance. As we will see, with an accurate prediction, one can beat the guarantee of $1/e$ from Theorem 1.1.3.

1.3 Overview of Results and Techniques in this Thesis

In this section, we give a more detailed overview of the results in this thesis and highlight differences to previous work as well as technical contributions. Additional notes on how the chapters differ from the papers they are based on as well as the contribution of the authors to the results are addressed later in Section 1.5.

Quality of Algorithms. Throughout this thesis, we measure the quality of all algorithms by comparing to an optimum benchmark. As we do usually not assume agents' valuations to be bounded, we aim for multiplicative guarantees. To this end, let us denote by ALG the (expected) objective value achieved by our algorithm, assume $\zeta \in [0, 1]$

and OPT denotes the (expected) optimal objective value[3].

I. *Loss due to limited information*: We say that an online algorithm is ζ-*competitive* compared to the (expected) offline optimum OPT if for any instance of the problem, we have ALG $\geq \zeta \cdot$ OPT.

II. *Loss due to limited computational power*: If an optimal solution to a problem cannot be found efficiently[4], we aim for polynomial-time approximation algorithms. We say that a polynomial-time algorithm is a ζ-*approximation* if for any instance of the problem, we have ALG $\geq \zeta \cdot$ OPT.

III. *Loss due to strategic behavior*: In settings where agents behave strategically and may misreport valuations in order to be assigned an allocation, we say that an allocation algorithm is a ζ-*approximation* if ALG $\geq \zeta \cdot$ OPT and OPT has access to the true valuations directly, i.e. it is not restricted by any strategic behavior.

Expectations are taken over inherent randomness of the algorithm, due to sampling valuations from distributions or random arrival times. To make this dependence more clear, we sometimes write $\mathbf{E}[\text{ALG}]$ and/or $\mathbf{E}[\text{OPT}]$ instead of ALG and OPT. As the benchmarks will be different across chapters in the thesis, we give more precise, formal definitions of what is meant in the preliminaries section in the respective chapters.

The key points of comparison are the following: The main statements in Chapter 3, Chapter 6 and Chapter 7 derive competitive ratios (I.) for online algorithms. In Chapter 4, we design and analyze a polynomial-time approximation algorithm (II.) and in Chapter 5, we compare the performance of a mechanism which is required to deal with strategic behavior to an optimum allocation (III.).

1.3.1 Chapter 3: Simplified Prophet Inequalities for Combinatorial Auctions and Matroids

In this chapter, we derive simplified proofs[5] for Prophet Inequalities in combinatorial auctions. To this end, we make use of a linear programming formulation capturing the existence of static item prices, inspired by an LP introduced in Dütting et al. [2020]. An application of strong LP duality allows us to change our perspective and move into dual space. The interpretation of the dual variables as probabilities over the presence and absence of items allows us to draw the conclusion: There exist static item prices which lead to the desired competitive ratios.

Comparing this to previous work, we can also interpret the static and anonymous item prices from Feldman et al. [2015] and Dütting et al. [2020] in the context of the linear program. Their prices are actually feasible solutions to our linear program. Still, the proofs for competitive ratios in the work of Feldman et al. [2015] and Dütting et al. [2020] are based on arguments about a specific choice of the price vector $(p_j)_{j \in M}$. In particular, they argue about prices and valuation functions and show that, on the one hand, prices are sufficiently high to cover the welfare loss induced by allocating an item.

[3]For simplicity, whenever appropriate in this thesis, we will overload notation and use ALG and OPT for the objective value as well as the algorithm itself and the allocation which achieves this value.

[4]Unless $\mathsf{P} = \mathsf{NP}$. By *efficient* we mean that the running time of the algorithm is polynomial in the input size of the problem instance.

[5]We highlight that most of the theorems in this chapter were already known in the literature before, but with different, more complicated proofs. In addition, Chapter 3 is the only chapter in this thesis in which the theorems were known already before.

On the other hand, they need to argue that prices are low enough such that agents are willing to buy. In contrast, our approach can avoid any argument on specific buyers' valuations at all.

Concerning competitive ratios, our LP based proof recovers the tight guarantees of $1/2$ for a single item from Theorem 1.1.2 as well as XOS[6] valuation functions [Feldman et al., 2015] and the currently best known $1/4k-2$ for MPH-k valuations[6] [Dütting et al., 2020]. As a side remark, for any $k \geq 2$, we also get a tiny improvement in the competitive ratio for MPH-k Prophet Inequalities from $1/4k-2$ to $1/\left(2k+2\sqrt{k(k-1)}-1\right)$ which is slightly better than $1/4k-2$.

Complementing this, we show that a related LP-based approach can also be used for Prophet Inequalities with matroid feasibility constraints. To this end, we give a proof for the existence of dynamic prices in Section 3.4 leading to a tight $1/2$-competitive algorithm, a statement initially shown by Kleinberg and Weinberg [2019].

1.3.2 Chapter 4: Approximating Optimum Online for Multi-Demand Buyers

Even though the proofs for the competitive ratios from Chapter 3 are simple, the approach is unfortunately not computationally efficient. Luckily, Dütting et al. [2020] argue that we can also achieve a polynomial time $1/2$-competitive algorithm against the expected offline optimum. Note that this directly implies a polynomial time $1/2$-approximation algorithm for the (computationally unrestricted) optimal online algorithm[7]. Still, by Example 1.1.1, there is no hope to get better-than-half approximation algorithms compared to the optimum online algorithm when using this implication. Hence, an immediate question is the following: Is there another way to get improved approximation algorithms for combinatorial auctions when comparing to the (computationally unrestricted) optimal online algorithm?

We study this problem for buyers with multi-demand valuation functions[6] and answer this question in the affirmative. In particular, we show the following theorem.

Theorem 1.3.1. *When buyers have multi-demand valuation functions, there exists a polynomial time* $(1/2 + \kappa)$*-approximation algorithm with respect to the expected social welfare of the optimal online algorithm, for a constant* $\kappa \geq 0.0115$.

Our approximation algorithm works in the following way: It rounds an LP relaxation online while introducing a controllable amount of positive correlation among items. For each buyer i, we apply two rounds of so-called *pivotal sampling* to the set of items. Doing so, we guarantee to never "over-allocate" items to buyer i beyond its remaining demand. In addition, we only randomly allocate a subset of these sampled groups to avoid large positive correlation between items while still ensuring that items are allocated with sufficiently high probability.

The LP relaxation uses an "online constraint" which separates online and offline algorithms, as in Papadimitriou et al. [2021], Torrico and Toriello [2022]. Also, from a high-level perspective, our algorithm template looks similar to the one in Papadimitriou et al. [2021] for the unit-demand case. Still, as we highlight in Chapter 4, for the

[6] We give formal definitions for all valuation functions in Chapter 2.

[7] The offline optimum is a stronger benchmark than the online optimum. To see this, simply note that one possible strategy for the offline optimum is to mimic the online optimum by not considering future realizations.

algorithm as well as its analysis, new ideas are required when considering more general multi-demand valuations.

Interestingly, this is a sharp contrast to Chapter 3: Through the lens of Prophet Inequalities (i.e. comparing to the offline optimum), the unit-demand and the general multi-demand variant of the problem behave nearly identically. These variants can all be handled by the same algorithmic template (for example using the one from Chapter 3) and techniques for the unit-demand case directly carry over. As we will discuss in Chapter 4, this is no longer true when considering the optimum online algorithm as a benchmark.

1.3.3 Chapter 5: Truthful Mechanisms for Two-Sided Markets via Prophet Inequalities

Given the discussion on resource allocation problems in Chapter 3 and Chapter 4 in which all items are available initially, in this chapter we shift our perspective towards a more strategic environment. Now, items are initially held by sellers who have a privately known value for items. In order to make an item available for buyers, we need to ensure the payment of a sufficiently high price to the seller holding an item. Still, suitable mechanisms are not allowed to subsidize potential trades. Hence, we end up facing the following challenge: Are there mechanisms which do not need to subsidize trades and still achieve nearly optimal social welfare?

To tackle this question, we design pricing-based mechanisms for two-sided markets. In all our mechanisms, truth-telling is a dominant strategy for every agent. More precisely, all mechanisms are dominant-strategy incentive-compatible (DSIC) for all buyers and sellers. Also, participation is not harmful for agents when using this dominant strategy, so our mechanisms are individually rational (IR). In addition, they fulfill different variants of budget balance.

In other words, we show the following theorem(s) which are further specified in Table 1.1.

Theorem 1.3.2. *There exists mechanisms for matroid/combinatorial/knapsack double auctions which are DSIC, IR for all buyers and sellers, budget balanced and whose social welfare are constant-factor approximations to the optimal (first best) social welfare.*

Our main results are for matroid double auctions (DA): Sellers hold identical items and the set of buyers who are allocated an item needs to be independent in a matroid. For these, we present two mechanisms. The first is strongly budget balanced and a $1/3$-approximation of the expected optimal (first best) social welfare (Section 5.3). Still, it requires agents to trade in a customized order which we compute in the ongoing mechanism. Our second mechanism is only weakly budget balanced, but achieves half of the expected optimal social welfare (Section 5.4). In addition, the second mechanism can handle agents trading in an arbitrary worst-case order.

We extend our techniques to combinatorial double auctions with heterogeneous items (Section 5.5) and to knapsack double auctions with homogeneous items and a knapsack constraint over the set of buyers (Section 5.6 and Section 5.7). A complete overview of the guarantees derived in Chapter 5 can be found below in Table 1.1.

In order to derive our mechanisms, we extend balanced prices [Kleinberg and Weinberg, 2019, Feldman et al., 2015, Dütting et al., 2020] from Prophet Inequalities to two-sided environments. When prices from the corresponding Prophet Inequality are

	Items	**Budget Bal.**	**Approx.**	**Previous Best**
Matroid DA:	Identical	Strong	1/3	1/16 [a] and
	Identical	Weak	1/2	$1/(3+\sqrt{3})$ [b]
Combinatorial DA:				
XOS + Unit-Supply	Heterog.	Strong	1/2	1/6 [c] and
Additive + Additive	Heterog.	Strong	1/2	1/3 [b]
Knapsack DA:	Identical	Strong	1/12	
	Identical	Weak	1/7	

Table 1.1: The approximation guarantees for mechanisms in matroid, combinatorial and knapsack double auctions. Concerning the previous best results, "a" can be found in Colini-Baldeschi et al. [2016], "b" in Dütting et al. [2021a] and "c" in Colini-Baldeschi et al. [2020]. We note that the setting in Dütting et al. [2021a] is different from the one in Chapter 5 as they construct mechanisms with only limited sample-based knowledge of the distributions. In contrast, we assume in Chapter 5 to have full knowledge of the distribution.

static and anonymous (as for example in Algorithm 1), results carry over in a straightforward manner. On the other hand, once a Prophet Inequality uses dynamic prices (as for matroid constraints in Section 3.4), we need to be much more careful. Here, our mechanism will consist of fine-grained trade proposals and requires a careful choice of the order in which we consider agents for trades. Also, the prices which we use will only be inspired by the ones in Kleinberg and Weinberg [2019], but are required to capture the more complex structure of the two-sided problem.

The analysis of the mechanisms is based on the standard base value and surplus decomposition as in the Prophet Inequality literature. Still, the surplus for buyers and sellers usually requires two different analyses. However, as a matter of fact, it will turn out that it does not play a key role which agent purchases or keeps which item — since any irrevocably allocated item ensures a sufficient contribution to welfare via its price.

1.3.4 Chapter 6: Asymptotically Optimal Welfare of Posted Pricing with MHR Distributions

In this chapter, we work in the one-sided environment (i.e. all items are available initially), but strengthen the assumption on the distribution and aim for stronger guarantees. In particular, we assume that the valuation for each buyer is an independent *and* identically distributed draw from the same distribution. In other words, there is a single distribution from which we draw an i.i.d. sample for the valuation function in every round. Note the contrast to Chapter 3, Chapter 4 and Chapter 5, where we worked in settings with the assumption that distributions are independent across buyers, but not necessarily identical.

For the main result, we focus on posted pricing mechanisms for unit-demand combinatorial auctions (a.k.a. bipartite matching) with the assumption that different item values are independent. To stress this point once more: Across buyers, distributions are independent and identical; across items, distributions are independent, but not necessarily identical. For example, with two items, every buyers' value for the first could be an i.i.d. draw from a uniform distribution, the value for the second one an i.i.d. draw from an exponential distribution.

We will use the key assumption that the marginal distributions are concentrated; more precisely that they have a monotone hazard rate (MHR).

Theorem 1.3.3. *When marginal distributions are independent across items and each marginal has a monotone hazard rate, there is a mechanism using dynamic prices whose expected social welfare ensures a* $(1 - O(1/\log n))$*-fraction of the expected optimal social welfare.*

Complementing this, for static prices, we obtain a $(1 - O(\log\log\log n/\log n))$-competitive algorithm with respect to the expected offline optimum. We also show that these guarantees are best possible, even in the case of only a single item. Note that the bounds are independent of the number of items m.

From a technical perspective, the mechanism on dynamic prices sets prices so that the offline optimum is mimicked. However, analyzing such a selling process is still difficult because items are incomparable and bounds for MHR distributions cannot be applied directly to draws from multiple distributions, which are not necessarily identical. To bypass this problem, we introduce a reduction that allows us to view item valuations not only as independent but also as identically distributed. In contrast, the idea of our mechanism using static prices is to set prices suitably high in order to bound the revenue of our mechanism with a sufficient fraction of the optimum.

We also demonstrate that our techniques are applicable beyond unit-demand settings by giving mechanisms for the more general class of subadditive valuation functions in Section 6.4. Our dynamic pricing mechanism is $1 - O(1+\log m/\log n)$-competitive for subadditive buyers. We complement this by a static pricing mechanism which is $1 - O(\log\log\log n/\log n + \log m/\log n)$-competitive. Both guarantees can be derived by showing that the revenue of the posted pricing mechanism is at least as high as the respective fraction of the optimal social welfare. As a consequence, these bounds directly imply the competitive ratios for welfare and revenue. For a constant number of items m, these bounds are again asymptotically tight by our optimality results. Obtaining tight bounds for large m remains an open problem.

1.3.5 Chapter 7: The Secretary Problem with Predicted Additive Gap

In all of the Chapters 3 to 6, we are using the key assumption that we have access to distributions upfront from which valuation functions are drawn. In contrast, in this chapter, we work in the Secretary model where weights are fixed adversarially, but arrive in random order. In Chapter 7, we go beyond the classical formulation of the Secretary problem. In addition to knowing that the adversarial weights arrive in random order, we are given a single piece of information upfront: An additive gap between the largest and k^{th} largest weight in the sequence. Following a standard assumption in the literature on algorithms with predictions (references can be found at Algorithms-with-Predictions), our algorithm must be able to deal with potentially inaccurate predictions for the additive gap.

Our contribution for this problem is threefold. First, we show that knowing an exact additive gap allows us to beat the competitive ratio of $1/e$ by a constant.

Theorem 1.3.4 (Theorem 7.2.1, simplified form)**.** *There exists an online algorithm which achieves an expected weight of* $\mathbf{E}[\text{ALG}] \geq 0.4 \cdot \text{OPT}$ *given access to a single additive gap* c_k *for* $c_k = w_1 - w_k$ *and some* k.

Still, as mentioned, getting to know an exact gap might be too much to expect. Hence, we next introduce a slight modification in the algorithm to make it robust with respect to errors in the predicted gap while simultaneously outperforming the prevalent competitive ratio of $1/e$ by a constant for accurate gaps.

Theorem 1.3.5 (Theorem 7.3.1, simplified form)**.** *There exists an online algorithm which uses a predicted additive gap and is simultaneously* $(1/e + O(1))$*-consistent and* $O(1)$*-robust.*

Theorem 1.3.5 does not assume any bounds on the error of the predicted additive gap used by our algorithm. In particular, the error of the prediction might be unbounded and our algorithm is still constant competitive. However, if we know that the error is bounded, we can do much better and only incur an additional additive loss with respect to the guarantee we achieved for exact predictions. In particular, incorporating this error bound ϵ in the algorithm allows to obtain an expected weight of $\mathbf{E}[\text{ALG}] \geq 0.4 \cdot w_1 - 2\epsilon$.

As a corollary of our main theorem, we show that we can beat the competitive ratio of $1/e$ even if we only know the gap $w_1 - w_k$ but do not get to know the index k. In particular, this proves that even an information like "there is a gap of c in the instance" is helpful to beat $1/e$, no matter which weights are in the sequence and which value c attains.

Our algorithms are inspired by the one for the classical Secretary problem, but additionally incorporate the gap: Wait for some time to get a flavor for the weights in the sequence, set a threshold based on the past observations and the gap, pick the first element exceeding the threshold. Even though this might not sound too promising at first glance, we will show that it allows to beat the prevalent bounds.

From a technical perspective, the analyses are based on a case distinction. We compare the weight w_1 to the weight of the element w_k to which we observe the predicted gap. If this gap is small, any element among $w_2, \ldots, w_k$ can sufficiently contribute to the algorithm's performance. If this gap is large, the gap itself is a good surrogate for the largest weight in the sequence and can hence help to exclude many elements from being selected.

1.4 Related Work

In order to embed the results mentioned in Section 1.3 in the existing literature, we give an overview of related work in Prophet Inequalities and the Secretary problem. Additional related work which is more topic-specific can be found in the respective chapters.

1.4.1 A Non-Exhaustive Overview on Prophet Inequalities

As mentioned in the introduction, the first statements like Theorem 1.1.2 already date back to the 1970s and 80s [Krengel and Sucheston, 1977, 1978, Samuel-Cahn, 1984]. Results of this type regained attention in the mechanism design community by the work of Hajiaghayi et al. [2007], Chawla et al. [2010] and Kleinberg and Weinberg [2012] due to the connection to sequential posted-pricing mechanisms. In other words, Algorithm 1 can be interpreted as a dominant-strategy incentive-compatible and individually rational mechanism: The allocation rule assigns the item to the buyer selected by the algorithm, her payment is the static price which we initially set in the algorithm. This connection established the growth of literature in more general environments: multiple (homo- or heterogeneous) items are available and buyers compete for subsets of them.

With k identical items for sale, Hajiaghayi et al. [2007] obtain a sequential posted-pricing mechanism which is $\left(1 - O(\sqrt{\ln k/k})\right)$-competitive. Chawla et al. [2010] argue that we cannot only obtain guarantees for welfare maximization via Prophet Inequalities, but also get bounds for revenue in single- and multi-parameter settings.

A few years later, Kleinberg and Weinberg [2012] introduce the first pricing-based algorithm for matroid constraints in Prophet Inequalities with a tight competitive ratio. That is, there are k identical items for sale and the set of buyers who are allocated an item needs to be independent in a matroid. Their algorithm based on dynamic prices obtains the tight $1/2$-competitive ratio for this problem.

Balanced Prices and Similar Approaches. Initiated by the work of Kleinberg and Weinberg [2012] for matroids and Feldman et al. [2015] for combinatorial auctions, there was a rise of pricing-based approaches for Prophet Inequalities. Feldman et al. [2015] give the first tight $1/2$-competitive algorithm when heterogeneous items are for sale and buyers have XOS valuation functions. In follow-up work, Dütting et al. [2020] unify and simplify these two approaches and obtain improved results in several other settings. All these algorithms are using so-called *balanced prices*: Prices which are low enough such that buyers are willing to purchase items, but are also high enough to cover the welfare loss due to allocating items. The results and techniques in Chapter 3 and Chapter 5 in this thesis are related to the balanced prices approach.

Using a very similar idea, Gravin and Wang [2019] get a $1/3$-competitive algorithm for edge-arrival bipartite matching using a posted-pricing algorithm with static prices. Edge-arrival bipartite matching can also be interpreted in the buyer-item-framework: items are nodes in the graph, every buyer corresponds to one edge and is interested in exactly two items, namely the ones which are incident to this edge. In other words, the buyer has a single-minded valuation with only non-negative value for the two items of interest. As a matter of fact, these single-minded valuation functions have complementarities — in contrast to the XOS valuations considered in Feldman et al. [2015] which are complement-free. Correa et al. [2022] extend the approach of Gravin and Wang [2019] to a broader class of valuation functions with complementarities.

Going beyond XOS valuation functions in the complement-free regime, Dütting et al. [2020] achieve a guarantee of $1/O(\log \log m)$ using static item prices when buyers have subadditive valuation functions. Obtaining a constant-factor Prophet Inequality for subadditive valuations was a long-standing open problem. Only recently, Correa and Cristi [2023] were able to introduce a $1/6$-competitive algorithm. Still, it is not pricing-based but uses a randomized allocation procedure. Obtaining a constant-factor Prophet Inequality for subadditive valuation functions which uses static and anonymous item prices remains an open problem.

There is also an excellent survey by Lucier [2017] which gives an overview of techniques and results on Prophet Inequalities through an economic lens.

Online Contention Resolution Schemes. In contrast to the pricing-based approaches discussed above, there is another line of work which derives Prophet Inequalities via online rounding. That is, we start with an ex-ante relaxation of the problem, obtain realizations of buyers in every iteration and round the respective variables of the relaxation online. This is done via so-called *online contention resolution schemes* (OCRS).

Alaei [2014] obtains a $(1 - O(1/\sqrt{k}))$-competitive algorithm when k identical items are for sale, improving the result from Hajiaghayi et al. [2007]. Alaei [2014] called his approach the magician problem. The term "online contention resolution scheme" was only introduced later by seminal work of Feldman et al. [2016] which derive Prophet Inequalities for several different feasibility families such as matroids or matching constraints. Ezra et al. [2020] use an OCRS to obtain guarantees for vertex- and edge-arrival matching in general graphs.

Combining pricing-based approaches and OCRS techniques, Lee and Singla [2018] show how to transform pricing-based Prophet Inequalities to OCRSes with the same competitive ratio. In addition, a long line of literature applied OCRS techniques in the last years in a variety of problems (see e.g. Zhang [2020], Pollner et al. [2022], Fu et al. [2022], Avadhanula et al. [2023], MacRury et al. [2023]). Our approach in Chapter 4 is inspired by the underlying ideas from the OCRS literature, and in particular the work of Ezra et al. [2020]. Still, it uses a different LP relaxation and also the algorithm is more evolved.

Matroids, Capacities and other Feasibility Constraints. Both of the approaches mentioned above, pricing-based ones as well as rounding algorithms, are widely applied to derive Prophet Inequalities when the set of buyers is restricted via a set of feasibility constraints. As mentioned, matroids are considered by Kleinberg and Weinberg [2019]. This approach was extended to polymatroids by Dütting and Kleinberg [2015]. Also Chawla et al. [2020] study Prophet Inequalities for matroids and restrict themselves to non-adaptive algorithms; a setting also studied by Pashkovich and Sayutina [2023].

Dütting et al. [2020] and Jiang et al. [2022] derive guarantees for knapsack (and more general packing) constraint families, Göbel et al. [2014], Rubinstein [2016], Baek and Ma [2019] take a look at more complex feasibility constraints such as independent sets in graphs or arbitrary downward-closed constraint families.

In another line of work, online allocation has also been studied in settings where offline nodes have capacities and can be allocated simultaneously in different rounds [Alaei et al., 2013, Alaei, 2014]. Chawla et al. [2017] also consider a combinatorial generalization with many item copies.

Prophet Inequalities beyond Adversarial Order. All of the results mentioned in this section so far hold for buyers arriving in adversarial order. A natural other direction is to study less pessimistic arrival orders. For a random arrival order, Esfandiari et al. [2017] and Correa et al. [2017] introduced the Prophet Secretary model and obtain a $1 - 1/\mathrm{e}$ when a single item is for sale. This result was later generalized to matroids and combinatorial auctions by Ehsani et al. [2018] using an algorithm which is based on dynamic prices. In contrast, Adamczyk and Włodarczyk [2018] and Lee and Singla [2018] use a rounding-based approach, known as random order contention resolution schemes. Going beyond the guarantee of $1 - 1/\mathrm{e}$, recent work of Azar et al. [2018] and Correa et al. [2019b] derives improved competitive ratios for the single item case.

With only two potential arrival orders, namely forward and backward, Arsenis et al. [2021] are able to prove a competitive ratio of $1/\phi$ for the single item problem, where ϕ is the golden ratio.

When the decision maker can choose the order, we enter the terrain of the free-order model. Abolhassani et al. [2017] show how to beat the guarantee of $1 - 1/\mathrm{e}$ when the decision maker can select the order, Yan [2011] derives a guarantee of $1 - 1/e$ for matroids.

As a special case, when assuming that all distributions are identical, we end up in the i.i.d. Prophet Inequality regime. Initially studied by Hill et al. [1982], Correa et al. [2017] proved the tight guarantee of approximately 0.745 for this problem. Interestingly, when not knowing the distribution, the best guarantee is $1/e$ [Correa et al., 2019a] which carries over from classical Secretary considerations, as e.g. in Theorem 1.1.3.

Changing the Benchmark: Approximating the Optimal Policy. So far, we only discussed literature which compares the performance of an online algorithm to the expected offline optimum which can see all realizations upfront and hence make optimal decisions. In recent line of work, the question arose what is possible when comparing to the less pessimistic optimal online algorithm (which has unbounded computational power, but does *not* see future realizations). For a single item with known adversarial order, computing the optimal policy can be done via backwards induction [Chow et al., 1971].

Anari et al. [2019] get a PTAS for a specific set of laminar matroid constraints. Dütting et al. [2023] run a similar comparison in the Prophet Secretary model.

Papadimitriou et al. [2021] argue that once we have multiple, heterogeneous items for sale and unit-demand valuation functions for buyers, it is PSPACE-hard to compute the optimal policy. They derive a 0.51-approximate algorithm, later improved to 0.52 [Saberi and Wajc, 2021], $1 - 1/e \approx 0.632$ [Braverman et al., 2022], and 0.652 [Naor et al., 2023]. In Chapter 4, we also compare to the online optimum as our benchmark. In contrast to the results mentioned here, we derive guarantees in a more general setting beyond unit-demand valuation functions.

Additional Literature on Prophet Inequalities. Complementing the results on full access to the distribution, Azar et al. [2014] and Rubinstein et al. [2020] study Prophet Inequalities when the decision maker only has access to a single sample from each distribution. Interestingly, this is sufficient to recover the guarantee from Theorem 1.1.2 by setting a price equal to the largest sample. Still, the analysis works in a different way. Finding a unified approach for this and the balanced prices techniques described above might lead to novel insights which could inspire future research.

Working in a similar direction beyond full access to the distributions, Dütting and Kesselheim [2019] study Prophet Inequalities when the decision maker has access to inaccurate prior distributions. They discuss previous approaches and give robustness guarantees depending on the inaccuracy with respect to distance measures of distributions, such as Levy, Wasserstein or Kolmogorov distance.

When having the goal of maximizing revenue instead of social welfare, one usually imposes the additional assumption that items are independent. This makes it possible to also apply Prophet Inequalities on the sequence of items rather than buyers and thus maximize revenue for unit-demand buyers via posted prices [Chawla et al., 2007, 2010]. Cai and Zhao [2017] consider more general XOS and subadditive valuations and apply a duality framework instead. They design a posted-prices mechanism with an entry fee that gives an $O(1)$ or $O(\log m)$ approximation to the optimal revenue. In Dütting et al. [2020], the approximation of the optimal revenue for subadditive valuations is improved to $O(\log \log m)$.

Beyond independence across buyers, there is work by Immorlica et al. [2020] who assume linear correlation across buyers as well as Rinott and Samuel-Cahn [1987] for negatively dependent values.

1.4.2 Some Additional Literature on the Secretary Problem

Since the introduction of the Secretary problem in the 1960s [Lindley, 1961, Dynkin, 1963], there have been a lot of extensions and generalizations of this problem with beautiful algorithmic ideas to solve them. For more details on the early stages of Secretary problems, there are nice surveys by Ferguson [1989] or Freeman [1983] giving a great overview on different variants of the problems and results.

Multi Selection and Feasibility Constraints. When having k identical items for sale, we can obtain a $1 - O(1/\sqrt{k})$-competitive algorithm [Kleinberg, 2005]. Beyond identical items, there are also results when the set of selected buyers needs to satisfy a knapsack constraint. Babaioff et al. [2007a] study this problem and give a $1/10e$-competitive algorithm.

Probably the most prominent problem among Secretary problems with feasibility constraints is Matroid Secretary. Initially studied by Babaioff et al. [2018b], currently the best guarantee is in the order of $1/\log\log(\text{rank})$, where "rank" refers to the rank of the matroid, in Lachish [2014] and follow-up work by Feldman et al. [2018]. Still, achieving a constant-factor for arbitrary matroids is a long-standing open problem. Along the way, a lot of special cases of matroids and other variants have been studied, for example in the ordinal Matroid Secretary problem [Soto et al., 2021].

Concerning the Secretary problem with more general constraints, among others, there are results for linear packing constraints [Kesselheim et al., 2018], independent set constraints [Göbel et al., 2014] and even arbitrary downward-closed feasibility constraints [Rubinstein, 2016].

Mechanism Design, Matching and Combinatorial Auctions. There is also literature on the application of Secretary problems in online mechanism design, for example in Hajiaghayi et al. [2004] and Babaioff et al. [2007b].

Being more precise, in the context of combinatorial auctions, when assuming unit-demand buyers arriving in random order, we end up with a bipartite matching problem. Introduced by Korula and Pál [2009], work by Mahdian and Yan [2011] tackled this max-weight bipartite matching problem with random arrival order of buyers via factor revealing LPs. The tight guarantee of $1/e$ for matching was finally obtained by Kesselheim et al. [2013], who also generalize their result to submodular combinatorial auctions, i.e. when buyers have submodular valuations functions over item bundles.

Additional Work on Secretary Problems. Complementing the assumption of uniform arrival orders, Kesselheim et al. [2015] study different distributions over permutations and give constant-factor guarantees. There are also variants where some of the weights are sampled upfront and the decision maker can observe these weights, e.g. in Kaplan et al. [2020] and Correa et al. [2021]. Here, some elements are revealed before the start of the sequence. The algorithm then tries to pick the best of the remaining weights. Guarantees are achieved with respect to the best remaining element in the sequence. Going beyond linear payoffs for the Secretary problem, Bateni et al. [2013]

and Feldman et al. [2011] study non-linear payoff functions. In particular, they derive results once the payoff functions are submodular. For a nice overview on random order models, there is a book chapter by Gupta and Singla [2020] giving a great overview of techniques and results in this field.

Bridging between the Secretary problem and the Prophet Inequality world, there is a recent line of work by Correa et al. [2023] or Correa et al. [2019a] and many more (see e.g. Bradac et al. [2020], Kesselheim and Molinaro [2020], Argue et al. [2022]) which interpolate between the two models in one way or another.

1.5 Bibliographic Notes and Contributions to the Chapters

The majority of the results which we discussed in Section 1.3 are contained in the following publications:

- *Simplified Prophet Inequalities for Combinatorial Auctions* [Braun and Kesselheim, 2023a], which is joint work with Thomas Kesselheim,
- *Approximating Optimum Online for Capacitated Resource Allocation* [Braun et al., 2024], which is joint work with Thomas Kesselheim, Tristan Pollner and Amin Saberi,
- *Truthful Mechanisms for Two-Sided Markets via Prophet Inequalities* [Braun and Kesselheim, 2021] (conference version) and [Braun and Kesselheim, 2023b] (journal version), which is joint work with Thomas Kesselheim,
- *Asymptotically Optimal Welfare of Posted Pricing for Multiple Items with MHR Distributions* [Braun et al., 2021], which is joint work with Matthias Buttkus and Thomas Kesselheim and
- *The Secretary Problem with Predicted Additive Gap* [Braun and Sarkar, 2023], which is joint work with Sherry Sarkar.

The remainder of this section gives an overview of the development of the respective results. As a matter of fact, in my research area, any publication is usually the joint work of several authors and it is very uncommon to split the contributions on a per-author level. In addition, results are mainly derived in a dynamic development process during which discussions of the authors play a pivotal role. That is also the reason why authors are listed alphabetically in all publications.

Chapter 3: Simplified Prophet Inequalities for Combinatorial Auctions and Matroids

The results on combinatorial auctions in Section 3.2 and Section 3.3 appeared in Braun and Kesselheim [2023a], which is joint work with Thomas Kesselheim. The simplified proof for matroids in Section 3.4 is unpublished work of myself.

Thomas Kesselheim and myself started studying this problem inspired by the work on subadditive combinatorial auctions by Dütting et al. [2020]. Still, using their LP directly only implies a competitive ratio of $1/4$, even for a single item. Hence, in a first step, we improved the LP to get the tight $1/2$-guarantee for a single item. Afterwards, the goal was to see in which generalizations one could make use of this LP: For subadditive combinatorial auctions, Dütting et al. [2020] were able to improve the competitive ratio from $1/O(\log m)$ to $1/O(\log \log m)$; making this a promising direction for Prophet Inequalities. While working on more general settings, Thomas Kesselheim and myself discovered the simplified proof for XOS valuation functions (Lemma 3.3.2) during several discussions, which I was able to generalize to MPH-k valuations (Lemma 3.3.3 and Lemma 3.3.4).

In addition, even though my proof for matroids in Section 3.4 is not as simple as the one for combinatorial auctions, it highlights that this approach might have more potential than we were able to use so far.

Chapter 4: Approximating Optimum Online for Multi-Demand Buyers

The results of this chapter are mainly contained in Braun et al. [2024], which is joint work with Thomas Kesselheim, Tristan Pollner and Amin Saberi. The paper is currently under review.

In the paper [Braun et al., 2024], we state the problem in a more general variant than in this thesis, namely with stochastic rewards. The analysis from Section 4.3 for the exponential time variant of the algorithm (Algorithm 4) is mainly the same as in the paper. The only step in Algorithm 4 which requires exponential time is to compute an expectation exactly. In the paper, we give a detailed proof on how to use samples to estimate this expectation with a bounded error and obtain a truly polynomial time algorithm. The sample-based analysis was done by Tristan Pollner and is not included in this thesis. In contrast, in this thesis, I complement the results from the paper by showing that beating the prevalent bound of $1/2$ by a constant is possible with an even simpler polynomial-time algorithm which does not use the mentioned expected value at all. This argument makes use of several lemmas from our paper, but differs in a few steps in the analysis.

We initiated this project while participating in the program on Data-Driven Decision Processes at the Simons Institute for the Theory of Computing at UC Berkeley, CA, USA. I went up to Tristan Pollner who had recently published a paper on unit-demand combinatorial auctions [Papadimitriou et al., 2021] comparing against the online benchmark. Since their techniques felt fairly related to online contention resolution schemes, I was interested in seeing if one could use their techniques also in a broader setting. In particular, since the OCRS approach works similar for unit-demand valuations as well as XOS valuations, I proposed discussing their approach for XOS valuations to see if results carry over. As a starting point during the semester, Tristan Pollner, Thomas Kesselheim and myself were considering a mix of additive and unit-demand valuations. As we figured out, a lot of properties from the purely unit-demand setting in Papadimitriou et al. [2021] and Braverman et al. [2022] do not carry over once we are required to allocate more than one item per round.

In weekly discussions after returning to our home universities, Tristan Pollner and myself continued studying the problem on additive and unit-demand valuations. Only several months later, we had a breakthrough, which was proposed by Tristan Pollner: When using the pivotal sampling subroutine, we can control the number of items being allocated. Adapting ideas from the two-proposal algorithm [Papadimitriou et al., 2021], we had the hope that this would allow us to solve the problem. In order to get the correlation bound to work, Tristan Pollner and myself tried several candidates for the function f used in the induction in Lemma 4.3.11. After several weeks, we ended up finding one to prove the main ingredient. Afterwards, I was able to extend the results beyond two-point distributions in Section 4.5 fairly easily. Also, coming up with the simplified variant in Section 4.4 took me only a couple of days afterwards.

Chapter 5: Truthful Mechanisms for Two-Sided Markets via Prophet Inequalities

The results of this chapter are contained in Braun and Kesselheim [2021] and Braun and Kesselheim [2023b], which is joint work with Thomas Kesselheim.

This work was started after attending a talk by Bart de Keijzer at the Day of Computational Game Theory 2020 in Enschede, Netherlands about bilateral trade and two-sided markets. After attending the talk, I discussed with Thomas Kesselheim that the techniques from the Prophet Inequality literature could easily be applied to get a (not tight) $1/2$-approximation for bilateral trade, as stated in Section 5.2.

To learn more about how to use Prophet Inequalities in two-sided environments, Thomas Kesselheim and myself started studying further directions. During this process, we figured out that the Prophet Inequality for XOS combinatorial auctions [Feldman et al., 2015, Dütting et al., 2020] could also be applied to two-sided markets, leading to the results in Section 5.5. Looking for other generalizations, I proposed studying matroid double auctions. As constant-factor Prophet Inequalities for matroids require dynamic prices, it was clear from the beginning that results could not be directly applied in a black-box manner. Still, we had the hope to make use of the tools from Prophet Inequalities to improve the previous guarantees for this problem mentioned in Section 1.3.3.

Thomas Kesselheim and myself made significant progress in this problem during independent work of each of us and joint discussions over a period of several months. After several months, the first breakthrough was my proof in Section 5.4 which uses a weaker variant of the budget balance constraint that we were initially hoping for. Having this proof, Thomas Kesselheim suggested a relaxation of the optimal (first best) social welfare which we use in Section 5.3. Given our previous insights from the last couple of months and this relaxation, I could finally prove the main result with the stronger notion of budget balance.

To make use of all the insights we had gained, I developed mechanisms and proved complementing theorems for knapsack double auctions with strong (Section 5.6) and weak budget balance (Section 5.7).

Chapter 6: Asymptotically Optimal Welfare of Posted Pricing with MHR Distributions

The results in this chapter are contained in Braun et al. [2021], which is joint work with Matthias Buttkus and Thomas Kesselheim.

This project was initiated before I started my PhD studies. In particular, Matthias Buttkus and Thomas Kesselheim had already worked on this project and were able to prove results on separable valuation functions (these results are not contained in this thesis, but appear in the paper) and the optimality results from Section 6.3.

Thomas Kesselheim and myself were then complementing this work by considering independent item valuations. As a starting point, we discussed several relaxations of the expected offline optimum. In the case that we have exactly as many items as buyers, we ended up using $\mathbf{E}[\mathrm{OPT}] \leq \mathbf{E}\left[\sum_j \max_{i \in [n]} v_{i,j}\right]$. After several months of independent work of Thomas Kesselheim and myself and joint discussions, Thomas Kesselheim proposed a simplified version of the quantile allocation rule from Algorithm 12 for which we were able to show the desired result. With this proof in mind, Thomas Kesselheim proposed to use the ex-ante relaxation for $\mathbf{E}[\mathrm{OPT}]$ instead of $\mathbf{E}\left[\sum_j \max_{i \in [n]} v_{i,j}\right]$. Having this, we were able to prove the main result for dynamic prices, leading to Section 6.2.1. On the way, we also discovered the bound for static prices and independent item values which can be found in Section 6.2.2.

With this in mind, we were able to introduce the notion of MHR marginals for subadditive valuations and I could extend the guarantees to more general valuation functions in Section 6.4.

Chapter 7: The Secretary Problem with Predicted Additive Gap

The results in this chapter are contained in Braun and Sarkar [2023], which is joint work with Sherry Sarkar. The paper is currently under review.

Sherry Sarkar and myself started working on this project during the program on Data-Driven Decision Processes at the Simons Institute for the Theory of Computing at UC Berkeley, CA, USA. In particular, Sherry Sarkar was interested in studying the Secretary problem when we have access to a piece of information which measures the spread of the weights in the sequence. During several discussions, we were considering, for example, the range of the values or the entropy as possible pieces of information. While for some of them we were able to prove improved guarantees fairly easily, for others it was sometimes even impossible due to impossibility results we could come up with.

In a joint discussion of Sherry Sarkar, Thomas Kesselheim and myself at the Simons Institute, we came up with the idea of considering an additive or multiplicative gap between the highest and second highest weight in the sequence. In follow-up discussions, Sherry Sarkar and myself were able to prove a better-than-$1/e$ guarantee for this setting.

After returning to our home universities, Sherry Sarkar and myself continued working on this problem and synchronized via weekly meetings. After several months of discussions, I was able to also prove a generalized variant of our statement which did hold for any gap and not only for the gap between the best and second best element, leading to the results in Section 7.2. Also, we could use the same approach and make the argument work for inaccurate gaps with bounded error, the results in Section 7.4.

Sherry Sarkar and myself also discussed several extensions and proved the robust and consistent variant of our algorithm which can now be found in Section 7.3. To complement results, I ran simulations on several classes of weights of potential interest whose outcome is stated in Section 7.5.

Chapter 2

A Few General Preliminaries

Before we dive into the technical part, we discuss a couple of preliminaries which are important throughout this thesis. Precise problem formulations, definitions and notation which is chapter-specific can be found in the preliminaries section in the respective chapter. We recap different classes of valuation functions, feasibility constraints and discuss how to model strategic behavior of agents.

Notation

For a natural number N, we denote by $[N]$ the set of natural numbers which are at most N, that is $[N] := \{1, \dots, N\}$. For a set X, we denote by 2^X the power set of X. By $\mathbb{R}_{\geq 0}$ we denote the set of non-negative real numbers.

By $(\cdot)^+$ we mean $\max(\cdot, 0)$. The indicator function $\mathbb{1}_{\mathrm{A}}$ is equal to one if A holds, otherwise it is zero. We use $\mathbb{1}_A$ and $\mathbb{1}[A]$ interchangeably.

For $q \in [0, 1]$, we denote by $\mathrm{Ber}(q)$ a Bernoulli random variable which is one with probability q and zero otherwise.

Valuation functions

For a set of m (potentially heterogeneous) items M, a valuation function is a mapping $v : 2^M \to \mathbb{R}_{\geq 0}$ which maps every subset of items to a non-negative real number. Usually, we assume that valuation functions are non-negative and finite for any bundle as well as monotone and normalized, i.e. $v(S) \leq v(S')$ for $S \subseteq S' \subseteq M$ and $v(\emptyset) = 0$. In this thesis, different classes of valuation functions will play an important role:

- *Additive*: A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is *additive* if and only if there are numbers $c_1, \dots, c_m \in \mathbb{R}_{\geq 0}$ such that for any $S \subseteq M$ we have $v(S) = \sum_{j \in S} c_j$.

- *Unit-Demand*: A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is *unit-demand* if and only if there are numbers $c_1, \dots, c_m \in \mathbb{R}_{\geq 0}$ such that for any $S \subseteq M$ we have $v(S) = \max_{j \in S} c_j$.

- *K-Demand/Multi-Demand*: A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is *K-demand*, also called *multi-demand*, if and only if there are numbers $c_1, \dots, c_m \in \mathbb{R}_{\geq 0}$ such that for any $S \subseteq M$ we have $v(S) = \max_{S' \subseteq S: |S'| \leq K} \sum_{j \in S'} c_j$.

We usually denote the numbers c_j by $v(\{j\})$ if the context is clear. Also, as we usually assume that valuation functions are buyer specific, we call the valuation function v_i and denote the respective $v_i(\{j\})$ by $v_{i,j}$ for shorter notation.

- *XOS*: A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is *XOS* (also called *fractionally subadditive*) if and only if there are additive functions $a_1, \dots, a_t$ such that for every $S \subseteq M$ we have $v(S) = \max_{\ell \in [t]} a_\ell(S)$.
- *Subadditive*: A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is *subadditive* if and only if for every two sets S and S', we have $v(S \cup S') \leq v(S) + v(S')$.

Subadditive valuation functions are also called complement-free. Interestingly, all of the above valuations are complement-free. In other words, we have

$$\text{Unit-Demand} \subseteq K\text{-Demand} \subseteq \text{XOS} \subseteq \text{Subadditive} \enspace ,$$

and of course also Additive $\subseteq$ XOS. A natural extension of XOS to functions which admit complementarities are MPH-k valuations [Feige et al., 2015].

- *MPH-k*: Consider a valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$. A hypergraph representation of the function v is a set function which satisfies $v(S) = \sum_{X \subseteq S} w(X)$. We call a set X with $w(X) \neq 0$ a hyperedge of w, and a positive-hyperedge if $w(X) > 0$. The *rank* of a hypergraph representation w is the cardinality of the largest hyperedge in w. If the hypergraph representation of v only contains non-negative hyperedges, we call this a positive-hyperedge-k function (PH-k) as introduced by Abraham et al. [2012]. The definition of the MPH-k hierarchy now represents a valuation function v as the maximum over a set of PH-k functions.

  Definition 2.0.1. *(Maximum-over-Positive-Hypergraph-k [Feige et al., 2015]) A monotone valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is MPH-k if there is a set $\{v_\ell\}_{\ell \in \mathcal{L}}$ of PH-k functions such that $v(S) = \max_{\ell \in \mathcal{L}} v_\ell(S)$ for any $S \subseteq M$ and arbitrary index set $\mathcal{L}$.*

This MPH-k hierarchy is containing very general valuation functions. In particular, any monotone valuation function is contained in one level of the MPH-k hierarchy. Also, note the inclusions

$$\text{XOS} = \text{MPH-1} \subseteq \text{MPH-2} \subseteq \cdots \subseteq \text{MPH-}k \subseteq \cdots \subseteq \text{MPH-}m = \text{All Monotone Valuations.}$$

As a side remark, there are also similar hierarchies which contain submodular or subadditive valuations in their lowest level instead of XOS valuations [Feige and Izsak, 2013, Chen et al., 2019].

Feasibility Constraints

In several chapters of this thesis, we assume that there is a constraint family over the set of buyers which we can accept. That is, for a set of buyers[1] $[n]$, there is a family $\mathcal{I} \subseteq 2^{[n]}$. Any set of buyers $X \subseteq [n]$ who are allocated (at least) one item needs to satisfy $X \in \mathcal{I}$. Most important for this thesis are matroid and knapsack constraints which are defined as follows.

[1] We note that in Chapter 5, we denote the set of buyers by $\mathcal{B}$ instead of $[n]$. Still, the following definitions also apply.

- *Matroid*: A *matroid* $([n], \mathcal{I})$ over ground set $[n]$ with non-empty set system $\mathcal{I} \subseteq 2^{[n]}$ is defined via the following properties. For two subsets $X \subseteq Y$ of $[n]$ with $Y \in \mathcal{I}$, also $X \in \mathcal{I}$. And for $X, Y \in \mathcal{I}$ with $|X| < |Y|$ there is a $y \in Y \setminus X$ such that $X \cup \{y\} \in \mathcal{I}$. We call sets in $\mathcal{I}$ *independent*.

- *Knapsack*: Without loss of generality, we assume a *knapsack* with capacity one. Each of the buyers in $[n]$ has a weight $w_i \in [0, 1]$. The set of buyers X who are allocated (at least) one item needs to satisfy $\sum_{i \in X} w_i \leq 1$.

For a matroid $([n], \mathcal{I})$, we define its *contraction* by $T \subseteq [n]$ as the following matroid with ground set $[n] \setminus T$: A set S is independent in the *contracted matroid* if $S \cup T_0$ is independent in the original matroid for any maximal independent subset $T_0 \in \mathcal{I}$ of T. For additional basic concepts concerning matroids, such as *rank* or *span*, we refer the reader to standard textbooks, for example Chapter 39 in Schrijver [2003].

Strategic Considerations

When buyers are behaving strategically, we aim to design allocation *mechanisms*. A (direct revelation) mechanism takes as input a vector of valuation functions which are reported by agents. Agents can report any possible valuation in their space of valuation functions V_i, not necessarily their true one. A mechanism outputs an allocation of items to agents $(X_1, \dots, X_n)$, as well as payments $(P_1, \dots, P_n)$. Agents are assumed to maximize *utility*. Fixing a valuation function v_i for buyer i, the (quasi-linear) utility of buyer i for being allocated bundle $X_i \subseteq M$ at price P_i is given by $u_i(X_i) = v_i(X_i) - P_i$.

Mechanisms are usually designed to fulfill the following desirable properties:

- *Dominant Strategy Incentive Compatibility* (DSIC): It is a dominant strategy for every agent to report her true valuation independent of the other agents' behavior. That is, no matter which valuation functions other agents report, it is always in the interest of an agent to report her true valuation.

- *Individual Rationality* (IR): When playing this dominant strategy, no agent decreases her utility by participating in the mechanism.

An important class of mechanisms which will be used multiple times in this thesis are *sequential posted prices mechanisms*. Every item $j \in M$ has a price $p_j^{(i)}$ which can either be *anonymous*, *static* or *dynamic*. In the latter case, prices may depend on the partial allocation at the arrival of buyer i. Prices which do not depend on the partial allocation are called *static*. Prices are called *anonymous* if they do not depend on the identity of the agent under consideration.

In sequential posted pricing mechanisms, buyers arrive sequentially and purchase the most desirable available bundle given the item prices. Buyers' payments are determined via the received bundle, so buyer i receiving item bundle X_i pays $\sum_{j \in X_i} p_j^{(i)}$. Posted prices mechanisms are DSIC and IR by design, so it is without loss to assume that buyers purchase their most preferred bundle with respect to their true valuation function v_i.

In the following pseudo code in Algorithm 3, we state the posted pricing mechanism with static item prices. In the variant with dynamic prices, p_j can depend in addition on the identity of buyer i as well as the set of previously allocated items.

Mechanisms like this will play a key role in Chapter 3 and Chapter 6.

Algorithm 3: Sequential Posted Pricing Mechanism

1 Compute item prices p_j based on $\mathcal{D}_1, \dots, \mathcal{D}_n$ before the first arrival
2 As buyer i arrives:
3   Buyer i buys bundle $X_i \subseteq M \setminus (\cup_{i'<i} X_{i'})$ among the available items maximizing her utility $v_i(X_i) - \sum_{j \in X_i} p_j$

Objective Functions

Throughout this thesis, our goal is to maximize a linear objective function. More precisely, we take the sum over the contribution of each individual agent and aim to

$$\text{maximize} \sum_i v_i(X_i) \ ,$$

where $(X_1, \dots, X_n)$ is the allocation computed by our algorithm (which may need to satisfy the feasibility constraints). In single-selection settings, this is simply the value of the agent who is allocated the item. There will be more detailed interpretations on a per-chapter basis in the respective preliminaries section.

As stated in Section 1.3, our goal is to derive bounds on the *competitive/approximation ratios* of our algorithms.

Chapter 3

Simplified Prophet Inequalities for Combinatorial Auctions and Matroids

Prophet Inequalities are an important tool to understand posted-pricing mechanisms for combinatorial auctions. In particular, the competitive ratio of a pricing-based Prophet Inequality directly implies a DSIC and IR posted-prices mechanism with the same approximation guarantee. In the past years, our understanding of pricing-based Prophet Inequalities rose constantly and significantly by a line of inspiring work, see e.g. Feldman et al. [2015], Dütting et al. [2020], Gravin and Wang [2019], Dütting et al. [2020] and Correa et al. [2022] among others.

For XOS and MPH-k combinatorial auctions [Feldman et al., 2015, Dütting et al., 2020], the common approach so far was to state a vector of static and anonymous item prices. Agents arrive one-by-one and can choose their most desired bundle among the remaining items. In order to derive a desirable competitive ratio with respect to the expected offline optimum, the algorithm designer is required to define prices carefully and tailored with respect to the class of valuation functions under consideration. In Feldman et al. [2015], item prices for XOS combinatorial auctions were set to be half of the expected contribution of an item to the expected offline optimum. Later, Dütting et al. [2020] discovered that it is sufficient to argue in the full information setting in contrast to dealing with random valuation profiles. Still, also their arguments are involved and require a deep understanding of the underlying habits. In contrast, easily accessible ideas concerning pricing-based Prophet Inequalities in combinatorial auctions are rare in the literature. Therefore, we try to advance our understanding of the following question:

What is the simplest way to prove pricing-based Prophet Inequalities?

In this chapter, in order to address this question, we derive simplified proofs for existing Prophet Inequalities in XOS and MPH-k combinatorial auctions which work as follows: First, we make use of a linear program with variables p_j for any item j corresponding to the static and anonymous item prices. For a variable assignment $(p_j)_{j \in M}$, the objective function of the LP will be non-negative if the posted-pricing mechanism with price vector $(p_j)_{j \in M}$ achieves a competitive ratio of $\zeta \in [0, 1]$ (Section 3.2.1). To show that there exists a solution to the LP whose corresponding prices achieve the competitive ratio, we want to argue that there always is a feasible primal solution with non-negative objective value. To prove this, we use strong LP duality and move into dual space. We interpret dual variables and constraints (Section 3.2.2) and show that every dual feasible solution has a non-negative objective value (Section 3.3). In particular, interpreting dual variables as probabilities over subsets of items, the dual constraints give a bound on the probability with which an item can be absent. With this interpretation in mind, we reformulate the dual objective to argue that it is non-negative for any dual feasible solution. As a consequence, at least the optimal primal solution also has to have a non-negative objective value. This implies that the corresponding prices $(p_j)_{j \in M}$ lead to a sequential posted-pricing mechanism with a competitive ratio ζ.

Using this technique, we can proof the following guarantees for XOS valuations. Originally, this statement was shown by Feldman et al. [2015].

Theorem 3.0.1. *If all buyers have XOS valuation functions, there exist static item prices $(p_j)_{j \in M}$ such that the sequential posted-pricing mechanism (Algorithm 3) satisfies*

$$\mathbf{E}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] \geq \frac{1}{2} \cdot \mathbf{E}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \quad ,$$

where $\mathrm{OPT}(\mathbf{v})$ *is the offline optimum allocation on valuation profile* $\mathbf{v}$*, i.e.* $\mathrm{OPT}(\mathbf{v}) := \arg\max_{(X_1,\ldots,X_n)} \sum_i v_i(X_i)$.

To allow complementarities across items in the valuations, if all buyers have MPH-k valuation functions, we can prove the following theorem for any $k \geq 2$ which slightly improves the previous best guarantee of $1/_{4k-2}$ by Dütting et al. [2020].

Theorem 3.0.2. *If all buyers have MPH-k valuation functions, there exist static item prices $(p_j)_{j \in M}$ such that the sequential posted-pricing mechanism (Algorithm 3) satisfies*

$$\mathbf{E}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] \geq \zeta \cdot \mathbf{E}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \quad ,$$

where $\zeta = 1/(2k+2\sqrt{k(k-1)}-1)$ *and* $\mathrm{OPT}(\mathbf{v})$ *is again the offline optimum allocation as defined before.*

Chapter Organization and Remarks

The part on combinatorial auctions in this chapter is based on *Simplified Prophet Inequalities for Combinatorial Auctions* [Braun and Kesselheim, 2023a], which is joint work with Thomas Kesselheim. Further bibliographic notes can be found in Section 1.5.

We discuss the framework for combinatorial auctions in Section 3.2 and Section 3.3. A proof for matroid feasibility constraints using LP duality is given afterwards in Section 3.4.

3.1 Notation and Preliminaries

In the first part of this chapter, we consider combinatorial auctions. Afterwards, in Section 3.4, we shift our perspective towards matroid feasibility constraints.

Combinatorial Auctions

Formally, in this chapter, we consider the following setting: There is a set of m heterogeneous items M and a sequence of n agents arriving online one-by-one. As agent i arrives, we get to know her valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ and agent i buys the bundle of (currently unassigned) items which maximizes her quasi-linear utility. We assume that each $v_i \sim \mathcal{D}_i$ is drawn independently from a publicly known, not necessarily identical distribution $\mathcal{D}_i$ and denote by $\mathbf{v} \sim \times_{i=1}^n \mathcal{D}_i$ the valuation profile of all agents.

An *allocation* $\mathbf{X} = (X_i)_{i\in[n]}$ is a vector of item bundles such that agent i is allocated bundle X_i and for two agents $i \neq i'$, we have $X_i \cap X_{i'} = \emptyset$. The *social welfare* of an allocation $\mathbf{X}$ given valuation profile $\mathbf{v}$ is defined as $\mathbf{v}(\mathbf{X}) := \sum_{i\in[n]} v_i(X_i)$. We compare the performance of the sequential posted-pricing mechanism to the expected offline optimal social welfare and aim for guarantees of the form $\mathbf{E}_{\mathbf{v}}[\mathbf{v}(\mathbf{X})] \geq \zeta \cdot \mathbf{E}_{\mathbf{v}}[\max_{\mathbf{X}^*} \mathbf{v}(\mathbf{X}^*)]$.

Matroids

In Section 3.4, we consider matroid Prophet Inequalities. That is, we have m homogeneous items which are all of the same kind. Every buyer has a unit-demand valuation function. As items are homogeneous, every unit-demand valuation function boils down to a single non-negative real number v_i which is buyer i's value for being allocated one of the items. Buyers in $[n]$ arrive one-by-one and the value v_i is drawn independently from a publicly known distribution $\mathcal{D}_i$.

In addition, we are given a matroid $([n], \mathcal{I})$ upfront and we can allocate items to buyers in $[n]$ as long as they form an independent set in the matroid. Our goal is to show the existence of *dynamic* prices such that the posted-pricing mechanism achieves a reasonable fraction of the expected offline optimal social welfare achievable in the matroid. In order words, we compare to the expected maximum weight basis in the matroid $\mathbf{E}_{\mathbf{v}}[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))] = \mathbf{E}_{\mathbf{v}}[\max_{X\in\mathcal{I}} \sum_{i\in X} v_i]$.

3.2 General Framework

In order to derive proofs for the competitive ratios in the combinatorial auction setting, we start with a lower bound on the expected social welfare of our algorithm. By $\mathrm{OPT}_i(\mathbf{v})$ we denote the (possibly empty) bundle of items which agent i gets in the optimal allocation on valuation profile $\mathbf{v}$.

Lemma 3.2.1. *For any combinatorial auction with monotone valuation functions, the social welfare of the sequential posted-prices mechanism with price vector $\mathbf{p} = (p_j)_{j\in M}$ fulfills for any $\beta \in [0,1]$*

$$\mathbf{E}_{\mathbf{v}}[\mathbf{v}(\mathrm{ALG}(\mathbf{v}))] \geq \min_{T\subseteq M} \left(\sum_{j\in T} p_j + \sum_{i=1}^n \mathbf{E}_{\mathbf{v}} \left[\sum_{S\subseteq M} \beta \left(v_i(S\setminus T) - \sum_{j\in S\setminus T} p_j \right) \mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})} \right] \right).$$

The proof of this lemma follows standard steps in the Prophet Inequality literature [Feldman et al., 2015, Dütting et al., 2020, Dütting et al., 2020, Gravin and Wang, 2019, Correa et al., 2022]. In particular, we split the social welfare into revenue and utility and bound each quantity separately. As a final step, we lower bound the social welfare by allowing an adversary to choose the set of allocated items T.

Proof of Lemma 3.2.1. We first split $\mathbf{E}_{\mathbf{v}}[\mathbf{v}(\mathrm{ALG}(\mathbf{v}))]$ into revenue and utility and bound each separately.

Let $T(\mathbf{v})$ denote the state of the set of allocated items T at the end of the allocation process when running the algorithm on valuation profile $\mathbf{v}$. In addition, let $T_i(\mathbf{v})$ denote the state of the set of allocated items T before the arrival of buyer i.

Revenue. The expected revenue of the algorithm is

$$\mathbf{E}_{\mathbf{v}}\left[\mathrm{rev}(\mathbf{v},\mathbf{p})\right] = \mathbf{E}_{\mathbf{v}}\left[\sum_{j\in T(\mathbf{v})} p_j\right] .$$

Utility. In order to lower bound the utility, consider an arbitrary buyer i. First, we consider an independent sample $\mathbf{v}'$ from the joint distribution $\times_{i=1}^{n}\mathcal{D}_i$. Note that buyer i could either buy nothing and hence obtains a utility which is non-negative. Another option is to buy the bundle $\mathrm{OPT}_i((v_i,\mathbf{v}'_{-i}))\setminus T((v'_i,\mathbf{v}_{-i}))$. To see that this is a feasible bundle, observe that the set of allocated items before the arrival of buyer i is independent of her valuation v_i. In particular, we have $T_i(\mathbf{v}) = T_i((v'_i,\mathbf{v}_{-i})) \subseteq T((v'_i,\mathbf{v}_{-i}))$, where the last inclusion is due to the fact that we only assign more items in the ongoing allocation process.

Taking the expectation over $\mathbf{v}'$, the utility of buyer i can be lower bounded by

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right]\\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{S\subseteq M}\left(v_i(S\setminus T((v'_i,\mathbf{v}_{-i}))) - \sum_{j\in S\setminus T((v'_i,\mathbf{v}_{-i}))} p_j\right)^{+} \mathbb{1}_{S=\mathrm{OPT}_i((v_i,\mathbf{v}'_{-i}))}\right]\\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{S\subseteq M}\beta\left(v_i(S\setminus T((v'_i,\mathbf{v}_{-i}))) - \sum_{j\in S\setminus T((v'_i,\mathbf{v}_{-i}))} p_j\right)^{+} \mathbb{1}_{S=\mathrm{OPT}_i((v_i,\mathbf{v}'_{-i}))}\right]\\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{S\subseteq M}\beta\left(v_i(S\setminus T((v'_i,\mathbf{v}_{-i}))) - \sum_{j\in S\setminus T((v'_i,\mathbf{v}_{-i}))} p_j\right) \mathbb{1}_{S=\mathrm{OPT}_i((v_i,\mathbf{v}'_{-i}))}\right]\\
&= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{S\subseteq M}\beta\left(v_i(S\setminus T(\mathbf{v}')) - \sum_{j\in S\setminus T(\mathbf{v}')} p_j\right) \mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right] .
\end{aligned}$$

Observe that in the second inequality we multiply a non-negative term by $\beta \in [0,1]$ and drop the $(\cdot)^+$ in the third inequality[1]. The last equality uses independence and the fact that $\mathbf{v}$ and $\mathbf{v}'$ are identically distributed.

Combination. As a consequence, summing over the lower bound of the utility for all

[1] Correa et al. [2022] use an equivalent line of arguments, but keep the $(\cdot)^+$ term in the utility instead of multiplying by β. Still, as a matter of fact, this will make the problem non-linear and hence, our arguments do not apply. Also, their arguments do not transfer to MPH-k functions, but require to bound the bundle size of requested items.

i and adding the revenue, we get

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] \\
&= \mathbf{E}_{\mathbf{v}}\left[\mathrm{rev}(\mathbf{v},\mathbf{p})\right] + \sum_{i=1}^{n} \mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right] \\
&\geq \mathbf{E}_{\mathbf{v}}\left[\sum_{j\in T(\mathbf{v})} p_j\right] + \sum_{i=1}^{n} \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{S\subseteq M} \beta\left(v_i(S\setminus T(\mathbf{v}')) - \sum_{j\in S\setminus T(\mathbf{v}')} p_j\right)\mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right] \\
&\geq \min_{T\subseteq M}\left(\sum_{j\in T} p_j + \sum_{i=1}^{n} \mathbf{E}_{\mathbf{v}}\left[\sum_{S\subseteq M} \beta\left(v_i(S\setminus T) - \sum_{j\in S\setminus T} p_j\right)\mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right]\right) \quad ,
\end{aligned}$$

where in the last inequality we lower bound the expectation by the worst possible choice for the set of allocated items T. □

3.2.1 An LP Formulation and its Dual

Having the lower bound on the social welfare obtained by the algorithm, we actually want to show that for any set $T \subseteq M$, the lower bound in Lemma 3.2.1 is at least as large as a ζ-fraction of the expected offline optimum. Interpreting this as a constraint for any set $T \subseteq M$, we can formulate an LP which has a non-negative objective value whenever the desired competitive ratio ζ can be achieved. Thus, in order to prove a competitive ratio, we only need to argue about the LP. For convenience in the remainder of this chapter, we define α such that $\zeta = \frac{1}{\alpha}$.

The LP has variables p_j for any item $j \in M$ where p_j corresponds to the static and anonymous item price for item j. In addition, there are slack variables ℓ_+ and ℓ_- which indicate if the desired competitive ratio can be achieved or not.

$$\begin{aligned}
\max \quad & \ell_+ - \ell_- \\
\text{s.t.} \quad & \sum_{j\in T} p_j + \sum_{i=1}^{n} \mathbf{E}_{\mathbf{v}}\left[\sum_{S\subseteq M} \beta\left(v_i(S\setminus T) - \sum_{j\in S\setminus T} p_j\right)\mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right] \\
& \qquad\qquad \geq \frac{1}{\alpha}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] + \ell_+ - \ell_- && \text{for all } T \subseteq M \\
& p_j \geq 0 && \text{for all } j \in M \\
& \ell_+, \ell_- \geq 0.
\end{aligned}$$

We note that Dütting et al. [2020] use a similar variant of this LP and its dual in order to show the existence of prices for subadditive combinatorial auctions[2]. Concerning the constraints for any set $T \subseteq M$, observe that they can be rearranged to

[2] In contrast to our LP, they draw the bundle S from some probability distribution whereas we set it equal to the bundle of items which agent i gets in the offline optimum on $\mathbf{v}$. In addition, they do subtract the prices for all items in S whereas in our formulation, it is essential to only consider prices of items in $S \setminus T$. Subtracting prices for any item in S will result in a worse competitive ratio already in the case of a single item.

$$\sum_{j\in M} p_j \left(\beta \sum_{i=1}^{n} \mathbf{Pr_v}\left[j \in \mathrm{OPT}_i(\mathbf{v})\right] \mathbb{1}_{j\notin T} - \mathbb{1}_{j\in T}\right) + \ell_+ - \ell_-$$
$$\leq \beta \sum_{i=1}^{n} \sum_{S\subseteq M} \mathbf{E_v}\left[v_i(S\setminus T)\mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right] - \frac{1}{\alpha}\mathbf{E_v}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right] \ .$$

In order to argue that the LP has a non-negative objective value, we can consider the dual program and use strong duality. In particular, we will argue that any feasible dual solution has an objective value which is non-negative. Via strong duality, this directly implies that at least the optimal primal solution has a non-negative objective value and hence, the corresponding prices lead to the desired competitive ratio.

The dual of the LP introduced above has variables $\mu_T \geq 0$ for every set $T \subseteq M$ and is given by

$$\begin{aligned}
\min \quad & \sum_T \mu_T \left(\sum_{i=1}^{n} \beta \mathbf{E_v}\left[\sum_{S\subseteq M} v_i(S\setminus T)\mathbb{1}_{S=\mathrm{OPT}_i(\mathbf{v})}\right] - \frac{1}{\alpha}\mathbf{E_v}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right]\right) \\
\text{s.t.} \quad & \sum_T \mu_T \left(\beta \sum_{i=1}^{n} \mathbf{Pr_v}\left[j \in \mathrm{OPT}_i(\mathbf{v})\right] \mathbb{1}_{j\notin T} - \mathbb{1}_{j\in T}\right) \geq 0 && \text{for all } j \in M \\
& \sum_T \mu_T = 1 \\
& \mu_T \geq 0 && \text{for all } T \subseteq M.
\end{aligned}$$

Having this, we are able to state the lemma which will simplify the proof of pricing-based Prophet Inequalities.

Lemma 3.2.2. *For any combinatorial auction with monotone valuation functions, there exists a sequential posted-prices mechanism with price vector $\mathbf{p} = (p_j)_{j\in M}$ which is $1/\alpha$-competitive with respect to the expected offline optimum if the objective value of the dual program is non-negative for any feasible dual solution.*

Observe that by the construction above, showing the existence of suitable prices boils down to arguing about the dual of a linear program.

3.2.2 Understanding the Dual Program

Before we continue to derive our competitive ratios, we will start by gaining a better understanding of the dual.

Dual Constraints. First, note that $\sum_T \mu_T = 1$ and $\mu_T \geq 0$ for any T. Hence, we can interpret the vector $(\mu_T)_T$ as a probability distribution over subsets $T \subseteq M$.

Second, by the monotonicity of valuations, we can without loss of generality assume that $\sum_{i=1}^{n} \mathbf{Pr_v}\left[j \in \mathrm{OPT}_i(\mathbf{v})\right] = 1$ as the optimum will always allocate any item in any realization $\mathbf{v}$. As a consequence, we can reformulate the dual constraints

$$\sum_T \mu_T \left(\beta \sum_{i=1}^{n} \mathbf{Pr_v}\left[j \in \mathrm{OPT}_i(\mathbf{v})\right] \mathbb{1}_{j\notin T} - \mathbb{1}_{j\in T}\right) \geq 0$$

for any item $j \in M$ as follows:

$$\begin{aligned}
&\sum_T \mu_T \left(\beta \sum_{i=1}^n \mathbf{Pr}_{\mathbf{v}} \left[j \in \mathrm{OPT}_i(\mathbf{v}) \right] \mathbb{1}_{j \notin T} - \mathbb{1}_{j \in T} \right) \\
&\quad = \sum_T \mu_T \left(\beta \mathbb{1}_{j \notin T} - \mathbb{1}_{j \in T} \right) \\
&\quad = \beta \sum_T \mu_T \mathbb{1}_{j \notin T} - \sum_T \mu_T \mathbb{1}_{j \in T} \\
&\quad = \beta \mathbf{Pr}_\mu \left[j \notin T \right] - \mathbf{Pr}_\mu \left[j \in T \right] \geq 0 \ ,
\end{aligned}$$

where we denote by $\mathbf{Pr}_\mu [j \notin T]$ the probability that some item j is not contained in a set $T \subseteq M$ sampled with respect to distribution $(\mu_T)_T$. In other words, $\mathbf{Pr}_\mu [j \notin T] := \mathbf{Pr}_{T \sim \mu} [j \notin T] := \sum_T \mu_T \mathbb{1}_{j \notin T}$.

Using that $\mathbf{Pr}_\mu [j \notin T] + \mathbf{Pr}_\mu [j \in T] = 1$, the dual constraints for any item $j \in M$ are equivalent to our *first key property*:

$$\mathbf{Pr}_\mu \left[j \notin T \right] \geq \frac{1}{1+\beta} \qquad \text{and} \qquad \mathbf{Pr}_\mu \left[j \in T \right] \leq \frac{\beta}{1+\beta} \ . \tag{3.1}$$

We will later set $\beta = 1$ for a single item and XOS functions. In this case, Inequality (3.1) simply states that any item j can only be in T with probability at most $\frac{1}{2}$.

Dual Objective. To make the dual objective more accessible, we change the order of summation:

$$\begin{aligned}
\text{dual obj.} &= \sum_T \mu_T \sum_{i=1}^n \beta \mathbf{E}_{\mathbf{v}} \left[\sum_{S \subseteq M} v_i(S \setminus T) \mathbb{1}_{S = \mathrm{OPT}_i(\mathbf{v})} \right] - \frac{1}{\alpha} \mathbf{E}_{\mathbf{v}} \left[\mathbf{v}(\mathrm{OPT}(\mathbf{v})) \right] \\
&= \sum_{i=1}^n \sum_{S \subseteq M} \mathbf{E}_{\mathbf{v}} \left[\mathbb{1}_{S = \mathrm{OPT}_i(\mathbf{v})} \beta \sum_T \mu_T v_i(S \setminus T) \right] - \frac{1}{\alpha} \mathbf{E}_{\mathbf{v}} \left[\sum_{i=1}^n \sum_{S \subseteq M} \mathbb{1}_{S = \mathrm{OPT}_i(\mathbf{v})} v_i(S) \right] \\
&= \sum_{i=1}^n \sum_{S \subseteq M} \mathbf{E}_{\mathbf{v}} \left[\mathbb{1}_{S = \mathrm{OPT}_i(\mathbf{v})} \left(\beta \sum_T \mu_T v_i(S \setminus T) - \frac{1}{\alpha} v_i(S) \right) \right]
\end{aligned}$$

Instead of arguing that the dual objective is non-negative, we will show the following for suitable choices of α and β: If a vector $(\mu_T)_T$ is feasible with respect to the dual, then the term

$$\beta \sum_T \mu_T v_i(S \setminus T) - \frac{1}{\alpha} v_i(S)$$

is non-negative for any v_i and S. In particular, we will show the following equivalent claim for the respective choices of α and β:

Claim 3.2.3. *For any set S, any XOS/MPH-k function v_i and any dual feasible solution $(\mu_T)_T$, the following holds:*

$$\beta \sum_T \mu_T v_i(S \setminus T) \geq \frac{1}{\alpha} v_i(S) \quad \text{or equivalently for } \beta > 0 \quad \sum_T \mu_T v_i(S \setminus T) \geq \frac{1}{\alpha\beta} v_i(S) \tag{3.2}$$

This will be our *second key ingredient*.
We can interpret it as follows: When drawing a set T with respect to distribution $(\mu_T)_T$, the value that remains from some fixed set S after removing T is still at least a $\frac{1}{\alpha\beta}$-fraction of the original value $v_i(S)$.
Having our two key ingredients, we are ready to derive the competitive ratios.

3.3 Deriving Competitive Ratios Easily

By the construction in Section 3.2, we are only required to argue that the dual objective is non-negative for the choices of α and β in the respective settings. As mentioned, we will show a stronger result, namely that Inequality (3.2) holds for any subset of items $S \subseteq M$ and any valuation function v_i which is XOS or MPH-k. We summarize the respective statements first and give proofs for each afterwards.

Lemma 3.3.1. *In the case of a single item, Inequality* (3.2) *in Claim 3.2.3 holds for* $\alpha = 2$ *and* $\beta = 1$.

Note that Lemma 3.3.1 implies a proof for Theorem 1.1.2 based on LP duality.

Lemma 3.3.2. *For XOS valuation functions, Inequality* (3.2) *in Claim 3.2.3 holds for* $\alpha = 2$ *and* $\beta = 1$.

Observe that these two lemmas directly correspond to the best known competitive ratios and are tight. In addition, note that the class of XOS valuation functions contains e.g. submodular functions and is equivalent to the class of MPH-1 valuations.

For MPH-k valuation functions, we first give a simplified proof of the competitive ratio of Dütting et al. [2020]. In particular, Dütting et al. [2020] introduced a reduction which allows to only argue about deterministic valuation functions instead of randomly drawn ones. Our proof will be even simpler: our argument only requires to consider the sizes of sets which are relevant in MPH-k valuations, the valuations as such do not play a role at all.

Lemma 3.3.3. *For MPH-k valuation functions, Inequality* (3.2) *in Claim 3.2.3 holds for* $\alpha = 4k - 2$ *and* $\beta = \frac{1}{2(k-1)}$ *when* $k \geq 2$.

Finally, for all $k \geq 2$, we show that we can get a tiny improvement in the competitive ratio compared to previously known results for MPH-k valuations as $2k + 2\sqrt{k(k-1)} - 1 < 4k - 2$.

Lemma 3.3.4. *For MPH-k valuation functions,* $\alpha = 2k + 2\sqrt{k(k-1)} - 1$ *and* $\beta = \sqrt{\frac{k}{k-1}} - 1$, *Inequality* (3.2) *in Claim 3.2.3 holds when* $k \geq 2$.

3.3.1 Warm-Up: A Single Item

Observe that in the case of a single item, the vector $(\mu_T)_T$ only has two entries $\mu_\emptyset$ and $\mu_{\{\text{item}\}}$.

Proof of Lemma 3.3.1. We overload notation and denote by v_i the value of agent i for the item in order to rewrite the left-hand side of Inequality (3.2) as follows:

$$\sum_T \mu_T v_i(S \setminus T) = \sum_T \mu_T \left(v_i \mathbb{1}_{S=\{\text{item}\}} \mathbb{1}_{T=\emptyset} \right) = \left(\sum_T \mu_T \mathbb{1}_{T=\emptyset} \right) v_i \mathbb{1}_{S=\{\text{item}\}}$$
$$= \mu_\emptyset \cdot v_i(S) = \mathbf{Pr}_\mu \left[T = \emptyset \right] \cdot v_i(S) \ .$$

Hence, what remains is to show that $\mathbf{Pr}_\mu [T = \emptyset] \geq \frac{1}{2}$ for any dual feasible solution. Indeed, observe that by Inequality (3.1) for $\beta = 1$, we have

$$\mathbf{Pr}_\mu \left[T = \emptyset \right] = \mathbf{Pr}_\mu \left[\text{item} \notin T \right] \geq \frac{1}{2} \ .$$
□

Note that we did only argue about the value of the dual variables μ_T and did not at all need to take a specific value v_i into account. In particular, it is only important that the probability distribution $(\mu_T)_T$ does not put too much mass on $T = \{\text{item}\}$ or in other words, $\mu_\emptyset$ is sufficiently large. The remarkable thing is that this argument nicely extends to XOS and MPH-k valuation functions without becoming much more involved. In the more general settings, we will consider (subsets of) items separately and argue in an equivalent way about $(\mu_T)_T$.

3.3.2 XOS valuation functions

We proceed in a similar way as in the proof of Lemma 3.3.1. Still, we need to take the combinatorial structure of the valuation functions into account.

Proof of Lemma 3.3.2. First, observe that for any XOS valuation function v_i, when fixing one of the additive supporting functions, we obtain a lower bound on the function's value. In particular, we can bound the value of $v_i(S \setminus T)$ from below by considering the additive function with which buyer i would evaluate the set S. We denote this function by w_i^S. Using this property, we can lower bound the left-hand side of Inequality (3.2) as follows:

$$\sum_T \mu_T v_i(S \setminus T) \geq \sum_T \mu_T w_i^S(S \setminus T) = \sum_T \mu_T \sum_{j \in S} w_i^S(\{j\}) \mathbb{1}_{j \notin T}$$
$$= \sum_{j \in S} w_i^S(\{j\}) \sum_T \mu_T \mathbb{1}_{j \notin T} = \sum_{j \in S} w_i^S(\{j\}) \mathbf{Pr}_\mu \left[j \notin T \right]$$
$$\geq \frac{1}{2} \sum_{j \in S} w_i^S(\{j\}) = \frac{1}{2} v_i(S) \ .$$

The first inequality holds due to the XOS property of v_i, the second step exploits that the function w_i^S is additive. In the last inequality, we made use of Inequality (3.1) for $\beta = 1$ to see that $\mathbf{Pr}_\mu [j \notin T] \geq \frac{1}{2}$. □

3.3.3 MPH-k valuation functions

Also the proof for MPH-k valuations follows a similar template. In contrast to XOS valuation functions, we have to take into account that items can complement each other. Further, we give a combined proof for Lemmas 3.3.3 and 3.3.4 with general $\alpha \geq 1$ and $\beta \in (0, 1]$ and use two observations afterwards to show the desired competitive ratios.

Proof of Lemmas 3.3.3 and 3.3.4. As in the case of XOS valuation functions, also for MPH-k valuations we can fix one of the supporting PH-k functions to obtain a lower bound on the value. In particular, we can bound the value of $v_i(S \setminus T)$ from below by only considering the PH-k function with which buyer i would evaluate the set S, denoted by $v_i^S(\cdot)$ with corresponding weights on hyperedges denoted by $w_i^S(\cdot)$. This implies the following lower bound on the left-hand side of Inequality (3.2):

$$\begin{aligned}\sum_T \mu_T v_i(S \setminus T) &\geq \sum_T \mu_T v_i^S(S \setminus T) \\ &= \sum_T \mu_T \sum_{X \subseteq S} w_i^S(X) \mathbb{1}_{X \cap T = \emptyset} \\ &= \sum_{X \subseteq S} w_i^S(X) \left(\sum_T \mu_T \mathbb{1}_{X \cap T = \emptyset} \right) \\ &= \sum_{X \subseteq S} w_i^S(X) \mathbf{Pr}_\mu \left[X \cap T = \emptyset \right] \quad .\end{aligned}$$

Next, we argue that for the respective choices of $\alpha \geq 1$ and $\beta \in (0,1]$, the term $\mathbf{Pr}_\mu [X \cap T = \emptyset]$ is at least as large as $\frac{1}{\alpha\beta}$ for any X with $|X| \leq k$ and any feasible dual solution.

To this end, first note that by the union bound and Inequality (3.1)

$$\mathbf{Pr}_\mu [\exists\, j \in X : j \in T] \leq \sum_{j \in X} \mathbf{Pr}_\mu [j \in T] \leq \sum_{j \in X} \frac{\beta}{1+\beta} \leq \frac{k\beta}{1+\beta} \quad ,$$

where in the last inequality, we used that $|X| \leq k$ via the MPH-k property. Hence, we can lower bound $\mathbf{Pr}_\mu [X \cap T = \emptyset]$ via

$$\mathbf{Pr}_\mu [X \cap T = \emptyset] = \mathbf{Pr}_\mu [\forall\, j \in X : j \notin T] = 1 - \mathbf{Pr}_\mu [\exists\, j \in X : j \in T] \geq 1 - \frac{k\beta}{1+\beta} \quad .$$

As a consequence, we obtain that

$$\sum_T \mu_T v_i(S \setminus T) \geq \sum_{X \subseteq S} w_i^S(X) \left(1 - \frac{k\beta}{1+\beta}\right) = \left(1 - \frac{k\beta}{1+\beta}\right) v_i(S) \quad .$$

We can conclude the proof of Lemma 3.3.3 by observing that for $\alpha = 4k - 2$ and $\beta = \frac{1}{2(k-1)}$, we have $1 - \frac{k\beta}{1+\beta} = \frac{1}{\alpha\beta}$.

The proof of Lemma 3.3.4 follows by checking that for $\alpha = 2k + 2\sqrt{k(k-1)} - 1$ and $\beta = \sqrt{\frac{k}{k-1}} - 1$, also $1 - \frac{k\beta}{1+\beta} = \frac{1}{\alpha\beta}$ which again leads to the desired result. □

One might wonder what is really going on. As a matter of fact, we made use of the property that the probability distribution $(\mu_T)_T$ cannot put too much mass on any item. In other words, the dual constraints ensure the following: When drawing a set T with respect to $(\mu_T)_T$, any item $j \in M$ is ensured to be not in T with a reasonably high probability. This implies for XOS and MPH-k functions that the values of $v_i(S \setminus T)$ and $v_i(S)$ are sufficiently close to each other.

3.4 Matroid Prophet Inequality via LP Duality

Let us now consider the setting of matroid Prophet Inequalities as studied in Kleinberg and Weinberg [2019]. Recall from Section 3.1 that we assume items to be identical and hence, buyers' valuations boil down to a single, non-negative, real number for being allocated one of the items. We denote by T the set of chosen buyers who are allocated an item and by T_i the set of chosen buyers before the arrival of buyer i. Our algorithm is based on dynamic prices for buyers. As buyer i arrives, if she can be feasibly added to T_i with respect to the matroid constraint $\mathcal{I}$, we offer a price $p_i(T_i)$. As prices are dynamic, the price $p_i(T_i)$ will depend on T_i.

Now, $\mathrm{OPT}(\mathbf{v})$ is a maximum weight basis in the matroid with weights $\mathbf{v}$. In addition, $\mathrm{OPT}(\mathbf{v} \mid T)$ denotes the maximum weight basis in the matroid contracted by T. Observe, that we need to ensure that $T \in \mathcal{I}$.

Theorem 3.4.1. *Let $\mathcal{M} = ([n], \mathcal{I})$ be a matroid. There exist prices $p_i(T_i)$ for any i and $T_i \subseteq [i-1]$ with $T_i \in \mathcal{I}$ which ensure*

$$\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] \geq \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \enspace .$$

Note that the competitive ratio of $1/2$ is tight here, as it is already tight in the single-item case (which corresponds to the 1-uniform matroid) by Example 1.1.1. Theorem 3.4.1 was originally shown by Kleinberg and Weinberg [2019]. Their original proof did not use LP duality but proceeded in a different way: They state a particular choice of prices $p_i(T_i)$ and show that these prices lead to the desired competitive ratios. Even though the following proof differs from their original one, we will later in Lemma 3.4.3 make use of a desirable property for matroids from Kleinberg and Weinberg [2019].

For the proof via LP duality, we start by proving the following lemma.

Lemma 3.4.2. *The value achieved by the algorithm with prices $p_i(T_i)$ for any i and $T_i \subseteq [i-1]$ with $T_i \in \mathcal{I}$ ensures*

$$\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] \geq \min_{T \in \mathcal{I}} \left(\sum_{i \in T} p_i(T_i) + \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v} \mid T)} (v_i - p_i(T_i)) \right] \right) \enspace .$$

Proof. In order to prove this theorem, we split $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right]$ into revenue and utility and consider each quantity separately.

Revenue. The expected revenue of the algorithm is

$$\mathbf{E}_{\mathbf{v}}\left[\mathrm{rev}(\mathbf{v}, \mathbf{p})\right] = \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in T(\mathbf{v})} p_i(T_i(\mathbf{v})) \right] \enspace ,$$

where $T(\mathbf{v})$ (and $T_i(\mathbf{v})$ respectively) denotes the set of chosen agents by the algorithm when the valuation profile is $\mathbf{v}$.

Utility. In order to lower bound the utility, consider an arbitrary buyer i. Note that we could add i to our solution, if her value exceeds her price $p_i(T_i(\mathbf{v}))$ and she can be feasibly added to T_i with respect to the matroid constraint $\mathcal{I}$. In addition, observe that T_i is independent of v_i, because only agents $1, \ldots, i-1$ are taken into account for T_i.

In particular, for any v_i', we have $T_i(\mathbf{v}) = T_i((v_i', \mathbf{v}_{-i}))$. As a consequence, we get the following bound:

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right] &\geq \mathbf{E}_{\mathbf{v}}\left[(v_i - p_i(T_i(\mathbf{v})))^+ \mathbb{1}_{i \notin \operatorname{span}(T_i(\mathbf{v}))}\right] \\
&= \mathbf{E}_{\mathbf{v}}\left[(v_i - p_i(T_i((v_i', \mathbf{v}_{-i}))))^+ \mathbb{1}_{i \notin \operatorname{span}\left(T_i((v_i', \mathbf{v}_{-i}))\right)}\right] \\
&\geq \mathbf{E}_{\mathbf{v}}\left[(v_i - p_i(T_i((v_i', \mathbf{v}_{-i})))) \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}_{-i}') \mid T(v_i', \mathbf{v}_{-i})\right)}\right] \quad .
\end{aligned}$$

In the last inequality we used that for any valuation profile $\tilde{\mathbf{v}}$, the following holds: if $i \in \mathrm{OPT}\left(\tilde{\mathbf{v}} \mid T(v_i', \mathbf{v}_{-i})\right)$, then i can be added to $T(v_i', \mathbf{v}_{-i})$ without violating the matroid constraint. In particular, as $T(v_i', \mathbf{v}_{-i}) \supseteq T_i(v_i', \mathbf{v}_{-i})$, the inequality holds.

Choosing $\mathbf{v}' \sim \times_{i=1}^n \mathcal{D}_i$ to be an independent and identically distributed sample from distribution $\times_{i=1}^n \mathcal{D}_i$, we get

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v_i - p_i(T_i((v_i', \mathbf{v}_{-i})))) \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}_{-i}') \mid T(v_i', \mathbf{v}_{-i})\right)}\right] \\
&= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v_i - p_i(T_i(\mathbf{v}'))) \mathbb{1}_{i \in \mathrm{OPT}(\mathbf{v} \mid T(\mathbf{v}'))}\right]
\end{aligned}$$

Summing over all i, we get

$$\sum_{i=1}^n \mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right] \geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v} \mid T(\mathbf{v}'))} (v_i - p_i(T_i(\mathbf{v}')))\right]$$

Combination. Combining revenue and utility, we get

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{ALG}(\mathbf{v})\right)\right] &= \mathbf{E}_{\mathbf{v}}\left[\mathrm{rev}(\mathbf{v},\mathbf{p})\right] + \sum_{i=1}^n \mathbf{E}_{\mathbf{v}}\left[u_i(\mathbf{v},\mathbf{p})\right] \\
&\geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in T(\mathbf{v})} p_i(T_i(\mathbf{v}))\right] + \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v} \mid T(\mathbf{v}'))} (v_i - p_i(T_i(\mathbf{v}')))\right] \\
&\geq \min_{T \in \mathcal{I}}\left(\sum_{i \in T} p_i(T_i) + \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v} \mid T)} (v_i - p_i(T_i))\right]\right)
\end{aligned}$$

□

Our task now is to show that for any set $T \in \mathcal{I}$, we have

$$\sum_{i \in T} p_i(T_i) + \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v} \mid T)} (v_i - p_i(T_i))\right] \geq \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \quad .$$

In particular, we can again formulate a linear program which has a non-negative objective value if there are prices $p_i(T_i)$ leading to the desired competitive ratio. To this end, we state the following LP with variables $p_i(T_i)$ for every i and $T_i \in \mathcal{I}$ such that $T_i \subseteq \{1, \ldots, i-1\}$ and slack variables ℓ_+ and ℓ_-.

$$
\begin{aligned}
\max \quad & \ell_+ - \ell_- \\
\text{s.t.} \quad & \sum_{i \in T} p_i(T_i) + \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}|T)} (v_i - p_i(T_i))\right] \\
& \qquad \geq \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] + \ell_+ - \ell_- && \text{for all } T \in \mathcal{I} \\
& p_i(T_i) \geq 0 && \text{for all } i \text{ and } T_i \\
& \ell_+ \geq 0 \\
& \ell_- \geq 0.
\end{aligned}
$$

In order to derive the dual, let us reformulate the constraints to

$$
\begin{aligned}
& \sum_{i=1}^{n} p_i(T_i) \cdot \left(\mathbf{Pr}_{\mathbf{v}}\left[i \in \mathrm{OPT}\left(\mathbf{v} \mid T\right)\right] - \mathbb{1}_{i \in T}\right) + \ell_+ - \ell_- \\
& \qquad \leq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}|T)} v_i\right] - \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right]
\end{aligned}
$$

Observe that if the optimal solution to this LP has a non-negative objective value, then the corresponding prices ensure a 1/2-competitive algorithm. In order to show this, by strong LP duality, we can argue in dual space in the same way as we did in Section 3.2.1. In particular, using strong duality, having a non-negative objective is equivalent to showing that any feasible dual solution has a non-negative dual objective. To this end, we consider the dual LP with variables μ_T for $T \subseteq [n]$ with $T \in \mathcal{I}$.

$$
\begin{aligned}
\min \quad & \sum_{T : T \in \mathcal{I}} \mu_T \left(\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}|T)} v_i\right] - \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}\left(\mathbf{v}\right)\right)\right]\right) \\
\text{s.t.} \quad & \sum_{T : T \in \mathcal{I}} \mu_T = 1 \\
& \sum_{\substack{T : T \in \mathcal{I} \text{ and} \\ T' = T \cap [i-1]}} \mu_T \left(\mathbf{Pr}_{\mathbf{v}}\left[i \in \mathrm{OPT}\left(\mathbf{v} \mid T\right)\right] - \mathbb{1}_{i \in T}\right) \geq 0 && \text{for all } i,\, T' \subseteq [i-1] \text{ with } T' \in \mathcal{I} \\
& \mu_T \geq 0 && \text{for all } T \text{ with } T \in \mathcal{I}.
\end{aligned}
$$

In order to show that the dual objective is always non-negative, we will next present our three key ingredients.

Dual constraints. We can rewrite the second set of constraints for all buyers i and $T' \subseteq [i-1]$ with $T' \in \mathcal{I}$ in a more useful way:

$$
\sum_{\substack{T : T \in \mathcal{I} \text{ and} \\ T' = T \cap [i-1]}} \mu_T \cdot \mathbf{Pr}_{\mathbf{v}}\left[i \in \mathrm{OPT}(\mathbf{v} \mid T)\right] \geq \sum_{\substack{T : T \in \mathcal{I} \text{ and} \\ T' = T \cap [i-1]}} \mu_T \cdot \mathbb{1}_{i \in T} \quad . \tag{3.3}
$$

Dual objective. We use that $\sum_{T:T\in\mathcal{I}} \mu_T = 1$ as well as linearity of expectations to rearrange the terms in the dual objective:

$$\begin{aligned}\text{dual obj.} &= \sum_{T:T\in\mathcal{I}} \mu_T \left(\mathbf{E}_{\mathbf{v}} \left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right)\right] - \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right]\right) \\ &= \mathbf{E}_{\mathbf{v}}\left[\sum_{T:T\in\mathcal{I}} \mu_T \ \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right)\right] - \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \\ &= \mathbf{E}_{\mathbf{v}}\left[\sum_{T:T\in\mathcal{I}} \mu_T \ \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) - \frac{1}{2}\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)\right] \ .\end{aligned}$$

In order to argue that the dual objective is always non-negative, we will show that the following inequality holds pointwise for any $\mathbf{v}$:

$$\sum_{T:T\in\mathcal{I}} \mu_T \ \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \geq \frac{1}{2}\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right) \ . \tag{3.4}$$

Observe that instead of arguing in the incomplete information setting, we can now restrict to the full information setting with a known valuation profile $\mathbf{v}$.

Further useful ingredients. We cite a useful statement from Kleinberg and Weinberg [2012] concerning matroids which we will apply in the following.

Lemma 3.4.3. [Kleinberg and Weinberg, 2012, Proposition 2] *Fix valuation profile* $\mathbf{v}$ *and let* $T \in \mathcal{I}$*. For any* $W \in \mathcal{I}$ *with* $T \cap W = \emptyset$ *and* $T \cup W \in \mathcal{I}$*, it holds*

$$\mathbf{v}\left(\mathrm{OPT}\left(\mathbf{v} \mid T\right)\right) \geq \sum_{i\in W} \left(\mathbf{v}\left(\mathrm{OPT}\left(\mathbf{v} \mid T_i\right)\right) - \mathbf{v}\left(\mathrm{OPT}\left(\mathbf{v} \mid T_i \cup \{i\}\right)\right)\right) \ .$$

Having this, we are ready to state our final claim which proves Theorem 3.4.1.

Claim 3.4.4. $\sum_{T:T\in\mathcal{I}} \mu_T \ \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \geq \frac{1}{2}\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)$

Once we have shown this claim, as an immediate consequence, we get that the dual objective is always non-negative by the above considerations.

Proof of Claim 3.4.4. Observe that we can set $W := \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)$ in Lemma 3.4.3 for some valuation profile $\tilde{\mathbf{v}}$. Hence, we get

$$\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \geq \sum_{i\in \mathrm{OPT}(\tilde{\mathbf{v}}\mid T)} \left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)\right) \ .$$

If we assume that $\tilde{\mathbf{v}} \sim \times_{i=1}^n \mathcal{D}_i$ is an independent fresh draw from $\times_{i=1}^n \mathcal{D}_i$, the lemma still holds pointwise for any realization. Taking expectations over $\tilde{\mathbf{v}}$ yields the following sequence of inequalities:

$$\begin{aligned}&\sum_{T:T\in\mathcal{I}} \mu_T \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \\ &\geq \sum_{T:T\in\mathcal{I}} \mu_T \mathbf{E}_{\tilde{\mathbf{v}}}\left[\sum_{i\in \mathrm{OPT}(\tilde{\mathbf{v}}\mid T)} \left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)\right)\right] \\ &= \sum_{i} \sum_{T:T\in\mathcal{I}} \mu_T \mathbf{Pr}_{\tilde{\mathbf{v}}}\left[i \in \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)\right] \left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)\right) \ .\end{aligned}$$

Next, we will reorder the sums in order to make use of the dual constraints as stated in Equation (3.3). To simplify notation, we define $\mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) := \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)$ as the marginal loss in $\mathrm{OPT}(\mathbf{v} \mid T_i)$ when adding i to T_i. Hence, we get

$$
\begin{aligned}
&\sum_i \sum_{T:T\in\mathcal{I}} \mu_T \mathbf{Pr}_{\tilde{\mathbf{v}}}\left[i \in \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)\right]\left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)\right) \\
&= \sum_i \sum_{T:T\in\mathcal{I}} \mu_T \cdot \mathbf{Pr}_{\tilde{\mathbf{v}}}\left[i \in \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)\right] \cdot \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \\
&= \sum_i \sum_{\substack{T_i:T_i\in\mathcal{I} \\ T_i\subseteq[i-1]}} \sum_{\substack{T:T\in\mathcal{I} \\ T_i=T\cap[i-1]}} \mu_T \cdot \mathbf{Pr}_{\tilde{\mathbf{v}}}\left[i \in \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)\right] \cdot \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \\
&= \sum_i \sum_{\substack{T_i:T_i\in\mathcal{I} \text{ and} \\ T_i\subseteq[i-1]}} \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \left(\sum_{\substack{T:T\in\mathcal{I} \text{ and} \\ T_i=T\cap[i-1]}} \mu_T \cdot \mathbf{Pr}_{\tilde{\mathbf{v}}}\left[i \in \mathrm{OPT}(\tilde{\mathbf{v}} \mid T)\right] \right) \\
&\geq \sum_i \sum_{\substack{T_i:T_i\in\mathcal{I} \text{ and} \\ T_i\subseteq[i-1]}} \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \left(\sum_{\substack{T:T\in\mathcal{I} \text{ and} \\ T_i=T\cap[i-1]}} \mu_T \cdot \mathbb{1}_{i\in T} \right) \\
&= \sum_i \sum_{T:T\in\mathcal{I}} \mu_T \cdot \mathbb{1}_{i\in T} \cdot \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \quad ,
\end{aligned}
$$

where the inequality simply uses Constraints (3.3). Finally, by reordering the sums, we make use of a telescopic sum argument:

$$
\begin{aligned}
&\sum_i \sum_{T:T\in\mathcal{I}} \mu_T \cdot \mathbb{1}_{i\in T} \cdot \mathrm{COST}\left(\mathbf{v}, i \mid T_i\right) \\
&= \sum_{T:T\in\mathcal{I}} \mu_T \sum_{i\in T} \left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i)\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T_i \cup \{i\})\right)\right) \\
&= \sum_{T:T\in\mathcal{I}} \mu_T \left(\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right) - \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right)\right) \\
&= \mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right) - \sum_{T:T\in\mathcal{I}} \mu_T \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \quad ,
\end{aligned}
$$

where in the last equality we made use of the fact that $\sum_{T:T\in\mathcal{I}} \mu_T = 1$ by the first dual constraint. Combining all of this, we get

$$
\sum_{T:T\in\mathcal{I}} \mu_T \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \geq \mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right) - \sum_{T:T\in\mathcal{I}} \mu_T \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right)
$$

or equivalently

$$
\sum_{T:T\in\mathcal{I}} \mu_T \mathbf{v}\left(\mathrm{OPT}(\mathbf{v} \mid T)\right) \geq \frac{1}{2} \mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right)
$$

which proves the desired statement. □

Chapter 4

Approximating Optimum Online for Multi-Demand Buyers

In the previous chapter, we compared the performance of pricing-based online algorithms to the expected offline optimum. For example, Theorem 3.0.1 from Chapter 3 shows the existence of static item prices which imply a $1/2$-competitive algorithm for combinatorial auctions when buyers have XOS valuation functions. Using the results from Dütting et al. [2020], there are even prices which can be computed in polynomial time.

Still, comparing to the expected offline optimum as a benchmark might be too pessimistic in several applications. Its access to future realizations is unattainable for online algorithms. Therefore, a recent line of work (e.g. Papadimitriou et al. [2021], Anari et al. [2019], Braverman et al. [2022] or Dütting et al. [2023]) has started to complement previous results by shifting attention towards the following question:

How well can we approximate the optimal (computationally unbounded) online algorithm in polynomial time?

In other words, how much must we lose when restricting to efficient algorithms instead of solving the optimal dynamic program? On the one hand, even for buyers with unit-demand valuations it is PSPACE-hard to approximate the optimum online algorithm within some absolute constant $1 - \epsilon$ [Papadimitriou et al., 2021]. Luckily, approximations strictly better than $1/2$ exist for unit-demand buyers: Papadimitriou et al. [2021] gave a 0.51-approximate algorithm, later improved to 0.52 [Saberi and Wajc, 2021], $1 - 1/e \approx 0.632$ [Braverman et al., 2022], and 0.652 [Naor et al., 2023]. Motivated by this, we ask:

Can we obtain a better than $1/2$-approximate algorithm to the optimal (computationally unbounded) online algorithm beyond unit-demand valuation functions?

As we will see in this chapter, this is indeed possible. In particular, we show the following theorem for multi-demand buyers[1].

Theorem 1.3.1. *When buyers have multi-demand valuation functions, there exists a polynomial time $(1/2 + \kappa)$-approximation algorithm with respect to the expected social welfare of the optimal online algorithm, for a constant $\kappa \geq 0.0115$.*

[1] We assume that every agent has a k_i-demand valuation function as defined in Chapter 2, where k_i denotes the buyer-specific demand size of agent i. This class of valuation functions is studied with the name "multi-demand buyers", e.g. by Berger et al. [2020]. We will sometimes refer to multi-demand buyers as k_i-demand buyers to make the demand size more clear.

Our algorithm for Theorem 1.3.1 is based on an LP relaxation of the welfare obtained by the optimal online algorithm. This LP relaxation is rounded online given the observed valuations. Inspired by the work of Papadimitriou et al. [2021], we make use of a two-proposal approach for the rounding procedure: We first propose a set of items to buyer i from which a random subset is allocated. Given the remaining space in the demand k_i, we fill this space with another random proposal while ensuring that we never allocate more than k_i items to buyer i.

It is crucial that the algorithm has access to the exact valuations of all buyers. Still, using recent work of Banihashem et al. [2024], we are also able to deal with strategic buyers who may misreport their own valuations. In particular, we can argue that our algorithm can be made dominant-strategy incentive-compatibility (DISC) if we bound the demand size k_i of buyers by a constant. In particular, Theorem 1.3.1 implies the following result.

Theorem 4.0.1. *If every buyer i has a k_i-demand valuation function and k_i is upper bounded by a constant, then there exists a polynomial-time DSIC mechanism giving a $(1/2 + \kappa)$-approximation to the social welfare of the optimal online algorithm.*

Proof. We apply Theorem 19 of Banihashem et al. [2024], as our problem with k_i-demand valuation functions can be viewed exactly as what they call a "Prophet Inequalities problem" in their paper. Using their notation, we take $\mathcal{A}^{\text{inp}}$ to be our algorithm which we use to prove Theorem 1.3.1, with expected social welfare $\mathbf{E}\left[\mathbf{v}(\mathcal{A}^{\text{inp}})\right]$. Note that our algorithm is what Banihashem et al. [2024] call "past-valuation-independent" as its allocation decision for buyer i depends only on the set of available items, the arriving buyer's valuation $v_i(\cdot)$, and knowledge of the input distributions. Note also that for each buyer i, the outcome space is of size at most $\binom{n}{k_i} = \text{poly}(n)$, because k_i is upper bounded by a constant. Finally, although our distribution over $v_i(\cdot)$ is not continuous, it is not hard to satisfy the required assumption by adding a tie-breaking coordinate (as mentioned in Banihashem et al. [2024]).

Hence, there is a pricing-based algorithm $\mathcal{A}^{\text{out}}$ which uses $\text{poly}(n, \binom{n}{\max_i k_i}, 1/\epsilon)$ many samples, runs in time $\text{poly}(n, \binom{n}{\max_i k_i}, 1/\epsilon)$ and whose expected social welfare satisfies

$$\mathbf{E}\left[\mathbf{v}(\mathcal{A}^{\text{out}})\right] \geq (1-\epsilon) \cdot \mathbf{E}\left[\mathbf{v}(\mathcal{A}^{\text{in}})\right] \quad .$$

□

Note that our main Theorem 1.3.1 does not require any upper bounds on the demand size k_i. In particular, the demand k_i can be as large as the number of items (which corresponds to an additive valuation function). The upper bound on k_i in Theorem 4.0.1 is only required such that the algorithm from Banihashem et al. [2024] runs in polynomial time.

Concerning the analysis of our algorithm, we distinguish for each buyer-item pair (i, j) whether it is assigned with sufficiently large probability already from the first proposal, or if it requires the second proposal. In the first case, the analysis proceeds in a straightforward way similar to the OCRS literature (see e.g. Ezra et al. [2020]). In the second case, we are required to be much more careful. In particular, the number of allocated items via the first proposal is affecting our second proposal. To this end, we need to consider the correlation between the presence and absence of items after the first proposal. Controlling this correlation will be the major technical challenge in this chapter.

As a side remark, in previous work with unit-demand valuations, this challenge was readily handled by using negative correlation. Unfortunately, simple examples show that in our case positive correlation is *required* to go beyond an approximation ratio of $1/2$.

Observation 4.0.2. *Any algorithm for buyers with k_i-demand valuations which has an approximation ratio better than $1/2$ with respect to the LP relaxation $\mathrm{LP_{on}}$ (to be defined in Section 4.1) must create positive correlation between the events of items being available.*

Proof. Let $F_{i,j}$ denote the event of item j being free just before the arrival of buyer i. Consider buyer i with demand size $k_i = 2$ arriving with probability ϵ and having unit values only for two items j and j'. With probability $1-\epsilon$, buyer i is not arriving (or equivalently, buyer i samples a valuation function which is zero for any subset of items). Imagine the LP sets a value of ϵ on each pair (i,j) and (i,j'). To achieve an approximation factor of $(0.5+\kappa)$ against the LP, we are required to have that the expected number of items assigned to i is at least $(0.5+\kappa)\cdot 2\epsilon$. Equivalently, we must have

$$\mathbf{Pr}\left[F_{i+1,j}\right] + \mathbf{Pr}\left[F_{i+1,j'}\right] < 2 - (0.5+\kappa)\cdot 2\epsilon$$

implying

$$\mathbf{Pr}\left[F_{i+1,j}\right] \cdot \mathbf{Pr}\left[F_{i+1,j'}\right] < (1-(0.5+\kappa)\cdot\epsilon)^2 = 1-(1+2\kappa)\epsilon + O(\epsilon^2).$$

However, because j and j' can only be allocated if i arrives, we have

$$\mathbf{Pr}\left[F_{i+1,j} \wedge F_{i+1,j'}\right] \geq 1-\epsilon > \mathbf{Pr}\left[F_{i+1,j}\right] \cdot \mathbf{Pr}\left[F_{i+1,j'}\right] \quad ,$$

where the final inequality holds for sufficiently small ϵ. □

While this unfortunately shows that it is not possible to obtain negative correlation for general k_i-demand valuation functions, we show that our algorithm obtains a good approximation if it can just avoid introducing a "large" amount of positive correlation.

Chapter Organization and Remarks

This chapter is based on *Approximating Optimum Online for Capacitated Resource Allocation* [Braun et al., 2024], which is joint work with Thomas Kesselheim, Tristan Pollner and Amin Saberi. Further bibliographic notes can be found in Section 1.5.

We start with some preliminaries and a formal problem definition in Section 4.1. Afterwards, the algorithm is stated for two-point distributions in Section 4.2 and analyzed in Section 4.3. The algorithm as stated in Section 4.2 is not polynomial time. While we can estimate the crucial quantity via samples as shown in our paper [Braun et al., 2024], we follow a different path in this thesis: We show in Section 4.4 that it is not even necessary to use these estimates as even a simpler algorithm allows us to beat the approximation ratio of half. In Section 4.5, we finally discuss how to extend the algorithm and its analysis towards general distributions.

4.1 Formal Problem Statement and Preliminaries

There is a set of m heterogeneous items M and a sequence of buyers $[n]$ arrives online in known order[2]. In the special case of Bernoulli distributions, we are given a k_i-demand

[2] We highlight that knowing the arrival order is important – in contrast to the results in Chapter 3, where results hold also against an adaptive adversary.

valuation function v_i for every buyer i upfront. To simplify notation, we denote buyer i's value for item j by $v_{i,j} \geq 0$. As a consequence, the valuation function of buyer i can be expressed as

$$v_i(S) = \max_{S' \subseteq S : |S'| \leq k_i} \sum_{j \in S'} v_{i,j} \quad .$$

In step i, buyer i *arrives* (also noted as *active*) independently with known probability q_i and does not arrive otherwise.

Every item can be allocated to at most one buyer; any buyer is interested in up to k_i many items. We call k_i the *demand size* or *capacity* of buyer i and emphasize that k_i can be buyer-specific, i.e. we allow different buyers to have different demand sizes. Upon the arrival of buyer i, we observe the random realization if the buyer is active, and can choose which items $X_i \subseteq M$ (if any) we would like to allocate, subject to the constraints that each item can be assigned to at most one buyer and $|X_i| \leq k_i$. If buyer i does not arrive, for convenience, we take $X_i = \emptyset$.

Our objective is to maximize the *expected social welfare*, defined as $\mathbf{E}[\text{ALG}] := \mathbf{E}\left[\sum_i \sum_{j \in X_i} v_{i,j}\right]$, where ALG is required to run in polynomial time.

Benchmark. As mentioned before, we want to compare to the optimal (computationally unbounded) online algorithm. In other words, we aim for a statement of the form $\mathbf{E}[\text{ALG}] \geq \zeta \cdot \text{OPT}_{\text{on}}$, where OPT_{on} is the expected welfare achieved by the optimal online algorithm. The optimal online algorithm has unlimited computational power and also knows all distributions upfront, but only observes realizations one at a time and needs to make an irrevocable decision before observing the next realization. Formally, we can define OPT_{on} via a Bellman equation. To this end, let $\text{OPT}_{\text{on}}(i, J)$ denote the optimum gain achievable from buyers $\{i, i+1, \ldots, n\}$ with items $J \subseteq M$ available. Then, recursively we have

$$\begin{aligned} \text{OPT}_{\text{on}}(i, J) := {} & (1 - q_i) \cdot \text{OPT}_{\text{on}}(i+1, J) \\ & + q_i \cdot \max_{J' \subseteq J, |J'| \leq k_i} \left(\sum_{j \in J'} v_{i,j} + \text{OPT}_{\text{on}}(i+1, J \setminus J') \right) \quad . \end{aligned}$$

We recall that even in the special case of unit-demand valuations, it is PSPACE-hard to approximate OPT_{on} within a $(1 - \epsilon)$-factor [Papadimitriou et al., 2021].

LP Relaxation. We will use an LP relaxation of the optimum online algorithm which generalizes the one from the unit-demand case [Papadimitriou et al., 2021, Braverman et al., 2022, Torrico and Toriello, 2022]. It has a variable $x_{i,j}$ for every buyer-item pair (i, j):

$$\begin{aligned}
\max \quad & \sum_{i,j} v_{i,j} \cdot x_{i,j} && && (\mathrm{LP_{on}}) \\
\text{s.t.} \quad & \sum_{i} x_{i,j} \leq 1 && \text{for all } j \in M && (4.1) \\
& \sum_{j} x_{i,j} \leq q_i \cdot k_i && \text{for all } i \in [n] && (4.2) \\
& 0 \leq x_{i,j} \leq q_i \cdot \left(1 - \sum_{i'<i} x_{i',j}\right) && \text{for all } j \in M, i \in [n] \ . && (4.3)
\end{aligned}$$

This LP indeed relaxes the optimal online algorithm, as we will see in the following observation.

Observation 4.1.1. *The optimum objective value of* $\mathrm{LP_{on}}$ *upper bounds the gain of optimum online. In other words,* $\mathrm{OPT}(\mathrm{LP_{on}}) \geq \mathrm{OPT_{on}}$.

Proof. Define an indicator random variable $X_{i,j}$ for every pair (i,j) which is one if and only if the optimum online algorithm allocates item j to buyer i. Denote by $x^*_{i,j} = \mathbf{E}\left[X_{i,j}\right]$ its expectation.

First, note that the welfare achieved by the optimum online algorithm is

$$\mathrm{OPT_{on}} = \mathbf{E}\left[\sum_{i,j} v_{i,j} X_{i,j}\right] = \sum_{i,j} v_{i,j} \cdot x^*_{i,j},$$

coinciding with the objective of $\mathrm{LP_{on}}$.

Next, we observe that any algorithm can allocate items at most once, hence for any item $j \in M$, we have

$$\sum_{i} x^*_{i,j} = \sum_{i} \mathbf{E}\left[X_{i,j}\right] = \mathbf{E}\left[\sum_{i} X_{i,j}\right] \leq 1 \ .$$

Also, note that for any buyer i, we have $\sum_j X_{i,j} = 0$ if the buyer does not arrive, and $\sum_j X_{i,j} \leq k_i$ if the buyer arrives, as the optimum online algorithm will assign at most k_i items to buyer i if the buyer arrives. Hence

$$\sum_{j} x^*_{i,j} = \mathbf{E}\left[\sum_{j} X_{i,j}\right] = \mathbf{Pr}\left[i \text{ arrives}\right] \cdot \mathbf{E}\left[\sum_{j} X_{i,j} \,\middle|\, i \text{ arrives}\right] \leq q_i \cdot k_i.$$

Finally, observe that if buyer i arrives, the optimum online algorithm can only allocate item j if it is available. Item j is available if it was not allocated to some previous buyer $i' < i$. Crucially, for any online algorithm, the event that item j is available at time i is independent of the arrival of buyer i (this does not hold for an offline algorithm). Hence, we obtain

$$\begin{aligned}
x^*_{i,j} &= \mathbf{E}\left[X_{i,j}\right] = \mathbf{Pr}\left[i \text{ arrives}\right] \cdot \mathbf{E}\left[X_{i,j} \mid i \text{ arrives}\right] \\
&\leq q_i \cdot \mathbf{E}\left[1 - \sum_{i'<i} X_{i',j} \,\middle|\, i \text{ arrives}\right] \\
&= q_i \cdot \mathbf{E}\left[1 - \sum_{i'<i} X_{i',j}\right] = q_i \cdot \left(1 - \sum_{i'<i} x^*_{i',j}\right) .
\end{aligned}$$

As a consequence, $\{x^*_{i,j}\}_{i,j}$ is a feasible solution to $\mathrm{LP}_{\mathrm{on}}$ and hence, $\mathrm{OPT}(\mathrm{LP}_{\mathrm{on}}) \geq \mathrm{OPT}_{\mathrm{on}}$. □

Generalized Problem Definition. In the above problem definition, we made the simplifying assumption that the buyer arriving at time i has a simple Bernoulli distribution determining if she is active or not. In the general model, in every round, a buyer randomly realizes one of many possible valuation functions. More formally, buyer i realizes one of L possible demand sizes $k_{i,\ell}$ together with a vector of values $(v_{i,j,\ell})_j$, where each realization ℓ is sampled with probability $q_{i,\ell}$. We highlight that demand sizes and values for a single buyer i can be arbitrarily correlated, although across different rounds we assume independence. In Section 4.5 we argue that our LP, algorithm, and analysis can be extended to this general setting.

Pivotal Sampling

As a part of our online algorithm we invoke the randomized offline rounding framework of *pivotal sampling* (also called *Srinivasan rounding* or *dependent rounding*) [Srinivasan, 2001, Gandhi et al., 2006]. To simplify the description, denote the set of items $M = \{1, \ldots, m\}$.

Imagine we are given marginals $x_1, \ldots, x_m$ with each $x_j \in [0, 1]$ and $\sum_j x_j \leq K$ for some positive integer K. We would like to randomly select at most K indices from $\{1, 2, \ldots, m\}$ such that j is selected with probability x_j. Pivotal sampling selects such a subset while also guaranteeing strong negative correlation properties between individual indices. It does so by sequentially choosing a pair of fractional marginals, and applying a randomized "pivot" operation that makes at least one integral. We formally state some of the properties of the algorithm below which will be used in our analysis in Section 4.3.

Theorem 4.1.2 (as in Srinivasan [2001])**.** *The pivotal sampling algorithm with input $(x_j)_{j=1}^m$ where $\sum_j x_j \leq K$ efficiently produces a random subset $\mathsf{PS}(x_1, \ldots, x_m) \subseteq M$ with the following properties:*

(P1) For every $j \in M$, we have $\mathbf{Pr}\left[j \in \mathsf{PS}(x_1, \ldots, x_m)\right] = x_j$.

(P2) The number of elements in $\mathsf{PS}(x_1, \ldots, x_m)$ is always at most K.

(P3) (Negative cylinder dependence) For any $J \subseteq M$, we have

$$\mathbf{Pr}\left[\bigwedge_{j \in J} j \in \mathsf{PS}(x_1, \ldots, x_m)\right] \leq \prod_{j \in J} \mathbf{Pr}\left[j \in \mathsf{PS}(x_1, \ldots, x_m)\right] \quad ,$$

and

$$\mathbf{Pr}\left[\bigwedge_{j \in J} j \notin \mathsf{PS}(x_1, \ldots, x_m)\right] \leq \prod_{j \in J} \mathbf{Pr}\left[j \notin \mathsf{PS}(x_1, \ldots, x_m)\right] \quad .$$

4.2 The Algorithm: A Two-Step Approach

We begin with a short description of our algorithm before presenting the pseudo code in Algorithm 4. To this end, we say item j is "free at i" or "available at i" (or "free"/"available", if the context is clear) if just before the arrival of buyer i, item j has not yet been allocated to any previous buyer.

Our algorithm uses an optimal solution $\{x_{i,j}\}_{i,j}$ to LP_{on} as input. After observing if buyer i arrives, if so, we sample a set of at most k_i items FP_i (denoting the **f**irst **p**roposal for i) using pivotal sampling, such that each item j is included in FP_i with marginal probability $x_{i,j}/q_i$. For every item $j \in \mathsf{FP}_i$, if j is still available, we toss a coin independently with probability $\alpha_{i,j} := \min\left(1, \frac{0.5+\kappa}{1-(0.5+\kappa)\cdot\sum_{i'<i} x_{i',j}}\right)$, and allocate item j to buyer i if this coin toss is successful.

After this procedure, we have a number A_i of items allocated to buyer i, where A_i is a random variable which can take values in $\{0, \ldots, k_i\}$. In order to make use of the remaining space in the demand size of buyer i, we make a second proposal. Again via the pivotal sampling subroutine, this time with a reduced marginal probability of $(1 - \frac{A_i}{k_i}) \cdot x_{i,j}/q_i$ for every item j, we sample a set of items SP_i, denoting the **s**econd **p**roposal with size at most $k_i - A_i$. Among these items, we consider only those items j for which $\alpha_{i,j} = 1$, j was free at i, and j was not allocated to i via a first proposal. For each such item j, we allocate j to i with probability $\beta_{i,j}$. The factor $\beta_{i,j}$ is chosen in a way to ensure that $\mathbf{Pr}\left[j \text{ allocated to } i\right] = (0.5 + \kappa) \cdot x_{i,j}$.

Algorithm 4: Allocation Algorithm

1 $\kappa \leftarrow 0.0115$
2 Solve LP_{on} for $\{x_{i,j}\}_{i,j}$
3 **for** *each buyer* i*, if* i *arrives* **do**
4   Define items $\mathsf{FP}_i := \mathsf{PS}((x_{i,j}/q_i)_j)$
5   **for** *each item* $j \in \mathsf{FP}_i$ **do**
6   **if** j *is available* **then**
7   Allocate j to i with probability $\alpha_{i,j} := \min\left(1, \frac{0.5+\kappa}{1-(0.5+\kappa)\cdot\sum_{i'<i} x_{i',j}}\right)$
8     $A_i \leftarrow$ number of items allocated to i thus far
9   Define items $\mathsf{SP}_i := \mathsf{PS}(((1 - \frac{A_i}{k_i}) \cdot x_{i,j}/q_i)_j)$
10   **for** *each item* $j \in \mathsf{SP}_i$ *with* $\alpha_{i,j} = 1$ **do**
11   **if** j *is available* **then**
12   Compute $\rho_{i,j} := \mathbf{E}\left[\mathbb{1}[j \text{ available after Line } 8] \cdot (1 - \frac{A_i}{k_i}) \mid i \text{ arrived}\right]$
13             $\beta_{i,j} \leftarrow \min\left(1, ((0.5+\kappa)\cdot\sum_{i'<i} x_{i',j} - (0.5-\kappa)) \cdot \frac{1}{\rho_{i,j}}\right)$
14   Allocate j to i with prob. $\beta_{i,j}$

Concerning the definition of $\rho_{i,j}$, we note that the expectation is over the randomness in the arrivals and the algorithm up to when Algorithm 4 reaches Line 8 for arrival i (in particular, we consider "re-running" the algorithm as defined thus far on a fresh instance). The indicator $\mathbb{1}[j \text{ available after Line } 8]$ refers to the event that j was not allocated to some $i' < i$ and is also not allocated to i via a first proposal. This indicator is potentially correlated with the number of allocated items A_i.

The $\min(1, \cdot)$ in the definition of $\beta_{i,j}$ is for convenience only; in particular, it is thus easy to see that the algorithm is well-defined. As a crux of our analysis, we will show that using $\kappa = 0.0115$ ensures that the $\min(1, \cdot)$ in the definition of $\beta_{i,j}$ is actually redundant.

In the remainder of this section, we will argue that Algorithm 4 is well-defined and guarantees to respect the demand constraints of buyers.

Observation 4.2.1. *Algorithm 4 is well-defined.*

Proof. Note first that in Line 4, our call to the pivotal sampling algorithm $\mathsf{PS}(\cdot)$ is well-defined as each marginal $x_{i,j}/q_i$ is in $[0,1]$ by $\mathrm{LP_{on}}$-Constraint (4.3). Each $\alpha_{i,j}$ as defined in Line 7 is clearly a probability by construction. Our second call to $\mathsf{PS}(\cdot)$ is similarly well-defined. Note that $\beta_{i,j}$ is always a probability — if $\alpha_{i,j} = 1$, it implies that $(0.5+\kappa)\cdot\sum_{i'<i} x_{i',j} \geq (0.5-\kappa)$ by definition. This in turn shows that $\beta_{i,j}$ is always in the interval $[0,1]$.

Finally, note that item j is allocated only if available, and hence never allocated to two different buyers (or to the same buyer twice). ☐

We also have that our algorithm respects the demand sizes for each buyer arriving online.

Observation 4.2.2. *The number of items allocated to buyer i by Algorithm 4 is always at most k_i.*

Proof. By Property (P2) of pivotal sampling, the size of FP_i is never larger than k_i as $\sum_j \frac{x_{i,j}}{q_i} \leq k_i$ by Constraint (4.2). In addition, as we scale the marginals down for the second proposal set SP_i, we are guaranteed that buyer i is only allocated at most $k_i - A_i$ many items during the second proposal. ☐

We also note that every line except Line 12 can be implemented in polynomial time. Indeed, note Line 2 can be run efficiently as $\mathrm{LP_{on}}$ has polynomial size. Also, our calls to pivotal sampling can be implemented efficiently [Srinivasan, 2001].

Line 12 requires exponential time as written, and for ease of presentation, first, we analyze the above exponential time algorithm. In Braun et al. [2024], we show that we can replace this computation with a sample average and appeal to concentration bounds, while only losing an arbitrarily small ϵ in the approximation ratio. The main point of care is to argue that $\rho_{i,j}$ is bounded away from 0 so that we can get an approximation with good *multiplicative* error. This will be a consequence of our analysis of the approximation ratio in Section 4.3.

In this thesis, we complement this result by pursuing a different path. As we will see, the subsampling with $\beta_{i,j}$ is *not* required to beat the approximation ratio of $1/2$. In particular, in Section 4.4, we show that we can set $\beta_{i,j} = 1$ for any pair (i,j) and still outperform the prevalent bound of half by a constant. Hence, any computation or estimation of the $\rho_{i,j}$ can be avoided while still getting an improved guarantee.

4.3 Analysis: Beating a 1/2-Approximation

Our main result is as follows.

Theorem 4.3.1. *For $\kappa = 0.0115$, the social welfare achieved by Algorithm 4 satisfies*

$$\mathbf{E}[\mathrm{ALG}] \geq (0.5+\kappa)\cdot \mathrm{OPT_{on}} \ .$$

This section is dedicated to the proof of this theorem. As mentioned before, in this section we analyze the algorithm which has access to the expectation $\rho_{i,j}$ exactly and recall that we can use samples to estimate the quantity in Line 12 [Braun et al., 2024].

Outline. Before diving into details we outline the ingredients in our proof of Theorem 4.3.1. First, we note that by Observation 4.2.2, the size of the set of items allocated to i (denote this set by X_i) is always at most k_i, so

$$\mathbf{E}\left[\text{ALG}\right] = \mathbf{E}\left[\sum_i \max_{S \subseteq X_i, |S| \le k_i} \left(\sum_{j \in S} v_{i,j}\right)\right] = \sum_{i,j} v_{i,j} \cdot \mathbf{Pr}\left[j \text{ allocated to } i\right].$$

Next, observe that bounding the term $\mathbf{Pr}\left[j \text{ allocated to } i\right]$ naturally brings us into one of two cases. If (i,j) is such that $\alpha_{i,j} < 1$, the allocation of j to i can only happen in Line 7 of our algorithm, and consequently it is straightforward to bound the resulting welfare (which we do in Observation 4.3.4). We then turn our perspective towards pairs (i,j) with a subsampling probability $\alpha_{i,j} = 1$; for these, the analysis requires much more care. Again, we start by considering the contribution of allocating via a first proposal in Lemma 4.3.5 (i). Here the first proposal alone is not sufficient, and we are required to compensate for this via a suitable bound on the allocation probability via a second proposal. We do so by proving Lemma 4.3.5 (ii) which gives a sufficient bound of the contribution from a second proposal. This is the main technical contribution and will use lemmas analyzing the evolution of the correlation between offline items in Section 4.3.3.

Notation. For convenience, we let $y_{i,j} := \sum_{i' < i} x_{i',j}$. Note that $\alpha_{i,j} < 1$ exactly when $y_{i,j} < (0.5 - \kappa)/(0.5 + \kappa)$. We hence define $\tau := (0.5 - \kappa)/(0.5 + \kappa)$ as this threshold for $y_{i,j}$ after which the subsampling probability $\alpha_{i,j}$ becomes one. If for buyer i and item j we have $y_{i,j} \le \tau$, then we call the pair (i,j) *early*. Otherwise, we call the pair (i,j) *late*. In addition, we define $\mathcal{A}_1$ as the set of all pairs (i,j) such that item j was allocated to buyer i in Line 7, and $\mathcal{A}_2$ as the set of all pairs (i,j) such that j was allocated to i in Line 14.

As j is not allocated more than once in our algorithm, we quickly observe the following claim.

Observation 4.3.2. *For any buyer i, we have*

$$\mathbf{E}\left[\max_{S \subseteq X_i, |S| \le k_i} \left(\sum_{j \in S} v_{i,j}\right)\right] = \sum_j v_{i,j} \cdot \left(\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] + \mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right]\right).$$

To analyze the probabilities $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right]$ and $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right]$, we consider two separate cases based on whether (i,j) is early (Section 4.3.1) or late (Section 4.3.2).

4.3.1 Analysis for Early Pairs

It will be crucial to bound the probability of an item j being free at time i. We denote the event that item j is *free* or *available* (i.e., not allocated) at the arrival of buyer i by $F_{i,j}$. The following observation gives an expression of the probability with respect to the LP variables. As an important remark, note that if a pair (i,j) is early, so is every pair (i',j) with $i' < i$.

Observation 4.3.3. *For early pairs (i,j), we have* $\mathbf{Pr}\left[F_{i,j}\right] = 1 - (0.5 + \kappa) \cdot y_{i,j}$.

Proof. We proceed via induction on i. Before the arrival of the first buyer, the claim is trivially true, as all items are available with probability one. Afterwards, note that

$$\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] = q_i \cdot \mathbf{Pr}\left[j \in \mathsf{FP}_i\right] \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \alpha_{i,j} \tag{4.4}$$

as i's arrival, j being included in FP_i, $F_{i,j}$ and the algorithm's $\text{Ber}(\alpha_{i,j})$ coin flip are mutually independent events. As (i,j) is early, then $\alpha_{i,j} = \frac{0.5+\kappa}{1-(0.5+\kappa)\cdot y_{i,j}}$, so we have

$$\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] = q_i \cdot \frac{x_{i,j}}{q_i} \cdot (1 - (0.5+\kappa) \cdot y_{i,j}) \cdot \frac{0.5+\kappa}{1-(0.5+\kappa)\cdot y_{i,j}} = (0.5+\kappa) \cdot x_{i,j},$$

where we also use the induction hypothesis for the probability of item j being free at the arrival of buyer i. For early (i,j), we also clearly have $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] = 0$, so

$$\mathbf{Pr}\left[F_{i+1,j}\right] = \mathbf{Pr}\left[F_{i,j}\right] - \mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] = 1 - (0.5+\kappa) \cdot y_{i+1,j}. \qquad \square$$

As a consequence we can bound the contribution of an early pair (i,j) to $\mathcal{A}_1$ and $\mathcal{A}_2$, as follows.

Observation 4.3.4. *For early pairs (i,j), we have*

(i) $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] = (0.5+\kappa) \cdot x_{i,j}$ *and*

(ii) $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] = 0$.

Thus for early pairs (i,j), our algorithm achieves the desired allocation probability.

4.3.2 Analysis for Late Pairs implies Theorem 4.3.1

For late pairs, we show the following lemma which will be sufficient to prove our main Theorem 4.3.1.

Lemma 4.3.5. *For late pairs (i,j), the following two statements hold:*

(i) $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] = (1 - (0.5+\kappa) \cdot y_{i,j}) \cdot x_{i,j}$, *and*

(ii) $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] = ((0.5+\kappa) \cdot y_{i,j} - 0.5 + \kappa) \cdot x_{i,j}$.

We note that this immediately implies our main result.

Proof of Theorem 4.3.1. We have $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] + \mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] = (0.5+\kappa) \cdot x_{i,j}$ for any pair (i,j) by Observation 4.3.4 and Lemma 4.3.5. Hence, using the decomposition in Observation 4.3.2, we have

$$\begin{aligned}
\mathbf{E}\left[\text{ALG}\right] &= \mathbf{E}\left[\sum_{i,j} v_{i,j} \mathbb{1}_{j \in X_i}\right] \\
&= \sum_i \sum_j v_{i,j} \cdot \left(\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1\right] + \mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right]\right) \\
&= \sum_i \sum_j v_{i,j} \cdot (0.5+\kappa) \cdot x_{i,j} \\
&= (0.5+\kappa) \cdot \text{OPT}(\text{LP}_{\text{on}}) \geq (0.5+\kappa) \cdot \text{OPT}_{\text{on}} \ .
\end{aligned}$$

$\square$

Thus, it remains to prove Lemma 4.3.5. Our analysis here requires significantly more care as it must bound the gain from the second proposal. As the second proposal's marginal probabilities are dependent on which items were allocated in the first proposal, a complete analysis must consider the correlation introduced.

Proof of Lemma 4.3.5 (i)

As for early pairs, the remainder of our proof will proceed by induction on i. Thus, for every late pair (i', j) with $i' < i$, by the inductive hypothesis we have $\mathbf{Pr}[(i', j) \in \mathcal{A}_1] + \mathbf{Pr}[(i', j) \in \mathcal{A}_2] = (0.5 + \kappa) \cdot x_{i',j}$. Recall also that for every early pair (i', j) we know from Observation 4.3.4 that $\mathbf{Pr}[(i', j) \in \mathcal{A}_1] + \mathbf{Pr}[(i', j) \in \mathcal{A}_2] = (0.5 + \kappa) \cdot x_{i',j}$. Thus, we may assume that for the late pair (i, j) we have

$$\mathbf{Pr}[F_{i,j}] = 1 - (0.5 + \kappa) \cdot y_{i,j}. \tag{4.5}$$

With this, bounding the probability of allocation along a first proposal is very straightforward.

Proof of Lemma 4.3.5 (i). Note that

$$\begin{aligned}
\mathbf{Pr}[(i, j) \in \mathcal{A}_1] &= q_i \cdot \mathbf{Pr}[F_{i,j}] \cdot \mathbf{Pr}[j \in \mathsf{FP}_i] \cdot \alpha_{i,j} && \text{(via Equation (4.4))} \\
&= q_i \cdot (1 - (0.5 + \kappa) \cdot y_{i,j}) \cdot \frac{x_{i,j}}{q_i} \cdot 1 && \text{(via Equation (4.5))} \\
&= (1 - (0.5 + \kappa) \cdot y_{i,j}) \cdot x_{i,j}.
\end{aligned}$$

□

This completes the proof of Lemma 4.3.5 (i), and the remainder of this section is dedicated to the proof of Lemma 4.3.5 (ii).

Proof of Lemma 4.3.5 (ii)

We begin by bounding $\mathbf{Pr}[(i, j) \in \mathcal{A}_2]$ for late pairs (i, j), in the natural way which depends on the number of allocated items during the first proposal in Line 7. Recall that this is because for second proposals, we reduce the marginal probabilities for the pivotal sampling algorithm by a factor of $1 - A_i/k_i$. Note that for (i, j) to be assigned via a second proposal we need all of the following to happen: (i) i should arrive, (ii) j must be available after Line 8, and included as a second proposal, and (iii) the pair (i, j) should survive the final downsampling by $\beta_{i,j}$. This lets us observe

$$\begin{aligned}
&\mathbf{Pr}[(i, j) \in \mathcal{A}_2] \\
&= q_i \cdot \mathbf{Pr}[j \text{ available after Line } 8 \wedge j \in \mathsf{SP}_i \mid i \text{ arrived}] \cdot \beta_{i,j} \\
&= q_i \cdot \mathbf{E}\left[\mathbb{1}[j \text{ available after Line } 8] \cdot \left(1 - \frac{A_i}{k_i}\right) \cdot \frac{x_{i,j}}{q_i} \;\middle|\; i \text{ arrived}\right] \cdot \beta_{i,j} \\
&= x_{i,j} \cdot \rho_{i,j} \cdot \beta_{i,j} \ .
\end{aligned} \tag{4.6}$$

Note that for the second equality, we relied on Property (P1) of pivotal sampling, which guarantees that individual elements are sampled with exactly their marginal probability. Note that the marginal probability is random. In addition, the indicator $\mathbb{1}[j \text{ available after Line } 8]$ and the marginal probability for the pivotal sampling are correlated.

Recall that $\beta_{i,j} := \min\left(1, ((0.5 + \kappa) \cdot y_{i,j} - (0.5 - \kappa)) \cdot \frac{1}{\rho_{i,j}}\right)$. If the $\min(1, \cdot)$ here is redundant, we are immediately done; this is concretized in the following observation.

Observation 4.3.6. *If* $\rho_{i,j} \geq (0.5+\kappa)y_{i,j} - (0.5-\kappa)$*, then*

$$\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] = x_{i,j} \cdot ((0.5+\kappa)\cdot y_{i,j} - (0.5-\kappa)) \,.$$

Thus it suffices to show that the hypothesis of this observation holds. In other words, in order to conclude the proof, the only thing we need to show is the following proposition.

Proposition 4.3.7. *For any late pair* (i,j)*, we have* $\rho_{i,j} \geq (0.5+\kappa)y_{i,j} - (0.5-\kappa)$*.*

As a first step, we start with the following lower bound on $\rho_{i,j}$.

Lemma 4.3.8. *For late pairs* (i,j)*,*

$$\rho_{i,j} \geq (1-(0.5+\kappa)\cdot y_{i,j}) \cdot \left(\tau - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i}\right) .$$

Proof of Lemma 4.3.8. We can expand

$$\begin{aligned}
\rho_{i,j} &= \mathbf{E}\left[\mathbb{1}[j \text{ available after Line 8}] \cdot \left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}\right] \\
&= \mathbf{Pr}\left[F_{i,j} \mid i \text{ arrived}\right] \cdot \mathbf{E}\left[\mathbb{1}[j \text{ not allocated in Line 7}] \cdot \left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}\right] \\
&= \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{E}\left[\mathbb{1}[j \text{ not allocated in Line 7}] \cdot \left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}\right] \quad .
\end{aligned}$$

As the pair (i,j) is late, we have $\alpha_{i,j} = 1$. Hence, conditioned on being free and the arrival of buyer i, item j is not allocated in Line 7 if and only if it is not contained in the set FP_i. This allows us to bound

$$\begin{aligned}
&\mathbf{E}\left[\mathbb{1}[j \text{ not allocated in Line 7}] \cdot \left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}\right] \\
&= \mathbf{E}\left[\mathbb{1}[j \notin \mathsf{FP}_i] \cdot \left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}\right] \\
&= \left(1-\frac{x_{i,j}}{q_i}\right) \cdot \mathbf{E}\left[\left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}, j \notin \mathsf{FP}_i\right] \\
&\geq \tau \cdot \mathbf{E}\left[\left(1-\frac{A_i}{k_i}\right) \,\middle|\, i \text{ arrived}, F_{i,j}, j \notin \mathsf{FP}_i\right] .
\end{aligned}$$

To argue about the resulting expectation, we apply the following bound to remove the conditioning on $j \notin \mathsf{FP}_i$:

$$\begin{aligned}
\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}, j \notin \mathsf{FP}_i\right] &= \frac{\mathbf{E}\left[A_i \cdot \mathbb{1}_{j \notin \mathsf{FP}_i} \mid i \text{ arrived}, F_{i,j}\right]}{\mathbf{Pr}\left[j \notin \mathsf{FP}_i \mid i \text{ arrived}, F_{i,j}\right]} \\
&\leq \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{\mathbf{Pr}\left[j \notin \mathsf{FP}_i \mid i \text{ arrived}, F_{i,j}\right]} .
\end{aligned}$$

In addition, note that $\mathbf{Pr}\left[j \notin \mathsf{FP}_i \mid i \text{ arrived}, F_{i,j}\right] = 1 - \frac{x_{i,j}}{q_i} \geq y_{i,j} \geq \tau$ as pair (i,j) is late. Thus we get

$$\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}, j \notin \mathsf{FP}_i\right] \leq \frac{1}{\tau} \cdot \mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] \quad .$$

By substitution and using Equation (4.5), we directly conclude

$$\rho_{i,j} \geq \mathbf{Pr}\left[F_{i,j}\right] \cdot \tau \cdot \left(1 - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i} \cdot \frac{1}{\tau}\right) \tag{4.7}$$

$$= (1 - (0.5 + \kappa) \cdot y_{i,j}) \cdot \left(\tau - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i}\right) \qquad \text{(via Equation (4.5))}$$

as claimed. □

In order to exploit the bound obtained in Lemma 4.3.8, we need to control the conditional expectation $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$. In particular, our goal is to show that $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$ is bounded away from k_i by a multiplicative constant smaller than one. If there was no conditioning on $F_{i,j}$, it is easy to check that

$$\mathbf{E}\left[A_i \mid i \text{ arrived}\right] = \sum_{j'} \mathbf{Pr}\left[F_{i,j'}\right] \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j'} \leq (0.5 + \kappa) \cdot k_i \ .$$

The conditioning could however lead us into trouble in the following way: When facing the conditioning, we end up with the expression

$$\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] = \sum_{j'} \mathbf{Pr}\left[F_{i,j'} \mid F_{i,j}\right] \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j'} \ .$$

If $F_{i,j}$ implies $F_{i,j'}$ for every other item $j' \neq j$, and $\alpha_{i,j'} \approx 1$ for every $j' \neq j$, then $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] \approx \sum_{j'} 1 \cdot \frac{x_{i,j'}}{q_i} \cdot 1$ where the right-hand side could equal k_i. This, in particular, would make the second proposal in our algorithm completely useless as we would reduce the marginal probabilities for the pivotal sampling in Line 9 to (almost) zero. The most crucial part of our analysis is to demonstrate that this cannot happen, by bounding the positive correlation introduced between items.

Lemma 4.3.9. *For any distinct items j and j', and $\Delta_\kappa := \left(1 + \frac{(0.5+\kappa)^2}{0.5-\kappa}\right) \cdot \left(\frac{0.5+\kappa}{0.5-\kappa}\right)^2$, for any i we have*

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq \Delta_\kappa \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \ .$$

The proof of Lemma 4.3.9 is deferred to Section 4.3.3; in the remainder of this section we demonstrate why it implies our bound on the approximation ratio. We note that for $\kappa = 0.0115$ (the value we choose in Algorithm 4), we have $\Delta_\kappa \approx 1.68$. As a concrete example, note that if (i, j) and (i, j') are both late with $\mathbf{Pr}\left[F_{i,j}\right] \approx \mathbf{Pr}\left[F_{i,j'}\right] \approx 1/2$, this bound quantifies that we avoid perfect positive correlation between $F_{i,j}$ and $F_{i,j'}$.

Having Lemma 4.3.9, we can prove the bound on $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$ which we state formally in Corollary 4.3.10 via

$$\begin{aligned} \mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] &= \sum_{j'} \mathbf{Pr}\left[F_{i,j'} \mid F_{i,j}\right] \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j'} \\ &= \frac{x_{i,j}}{q_i} + \sum_{j' \neq j} \frac{\mathbf{Pr}\left[F_{i,j'} \wedge F_{i,j}\right]}{\mathbf{Pr}\left[F_{i,j}\right]} \cdot \frac{x_{i,j'}}{q_i} \cdot \alpha_{i,j'} \qquad (4.8) \\ &\leq \frac{x_{i,j}}{q_i} + \sum_{j' \neq j} \Delta_\kappa \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot \frac{x_{i,j'}}{q_i} \cdot \alpha_{i,j'} \\ &\leq \frac{x_{i,j}}{q_i} + \Delta_\kappa \cdot (0.5 + \kappa) \cdot k_i. \end{aligned}$$

The last inequality uses the fact that $\mathbf{Pr}[F_{i,j'}] \cdot \alpha_{i,j'} \leq 0.5 + \kappa$ and upper bounds $\sum_{j' \neq j} \frac{x_{i,j'}}{q_i}$ by k_i. By Constraint (4.3) and the property that $y_{i,j} > \tau$ for late pairs (i, j), we have that $\frac{x_{i,j}}{q_i} \leq 1 - \tau$. Hence, we can conclude that

$$\begin{aligned} \mathbf{E}[A_i \mid i \text{ arrived}, F_{i,j}] &\leq 1 - \tau + \Delta_\kappa \cdot (0.5 + \kappa) \cdot k_i \\ &\leq (1 - \tau + \Delta_\kappa \cdot (0.5 + \kappa)) \cdot k_i. \end{aligned} \tag{4.9}$$

Although the last inequality appears quite loose if k_i is larger than 1, in Section 4.3.4 we show that a fine-grained bound in terms of $\min_i k_i$ only results in limited improvements in the analysis. So, we use in the remainder that Equation (4.9) implies the following corollary of our correlation bound.

Corollary 4.3.10. *Let* $\Delta_\kappa := \left(1 + \frac{(0.5+\kappa)^2}{0.5-\kappa}\right) \cdot \left(\frac{0.5+\kappa}{0.5-\kappa}\right)^2$. *For any late* (i, j) *we have*

$$\mathbf{E}[A_i \mid i \text{ arrived}, F_{i,j}] \leq (1 - \tau + \Delta_\kappa \cdot (0.5 + \kappa)) \cdot k_i.$$

We are now able to conclude the proof of Lemma 4.3.5 (ii), as follows.

Proof of Lemma 4.3.5 (ii). By Observation 4.3.6, it suffices to show that $\rho_{i,j} \geq (0.5 + \kappa)y_{i,j} - (0.5 - \kappa)$. Combining the bound of Lemma 4.3.8 with Corollary 4.3.10 implies

$$\rho_{i,j} \geq (1 - (0.5 + \kappa) \cdot y_{i,j}) \cdot (\tau - (1 - \tau + \Delta_\kappa \cdot (0.5 + \kappa))) \ . \tag{4.10}$$

For convenience let $g(\kappa) := 2\tau - 1 - \Delta_\kappa \cdot (0.5 + \kappa)$, recalling that τ is a function of κ. Then, it suffices to show $(1 - (0.5 + \kappa) \cdot y_{i,j}) \cdot g(\kappa) \geq (0.5 + \kappa)y_{i,j} - (0.5 - \kappa)$, or equivalently

$$g(\kappa) + 0.5 - \kappa \geq (0.5 + \kappa + (0.5 + \kappa)g(\kappa)) \cdot y_{i,j} \ .$$

For $\kappa = 0.0115$, we can confirm that the coefficient of $y_{i,j}$ on the right-hand side is positive, and hence it suffices to show this inequality when $y_{i,j} = 1$. This reduces to

$$g(\kappa) \geq \frac{2\kappa}{0.5 - \kappa}$$

which is readily confirmed by direct computation at $\kappa = 0.0115$.

□

As a side remark, using Equation (4.10), we can observe that for our choice of $\kappa = 0.0115$, the expectation $\rho_{i,j}$ is bounded away from zero by a constant. In particular, for $\kappa = 0.0115$, we have that $\rho_{i,j} \geq 0.02389$. This can be used to estimate $\rho_{i,j}$ via sampling with small multiplicative error using standard Chernoff-Hoeffding-bounds.

In order to finalize our proof of Lemma 4.3.5 (ii), it only remains to prove our bound on the correlation introduced between items, which we do in the following section.

4.3.3 Bounding the Correlation — Proof of Lemma 4.3.9

To conclude the proof of Theorem 4.3.1, we need to control the correlation of the events that two items j and j' are free simultaneously. In particular, our goal is to prove the bound from Lemma 4.3.9. To this end, we first state and prove Lemma 4.3.11 which uses the assumption that $y_{i-1,j}$ and $y_{i-1,j'}$ are at most τ. Afterwards, we discuss its implications towards Lemma 4.3.9 with no restrictions on $y_{i-1,j}$ and $y_{i-1,j'}$.

Lemma 4.3.11. *Define* $\gamma_\kappa := 1 + \frac{(0.5+\kappa)^2}{0.5-\kappa}$. *For any distinct items* j *and* j'*, and any time* i *such that* $y_{i-1,j}, y_{i-1,j'} \leq \tau$*, we have*

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq \gamma_\kappa \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right].$$

In order to prove this, we consider the function

$$f(z) := 1 + z \cdot \left(\frac{(0.5+\kappa)^2}{1 - z \cdot (0.5+\kappa)}\right),$$

which depends on our choice of κ. Also, note that $\gamma_\kappa = f(1)$. For this function, we can prove the following claim.

Claim 4.3.12. *For any distinct items* j *and* j'*, and any buyer* i *with* $y_{i-1,j}, y_{i-1,j'} \leq \tau$*, we have*

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right],$$

where $f(z) := 1 + z \cdot \left(\frac{(0.5+\kappa)^2}{1-z\cdot(0.5+\kappa)}\right)$.

In order to prove Lemma 4.3.11 from Claim 4.3.12, it suffices to note that f is monotonically increasing in $[0,1]$, and hence, $f(z) \leq f(1) = \gamma_\kappa$ for all $z \in [0,1]$.

Proof of Claim 4.3.12. We give a proof by induction. As $f(0) = 1$ and all items are available initially, the base case is clear. Assuming the claim is true for fixed i, we will prove it for $i+1$ with the assumption $y_{i,j}, y_{i,j'} \leq \tau$.

Proof outline for the inductive step. Our proof proceeds with the following steps:

(S1) We find an upper bound for the probability that both j and j' are not assigned to i via a first proposal conditioned on being free.

(S2) We compute $\mathbf{Pr}\left[F_{i+1,j}\right] / \mathbf{Pr}\left[F_{i,j}\right]$, in order to apply the induction hypothesis.

(S3) We apply the induction hypothesis, and use Step (S2) to rewrite our bound in terms of $\mathbf{Pr}\left[F_{i+1,j}\right]$ and $\mathbf{Pr}\left[F_{i+1,j'}\right]$.

(S4) We argue that we can upper bound the coefficient in front of $\mathbf{Pr}\left[F_{i+1,j}\right] \cdot \mathbf{Pr}\left[F_{i+1,j'}\right]$ with $f(y_{i+1,j})$.

Step (S1): Bounding the probability of not assigning both items via a first proposal. As $y_{i,j}, y_{i,j'} \leq \tau$, items can only be assigned as first proposals; hence the probability that both j and j' are free at time $i+1$ is

$$\mathbf{Pr}\left[F_{i+1,j} \wedge F_{i+1,j'}\right] = \mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \cdot \underbrace{\mathbf{Pr}\left[(i,j) \notin \mathcal{A}_1 \wedge (i,j') \notin \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right]}_{(\star)}. \tag{4.11}$$

The first term on the right-hand side of Equation (4.11) will later be bounded via the induction hypothesis. The second term $(\star) := \mathbf{Pr}\left[(i,j) \notin \mathcal{A}_1 \wedge (i,j') \notin \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right]$ can be equivalently written as

$$\begin{aligned}(\star) = 1 &- \mathbf{Pr}\left[(i,j) \in \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] - \mathbf{Pr}\left[(i,j') \in \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] \\ &+ \mathbf{Pr}\left[(i,j) \in \mathcal{A}_1 \wedge (i,j') \in \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] .\end{aligned} \tag{4.12}$$

Now, observe that $\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] = q_i \cdot \mathbf{Pr}\left[j \in \mathsf{FP}_i\right] \cdot \alpha_{i,j} = x_{i,j} \cdot \alpha_{i,j}$. The analogous equality holds for j'. Hence, it remains to get a suitable bound on the joint probability that both items j and j' are assigned via a first proposal given they were both free. To this end, we make use of the negative cylinder dependence in pivotal sampling, observing

$$\begin{aligned}
&\mathbf{Pr}\left[(i,j) \in \mathcal{A}_1 \wedge (i,j') \in \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] \\
&= q_i \cdot \mathbf{Pr}\left[j \in \mathsf{FP}_i \wedge j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j} \cdot \alpha_{i,j'} \\
&\leq q_i \cdot \mathbf{Pr}\left[j \in \mathsf{FP}_i\right] \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j} \cdot \alpha_{i,j'} \qquad \text{(Pivotal Sampling Property (P3))} \\
&= q_i \cdot \frac{x_{i,j} \cdot x_{i,j'}}{q_i^2} \cdot \alpha_{i,j} \cdot \alpha_{i,j'} \\
&= \frac{x_{i,j} \cdot x_{i,j'}}{q_i} \cdot \alpha_{i,j} \cdot \alpha_{i,j'} \ .
\end{aligned}$$

Combining all of the above, we can bound the conditional probability that neither j nor j' is allocated to i via a first proposal. In other words, the left-hand side of Equation (4.12) is at most

$$\begin{aligned}
&\mathbf{Pr}\left[(i,j) \notin \mathcal{A}_1 \wedge (i,j') \notin \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] \\
&\leq 1 - x_{i,j}\alpha_{i,j} - x_{i,j'}\alpha_{i,j'} + \frac{1}{q_i} \cdot x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'} \qquad (4.13) \\
&= (1 - x_{i,j}\alpha_{i,j})(1 - x_{i,j'}\alpha_{i,j'}) + \left(\frac{1}{q_i} - 1\right) x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'} \ .
\end{aligned}$$

Step (S2): Comparing $\mathbf{Pr}\left[F_{i+1,j}\right]$ to $\mathbf{Pr}\left[F_{i,j}\right]$. To prepare for the use of the inductive hypothesis, we compute $\mathbf{Pr}\left[F_{i+1,j}\right] / \mathbf{Pr}\left[F_{i,j}\right]$ via a straightforward calculation:

$$\begin{aligned}
\mathbf{Pr}\left[F_{i+1,j}\right] = 1 - (0.5 + \kappa) \cdot y_{i+1,j} &= \mathbf{Pr}\left[F_{i,j}\right] \cdot \frac{1 - (0.5 + \kappa) \cdot y_{i+1,j}}{1 - (0.5 + \kappa) \cdot y_{i,j}} \\
&= \mathbf{Pr}\left[F_{i,j}\right] \cdot (1 - x_{i,j} \cdot \alpha_{i,j}) \, . \qquad (4.14)
\end{aligned}$$

In the final line, we used that (i,j) is early. For j', we analogously have

$$\mathbf{Pr}\left[F_{i+1,j'}\right] = \mathbf{Pr}\left[F_{i,j'}\right] \cdot (1 - x_{i,j'} \cdot \alpha_{i,j'}) \, .$$

Step (S3): Applying the induction hypothesis. Applying the induction hypothesis to Equation (4.11), plugging in Inequality (4.13) and using Equation (4.14), we can bound

$$\begin{aligned}
&\mathbf{Pr}\left[F_{i+1,j} \wedge F_{i+1,j'}\right] \\
&= \mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \cdot \mathbf{Pr}\left[(i,j) \notin \mathcal{A}_1 \wedge (i,j') \notin \mathcal{A}_1 \mid F_{i,j} \wedge F_{i,j'}\right] \\
&\leq \mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \cdot \left((1 - x_{i,j}\alpha_{i,j})(1 - x_{i,j'}\alpha_{i,j'}) + \left(\frac{1}{q_i} - 1\right) x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'}\right) \\
&\leq f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot \left((1 - x_{i,j}\alpha_{i,j})(1 - x_{i,j'}\alpha_{i,j'}) + \left(\frac{1}{q_i} - 1\right) x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'}\right) \\
&= f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i+1,j}\right] \cdot \mathbf{Pr}\left[F_{i+1,j'}\right] + f(y_{i,j}) \left(\frac{1}{q_i} - 1\right) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'} .
\end{aligned}$$

Here, the first inequality uses Inequality (4.13) from Step (S1), i.e. the upper bound on the probability of both items not being allocated via a first proposal. The second inequality applies the induction hypothesis for $\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right]$, the last equality uses Equation (4.14) from Step (S2) for both items j and j' and rearranges terms.

Having this, we pause for a moment to bound the second summand of the expression above, i.e., $f(y_{i,j})\left(\frac{1}{q_i}-1\right) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'}$. To this end, we will use the following inequality.

Fact 4.3.13. *For any (i,j) we have*

$$x_{i,j} \cdot \alpha_{i,j} \cdot \left(\frac{1}{q_i}-1\right) \leq (0.5+\kappa) \cdot (1 - x_{i,j}\alpha_{i,j})\,. \tag{4.15}$$

Proof. By Constraint (4.3) of the LP, we have that $\frac{1}{q_i} \leq \frac{1-y_{i,j}}{x_{i,j}}$. Thus it suffices to show that

$$\alpha_{i,j}(1-y_{i,j}) - x_{i,j}\alpha_{i,j} \leq (0.5+\kappa)(1-x_{i,j}\alpha_{i,j})$$

which is equivalent to

$$\alpha_{i,j}(1-y_{i,j}-(0.5-\kappa)x_{i,j}) \leq 0.5+\kappa.$$

As $\alpha_{i,j} \leq \frac{0.5+\kappa}{1-(0.5+\kappa)y_{i,j}}$, the claim follows. □

We can apply Fact 4.3.13 to item j' and combine it with Equation (4.14) in order to bound the second summand via

$$\begin{aligned}
&f(y_{i,j})\left(\frac{1}{q_i}-1\right) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot x_{i,j}\alpha_{i,j}x_{i,j'}\alpha_{i,j'} \\
&\leq f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot (0.5+\kappa) \cdot (1-x_{i,j'}\alpha_{i,j'}) \cdot x_{i,j}\alpha_{i,j} \\
&= (0.5+\kappa) \cdot f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i+1,j'}\right] \cdot x_{i,j}\alpha_{i,j} \cdot \mathbf{Pr}\left[F_{i+1,j}\right] \cdot (1-x_{i,j} \cdot \alpha_{i,j})^{-1}.
\end{aligned}$$

Overall, we thus have

$$\mathbf{Pr}\left[F_{i+1,j} \wedge F_{i+1,j'}\right] \leq f(y_{i,j}) \cdot \mathbf{Pr}\left[F_{i+1,j}\right] \cdot \mathbf{Pr}\left[F_{i+1,j'}\right] \cdot \left(1 + (0.5+\kappa) \cdot \frac{x_{i,j}\alpha_{i,j}}{1-x_{i,j}\alpha_{i,j}}\right).$$

Step (S4): Upper bounding the coefficient by $f(y_{i+1,j})$. In order to complete the inductive step, we would like to show that

$$f(y_{i,j}) \cdot \left(1 + (0.5+\kappa) \cdot \frac{x_{i,j}\alpha_{i,j}}{1-x_{i,j}\alpha_{i,j}}\right) \leq f(y_{i,j}+x_{i,j}).$$

First, note that as we only consider early pairs, $\alpha_{i,j}$ is always equal to $\frac{0.5+\kappa}{1-(0.5+\kappa)\cdot y_{i,j}}$, so we know

$$\frac{x_{i,j}\alpha_{i,j}}{1-x_{i,j}\alpha_{i,j}} = \frac{(0.5+\kappa) \cdot x_{i,j}}{1-(0.5+\kappa) \cdot (x_{i,j}+y_{i,j})}\ .$$

Thus to conclude the proof, we are required to show that

$$f(y_{i,j}) \cdot \left(1 + (0.5+\kappa) \cdot \frac{(0.5+\kappa) \cdot x_{i,j}}{1-(0.5+\kappa) \cdot (x_{i,j}+y_{i,j})}\right) \leq f(y_{i,j}+x_{i,j}).$$

In fact, this property is a consequence of our definition of

$$f(z) := 1 + z \cdot \left(\frac{(0.5+\kappa)^2}{1 - z \cdot (0.5+\kappa)} \right).$$

In particular the following claim completes the inductive step by providing exactly the required property.

Claim 4.3.14. *For any $x, y \in [0,1]$ with $x + y \leq 1$ and $f(\cdot)$ as stated above, we have*

$$f(y) \cdot \left(1 + (0.5+\kappa) \cdot \frac{(0.5+\kappa) \cdot x}{1 - (0.5+\kappa) \cdot (x+y)} \right) \leq f(y+x).$$

Proof of Claim 4.3.14. Plugging in the definition of $f(z) = 1 + z \cdot \left(\frac{(0.5+\kappa)^2}{1-z\cdot(0.5+\kappa)} \right)$, the claim is equivalent to

$$\left(1 + \frac{(0.5+\kappa)^2 y}{1-(0.5+\kappa)y} \right) \left(1 + \frac{(0.5+\kappa)^2 x}{1-(0.5+\kappa)(x+y)} \right) \leq 1 + \frac{(0.5+\kappa)^2 (x+y)}{1-(0.5+\kappa)(x+y)}.$$

Multiplying out the left-hand side and subtracting $1 + \frac{(0.5+\kappa)^2 x}{1-(0.5+\kappa)(x+y)}$ on both sides, this is equivalent to

$$\frac{(0.5+\kappa)^2 y}{1-(0.5+\kappa)y} + \frac{(0.5+\kappa)^2 y}{1-(0.5+\kappa)y} \cdot \frac{(0.5+\kappa)^2 x}{1-(0.5+\kappa)(x+y)} \leq \frac{(0.5+\kappa)^2 y}{1-(0.5+\kappa)(x+y)}.$$

If $y = 0$, the claim is trivially true. If $y > 0$, we can divide both sides by $(0.5+\kappa)^2 y$ to get

$$\frac{1}{1-(0.5+\kappa)y} + \frac{1}{1-(0.5+\kappa)y} \cdot \frac{(0.5+\kappa)^2 x}{1-(0.5+\kappa)(x+y)} \leq \frac{1}{1-(0.5+\kappa)(x+y)}.$$

Multiplying both sides by $(1-(0.5+\kappa)y) \cdot (1-(0.5+\kappa)(x+y))$, we get

$$1 - (0.5+\kappa)(x+y) + (0.5+\kappa)^2 x \leq 1 - (0.5+\kappa)y.$$

Subtracting $1 - (0.5+\kappa)y$ on both sides, we finally end up with

$$-(0.5+\kappa)x + (0.5+\kappa)^2 x \leq 0.$$

If $x = 0$, this is trivially true; for $x > 0$, we can divide by $(0.5+\kappa)x$ to see that the claim follows.

□

This concludes the proof of Claim 4.3.12.

□

Now, we can finally prove Lemma 4.3.9 which is the last ingredient in the proof of Theorem 4.3.1. Recall Lemma 4.3.9:

Lemma 4.3.9. *For any distinct items j and j', and $\Delta_\kappa := \left(1 + \frac{(0.5+\kappa)^2}{0.5-\kappa}\right) \cdot \left(\frac{0.5+\kappa}{0.5-\kappa}\right)^2$, for any i we have*

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq \Delta_\kappa \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \ .$$

Proof of Lemma 4.3.9. We assume that both $y_{i,j} > \tau$ and $y_{i,j'} > \tau$; if neither inequality holds the result is clear and follows directly from Lemma 4.3.11 while if just one holds the proof proceeds nearly identically with a slightly better guarantee.

Let i^j denote the latest buyer in $[n]$ such that $y_{i^j-1,j} \leq \tau$ and $y_{i^j,j} > \tau$ and similarly let $i^{j'}$ denote the latest buyer in $[n]$ such that $y_{i^{j'}-1,j'} \leq \tau$ and $y_{i^{j'},j'} > \tau$.

Let A_j denote the event that j is allocated to some arrival in $[i^j, i-1]$ and let $A_{j'}$ denote the event that j' is allocated to some arrival in $[i^{j'}, i-1]$. By the hypothesis that $\mathbf{Pr}\left[(i',j) \in \mathcal{A}_1\right] + \mathbf{Pr}\left[(i',j) \in \mathcal{A}_2\right] = (0.5+\kappa) \cdot x_{i',j}$ for all $i' < i$, we have

$$\mathbf{Pr}\left[A_j\right] = \sum_{i' \in [i^j, i-1]} (0.5+\kappa) \cdot x_{i',j} = (0.5+\kappa) \cdot (y_{i,j} - y_{i^j,j}) \leq (0.5+\kappa) \cdot (1-\tau) = 2\kappa.$$

An analogous upper bound holds for $\mathbf{Pr}\left[A_{j'}\right]$.

To simplify notation, let us assume for a moment that $i^j \leq i^{j'}$ (if $i^j > i^{j'}$, simply swap the roles of j and j' in the following line). We apply Lemma 4.3.11 to get

$$\begin{aligned}
\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] &\leq \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'}\right] \\
&= \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^j,j'}\right] \cdot \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'} \;\middle|\; F_{i^j,j} \wedge F_{i^j,j'}\right] \\
&\leq \gamma_\kappa \cdot \mathbf{Pr}\left[F_{i^j,j}\right] \cdot \mathbf{Pr}\left[F_{i^j,j'}\right] \cdot \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'} \;\middle|\; F_{i^j,j} \wedge F_{i^j,j'}\right] \quad \text{(via Lemma 4.3.11)} \\
&= \gamma_\kappa \cdot \mathbf{Pr}\left[F_{i^j,j}\right] \cdot \mathbf{Pr}\left[F_{i^j,j'}\right] \cdot \mathbf{Pr}\left[F_{i^{j'},j'} \;\middle|\; F_{i^j,j} \wedge F_{i^j,j'}\right] .
\end{aligned}$$

In this expression, we aim to combine the last two factors concerning the events if item j' is free at some point in time. To this end, observe that

$$\begin{aligned}
\mathbf{Pr}\left[F_{i^j,j'}\right] \cdot \mathbf{Pr}\left[F_{i^{j'},j'} \;\middle|\; F_{i^j,j} \wedge F_{i^j,j'}\right] &= \mathbf{Pr}\left[F_{i^j,j'}\right] \cdot \prod_{i'=i^j}^{i^{j'}-1} \left(1 - q_{i'} \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_{i'}\right] \cdot \alpha_{i',j'}\right) \\
&= \mathbf{Pr}\left[F_{i^j,j'}\right] \cdot \prod_{i'=i^j}^{i^{j'}-1} \left(1 - x_{i',j'} \cdot \alpha_{i',j'}\right) \\
&= \mathbf{Pr}\left[F_{i^{j'},j'}\right] \quad ,
\end{aligned}$$

where the last equality uses the reasoning from Step (S2) in the proof of Claim 4.3.12. So, overall, we have

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'}\right] \leq \gamma_\kappa \cdot \mathbf{Pr}\left[F_{i^j,j}\right] \cdot \mathbf{Pr}\left[F_{i^{j'},j'}\right] \quad . \tag{4.16}$$

With this in mind, we are ready to prove the final statement as

$$\begin{aligned}
\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] &\leq \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'}\right] \\
&\leq \gamma_\kappa \cdot \mathbf{Pr}\left[F_{i^j,j}\right] \cdot \mathbf{Pr}\left[F_{i^{j'},j'}\right] && \text{(via Equation (4.16))} \\
&= \gamma_\kappa \cdot \left(\mathbf{Pr}\left[F_{i,j}\right] + \mathbf{Pr}\left[A_j\right]\right) \cdot \left(\mathbf{Pr}\left[F_{i,j'}\right] + \mathbf{Pr}\left[A_{j'}\right]\right) \\
&\leq \gamma_\kappa \cdot \left(\mathbf{Pr}\left[F_{i,j}\right] + 2\kappa\right) \cdot \left(\mathbf{Pr}\left[F_{i,j'}\right] + 2\kappa\right) \\
&\leq \gamma_\kappa \cdot \left(1 + \frac{4\kappa}{0.5-\kappa} + \frac{4\kappa^2}{(0.5-\kappa)^2}\right) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \\
&= \gamma_\kappa \cdot \left(\frac{0.5+\kappa}{0.5-\kappa}\right)^2 \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \\
&= \Delta_\kappa \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right],
\end{aligned}$$

where in the last inequality we used $\mathbf{Pr}\left[F_{i,j}\right], \mathbf{Pr}\left[F_{i,j'}\right] \geq 0.5-\kappa$ and the last equality applies $\gamma_\kappa := 1 + {}^{(0.5+\kappa)^2}/_{0.5-\kappa}$.

□

Having this, we concluded the proof of the approximation guarantee in Theorem 4.3.1.

4.3.4 A Bound Depending on $\min_i k_i$.

As mentioned in Section 4.3.2, the bound from Equation (4.9) in the proof of Corollary 4.3.10 is not tight if all k_i are strictly greater than one. At first glance, this step looks quite lossy. Still, we are not losing much in our analysis by replacing $\min_i k_i$ with one. To see this, consider replacing the last inequality in the proof of Corollary 4.3.10 with a bound depending on $\min_i k_i$. Doing so, we get

$$\begin{aligned}
\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] &\leq 1 - \tau + \Delta_\kappa \cdot (0.5+\kappa) \cdot k_i \\
&= \left(\frac{1-\tau}{k_i} + \Delta_\kappa \cdot (0.5+\kappa)\right) \cdot k_i \\
&\leq \left(\frac{1-\tau}{\min_{i'} k_{i'}} + \Delta_\kappa \cdot (0.5+\kappa)\right) \cdot k_i \quad . \qquad (4.17)
\end{aligned}$$

As a consequence, in order to show the desired lower bound on $\rho_{i,j}$, we can use the same reasoning as before, but apply Inequality (4.17) instead:

$$\rho_{i,j} \geq (1-(0.5+\kappa)\cdot y_{i,j}) \cdot \left(\tau - \left(\frac{1-\tau}{\min_{i'} k_{i'}} + \Delta_\kappa \cdot (0.5+\kappa)\right)\right) \quad .$$

Thus, the right-hand side needs to be at least as large as $(0.5+\kappa)y_{i,j} - (0.5-\kappa)$. In other words, we are required to show that

$$(1-(0.5+\kappa)\cdot y_{i,j}) \cdot \left(\tau - \left(\frac{1-\tau}{\min_{i'} k_{i'}} + \Delta_\kappa \cdot (0.5+\kappa)\right)\right) \geq (0.5+\kappa)y_{i,j} - (0.5-\kappa) \quad .$$

Hence we can choose any κ such that

$$\tau - \left(\frac{1-\tau}{\min_{i'} k_{i'}} + \Delta_\kappa \cdot (0.5+\kappa)\right) \geq \frac{2\kappa}{0.5-\kappa} \quad . \qquad (4.18)$$

$\min_i k_i$	1	2	3	4	5	6	7	8
κ	0.0115	0.0126	0.0131	0.0133	0.0134	0.0135	0.01362	0.01367

Table 4.1: Values of κ depending on $\min_i k_i$

When solving Equation (4.18) for κ, we get some small improvement upon the constant of 0.0115 as a function of $\min_i k_i$. In Table 4.1, we state these constants for $\min_i k_i \in \{2, \dots, 8\}$, demonstrating that there is little loss in our analysis of Algorithm 4 when replacing $\min_i k_i$ with one.

Also, even for large values of $\min_i k_i$, this approach is limited. In particular, for $\min_i k_i \to \infty$, the best κ one can hope for is approximately 0.01402 when being restricted to Equation (4.18). Therefore, in order to get a significant improvement, a more fine grained technique is required.

4.4 Beating Half with a Simplified Algorithm

We state a slightly modified variant of Algorithm 4 which can avoid sampling estimates to run in polynomial time. We remove the subsampling step with the $\beta_{i,j}$, or in other words, we set $\beta_{i,j} = 1$ for all j and i. As we will see, this is still sufficient to improve the approximation guarantee of $1/2$ by a constant. The algorithm is stated in Algorithm 5 below and is clearly running in polynomial time.

Algorithm 5: Simplified Allocation Algorithm

1 $\tilde{\kappa} \leftarrow 0.0076$
2 Solve LP_{on} for $\{x_{i,j}\}_{i,j}$
3 **for** *each buyer i, if i arrives* **do**
4   Define items $\mathsf{FP}_i := \mathsf{PS}((x_{i,j}/q_i)_j)$
5   **for** *each item $j \in \mathsf{FP}_i$* **do**
6   **if** *j is available* **then**
7   Allocate j to i with probability $\alpha_{i,j} := \min\left(1, \frac{0.5+\tilde{\kappa}}{1-(0.5+\tilde{\kappa})\cdot\sum_{i'<i} x_{i',j}}\right)$
8     $A_i \leftarrow$ number of items allocated to i thus far
9   Define items $\mathsf{SP}_i := \mathsf{PS}(((1 - \frac{A_i}{k_i}) \cdot x_{i,j}/q_i)_j)$
10   **for** *each item $j \in \mathsf{SP}_i$ with $\alpha_{i,j} = 1$* **do**
11   **if** *j is available* **then**
12   Allocate j to i

In addition, we can state the following theorem.

Theorem 4.4.1. *Denote by X_i the set of items which are allocated to buyer i via Algorithm 5. Then, for $\tilde{\kappa} = 0.0076$, we have*

$$\mathbf{E}\left[\sum_i \max_{S \subseteq X_i, |S| \le k_i} \left(\sum_{j \in S} v_{i,j}\right)\right] \ge (0.5 + \tilde{\kappa}) \cdot \text{OPT}_{\text{on}}.$$

The proof of this theorem proceeds fairly similar to the one of Theorem 4.3.1. Still, we need to be more careful as it is not always guaranteed that $\mathbf{Pr}\left[F_{i,j}\right] = 1 - (0.5 + \tilde{\kappa}) y_{i,j}$ due to possibly allocating with too high probability in Line 12 of Algorithm 5.

We discuss the modifications in the key observations and steps from the analysis of Algorithm 4 in the following section and show how extend it to prove Theorem 4.4.1.

4.4.1 Beating 1/2 via the Analysis from Section 4.3

Again, the proof proceeds by induction on i. We start with a few observations concerning the modified algorithm. First, we mention that τ now depends on $\tilde{\kappa}$ instead of κ, so does $\alpha_{i,j}$. Then, note that any early pair (i, j) is not affected by the replacement of $\beta_{i,j}$. The reason for this is that for early pairs (i, j), we have $\alpha_{i,j} < 1$ and hence, we never allocate these via a second proposal. As a consequence, Observation 4.3.3 and Observation 4.3.4 directly carry over (as they did not use any assumption on the exact choice of κ). These observations imply that the probability for any early pair (i, j) to be assigned is exactly $(0.5 + \tilde{\kappa}) \cdot x_{i,j}$.

So, using Observation 4.3.2 which also still holds, we only need to argue about the late pairs (i, j) and in particular, we need to prove a modified variant of Lemma 4.3.5. Once we have that for late pairs (i, j), the probability of being assigned is at least $(0.5 + \tilde{\kappa}) \cdot x_{i,j}$ via first and second proposals together, we can conclude as in the proof of Theorem 4.3.1.

So, in order to argue that Algorithm 5 allows to outperform 1/2 by a constant, we are going to prove the following modified variant of Lemma 4.3.5.

Lemma 4.4.2. *For late pairs (i, j) in Algorithm 5, the following two statements hold:*

(j) $\mathbf{Pr}\left[(i, j) \in \mathcal{A}_1\right] \geq (0.5 - 3\tilde{\kappa}) \cdot x_{i,j}$ *and*

(ii) $\mathbf{Pr}\left[(i, j) \in \mathcal{A}_2\right] \geq 4\tilde{\kappa} \cdot x_{i,j}$.

Note that if we sum the contribution of (i) and (ii) in the lemma above, we get that the probability of assigning the pair (i, j) either via a first or second proposal is at least $(0.5 + \tilde{\kappa}) \cdot x_{i,j}$ as desired. In order to prove this lemma, we will make use of the following observations.

Observation 4.4.3. *For any late pair (i, j), when using Algorithm 5, we have*

$$\mathbf{Pr}\left[F_{i,j}\right] \leq 0.5 + \tilde{\kappa} \ .$$

Proof. Note that for any late (i, j), we have that $y_{i,j} > \tau$. Therefore, we can bound

$$\begin{aligned}
\mathbf{Pr}\left[F_{i,j}\right] &\leq 1 - \mathbf{Pr}\left[j \text{ assigned to some } i' < i \text{ where } (i', j) \text{ was early}\right] \\
&= 1 - (0.5 + \tilde{\kappa}) \sum_{i' < i : (i', j) \text{ early}} x_{i',j} \\
&\leq 1 - (0.5 + \tilde{\kappa}) \cdot \tau = 1 - (0.5 + \tilde{\kappa}) \cdot \frac{0.5 - \tilde{\kappa}}{0.5 + \tilde{\kappa}} \\
&= 0.5 + \tilde{\kappa} \ ,
\end{aligned}$$

where the first equality uses that for early pairs (i', j), by Observation 4.3.4, the probability of assigning is exactly $(0.5 + \tilde{\kappa}) \cdot x_{i',j}$. For the second inequality, as (i, j) is late, the sum of the LP variables of early pairs (i', j) with $i' < i$ exceeds τ, so $\sum_{i' < i : (i', j) \text{ early}} x_{i',j} \geq \tau$. □

In addition, we can also find a lower bound on the probability of item j being free at the arrival of buyer i.

Observation 4.4.4. *For any late pair* (i, j)*, when using Algorithm 5, we have*

$$\mathbf{Pr}\left[F_{i,j}\right] \geq 0.5 - 3\tilde{\kappa} \ .$$

Proof. For any late pair (i', j) with $i' < i$, the probability of assigning j to i' via a first proposal equals

$$\mathbf{Pr}\left[(i', j) \in \mathcal{A}_1\right] = q_{i'} \cdot \mathbf{Pr}\left[F_{i',j}\right] \cdot \frac{x_{i',j}}{q_{i'}} = \mathbf{Pr}\left[F_{i',j}\right] \cdot x_{i',j} \leq (0.5 + \tilde{\kappa}) \cdot x_{i',j} \ ,$$

where the last inequality uses Observation 4.4.3.

In a more rude manner, the probability of assigning j to i' via a second proposal can be upper bounded by $q_{i'} \cdot \mathbf{Pr}\left[F_{i',j}\right] \cdot {}^{x_{i',j}}/_{q_{i'}}$ as well. To see this, observe that in order to assign pair (i', j) via a second proposal, first of all j needs to be free before the arrival of i', i' needs to arrive, the first pivotal sampling cannot sample j and the second pivotal sampling with reduced marginals needs to pick j in the second proposal set. Of course, the reduced marginals for the second proposal are at most ${}^{x_{i',j}}/_{q_{i'}}$. Therefore, we can bound

$$\mathbf{Pr}\left[(i', j) \in \mathcal{A}_2\right] \leq \mathbf{Pr}\left[F_{i',j}\right] \cdot q_{i'} \cdot {}^{x_{i',j}}/_{q_{i'}} = \mathbf{Pr}\left[F_{i',j}\right] \cdot x_{i',j} \leq (0.5 + \tilde{\kappa}) \cdot x_{i',j} \ ,$$

where the last inequality uses Observation 4.4.3 again.

Any early pair (i', j) is assigned via a first proposal exactly with probability $(0.5 + \tilde{\kappa}) \cdot x_{i',j}$ via Observation 4.3.4. As a consequence, we get that for a late pair (i, j),

$$\begin{aligned}
\mathbf{Pr}\left[F_{i,j}\right] &= 1 - \sum_{i'<i} \mathbf{Pr}\left[(i', j) \in \mathcal{A}_1\right] - \sum_{i'<i} \mathbf{Pr}\left[(i', j) \in \mathcal{A}_2\right] \\
&\geq 1 - \sum_{i'<i} (0.5 + \tilde{\kappa}) \cdot x_{i',j} - \sum_{i'<i:(i',j) \text{ late}} (0.5 + \tilde{\kappa}) \cdot x_{i',j} \\
&\geq 1 - (0.5 + \tilde{\kappa}) - (0.5 + \tilde{\kappa}) \left(1 - \sum_{i'<i:(i',j) \text{ early}} x_{i',j}\right) \\
&\geq 1 - (0.5 + \tilde{\kappa}) - (0.5 + \tilde{\kappa})(1 - \tau) \\
&= 0.5 - 3\tilde{\kappa} \ .
\end{aligned}$$

Here, the first inequality uses the bounds which we just derived before and the second inequality exploits constraints from the LP on the sum of LP variables. The last inequality is based on the fact that (i, j) is late and hence, the sum of LP variables of previously arrived early pairs (i', j) is at least τ. □

Having these two observations, we can prove Lemma 4.4.2 (i) fairly easily.

Proof of Lemma 4.4.2 (i). Note that if pair (i, j) is late, $\alpha_{i,j} = 1$. Hence, by the use of Observation 4.4.4

$$\begin{aligned}
\mathbf{Pr}\left[(i, j) \in \mathcal{A}_1\right] &= q_i \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[j \in \mathsf{FP}_i\right] \cdot \alpha_{i,j} \\
&= q_i \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \frac{x_{i,j}}{q_i} \cdot 1 \\
&= \mathbf{Pr}\left[F_{i,j}\right] \cdot x_{i,j} \\
&\geq (0.5 - 3\tilde{\kappa}) \cdot x_{i,j}.
\end{aligned}$$

□

So what remains is to prove Lemma 4.4.2 (ii) in order to conclude the proof of Theorem 4.4.1. To this end, we proceed similarly to Section 4.3.2.

Lemma 4.4.5. *For late pairs* (i, j)*, we have*

$$\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] \geq x_{i,j} \cdot (0.5 - 3\tilde{\kappa}) \cdot \left(\tau - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i}\right) \quad .$$

Proof. The proof proceeds equivalently to the one in Section 4.3.2. Still, we need to use Observation 4.4.4 when lower bounding the probability of j being free at time i.

Therefore, we get

$$\begin{aligned}\mathbf{Pr}\left[(i,j) \in \mathcal{A}_2\right] &= q_i \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[j \text{ not allocated in Line 7 } \wedge j \in \mathsf{SP}_i \mid i \text{ arrived}, F_{i,j}\right] \\ &\geq q_i \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \tau \cdot \left(1 - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i} \cdot \frac{1}{\tau}\right) \cdot \frac{x_{i,j}}{q_i} \\ &\geq x_{i,j} \cdot (0.5 - 3\tilde{\kappa}) \cdot \left(\tau - \frac{\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]}{k_i}\right) \quad ,\end{aligned}$$

where the first inequality uses exactly the same reasoning as the proof of Lemma 4.3.8 and the second inequality applies Observation 4.4.4. □

So, in order to bound $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$, we are required to adapt the bound on the correlation with respect to the modified algorithm. To this end, first observe that Lemma 4.3.11 is considering early pairs only and hence holds identically also for Algorithm 5. Using this, we can obtain a bound for the joint probability of two items j and j' being available at time i in Algorithm 5.

Lemma 4.4.6. *For any distinct items* j *and* j'*, and any time* i*, we have*

$$\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] \leq \gamma_{\tilde{\kappa}} \cdot \left(\frac{0.5 + \tilde{\kappa}}{0.5 - 3\tilde{\kappa}}\right)^2 \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \quad .$$

Proof. Again, let us assume that both $y_{i,j} > \tau$ and $y_{i,j'} > \tau$; if neither inequality holds the result is clear and follows directly while if just one holds the proof proceeds nearly identically.

Let i^j denote the latest buyer in $[n]$ such that $y_{i^j-1,j} \leq \tau$ and $y_{i^j,j} > \tau$ and similarly let $i^{j'}$ denote the latest buyer in $[n]$ such that $y_{i^{j'}-1,j'} \leq \tau$ and $y_{i^{j'},j'} > \tau$. Let A_j denote the event that j is allocated to some buyer in $[i^j, i-1]$ and let $A_{j'}$ denote the event that j' is allocated to some buyer in $[i^{j'}, i-1]$.

Now, observe that

$$\begin{aligned}\mathbf{Pr}\left[A_j\right] &= \sum_{i'<i:(i',j) \text{ late}} \left(\mathbf{Pr}\left[(i',j) \in \mathcal{A}_1\right] + \mathbf{Pr}\left[(i',j) \in \mathcal{A}_2\right]\right) \\ &\leq \sum_{i'<i:(i',j) \text{ late}} 2 \cdot (0.5 + \tilde{\kappa}) \cdot x_{i',j} \\ &\leq 2 \cdot (0.5 + \tilde{\kappa}) \cdot (1 - \tau) \\ &= 4\tilde{\kappa} \quad ,\end{aligned}$$

where the first inequality holds due to the same arguments as in the proof of Observation 4.4.4. An analogous upper bound can be stated for $\mathbf{Pr}\left[A_{j'}\right]$.

With this, we can use the same reasoning as in Lemma 4.3.9 by an application of Lemma 4.3.11 and Equation (4.16) to bound

$$\begin{aligned}
\mathbf{Pr}\left[F_{i,j} \wedge F_{i,j'}\right] &\le \mathbf{Pr}\left[F_{i^j,j} \wedge F_{i^{j'},j'}\right] \\
&\le \gamma_{\tilde{\kappa}} \cdot \mathbf{Pr}\left[F_{i^j,j}\right] \cdot \mathbf{Pr}\left[F_{i^{j'},j'}\right] && \text{(via Equation (4.16))} \\
&= \gamma_{\tilde{\kappa}} \cdot \left(\mathbf{Pr}\left[F_{i,j}\right] + \mathbf{Pr}\left[A_j\right]\right) \cdot \left(\mathbf{Pr}\left[F_{i,j'}\right] + \mathbf{Pr}\left[A_{j'}\right]\right) \\
&\le \gamma_{\tilde{\kappa}} \cdot \left(\mathbf{Pr}\left[F_{i,j}\right] + 4\tilde{\kappa}\right) \cdot \left(\mathbf{Pr}\left[F_{i,j'}\right] + 4\tilde{\kappa}\right) \\
&\le \gamma_{\tilde{\kappa}} \cdot \left(1 + \frac{8\tilde{\kappa}}{0.5 - 3\tilde{\kappa}} + \frac{16\tilde{\kappa}^2}{(0.5 - 3\tilde{\kappa})^2}\right) \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \\
&= \gamma_{\tilde{\kappa}} \cdot \left(\frac{0.5 + \tilde{\kappa}}{0.5 - 3\tilde{\kappa}}\right)^2 \cdot \mathbf{Pr}\left[F_{i,j}\right] \cdot \mathbf{Pr}\left[F_{i,j'}\right] \quad ,
\end{aligned}$$

where in the last inequality we used Observation 4.4.4 for both j and j'. ☐

Having this, we can now turn towards bounding $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$.

Lemma 4.4.7. *Let* $\gamma_{\tilde{\kappa}} := 1 + \frac{(0.5+\tilde{\kappa})^2}{0.5-\tilde{\kappa}}$. *For any buyer i, we have*

$$\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] \le \left(1 - \tau + \gamma_{\tilde{\kappa}} \cdot \frac{(0.5 + \tilde{\kappa})^3}{(0.5 - 3\tilde{\kappa})^2}\right) \cdot k_i \quad .$$

Proof. Applying Lemma 4.4.6, we can bound the probability of item j' being free conditioned on item j being also free via

$$\mathbf{Pr}\left[F_{i,j'} \mid F_{i,j}\right] = \frac{\mathbf{Pr}\left[F_{i,j'} \wedge F_{i,j}\right]}{\mathbf{Pr}\left[F_{i,j}\right]} \le \gamma_{\tilde{\kappa}} \cdot \left(\frac{0.5 + \tilde{\kappa}}{0.5 - 3\tilde{\kappa}}\right)^2 \cdot \mathbf{Pr}\left[F_{i,j'}\right] \quad .$$

Using this for all items $j' \neq j$, we can compute

$$\begin{aligned}
\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] &= \sum_{j'} \mathbf{Pr}\left[F_{i,j'} \mid F_{i,j}\right] \cdot \mathbf{Pr}\left[j' \in \mathsf{FP}_i\right] \cdot \alpha_{i,j'} \\
&\le \frac{x_{i,j}}{q_i} + \sum_{j' \neq j} \gamma_{\tilde{\kappa}} \cdot \left(\frac{0.5 + \tilde{\kappa}}{0.5 - 3\tilde{\kappa}}\right)^2 \cdot \mathbf{Pr}\left[F_{i,j'}\right] \cdot \frac{x_{i,j'}}{q_i} \cdot \alpha_{i,j'} \\
&\le \frac{x_{i,j}}{q_i} + \gamma_{\tilde{\kappa}} \cdot \left(\frac{0.5 + \tilde{\kappa}}{0.5 - 3\tilde{\kappa}}\right)^2 \cdot (0.5 + \tilde{\kappa}) \cdot k_i \quad .
\end{aligned}$$

The second inequality first uses the fact that $\mathbf{Pr}\left[F_{i,j'}\right] \cdot \alpha_{i,j'} \le 0.5 + \tilde{\kappa}$: For early (i, j'), this is indeed an equality, for late (i, j'), we can use that $\alpha_{i,j'} = 1$ together with Observation 4.4.3. Afterwards, we upper bound $\sum_{j' \neq j} \frac{x_{i,j'}}{q_i}$ by k_i. Using that $x_{i,j}/q_i \le 1 - \tau$ for late pairs (i, j), we can conclude that

$$\begin{aligned}
\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right] &\le 1 - \tau + \gamma_{\tilde{\kappa}} \cdot \frac{(0.5 + \tilde{\kappa})^3}{(0.5 - 3\tilde{\kappa})^2} \cdot k_i \\
&\le \left(1 - \tau + \gamma_{\tilde{\kappa}} \cdot \frac{(0.5 + \tilde{\kappa})^3}{(0.5 - 3\tilde{\kappa})^2}\right) \cdot k_i \quad .
\end{aligned}$$

☐

Having all of this, we are now able to conclude the proof of Lemma 4.4.2 (ii).

Proof of Lemma 4.4.2 (ii). We start with the lower bound on $\mathbf{Pr}\left[(i,j)\in\mathcal{A}_2\right]$ from Lemma 4.4.5. Combined with the bound on $\mathbf{E}\left[A_i \mid i \text{ arrived}, F_{i,j}\right]$ from Lemma 4.4.7, we get

$$\begin{aligned}\mathbf{Pr}\left[(i,j)\in\mathcal{A}_2\right] &\geq x_{i,j}\cdot(0.5-3\tilde{\kappa})\cdot\left(\tau-\frac{\mathbf{E}\left[A_i\mid i\text{ arrived}, F_{i,j}\right]}{k_i}\right)\\ &\geq x_{i,j}\cdot(0.5-3\tilde{\kappa})\left(\tau-\left(1-\tau+\gamma_{\tilde{\kappa}}\cdot\frac{(0.5+\tilde{\kappa})^3}{(0.5-3\tilde{\kappa})^2}\right)\right)\\ &= x_{i,j}\cdot(0.5-3\tilde{\kappa})\left(2\tau-1-\gamma_{\tilde{\kappa}}\cdot\frac{(0.5+\tilde{\kappa})^3}{(0.5-3\tilde{\kappa})^2}\right)\quad.\end{aligned}$$

Plugging in our choice of $\tilde{\kappa}=0.0076$, we obtain

$$(0.5-3\tilde{\kappa})\left(2\tau-1-\gamma_{\tilde{\kappa}}\cdot\frac{(0.5+\tilde{\kappa})^3}{(0.5-3\tilde{\kappa})^2}\right)\ \geq\ 4\tilde{\kappa}\ ,$$

which implies our desired bound. □

4.5 Beyond Two-Point Distributions

When not restricting the model to two-point distributions for the arrivals, for every buyer i, there is a known distribution $\{q_{i,\ell}\}_\ell$ over valuation functions $v_{i,1},\dots,v_{i,L}$. Equivalently, we can also think about distributions over vectors $\{v_{i,j,\ell}\}_j$ and demand sizes $k_{i,\ell}$. In this section, we do the latter and hence, we have $\sum_\ell q_{i,\ell}=1$ and each pair $(\{v_{i,j,\ell}\}_j, k_{i,\ell})$ is sampled in step i with probability $q_{i,\ell}$.

That said, upon the arrival of buyer i, she samples one index $\ell\in\{1,\dots,L\}$ with probability $q_{i,\ell}$[3] and realizes demand size $k_{i,\ell}$ and values $\{v_{i,j,\ell}\}_j$ over items.

Generalized LP. We generalize LP_{on} as follows.

$$\begin{aligned}\max\ & \sum_{i,j,\ell} v_{i,j,\ell}\cdot x_{i,j,\ell} && && (\text{General-LP}_{\text{on}})\\ \text{s.t.}\ & \sum_i\sum_\ell x_{i,j,\ell}\leq 1 && \text{for all } j\in M && (4.19)\\ & \sum_j x_{i,j,\ell}\leq q_{i,\ell}\cdot k_{i,\ell} && \text{for all } i\in[n],\ell\in[L] && (4.20)\\ & 0\leq x_{i,j,\ell}\leq q_{i,\ell}\cdot\left(1-\sum_{i'<i}\sum_{\ell'}x_{i',j,\ell'}\right) && \text{for all } j\in M, i\in[n],\ell\in[L] && (4.21)\end{aligned}$$

As in Observation 4.1.1, it is easy to see that $\text{General-LP}_{\text{on}}$ is a relaxation of the expected welfare achieved by the optimum online algorithm for general distributions, so $\text{OPT}(\text{General-LP}_{\text{on}})\geq\text{OPT}_{\text{on}}$.

[3] We assume without loss of generality that all buyers share the same space of valuation functions as we can set $q_{i,\ell}=0$ if realization ℓ is not feasible for buyer i. Also, we assume that buyers always arrive by adding a valuation function evaluating any set of items by zero which is sampled with the probability of buyer i not arriving.

Generalized Algorithm. In order to round any fractional LP solution to an integral one in an online fashion, we extend our Algorithm 4 as follows: In round i, we see the realization of index ℓ. We replace all previous LP variables with the ones from the generalized LP for index ℓ and run the slightly modified Algorithm 6.

Algorithm 6: Allocation Algorithm for General Distributions

---

1 $\kappa \leftarrow 0.0115$
2 Solve General-LP$_{\text{on}}$ for $\{x_{i,j,\ell}\}_{i,j,\ell}$
3 **for** *each buyer* i **do**
4   Observe index ℓ sampled from $(q_{i,\ell})_\ell$
5   Define items $\mathsf{FP}_{i,\ell} := \mathsf{PS}((x_{i,j,\ell}/q_{i,\ell})_j)$
6   **for** *each item* $j \in \mathsf{FP}_{i,\ell}$ **do**
7   **if** j *is available* **then**
8   Allocate j to i with probability $\alpha_{i,j} := \min\left(1, \frac{0.5+\kappa}{1-(0.5+\kappa)\cdot\sum_{i'<i}\sum_{\ell'} x_{i',j,\ell'}}\right)$
9     $A_{i,\ell} \leftarrow$ number of items allocated to i with sampled index ℓ thus far
10   Define items $\mathsf{SP}_{i,\ell} := \mathsf{PS}\left(\left(\left(1 - \frac{A_{i,\ell}}{k_{i,\ell}}\right) \cdot x_{i,j,\ell}/q_{i,\ell}\right)_j\right)$
11   **for** *each item* $j \in \mathsf{SP}_{i,\ell}$ *with* $\alpha_{i,j} = 1$ **do**
12   **if** j *is available* **then**
13   Compute $\rho_{i,j,\ell} :=$
               $\mathbf{E}\left[\mathbb{1}[j \text{ available after Line 9}] \cdot \left(1 - \frac{A_{i,\ell}}{k_{i,\ell}}\right) \mid i \text{ sampled index } \ell\right]$
14             $\beta_{i,j,\ell} \leftarrow \min\left(1, \left((0.5+\kappa)\cdot\sum_{i'<i}\sum_{\ell'} x_{i',j,\ell'} - (0.5-\kappa)\right)\cdot\frac{1}{\rho_{i,j,\ell}}\right)$
15   Allocate j to i with prob. $\beta_{i,j,\ell}$

As in the previously studied Bernoulli case, observe that we choose $\beta_{i,j,\ell}$ in a way so that the following holds: $\mathbf{Pr}\left[(i,j) \text{ assigned with sampled index } \ell\right] = (0.5+\kappa)\cdot x_{i,j,\ell}$. Also, note that this algorithm can be implemented in polynomial time in the number of buyers and items and the size of the support of the distributions. Concerning the computation of $\rho_{i,j,\ell}$, we can observe that for our choice of $\kappa = 0.0115$, the generalized analysis also shows that any $\rho_{i,j,\ell}$ is lower bounded by a constant; equivalently to the Bernoullli case. As before, this can be used to estimate $\rho_{i,j,\ell}$ via samples with a multiplicative error as small as desired, implying a $(0.5+\kappa-\epsilon)$-approximate algorithm. Also, the simplified version where we set $\beta_{i,j,\ell} = 1$ for all i, j and ℓ carries over while suffering a slightly worse improvement than $\kappa = 0.0115$. In the following paragraph, we restrict to a sketch on how to extend the analysis for Algorithm 4 to the more general case beyond Bernoulli distributions.

A Sketch on a Generalized Analysis. In order to prove the generalization of Theorem 4.3.1, the major work is to change the syntax of the lemmas. We do not give details for all lemmas but rather provide the key steps on what to change and how to overcome obstacles.

First, we extend and change several definitions such as $y_{i,j} := \sum_{i'<i}\sum_{\ell} x_{i',j,\ell}$ or $\mathcal{A}_1^\ell$, $\mathcal{A}_2^\ell$ as the set of assignments (i,j) if the realized index is ℓ via a first or second proposal. The lemmas, observations and statements which referred to "i arriving" are now with respect to the event "i realizes index ℓ". For example, when talking about assigning j to

i via a first proposal, we replace this by saying that we assign j to i via a first proposal when i realized the valuation function with index ℓ.

The proofs for the analysis of early pairs directly carry over after adapting the syntax. For late pairs, the generalization of the proof of Lemma 4.3.5 (i) is also straightforward, as is the combination of both cases at the end.

We need to take some care in generalizing the proof of Lemma 4.3.5 (ii). The majority of the steps can be extended directly via a syntactic generalization from Section 4.3. In contrast, the proof of the generalized version of the correlation bound from Section 4.3.3, and in particular Claim 4.3.12 need some updates. Note however that as Claim 4.3.12 only concerns early pairs, it is not affected by the choice of $\beta_{i,j,\ell}$.

To see why Claim 4.3.12 also holds in the more general variant, we go through its proof steps one-by-one. Concerning the generalization of Step (S1) we note that the probability of both items being free after time $i+1$ can still be decomposed as the product of the probability of both being free before buyer $i+1$ times the conditional probability of assigning neither via a first proposal (as in Equation (4.11)). Still, we are required to sum the latter conditional probabilities for all possible realizations of ℓ. Doing so, we first follow Steps (S1) and (S2) from the Bernoulli case. During Step (S3), we need to show that for two distinct items j, j' and buyer i, the following inequality holds:

$$\begin{aligned} \alpha_{i,j}\alpha_{i,j'}\mathbf{Pr}\left[F_{i,j}\right]\mathbf{Pr}\left[F_{i,j'}\right]\left(\sum_{\ell}\frac{x_{i,j,\ell}x_{i,j',\ell}}{q_{i,\ell}} - \left(\sum_{\ell}x_{i,j,\ell}\right)\left(\sum_{\ell}x_{i,j',\ell}\right)\right) \\ \leq \mathbf{Pr}\left[F_{i,j}\right]\mathbf{Pr}\left[F_{i,j'}\right](0.5+\kappa)\left(1-\alpha_{i,j'}\sum_{\ell}x_{i,j',\ell}\right)\alpha_{i,j}\left(\sum_{\ell}x_{i,j,\ell}\right). \end{aligned} \tag{4.22}$$

In order to argue that this inequality is indeed true, we depart from the proof of the Bernoulli case by controlling the term $\sum_{\ell}\frac{x_{i,j,\ell}x_{i,j',\ell}}{q_{i,\ell}}$ via the online constraint for item j'. By Constraint (4.21), we know that

$$\frac{x_{i,j',\ell}}{q_{i,\ell}} \leq 1 - y_{i,j'}.$$

Using this, we can bound

$$\sum_{\ell}\frac{x_{i,j,\ell}x_{i,j',\ell}}{q_{i,\ell}} \leq (1-y_{i,j'})\sum_{\ell}x_{i,j,\ell}.$$

Plugging this into the left-hand side of Equation (4.22) and rearranging terms, we can conclude in a similar way as we did using Fact 4.3.13 in the Bernoulli case. Afterwards, Step (S4) of the correlation bound can again proceed via a syntactic generalization.

Chapter 5

Truthful Mechanisms for Two-Sided Markets via Prophet Inequalities

The algorithms in Chapter 3 and Chapter 4 can be interpreted (or imply) pricing-based mechanisms in one-sided markets: All items are initially held by the auctioneer who does not have any value for any of them and strategic buyers are willing to purchase bundles. For these types of markets, various different other auction formats and allocation procedures have been developed, such as VCG [Groves, 1973, Vickrey, 1961, Clarke, 1971], posted-prices mechanisms (as e.g. in Chawla et al. [2010], Dütting et al. [2020]) and many more (see e.g. Dobzinski et al. [2012], Dughmi et al. [2011], Assadi and Singla [2019], Assadi et al. [2021]).

Still, there are several applications, in which the assumption fails that all items are held by the auctioneer. Examples are widely spread, as to mention stock exchanges, ad auctions, online marketplaces such as eBay or ride sharing platforms.

In this chapter, we ask how the story from Chapter 3 and Chapter 4 changes once strategic sellers hold items initially. In other words, each seller brings a set of items to the market and has a valuation over her bundle. In addition, every seller acts strategically with the goal of maximizing her own utility. So, for example, a seller might keep an item for herself if the offered price is too low.

To deal with this challenge, our goal is to construct a mechanism which specifies trades between buyers and sellers and determines suitable prices for each trade with the objective of maximizing the overall social welfare.

Standard requirements for mechanisms are *individual rationality* (IR) and *dominant strategy incentive compatibility* (DSIC). Furthermore, as one cannot assume that there is a superior authority funding beneficial trades in two-sided markets, an additional natural requirement is *budget balance*. Its stronger version, *strong budget balance* (SBB), means that the mechanism can neither subsidize trades nor is allowed to extract money from trades. In other words, this requires that all money which is spent by buyers is transferred to sellers. The weaker form, *weak budget balance* (WBB), only requires the first property, namely that subsidizing trades is prohibited, but the mechanism is allowed to extract money from trades.

Unfortunately, in their seminal work in the 80s, Myerson and Satterthwaite [1983] showed that no mechanism can simultaneously be individually rational, incentive compatible, budget balanced and optimize social welfare[1]. This result is a sharp contrast to one-sided markets where optimal results are possible [Vickrey, 1961, Myerson, 1981]. As a consequence, *approximating* the optimal social welfare becomes the typical workaround. Going one step further, trying to approximate the optimal social welfare with a rather simple mechanism which can be easily understood by all participants may be an even more desirable goal.

Double Auction Formats. Probably the most fundamental problem in this field is *bilateral trade* (see e.g. Myerson and Satterthwaite [1983], Blumrosen and Dobzinski [2021], Kang and Vondrák [2018]), studied in Section 5.2: There is one seller holding one indivisible item and one buyer. In more general *double auctions*, there might be multiple buyers, multiple sellers, multiple items, and complex combinatorial constraints. In *matroid double auctions* (see e.g. Dütting et al. [2014], Colini-Baldeschi et al. [2016]) in Section 5.3 and Section 5.4, each seller initially holds one of m identical items, each buyer wants to purchase at most one of them and the set of buyers who receive an item needs to be an independent set in a matroid. In *combinatorial double auctions* (see e.g. Blumrosen and Dobzinski [2014], Colini-Baldeschi et al. [2020]) in Section 5.5, there are k sellers holding m heterogeneous items and the agents have combinatorial valuation functions over item bundles. For *knapsack double auctions* in Section 5.6 and Section 5.7, the setting is very similar as in matroid double auctions. Still, the matroid constraint over the set of buyers is replaced by a knapsack constraint (see e.g. Dütting et al. [2014]). That is, each buyer has a weight and we need to select buyers in a way such that the sum of weights does not exceed a certain capacity.

Mechanisms based on Balanced Prices. In order to derive the mechanisms for two-sided markets in this chapter, we extend the idea of balanced prices [Kleinberg and Weinberg, 2019, Feldman et al., 2015, Dütting et al., 2020] from Prophet Inequalities to two-sided environments. We use prices that are low enough and high enough at the same time: On the one hand, prices should be low enough so that agents may have values exceeding the prices and hence, either keep items (as sellers) or purchase items (as buyers). On the other hand, prices should be high enough so that we can cover the loss in social welfare once an item is allocated.

In the example of bilateral trade in Section 5.2, this extension is fairly easy. We can use Algorithm 1 from Chapter 1 which is based on a static and anonymous price for the item and apply an "interpret-seller-as-buyer"-argument: Interpret the seller as a buyer which is considered first and ask her if she want to keep or sell the item for a given price. If she agrees to try selling, ask the buyer if she wants to purchase for the same price. This directly implies a $1/2$-approximation which is strongly budget balanced by design.

Extending this idea to more complex settings can turn out to be relatively easy if the corresponding Prophet Inequality uses static and anonymous item prices (as e.g. using the Prophet Inequality of Feldman et al. [2015] for combinatorial double auctions in Section 5.5). In particular, the guarantees of strong budget balance and a reasonable approximation for the two-sided market problem directly follow again by a generalized version of the "interpret-seller-as-buyer"-argument described above. The reason is that

[1] The original result from Myerson and Satterthwaite is for bilateral trade instances, i.e. one seller holding one item and one buyer. They show that even individual rationality and (Bayesian) incentive compatibility cannot be combined with achieving the optimal ex-post social welfare.

due to prices being static and anonymous, we get strong budget balance in the two-sided environment for free.

Still, in settings where the Prophet Inequality is required to have dynamic or agent-specific prices (as e.g. the Prophet Inequality of Kleinberg and Weinberg [2019] for matroids or Dütting et al. [2020] for knapsacks), this straightforward extension does not work anymore. The reason is that prices heavily depend on the identity of an agent and her role in the corresponding feasibility constraint. To address the key difficulty of fulfilling incentive as well as budget balance constraints simultaneously, we will have to carefully select when to offer a trade to which pair of agents.

The proofs concerning the approximation guarantees mimic the spirit of revenue and utility based-ones in one-sided markets: we split the contribution to welfare of each agent into the *base value*, defined by the price of the proposed trade, and *surplus*, which is the amount by how much the agent's value exceeds the price. Afterwards, we bound each quantity separately. As a matter of fact, it does not play a key role which agent purchases or keeps which item — since any irrevocably allocated item ensures a sufficient contribution to welfare via its price.

Additional Related Work on Two-Sided Markets

Two-sided markets have been studied for a long time, including the mentioned impossibility result by Myerson and Satterthwaite [1983] and pioneering work on trade-reduction mechanisms and their generalizations as e.g. considered in McAfee [1992], Dütting et al. [2014], Babaioff and Walsh [2003], Babaioff and Nisan [2004]. Only much more recently, worst-case approximation ratios have been considered. There has been a lot of progress on improving the guarantees for bilateral trade ([Blumrosen and Dobzinski, 2014, 2021, Kang and Vondrák, 2018, Gerstgrasser et al., 2019] among others). However, determining the optimal guarantee is still an open problem.

Most relevant for the content of this chapter is the work of Colini-Baldeschi et al. [2016] and Colini-Baldeschi et al. [2020], which derive mechanisms for matroid and combinatorial double auctions in Bayesian settings. The focus of Colini-Baldeschi et al. [2016] is on matroid double auctions. Here, mechanisms are designed with pricing strategies based on quantiles, whereas our approach uses balanced prices. In Colini-Baldeschi et al. [2020], the authors consider combinatorial double auctions using very similar prices as the ones we use in this chapter. However, the analysis is different as their proofs rely on case distinctions while our proofs use charging arguments from balanced prices. Another important contribution of Colini-Baldeschi et al. [2020] is the introduction and discussion of direct-trade budget balance, which we also adopt in this paper. Dütting et al. [2021a] consider the same constraints. Besides giving improved approximation guarantees, they change the fundamental assumption of the Bayesian setting: They design mechanisms given only sample-based access to the underlying distribution.

Our mechanisms can also be viewed trough the lens of simplicity vs. optimality. There is related work in this direction e.g. by Deng et al. [2022a] or Niazadeh et al. [2014].

We focus on maximizing the overall social welfare of the mechanisms in this chapter. Complementing this direction, there is also a line of work using different objective functions in two-sided markets, most prominently *gain from trade* [Blumrosen and Mizrahi, 2016, Brustle et al., 2017, Colini-Baldeschi et al., 2017, Babaioff et al., 2018a, Segal-Halevi et al., 2016, Feldman and Gonen, 2018, Cai et al., 2021]. In this setting, only the *increase* in welfare by transferring items from sellers to buyers is measured. An α-approximation with respect to gain from trade is also an α-approximation with respect to social welfare but not vice versa. Indeed, Blumrosen and Dobzinski [2021] and Blumrosen and Mizrahi [2016] show that approximating the gain from trade is harder than social welfare: There is no DISC, IR and SBB mechanism which can achieve a constant factor approximation to the optimal gain from trade. Only recently, Deng et al. [2022a] could provide a breakthrough and prove the first constant factor approximation to the optimal (first best) gain-from-trade via a BIC mechanism. Babaioff et al. [2020] tackle the question by how many buyers and sellers the size of the two-sided market needs to be increased in order to recover the optimal gain from trade from the original market, mirroring the seminal work of Bulow and Klemperer [1996]. Another interesting objective function is the profit of the sellers in two-sided markets as considered by Cai and Zhao [2019].

Chapter Organization and Remarks

This chapter is based on *Truthful Mechanisms for Two-Sided Markets via Prophet Inequalities* [Braun and Kesselheim, 2021] (conference version) and [Braun and Kesselheim, 2023b] (journal version), which is joint work with Thomas Kesselheim. Further bibliographic notes can be found in Section 1.5.

In this chapter, we start by giving some specific preliminaries in Section 5.1. In Section 5.3, we state our mechanism for matroid double auctions which satisfies strong budget balance. In Section 5.4, we give the corresponding weakly budget balanced mechanism with an improved approximation guarantee. Section 5.5 addresses combinatorial double auctions. In Section 5.6, we consider knapsack double auctions with strong budget balance and give a proof for the approximation guarantee. In Section 5.7, we complement this by allowing weak budget balance in knapsack double auctions.

5.1 Formal Problem Statement and Preliminaries

We consider the following setup for two-sided markets: There is a set of n buyers[2] $\mathcal{B}$, a set of k sellers $\mathcal{S}$ and a set of m items M. We assume that $\mathcal{B} \cap \mathcal{S} = \emptyset$, so any agent can either act as a buyer or a seller. Before running any (reallocation) mechanism, the set of items is initially held by the sellers. We denote by I_l the set of items which is hold by seller l initially and call the vector $(I_1, \dots, I_k)$ the *initial allocation*. Note that the sets I_l are pairwise disjoint, i.e. for any two sellers $l \neq l'$ we have $I_l \cap I_{l'} = \emptyset$, and further all items are allocated to some seller before running our mechanism, i.e. $\bigcup_{l \in \mathcal{S}} I_l = M$.

Any agent $i \in \mathcal{B} \cup \mathcal{S}$ has a privately known valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$. Any seller l is assumed to have only positive value for items in her initial bundle I_l, i.e. for any seller $l \in \mathcal{S}$ and $T \subseteq M$ it holds that $v_l(T) = v_l(T \cap I_l)$.

We consider a Bayesian setting where each agent i's valuation function is drawn independently from a publicly known, not necessarily identical probability distribution $\mathcal{D}_i$. That is, $\mathcal{D}_i$ is a probability distribution over the *space of valuation functions* V_i. We denote by $\mathcal{D} = \times_{i \in \mathcal{B} \cup \mathcal{S}} \mathcal{D}_i$ the joint probability distribution of the space of all agents' valuation functions $\mathbf{V} = \times_{i \in \mathcal{B} \cup \mathcal{S}} V_i$ and we refer to $\mathbf{v}$ as a *valuation profile* which consists of one valuation function per agent.

[2] As the set of buyers is *not required* to arrive sequentially in two-sided markets, we differ from the notation in Chapter 3 and Chapter 4 for the set of buyers in order to highlight this difference.

An *allocation* $\mathbf{X} = (X_i)_{i \in \mathcal{B} \cup \mathcal{S}}$ is a vector of item bundles such that agent i is allocated bundle X_i and for two agents $i \neq i'$, we have $X_i \cap X_{i'} = \emptyset$. Our goal is to maximize the *social welfare* of an allocation $\mathbf{X}$ given valuation profile $\mathbf{v}$, defined as $\mathbf{v}(\mathbf{X}) := \sum_{i \in \mathcal{B} \cup \mathcal{S}} v_i(X_i)$. Concerning *feasibility*, as said, any seller $l \in \mathcal{S}$ can only receive items in her initial allocation, i.e. $X_l \subseteq I_l$ for any $l \in \mathcal{S}$. For buyers, we address feasibility constraints in the respective sections below.

Mechanisms in Two-Sided Markets and their properties

A (direct revelation) *mechanism* takes as input a vector of valuation functions which are reported by agents. Agents can report any possible valuation in their space of valuation functions V_i, not necessarily their true one. A mechanism outputs an allocation of items to agents $\mathbf{X}$ as well as payments $\mathbf{P}$. Buyers pay money to the mechanism whereas sellers receive money.

Agents are assumed to maximize *utility*. Fixing a valuation profile $\mathbf{v}$, an allocation $\mathbf{X}$ and payments $\mathbf{P}$, the (quasi-linear) utility of buyer i for being allocated bundle $X_i \subseteq M$ is given by $u_i(X_i) = v_i(X_i) - P_i$ whereas the utility for seller l who remains with bundle $X_l \subseteq I_l$ is given by $u_l(X_l) = v_l(X_l) + P_l$.

Mechanisms in this context are designed to fulfill the following desirable properties:

- *Dominant Strategy Incentive Compatibility* (DSIC): It is a dominant strategy for every agent to report her true valuation independent of the other agents' behavior.
- *Individual Rationality* (IR): When playing this dominant strategy, no agent decreases her utility by participating in the mechanism. So, for buyers $v_i(X_i) - P_i \geq 0$ and for sellers $v_l(X_l) + P_l \geq v_l(I_l)$.
- *Weak/Strong Budget Balance* (WBB/SBB): The money received by sellers is at most/equals the payments made by buyers, i.e. $\sum_{i \in \mathcal{B}} P_i \overset{(=)}{\geq} \sum_{l \in \mathcal{S}} P_l$.

Concerning budget balance, Colini-Baldeschi et al. [2020] showed a weakness in the original definition of strong budget balance: cross subsidizing trades with already received money is not prohibited as long as the sum of payments is equal for buyers and sellers[3]. The stronger notion of *direct-trade weak/strong budget balance* (DWBB/DSBB) requires that the outcome of the mechanism can be obtained by a composition of bilateral trades. In each trade, an item is reallocated from seller l to buyer i, payments are transferred from some buyer i to some seller l and each item may only be traded at most once [Colini-Baldeschi et al., 2020]. If the buyer's payment exceeds the seller's receiving for at least one of the trades, the mechanism is DWBB; if payments in each of these bilateral trades are equal for buyers and sellers, we refer to DSBB. Our mechanisms will be DSIC, IR and satisfy the respective DSBB/DWBB property.

As a benchmark, we compare the expected social welfare achieved by our mechanism to the expected optimal (first best) social welfare $\mathbf{E}_{\mathbf{v}}\left[\max_{\mathbf{X}^*} \mathbf{v}(\mathbf{X}^*)\right]$. We aim to design ζ-approximative mechanisms. In the case that our mechanism can deal with online arrivals of agents, we may use the wording ζ-approximation and ζ-competitive interchangeably in this chapter.

[3] Colini-Baldeschi et al. [2020] argue that turning a WBB into an SBB mechanism is rather easy with a small loss in the approximation guarantee as one can simply draw one seller uniformly at random and give all the surplus money in the WBB mechanism to this seller.

5.2 Warm-Up: Bilateral Trade via Balanced Prices

In order to gain a better understanding on how to apply Prophet Inequality techniques in two-sided markets, we start by considering the problem of bilateral trade. In bilateral trades, there is one seller, initially equipped with one item and one buyer.

For this setting, there is the following – almost trivial – mechanism: Let v_s denote the seller's value and v_b denote the buyer's value for the item. Fix a price p and trade the item if and only if $v_b \geq p \geq v_s$. Among others, Blumrosen and Dobzinski [2014, 2021] and Gerstgrasser et al. [2019] set p to be the median of the seller's distribution which recovers an expected welfare of at least $1/2 \cdot \mathbf{E}\left[\max\{v_s, v_b\}\right]$, so it is a $1/2$-approximation.

Our mechanism uses a different price. We set $p = 1/2 \cdot \mathbf{E}\left[\max\{v_s, v_b\}\right]$ and trade the item if and only if $v_b \geq p \geq v_s$. As an important remark, note that we can interpret this mechanism as a sequential posted-prices mechanism with price p: First, ask the seller s if she would like to keep or try selling the item for price p. Afterwards, buyer b may buy the item for price p if the seller herself wanted to sell the item. This mechanism is DSIC, IR and SBB by design. Concerning the approximation guarantee, we can state the following proposition.

Proposition 5.2.1. *The bilateral trade mechanism with* $p = \frac{1}{2} \cdot \mathbf{E}\left[\max\{v_s, v_b\}\right]$ *is a* $\frac{1}{2}$*-approximation to the optimal social welfare.*

We give a proof applying the ideas from Prophet Inequalities.

Proof. We distinguish several cases: the mechanism extracts welfare if the item is either allocated as $v_s \geq p$ (the seller initially keeps the item), $v_b \geq p > v_s$ (a trade occurs), or $v_b < p$ and $v_s < p$ (both agents' values do not exceed the price, so the item remains at the seller). Observe that the social welfare achieved by the first two cases is clearly a lower bound on the overall social welfare achieved by the mechanism. In other words, we only consider contributions to social welfare if at least one of the agents exceeds price p.

We begin by splitting the social welfare achieved by the mechanism in base value and surplus.

For the base value, observe that if there exists $i \in \{b, s\}$ with $v_i \geq p$, we get a contribution to social welfare of p (actually, we get $p + (v_i - p)$, but the second summand is considered in the surplus). Hence,

$$\mathbf{E}\left[\text{Base Value}(\mathbf{v})\right] = \mathbf{Pr}\left[\text{There exists } i \in \{b, s\} \text{ with } v_i \geq p\right] \cdot p \ .$$

For the surplus, we first argue about the contribution of the seller, afterwards about the buyer. Note that the seller may keep the item initially if $v_s \geq p$. As a consequence, we can extract a surplus of $v_s - p$ if this quantity is non-negative. In other words, we get $(v_s - p)^+$ as a surplus from seller s. Buyer b can buy the item if the seller initially agreed to try selling, i.e. $v_s < p$ and if her value $v_b \geq p$ exceeds the price. Hence, we get $(v_b - p)^+ \mathbb{1}_{v_s < p}$ as a surplus. Observe that v_s and v_b are independent and hence, taking the expectation over $\mathbf{v}$, we can bound the expected surplus of buyer b via

$$\mathbf{E}\left[\text{surplus}_b(\mathbf{v})\right] \geq \mathbf{E}\left[(v_b - p)^+ \mathbb{1}_{v_s < p}\right] = \mathbf{E}\left[(v_b - p)^+\right] \cdot \mathbf{Pr}\left[v_s < p\right] \ .$$

Now, observe that $1 \geq \mathbf{Pr}[v_s < p] \geq \mathbf{Pr}[v_s < p, v_b < p]$, which allows to bound the sum of the seller's and buyer's surplus as

$$\begin{aligned}\mathbf{E}[\text{Surplus}(\mathbf{v})] &= \mathbf{E}[\text{surplus}_s(\mathbf{v})] + \mathbf{E}[\text{surplus}_b(\mathbf{v})] \\ &\geq \left(\mathbf{E}\left[(v_s - p)^+\right] + \mathbf{E}\left[(v_b - p)^+\right]\right) \cdot \mathbf{Pr}[v_s < p, v_b < p] \quad .\end{aligned}$$

Next, we use that $(v_s - p)^+ + (v_b - p)^+ \geq \max_{i \in \{s,b\}} (v_i - p)^+ \geq \max\{v_s, v_b\} - p$. Further, by our choice of p, we get that

$$\mathbf{E}[\text{Surplus}(\mathbf{v})] \geq \mathbf{E}[\max\{v_s, v_b\} - p] \cdot \mathbf{Pr}[v_s < p, v_b < p] = p \cdot \mathbf{Pr}[v_s < p, v_b < p] \quad .$$

Combining base value and surplus, we get

$$\begin{aligned}&\mathbf{E}[\text{Base Value}(\mathbf{v})] + \mathbf{E}[\text{Surplus}(\mathbf{v})] \\ &\geq p \cdot (\mathbf{Pr}[\text{There exists } i \in \{b, s\} \text{ with } v_i \geq p] + \mathbf{Pr}[v_s < p, v_b < p]) \\ &= p \cdot 1 \\ &= \frac{1}{2} \cdot \mathbf{E}[\max\{v_s, v_b\}] \quad .\end{aligned}$$

□

As a remark, recall that we did not consider any contribution to social welfare from the case that neither agent has a value which exceeds the price, i.e. $v_s < p$ and $v_b < p$. Still, in these cases, seller s keeps the item, so we could extract the welfare contribution in addition to the ones considered above in order to improve the approximation guarantee for bilateral trade instances. Nonetheless, when considering generalizations to matroid, knapsack and combinatorial double auctions, restricting to social welfare contributions of agents who exceed prices will turn out to simplify the arguments.

5.3 Matroid Double Auctions with Strong Budget Balance

Our first mechanism is for double auctions where the set of n buyers $\mathcal{B}$ is equipped with a matroid constraint. That is, there is a matroid $\mathcal{M}_{\mathcal{B}} = (\mathcal{B}, \mathcal{I}_{\mathcal{B}})$ and the set of buyers who receive an item in the mechanism needs to be an independent set in the matroid $\mathcal{M}_{\mathcal{B}}$. For this section, we assume buyers to be unit-demand and sellers to be unit-supply, i.e. every seller initially holds a single, indivisible item and hence $k = m$. Items are identical, meaning that $v_i(T)$ is zero if $T = \emptyset$ and (as buyers are unit-demand) equal to some fixed value for any $T \neq \emptyset$. Our mechanism requires an offline setting in which buyers and sellers can trade in any order. We determine the order adaptively during the mechanism. In particular, we assume that we can pick one buyer and one seller in any step out of the remaining ones and offer a trade at some price to both agents. Further, we simplify notation in this section. In the setting of this section, a valuation profile $\mathbf{v}$ can be interpreted as a $|\mathcal{B} \cup \mathcal{S}|$-dimensional vector over the non-negative real numbers in which each entry v_i corresponds to the value of an agent for being allocated an item. Further, we denote any unit-supply seller and the corresponding item by j.

In the following, we start by a description of the mechanism in Section 5.3.1, the pricing is described afterwards in Section 5.3.2. In particular, the definition of the prices will be based on the order in which we offer trades in the mechanism. In Section 5.3.3, we discuss properties of our mechanism such as DSBB, DSIC and IR followed by proofs of the approximation guarantee. We give a proof of the approximation guarantee in the full information case first and only consider the general incomplete information case afterwards.

5.3.1 The Mechanism

We formally state our mechanism in Algorithm 7 and start with an informal description. Throughout the algorithm, we maintain a set of agents $A = A_{\mathcal{B}} \cup A_{\mathcal{S}}$ who are irrevocably allocated an item with $A_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}$. In addition, in the set M_{SELL} we store all sellers (or equivalently items) who may still be considered for a possible trade. In particular, this set contains all sellers which have neither rejected a price yet nor participated in a trade. Analogously, the set M_{BUY} denotes the set of buyers who have not been considered for a trade yet and can feasibly be added to the current set of accepted buyers.

We maintain buyer-specific thresholds p_i and seller-specific ones p_j. The exact description of these thresholds will be discussed afterwards. The price for a trade between seller j and buyer i will then be defined as $p_{i,j} = \text{constant} \cdot (p_i + p_j)$. In every iteration, among all available sellers $j \in M_{\text{SELL}}$, we consider the one with the smallest threshold[4] p_j. We try to match her to buyers $i \in M_{\text{BUY}}$ in decreasing order of thresholds p_i. To this end, we first ask seller j if she wants to sell or keep her item for a price of $p_{i,j}$. If she wants to keep her item, we remove seller j from the set of available sellers. Otherwise, if seller j considers selling her item, we ask buyer i if she wants to buy the item for price $p_{i,j}$. If buyer i agrees, the item is transferred from j to i, both are removed from the set of available agents, i is irrevocably allocated an item, j is irrevocably discarded for holding an item and i pays $p_{i,j}$ to seller j. Else, buyer i is removed from the set of available buyers and irrevocably discarded while seller j remains available. As a consequence, as soon as an agent turned down an offer, she is not considered in our mechanism anymore. Then we move to the next iteration, in which we consider a different pair for trading.

5.3.2 The Pricing

We start by giving an intuitive approach to prices and describe them formally afterwards. As mentioned already, the price for a trade between buyer i and seller j will be a scaled sum of thresholds p_i and p_j.

The intuition behind the thresholds p_i and p_j is as follows: We consider a relaxation of the expected optimal social welfare to $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\text{OPT}_{\mathcal{B}}(\mathbf{v})\right)\right] + \mathbf{E}_{\mathbf{v}}\left[\sum_{j \in \mathcal{S}} v_j\right]$, where $\text{OPT}_{\mathcal{B}}(\mathbf{v})$ is the optimal choice when restricting to the set of buyers. That is, in our relaxation, each item can be counted twice: It will contribute to the first term by being assigned to a buyer while in the second term it is assumed that the seller keeps it. The buyer-specific thresholds p_i will compensate for the loss in the first term once we allocate an item. The thresholds p_j will do so in the second term of the relaxation of the optimal social welfare. In particular, once a trade occurs, the price for this trade covers the loss of both, the seller and the buyer, in the relaxed optimal social welfare. In addition, the remaining share of the social welfare will be covered by the surplus. By the choice of the order in which trades are offered, we ensure that prices are monotone for a fixed seller which is crucial concerning truthfulness and budget balance.

[4] Break ties arbitrarily, but always in the same way.

Algorithm 7: Mechanism for Matroid Double Auctions with Strong Budget Balance

---

Result: Set A of agents to get an item with $A \cap \mathcal{B} \in \mathcal{I}_{\mathcal{B}}$ and $|A| = |\mathcal{S}|$

$A_{\mathcal{B}} \longleftarrow \emptyset$; $A_{\mathcal{S}} \longleftarrow \emptyset$; $r \longleftarrow |\mathcal{S}|$; $M_{\text{SELL}} \longleftarrow \mathcal{S}$; $M_{\text{BUY}} \longleftarrow \mathcal{B}$

while $M_{\text{BUY}} \neq \emptyset$ and $M_{\text{SELL}} \neq \emptyset$ **do**

  Recompute the thresholds $p_i(A_{\mathcal{B}}, r)$ and $p_j(A_{\mathcal{B}}, r)$ with respect to current $A_{\mathcal{B}}$, r, M_{SELL} and M_{BUY}

  $j \in \arg\min_{j' \in M_{\text{SELL}}} p_{j'}(A_{\mathcal{B}}, r)$; $i \in \arg\max_{i' \in M_{\text{BUY}}} p_{i'}(A_{\mathcal{B}}, r)$

  if $A_{\mathcal{B}} \cup \{i\} \notin \mathcal{I}_{\mathcal{B}}$ or $|A_{\mathcal{B}} \cup \{i\}| > r$ **then**

    $M_{\text{BUY}} \longleftarrow M_{\text{BUY}} \setminus \{i\}$

    go to next iteration

  $p \longleftarrow p_{i,j}\left(A_{\mathcal{B}}, r\right)$

  if $v_j > p$ **then**

    $A_{\mathcal{S}} \longleftarrow A_{\mathcal{S}} \cup \{j\}$; $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$; $r \longleftarrow r - 1$

  if $v_j \leq p$ **then**

    $M_{\text{BUY}} \longleftarrow M_{\text{BUY}} \setminus \{i\}$

    if $v_i > p$ **then**

      $A_{\mathcal{B}} \longleftarrow A_{\mathcal{B}} \cup \{i\}$; $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$

return $A := A_{\mathcal{B}} \cup A_{\mathcal{S}} \cup M_{\text{SELL}}$

In order to address feasibility issues first, by construction, our mechanism never offers trades to buyers who cannot be feasibly added to $A_{\mathcal{B}}$. Hence, the mechanism ensures that the set of buyers $A_{\mathcal{B}}$ who receive an item in our mechanism is an independent set in the matroid, i.e. $A_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}$. Additionally, we do not promise items to agents once all items are irrevocably allocated. The price for any feasible trade is calculated in an agent-specific way extending the method of balanced thresholds by Kleinberg and Weinberg [2019] and balanced prices by Dütting et al. [2020] to two-sided markets.

Recall that $A_{\mathcal{B}}$ contains all buyers who receive an item and $A_{\mathcal{S}}$ contains all sellers who irrevocably keep their item. By r we denote the number of items which may be allocated to buyers in total, i.e. which are not irrevocably kept by a seller, so $r = |\mathcal{S}| - |A_{\mathcal{S}}|$. Observe that r is decreasing in our mechanism every time a seller decides to irrevocably keep her item. Given the matroid over the set of buyers, we need to ensure that we do not pick more than r buyers in our mechanism.

Fixing a valuation profile $\mathbf{v}$, we let

$$\text{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r) \in \underset{\mathcal{B}' \subseteq \mathcal{B} \setminus A_{\mathcal{B}}, \mathcal{B}' \cup A_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}, |\mathcal{B}' \cup A_{\mathcal{B}}| \leq r}{\arg\max} \left(\sum_{i \in \mathcal{B}'} v_i \right) .$$

That is, $\text{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$ denotes the following allocation. Assume that we are only allowed to assign items to buyers (not to sellers) and we have already allocated items to buyers in $A_{\mathcal{B}}$ and at most r items can be allocated to buyers in total. Then $\text{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$ is the allocation that maximizes the welfare increase. The value of this partial allocation is denoted by $\mathbf{v}\left(\text{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r)\right)$. Further, we define $\text{OPT}_{\mathcal{B}}(\mathbf{v}) = \text{OPT}_{\mathcal{B}}(\mathbf{v} \mid \emptyset, |\mathcal{S}|)$ to be the optimal allocation of *all* items to buyers.

The threshold of buyer i is defined with respect to the current state of $A_{\mathcal{B}}$ and the number of items r. For a fixed valuation profile $\mathbf{v}$, let

$$p_i(A_{\mathcal{B}}, r, \mathbf{v}) = \mathbf{v}\left(\text{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)\right) - \mathbf{v}\left(\text{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}} \cup \{i\}, r)\right)$$

if $A_{\mathcal{B}} \cup \{i\} \in \mathcal{I}_{\mathcal{B}}$ and $|A_{\mathcal{B}} \cup \{i\}| \leq r$. So, $p_i(A_{\mathcal{B}}, r, \mathbf{v})$ is the difference in welfare which we can achieve by allocating r items to buyers given we have already allocated items to buyers in $A_{\mathcal{B}}$ and $A_{\mathcal{B}} \cup \{i\}$ respectively. To simplify notation, we define $p_i(A_{\mathcal{B}}, r, \mathbf{v}) = \infty$ if $A_{\mathcal{B}} \cup \{i\} \notin \mathcal{I}_{\mathcal{B}}$ or $|A_{\mathcal{B}} \cup \{i\}| > r$.

Based on this, define buyer i's threshold as $p_i(A_{\mathcal{B}}, r) = \mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[p_i(A_{\mathcal{B}}, r, \tilde{\mathbf{v}})\right]$. For a seller j, we set the seller-specific threshold to $p_j = \mathbf{E}_{\tilde{v}_j \sim \mathcal{D}_j}\left[\tilde{v}_j\right]$, which is simply the expected value of the distribution of seller j's value for an item. Now, fix a buyer-seller-pair (i, j) which is available for trading and denote the price for a trade between i and j by

$$p_{i,j}(A_{\mathcal{B}}, r) := \frac{1}{3}\left(p_i(A_{\mathcal{B}}, r) + p_j\right) := \frac{1}{3}\left(\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[p_i(A_{\mathcal{B}}, r, \tilde{\mathbf{v}})\right] + \mathbf{E}_{\tilde{v}_j \sim \mathcal{D}_j}\left[\tilde{v}_j\right]\right) \ .$$

5.3.3 Properties of Our Mechanism

Note that our mechanism ensures that the final allocation is a feasible solution with respect to the matroid constraint on the buyers' side as we discard any buyer who cannot be added feasibly to the set of allocated agents. Further, we do not allocate more than $|\mathcal{S}|$ items in total among all agents in our allocation process.

Theorem 5.3.1. *The mechanism for matroid double auctions is DSBB, DSIC and IR for all buyers and sellers and a $\frac{1}{3}$-approximation to the optimal social welfare.*

By construction, Mechanism 7 consists of several bilateral trades, where an item is transfered from seller j to buyer i and a price of $p_{i,j}$ is paid by buyer i, received by seller j, so the mechanism satisfies DSBB.
We offer any buyer the possibility to participate in a trade at most once, so DSIC and IR for buyers follows directly. Also IR for sellers is rather simple as we ask seller j every time if she would like to participate in a trade for a given price. In order to show DSIC for sellers, we have to exploit the order in which trades are offered. By this choice, prices offered to a fixed seller are only non-increasing as the mechanism evolves. As a consequence, selling the item as early as possible is only beneficial for a seller (if she would like to sell the item at all). Truthfulness follows as misreporting the value for an item might allow or block unfavorable trades.

Proof of the Approximation Guarantee of Theorem 5.3.1 in the full information setting.

In order to illustrate the proof concerning the approximation guarantee, we give a simplified proof in the full information setting first. That is, the value v_i of an agent is not a random variable anymore, but rather deterministic. The general case with incomplete information can be found afterwards. In the full information setting, the price for a feasible trade between buyer i and seller j simplifies to

$$\frac{1}{3}\left(\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B},ij}, r_{ij})) - \mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B},ij} \cup \{i\}, r_{ij})) + v_j\right) \ ,$$

where $A_{\mathcal{B},ij}$ and r_{ij} are the states of $A_{\mathcal{B}}$ and r as we consider buyer i and seller j for a trade. First, note that any agent who keeps or purchases an item has a value exceeding some price. So for any agent $i \in A$, there is a price P_i which agent i's value did exceed when we added i to A. For sellers to which we did not offer any trade in our mechanism, we set P_i to zero as they keep their items anyway; for buyers who cannot be feasibly added to our set of chosen agents, we set P_i to infinity. We split the social welfare achieved by our mechanism in two parts, calling them *base value* and *surplus*:

$$\sum_{i \in A} v_i = \sum_{i \in A} P_i + \sum_{i \in A} (v_i - P_i)$$

Now, we bound each of these quantities separately.

Base Value. When irrevocably allocating an item during the offer of a trade to buyer-seller-pair (i, j), either the seller keeps the item or the buyer purchases it. In the first case, we reduce r by one, in the second, we add i to $A_\mathcal{B}$. In order to bound the loss incurred by a seller keeping her item, observe that

$$\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij} \cup \{i\}, r_{ij})) \leq \mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij}, r_{ij} - 1)) \enspace .$$

This is true as any feasible choice of buyers for the left-hand side is also feasible for the right-hand side. It further directly implies that

$$\begin{aligned} &\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij}, r_{ij})) - \mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij} \cup \{i\}, r_{ij})) \\ \geq\ &\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij}, r_{ij})) - \mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_{\mathcal{B},ij}, r_{ij} - 1)) \enspace . \end{aligned}$$

As prices in the next iteration are computed with respect to $r_{ij} - 1$ and $A_{\mathcal{B},ij}$, the loss in the buyers' optimal welfare when allocating an item to a seller is bounded by the buyer's contribution to the price. Summing the prices which we offered to agents in $A_\mathcal{B} \cup A_\mathcal{S}$ combined with this bound leads to a telescopic sum over the buyers' thresholds in the prices. Therefore, we can derive a bound of

$$\sum_{i \in A} P_i \geq \frac{1}{3}\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v})) - \frac{1}{3}\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_\mathcal{B}, r)) + \frac{1}{3} \sum_{j \in \mathcal{S} \setminus M_{\mathrm{SELL}}} v_j \tag{5.1}$$

for the base value.

Surplus. Concerning the surplus, we consider buyers and sellers separately. For the sellers, note that any seller who remains in M_{SELL} after the mechanism keeps her item. Therefore, the contribution to the surplus is v_j for any $j \in M_{\mathrm{SELL}}$. In the incomplete information setting, this turns out to be much more involved and a more sophisticated argument needs to be applied. As a consequence, we are only able to bound the sellers' surplus from below via

$$\sum_{i \in A_\mathcal{S} \cup M_{\mathrm{SELL}}} (v_i - P_i) \geq \frac{2}{3} \sum_{j \in M_{\mathrm{SELL}}} v_j - \frac{1}{3}\mathbf{v}(\mathrm{OPT}_\mathcal{B}(\mathbf{v}|A_\mathcal{B}, r)) \enspace . \tag{5.2}$$

For the buyers, we note that the prices for a fixed buyer are only non-decreasing as the allocation process evolves (see Lemma 5.3.5 which is a generalized version of Lemma 3 in Kleinberg and Weinberg [2012]). Further, any buyer to which we offer a trade gets an item if her value exceeds her price. Using this, we can bound the surplus of any buyer i to which we proposed a trade via

$$\begin{aligned} (v_i - P_i)^+ &= (v_i - p_{i,j_i}(A_{\mathcal{B},ij_i}, r_{ij_i}))^+ \geq (v_i - p_{i,j_i}(A_\mathcal{B}, r))^+ \\ &\geq \left(v_i - \min_{j \in M_{\mathrm{SELL}}} p_{i,j}(A_\mathcal{B}, r)\right)^+ \enspace , \end{aligned}$$

where we denote by j_i the seller which is matched to buyer i. Now, we consider all buyers which are in $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$. Any of these buyers could have purchased an item if her value had exceeded the price. To see this, note that $A_{\mathcal{B}} \cup \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$ needs to be independent. Further, if $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r) \neq \emptyset$, we have that $r > 0$ and so there are still items available after running the mechanism. As a consequence, any agent $i \in \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$ has a surplus of $(v_i - P_i)^+$ which is positive only if $i \in A_{\mathcal{B}}$ after running the mechanism[5]. For a buyer who does not exceed her price, this is zero as is her contribution to the surplus. Summing the surplus of all these buyers implies a lower bound on the overall buyers' surplus of

$$\begin{aligned}\sum_{i \in A_{\mathcal{B}}} (v_i - P_i) &\geq \sum_{i \in \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)} (v_i - P_i)^+ \\ &\geq \sum_{i \in \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)} v_i - \sum_{i \in \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)} \min_{j \in M_{\mathrm{SELL}}} p_{i,j}(A_{\mathcal{B}}, r) \ .\end{aligned}$$

Having a closer look at the sum of prices, we can apply a proposition from Kleinberg and Weinberg [2012, Proposition 2] on the buyers' contribution in order to derive a suitable bound:

$$\begin{aligned}\sum_{i \in \mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)} \min_{j \in M_{\mathrm{SELL}}} p_{i,j}(A_{\mathcal{B}}, r) &\leq \frac{1}{3}\left(\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) + |M_{\mathrm{SELL}}| \cdot \min_{j \in M_{\mathrm{SELL}}} v_j\right) \\ &\leq \frac{1}{3}\left(\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) + \sum_{j \in M_{\mathrm{SELL}}} v_j\right) \ .\end{aligned}$$

And so we get

$$\sum_{i \in A_{\mathcal{B}}} (v_i - P_i) \geq \frac{2}{3}\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) - \frac{1}{3}\sum_{j \in M_{\mathrm{SELL}}} v_j \ .$$

Hence, in combination with (5.2), we can lower-bound the overall surplus of all agents via

$$\begin{aligned}\sum_{i \in A} (v_i - P_i) &\geq \frac{2}{3}\sum_{j \in M_{\mathrm{SELL}}} v_j - \frac{1}{3}\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) + \frac{2}{3}\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) - \frac{1}{3}\sum_{j \in M_{\mathrm{SELL}}} v_j \\ &= \frac{1}{3}\sum_{j \in M_{\mathrm{SELL}}} v_j + \frac{1}{3}\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)) \ . \end{aligned} \tag{5.3}$$

Adding base value (5.1) and surplus of all buyers and sellers (5.3) proves the claim as $\mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})) + \sum_{j \in \mathcal{S}} v_j \geq \mathbf{v}(\mathrm{OPT}(\mathbf{v}))$.

Proof in the incomplete information setting.

Now, we give a complete proof of the approximation guarantee of Theorem 5.3.1.

Lemma 5.3.2. *Mechanism 7 for matroid double auctions is a $\frac{1}{3}$-approximation of the optimal social welfare.*

[5] This argument might sound strange in the full information case as buyers can either be included in $A_{\mathcal{B}}$ or $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}|A_{\mathcal{B}}, r)$. It will make much more sense in the incomplete information case later.

Proof. We start by a quick reformulation of the prices. Assume, we introduced a counter t starting at zero which increases by one in every iteration of the while-loop as soon as a buyer or a seller accepts a price. Every time the counter increases, one item is allocated irrevocably: Either the sellers decides to keep the item or a trade occurs and the item is allocated to the current buyer. Denote by $A_{\mathcal{B},t}$ the state of set $A_{\mathcal{B}}$ (similarly with $A_{\mathcal{S},t}$ for $A_{\mathcal{S}}$ etc.) as the counter shows t (i.e. t items are already allocated) and as before, if $A_{\mathcal{B},t} \cup \{i\} \in \mathcal{I}_B$ and $|A_{\mathcal{B},t} \cup \{i\}| \leq r_t$, let

$$p_{i,j}(A_{\mathcal{B},t}, r_t) = \frac{1}{3}\left(\mathbf{E}_{\tilde{\mathbf{v}}\sim\mathcal{D}}\left[p_i(A_{\mathcal{B},t}, r_t, \tilde{\mathbf{v}})\right] + \mathbf{E}_{\tilde{v}_j\sim\mathcal{D}_j}\left[\tilde{v}_j\right]\right)$$

be the price for buyer-seller-pair (i, j). Otherwise, as already mentioned, we will not consider buyer i and set any price $p_{i,j}$ for trades offered to buyer i to infinity. Note that this formulation is equivalent to our initial definition of the prices but rather allows to refer to the t-th irrevocably allocated item.

The set of agents who receive an item A depends on $\mathbf{v}$, so we denote by $A(\mathbf{v})$ the set A under valuation profile $\mathbf{v}$ (the same for $A_{\mathcal{B},t}(\mathbf{v})$ and $A_{\mathcal{S},t}(\mathbf{v})$ etc.). We want to compare $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(A(\mathbf{v}))\right]$ to $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right]$. To this end, we split the welfare of our algorithm into two parts, the base value and the surplus, and bound each quantity separately. (When thinking about one-sided markets, this corresponds to revenue and utility of buyers.) The base value is hereby defined as follows: let agent i receive an item in our mechanism, i.e. $i \in A$. Any buyer who gets an item has paid some price for the item. Any seller who decided to keep her item was asked to keep it for some specific price. The part of agent i's value which is below this price is denoted the base value. The surplus is the part of agent i's value above this threshold if it exists, otherwise it is zero. There might be sellers who are left unconsidered in our mechanism, i.e. we did never ask them if they would like to participate in a trade. These sellers keep their item without any contribution to the base value in our calculations. All their value is considered in the surplus.

Base Value. As said, all buyers and sellers who are irrevocably allocated an item (i.e. which are in A before adding the remaining sellers) have a value which exceeds some price. For any agent i, denote this price by P_i. Further, every time the counter t increases, we are allocating an item irrevocably in our mechanism.

As a first step, we need to argue about the two different scenarios which can occur in our mechanism as an item is allocated after offering a trade to buyer i and seller j with counter t. On the one hand, a trade may occur and buyer i is allocated seller j's item. In this scenario, the prices in the next iteration(s) with counter $t+1$ are computed with respect to $A_{\mathcal{B},t+1} = A_{\mathcal{B},t} \cup \{i\}$ and $r_{t+1} = r_t$. In addition, seller j is not available for a trade anymore. On the other hand, seller j may keep the item, so we compute prices at counter $t+1$ with respect to $A_{\mathcal{B},t+1} = A_{\mathcal{B},t}$ and $r_{t+1} = r_t - 1$. Note that our prices are adapted to mirror the first scenario. Taking the expectation over the inequality from Lemma 5.3.3 (which is formally stated only after the proof), we see that the impact of the second scenario (a seller keeping the item) can be bounded by the loss of the first one concerning the optimal social welfare.

Fixing a valuation profile $\mathbf{v}$ and summing over all agents in $A(\mathbf{v})$ in the order that they were added to A is equivalent to summing over all steps in which we increased the counter t. Denote by i_t and j_t the buyer and seller considered in this particular time step.

This allows to use a telescopic sum argument to obtain

$$\sum_{i \in A_{\mathcal{B}}(\mathbf{v}) \cup A_{\mathcal{S}}(\mathbf{v})} P_i = \sum_t \frac{1}{3} \left(\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}} \left[p_{i_t}(A_{\mathcal{B},t}, r_t, \tilde{\mathbf{v}}) \right] + \mathbf{E}_{\tilde{v}_{j_t} \sim \mathcal{D}_{j_t}} \left[\tilde{v}_{j_t} \right] \right)$$
$$= \frac{1}{3} \sum_t \left(\mathbf{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t}, r_t) \right) - \tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t} \cup \{i_t\}, r_t) \right) \right] + \mathbf{E}_{\tilde{v}_{j_t} \sim \mathcal{D}_{j_t}} \left[\tilde{v}_{j_t} \right] \right)$$
$$\overset{(\star\star)}{\geq} \frac{1}{3} \left(\mathbf{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}}) \right] - \mathbf{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})) \right) \right] \right) + \frac{1}{3} \sum_t \mathbf{E}_{\tilde{v}_{j_t} \sim \mathcal{D}_{j_t}} \left[\tilde{v}_{j_t} \right] \ .$$

To see why the last inequality $(\star\star)$ holds, we use Lemma 5.3.3. Consider the step with counter t. If buyer i_t gets the item, we argued that $A_{\mathcal{B},t+1} = A_{\mathcal{B},t} \cup \{i_t\}$ and hence, the sum telescopes. On the other hand, if seller j_t decided to keep the item, we note that by Lemma 5.3.3,

$$\mathbf{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t}, r_t) \right) - \tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t} \cup \{i_t\}, r_t) \right) \right]$$
$$\geq \mathbf{E}_{\tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t}, r_t) \right) - \tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B},t}, r_t - 1) \right) \right]$$

and again, the sum telescopes since the prices in the next step are computed with respect to $r_t - 1$.

Further, denote by $M_{\mathrm{SELL}}(\mathbf{v})$ the set M_{SELL} after running our mechanism with valuation profile $\mathbf{v}$. Note that any seller who is not in $M_{\mathrm{SELL}}(\mathbf{v})$ either participated in a trade or irrevocably kept the item during our mechanism. Therefore,

$$\sum_t \mathbf{E}_{\tilde{v}_{j_t} \sim \mathcal{D}_{j_t}} \left[\tilde{v}_{j_t} \right] = \sum_{j \in \mathcal{S} \setminus M_{\mathrm{SELL}}(\mathbf{v})} \mathbf{E}_{\tilde{v}_j \sim \mathcal{D}_j} \left[\tilde{v}_j \right] \ .$$

Taking the expectation over all valuation profiles $\mathbf{v}$, exploiting linearity of expectation and using that $\tilde{\mathbf{v}} \sim \mathcal{D}$ is an independent and identically distributed draw from the same distribution, we get

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i \in A_{\mathcal{B}}(\mathbf{v}) \cup A_{\mathcal{S}}(\mathbf{v})} P_i \right]$$
$$\geq \frac{1}{3} \mathbf{E}_{\mathbf{v}} \left[\mathbf{v} \left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}) \right) \right] - \frac{1}{3} \mathbf{E}_{\mathbf{v}, \tilde{\mathbf{v}}} \left[\tilde{\mathbf{v}} \left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})) \right) \right] + \frac{1}{3} \mathbf{E}_{\mathbf{v}, \tilde{\mathbf{v}}} \left[\sum_{j \in \mathcal{S} \setminus M_{\mathrm{SELL}}(\mathbf{v})} \tilde{v}_j \right] \ .$$

Surplus. The part of the welfare which is not covered by the base value is captured in the surplus. In order to talk about the surplus of any agent who receives an item, we split the set of agents and consider buyers and sellers separately.

Sellers: Fix seller j. Note that any seller whose value exceeds a price which we offered keeps her item. By construction of our mechanism, seller j is matched to some buyer(s) in the mechanism and asked if she would like to keep or try selling the item for price $p_{i,j}$. Let i_j denote the first buyer to which j is matched in the mechanism. This matching is independent of seller j's actual valuation since it only depends on the prices for seller j and buyer i_j. In the case that i_j does not exist (i.e. seller j was never offered

a trade), we can simply set $i_j = \perp$ and $p_{i_j,j} = 0$ and apply the same argument. The last price offered to seller j is P_j (maybe 0 if seller j was never offered a trade) and let the counter show t at this point.
Note that the prices which we offered to seller j cannot have increased in the process. Hence, the last price P_j which we offered to seller j is clearly upper bounded by the first price $p_{i_j,j}$ which we offered to j. Further, by Lemma 5.3.4, the price for the trade between j and i_j is only non-decreasing compared to offering a trade between buyer i_j and seller j later in the process again. Therefore, we can bound the surplus of seller j as follows:

$$\begin{aligned}(v_j - P_j)^+ &\geq \left(v_j - p_{i_j,j}(A_{\mathcal{B},t}(\mathbf{v}), r_t)\right)^+ \geq \left(v_j - p_{i_j,j}(A_{\mathcal{B}}((v'_j, \mathbf{v}_{-j})), r((v'_j, \mathbf{v}_{-j})))\right)^+ \\ &\geq \left(v_j - p_{i_j,j}(A_{\mathcal{B}}((v'_j, \mathbf{v}_{-j})), r((v'_j, \mathbf{v}_{-j})))\right)^+ \cdot \mathbb{1}_{j \in M_{\text{SELL}}(v'_j, \mathbf{v}_{-j})} \ .\end{aligned}$$

Taking expectations on both sides and exploiting that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed allows to bound

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[(v_j - P_j)^+\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v_j - p_{i_j,j}(A_{\mathcal{B}}((v'_j, \mathbf{v}_{-j})), r((v'_j, \mathbf{v}_{-j})))\right)^+ \cdot \mathbb{1}_{j \in M_{\text{SELL}}(v'_j, \mathbf{v}_{-j})}\right] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v'_j - p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+ \cdot \mathbb{1}_{j \in M_{\text{SELL}}(\mathbf{v})}\right] \ .\end{aligned}$$

Next, we can sum over all sellers and use linearity of expectation to obtain

$$\begin{aligned}\sum_{j \in \mathcal{S}} \mathbf{E}_{\mathbf{v}}\left[(v_j - P_j)^+\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in \mathcal{S}} \left(v'_j - p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+ \cdot \mathbb{1}_{j \in M_{\text{SELL}}(\mathbf{v})}\right] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in M_{\text{SELL}}(\mathbf{v})} \left(v'_j - p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+\right] \\ &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in M_{\text{SELL}}(\mathbf{v})} v'_j\right] - \mathbf{E}_{\mathbf{v}}\left[\sum_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right] \ .\end{aligned}$$

Let us pause for a moment and consider the sum over the prices. First of all, note that by construction of our mechanism, at most one seller $j^* \in M_{\text{SELL}}(\mathbf{v})$ is offered (maybe multiple times) a trade at all. Therefore, any other seller j satisfies that $i_j = \perp$ and hence for all sellers except j^*, we can set $p_{i_j,j} = 0$.
Having a look at the seller $j^* \in M_{\text{SELL}}(\mathbf{v})$ who is offered a trade (if j^* exists), the price for a trade between j^* and i_{j^*} was well-defined in the iteration that j^* and i_{j^*} were considered for a trade. Note that $A_{\mathcal{B}}$ and r did not change after this iteration anymore, so if i_{j^*} could be feasibly added to $A_{\mathcal{B}}$ at the step we offered a trade, she also can be feasibly added to $A_{\mathcal{B}}$ after the mechanism. Therefore, combining the price given by

$$p_{i_{j^*},j^*}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})) = \frac{1}{3}\left(\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[p_{i_{j^*}}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}), \tilde{\mathbf{v}})\right] + \mathbf{E}_{\tilde{v}_{j^*} \sim \mathcal{D}_{j^*}}\left[\tilde{v}_{j^*}\right]\right)$$

with

$$\begin{aligned}p_{i_{j^*}}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}), \tilde{\mathbf{v}}) &= \tilde{\mathbf{v}}\left(\text{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right) - \tilde{\mathbf{v}}\left(\text{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}) \cup \{i_{j^*}\}, r(\mathbf{v}))\right) \\ &\leq \tilde{\mathbf{v}}\left(\text{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)\end{aligned}$$

allows to bound the sum of prices as follows:

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}}\left[\sum_{j \in M_{\mathrm{SELL}}(\mathbf{v})} p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right] \\
&\leq \frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)\right] + \frac{1}{3}\mathbf{E}_{\mathbf{v}}\left[\mathbf{E}_{\tilde{v}_{j^*} \sim \mathcal{D}_{j^*}}\left[\tilde{v}_{j^*}\right]\right] \\
&\leq \frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)\right] + \frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\sum_{j \in M_{\mathrm{SELL}}(\mathbf{v})} \tilde{v}_j\right] \quad .
\end{aligned}$$

Now, we use that $\mathbf{v}$, $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed. Therefore, we can bound the surplus of all sellers by the following expression:

$$\begin{aligned}
\sum_{j \in \mathcal{S}} \mathbf{E}_{\mathbf{v}}\left[(v_j - P_j)^+\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in M_{\mathrm{SELL}}(\mathbf{v})} v'_j\right] - \mathbf{E}_{\mathbf{v}}\left[\sum_{j \in M_{\mathrm{SELL}}(\mathbf{v})} p_{i_j,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right] \\
&\geq \frac{2}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in M_{\mathrm{SELL}}(\mathbf{v})} v'_j\right] - \frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)\right] \quad .
\end{aligned} \tag{5.4}$$

Buyers: First, observe that initially, all buyers can be feasibly added to $A_{\mathcal{B}}$. During the mechanism, buyers may become infeasible at some point in time. Once a buyer cannot be feasibly added anymore, this buyer will remain infeasible for the remainder of the mechanism. On the other hand, if a buyer can be feasibly added at some point in time, she could also be feasibly added at any time before. During our mechanism, we offer trades to all buyers except of those who did become infeasible on the way. Any of the buyers to which we offer a trade for a finite price gets an item if her value exceeds the offered price. As a consequence, we are allowed to consider $(v_i - P_i)^+$ as the contribution to the surplus for all buyers. We can define P_i for buyer i to be infinity if buyer i was not offered a trade in our mechanism due to the fact that buyer i became infeasible. In the same way, if $p_{i,j}(A_{\mathcal{B}}, r)$ is not well-defined for a buyer due to the fact that this buyer did become infeasible, we defined the price to be infinity. This directly implies a zero contribution to the surplus, so we do not need to focus on these buyers anymore in our considerations. Otherwise, as before, P_i denotes the price which we offered to buyer i.

Observe that by Lemma 5.3.5, the prices which are proposed to buyer i are non-decreasing as the allocation process proceeds. As said, any buyer who is offered a trade and exceeds her price gets an item in our mechanism. Note that the price which we offered to buyer i only depends on the sellers and all buyers which we did consider before i. In particular, being offered a trade and its price are independent of buyer i's value.

Having this, let j_t denote the seller which is matched to buyer i in round t, i.e. in the round in which buyer i receives an item (if she does). Hence, for all buyers which are offered trades, we are allowed to calculate

$$\begin{aligned}(v_i - P_i)^+ &= (v_i - p_{i,j_t}(A_{\mathcal{B},t}(\mathbf{v}), r_t))^+ \\ &\geq \left(v_i - \min_{j \in M_{\text{SELL}}((v_i', \mathbf{v}_{-i}))} p_{i,j}(A_{\mathcal{B}}((v_i', \mathbf{v}_{-i})), r((v_i', \mathbf{v}_{-i})))\right)^+ \\ &\geq \left(v_i - \min_{j \in M_{\text{SELL}}((v_i', \mathbf{v}_{-i}))} p_{i,j}(A_{\mathcal{B}}((v_i', \mathbf{v}_{-i})), r((v_i', \mathbf{v}_{-i})))\right)^+ \\ &\quad \cdot \mathbb{1}_{i \in \text{OPT}_{\mathcal{B}}\left((v_i, \mathbf{v}'_{-i}) \mid A_{\mathcal{B}}((v_i', \mathbf{v}_{-i})), r((v_i', \mathbf{v}_{-i}))\right)} \ .\end{aligned}$$

Note that if $M_{\text{SELL}}((v_i', \mathbf{v}_{-i}))$ is empty, then the minimum is taken over the empty set and we do not consider buyer i anymore as in this case buyer i cannot be feasibly added to $A_{\mathcal{B}}$. Taking expectations on both sides and exploiting that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed allows the following:

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[(v_i - P_i)^+\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\Bigg[\left(v_i - \min_{j \in M_{\text{SELL}}((v_i', \mathbf{v}_{-i}))} p_{i,j}(A_{\mathcal{B}}((v_i', \mathbf{v}_{-i})), r((v_i', \mathbf{v}_{-i})))\right)^+ \\ &\qquad \cdot \mathbb{1}_{i \in \text{OPT}_{\mathcal{B}}\left((v_i, \mathbf{v}'_{-i}) \mid A_{\mathcal{B}}((v_i', \mathbf{v}_{-i})), r((v_i', \mathbf{v}_{-i}))\right)}\Bigg] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v_i' - \min_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+ \cdot \mathbb{1}_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))}\right] \ .\end{aligned}$$

Now, taking the sum over all buyers, we get

$$\begin{aligned}&\sum_{i \in \mathcal{B}} \mathbf{E}_{\mathbf{v}}\left[(v_i - P_i)^+\right] \\ &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathcal{B}} \left(v_i' - \min_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+ \cdot \mathbb{1}_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))}\right] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))} \left(v_i' - \min_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right)^+\right] \\ &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))} v_i'\right] - \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))} \min_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\text{OPT}_{\mathcal{B}}\left(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})\right)\right)\right] - \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))} \min_{j \in M_{\text{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))\right] \ .\end{aligned}$$

Let us take a closer look at the sum over the prices. Using Lemma 5.3.6 (to be stated later), we can upper bound the prices as follows:

$$\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i\in\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}'|A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}))}\min_{j\in M_{\mathrm{SELL}}(\mathbf{v})} p_{i,j}(A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}))\right]$$
$$=\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i\in\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}'|A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}))}\mathbf{E}_{\tilde{\mathbf{v}}}\left[p_i(A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}),\tilde{\mathbf{v}})\right]\right]+\frac{1}{3}\mathbf{E}_{\mathbf{v}}\left[|M_{\mathrm{SELL}}(\mathbf{v})|\cdot\min_{j\in M_{\mathrm{SELL}}(\mathbf{v})}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{v}_j\right]\right]$$
$$\le\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}\left(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v})\right)\right)\right]+\frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}\tilde{v}_j\right].$$

Overall, the surplus of all buyers can be bounded as follows:

$$\sum_{i\in\mathcal{B}}\mathbf{E}_{\mathbf{v}}\left[(v_i-P_i)^+\right]\ge\frac{2}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}\left(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v})\right)\right)\right]-\frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}\tilde{v}_j\right]. \tag{5.5}$$

Combination: Having discussed the surplus of buyers and sellers separately, we combine the two bounds in order to bound the total surplus of our mechanism by summing over all buyers and sellers. Therefore, we sum inequalities (5.4) and (5.5) and use that $\mathbf{v}$, $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed:

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in\mathcal{B}\cup\mathcal{S}}(v_i-P_i)^+\right]\ge&\frac{2}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}v'_j\right]-\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}))\right)\right]\\&+\frac{2}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}\left(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v})\right)\right)\right]-\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}v'_j\right]\\&=\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}\left(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v})\right)\right)\right]+\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}v'_j\right].\end{aligned}$$

Combining Base Value and Surplus. Adding base value and surplus and again, using that $\mathbf{v}$, $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed, we can lower bound the social welfare of our mechanism by

$$\begin{aligned}&\mathbf{E}_{\mathbf{v}}\left[\text{Base Value}\right]+\mathbf{E}_{\mathbf{v}}\left[\text{Surplus}\right]\\&\ge\frac{1}{3}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})\right)\right]-\frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}}\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v}))\right)\right]\\&\quad+\frac{1}{3}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\sum_{j\in\mathcal{S}\setminus M_{\mathrm{SELL}}(\mathbf{v})}\tilde{v}_j\right]+\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}_{\mathcal{B}}\left(\mathbf{v}'\mid A_{\mathcal{B}}(\mathbf{v}),r(\mathbf{v})\right)\right)\right]\\&\quad+\frac{1}{3}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in M_{\mathrm{SELL}}(\mathbf{v})}v'_j\right]\\&=\frac{1}{3}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})\right)\right]+\frac{1}{3}\mathbf{E}_{\mathbf{v}}\left[\sum_{j\in\mathcal{S}}v_j\right].\end{aligned}$$

We can conclude as $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})\right)\right]+\mathbf{E}_{\mathbf{v}}\left[\sum_{j\in\mathcal{S}} v_j\right]=\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})\right)+\sum_{j\in\mathcal{S}} v_j\right]$ and for each valuation profile $\mathbf{v}$, we have that $\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v})\right)+\sum_{j\in\mathcal{S}} v_j$ is an upper bound on the optimal social welfare which can be achieved by allocating the items among all agents.

□

In order to conclude the proof of the approximation guarantee of Theorem 5.3.1, we show the remaining lemmas used in the proof. First, we aim for a bound of

$$p_i(A_{\mathcal{B}}, r, \mathbf{v}) = \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r)\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}} \cup \{i\}, r)\right)$$

with respect to a change in r instead of adding i to the set $A_{\mathcal{B}}$.

Lemma 5.3.3. *Fix some buyer i and a valuation profile $\mathbf{v}$. Let $A_{\mathcal{B}}$ and r be such that $A_{\mathcal{B}} \cup \{i\} \in \mathcal{I}_{\mathcal{B}}$ and $|A_{\mathcal{B}} \cup \{i\}| \leq r$. Then*

$$\begin{aligned}&\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r)\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}} \cup \{i\}, r)\right)\\ &\quad \geq \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r)\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r-1)\right) \enspace .\end{aligned}$$

Proof. We argue that $\mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}} \cup \{i\}, r)\right) \leq \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r-1)\right)$ which immediately proves the claim. As the key insight, note that any possible choice of agents for $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}} \cup \{i\}, r)$ is also a feasible choice for $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid A_{\mathcal{B}}, r-1)$ and hence, the claim follows. □

Second, we show that for a fixed buyer-seller-pair (i, j), the prices which we consider in our mechanism are only non-decreasing as the process evolves.

Lemma 5.3.4. *Fix buyer i and seller j. Let the price for trading between buyer i and seller j be $p_{i,j}(X, r)$ for some set of remaining sellers M_{SELL}. Then we have $p_{i,j}(X, r) \leq p_{i,j}(X', r')$ for any superset of allocated agents $X' \supseteq X$ and $r' \leq r$.*

Proof. First, if a buyer is infeasible with respect to X and r, or if she could be feasibly added with respect to X and r, but not to X' with r', the claim is trivial. Therefore, it remains to consider the case where both sides of the inequality are finite. So let us consider X, X' and r, r' such that i can feasibly be added and let $j \in M_{\mathrm{SELL}}$. By definition, $p_{i,j}(X, r) = \frac{1}{3}\left(\mathbf{E}_{\tilde{\mathbf{v}}\sim\mathcal{D}}\left[p_i(X, r, \tilde{\mathbf{v}})\right] + \mathbf{E}_{\tilde{v}_j\sim\mathcal{D}_j}\left[\tilde{v}_j\right]\right)$. We show the inequality pointwise for any $\tilde{\mathbf{v}}$ in the first summand (the second is independent of X and r anyway) and conclude by taking the expectation. Therefore, fix a valuation profile $\mathbf{v}$ and consider $p_i(X, r, \mathbf{v})$. We show that

$$\begin{aligned}p_i(X, r, \mathbf{v}) &= \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r)\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X \cup \{i\}, r)\right)\\ &\overset{(1)}{\leq} \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r')\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X \cup \{i\}, r')\right)\\ &\overset{(2)}{\leq} \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X', r')\right) - \mathbf{v}\left(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X' \cup \{i\}, r')\right) = p_i(X', r', \mathbf{v}) \enspace .\end{aligned}$$

To show Inequality (1), we first use that the basis $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r')$ can be chosen to be a subset of $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r)$. To see this, denote by $\{b_1, \ldots, b_m\}$ the basis $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r)$ in decreasing order of weights. We show that there is an m' such that $\{b_1, \ldots, b_{m'}\}$ is equal to $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r')$, where m' is chosen in a way that $|X \cup \{b_1, \ldots, b_{m'}\}| \leq r'$ and that $X \cup \{b_1, \ldots, b_{m'}\}$ has maximum size with respect to this property (i.e. either

we have equality or r' is larger than the cardinality of any independent set — in the latter case, we can just choose a basis without considering r').
Assume there is a set $\{b'_1, \dots, b'_{m'}\}$ such that $\sum_{k=1}^{m'} v_{b'_k} > \sum_{k=1}^{m'} v_{b_k}$, so $\{b_1, \dots, b_{m'}\}$ would not be a maximum weight basis with respect to X and r'. We know that $\{b'_1, \dots, b'_{m'}\}$ also needs to be independent with respect to X and r and further $m' \leq m$. Therefore, there are $m - m'$ elements in $\{b_1, \dots, b_m\}$ which we can add to $\{b'_1, \dots, b'_{m'}\}$ in order to get a basis in the matroid with respect to X and r. Denote these $m - m'$ elements with $b_{\pi_1}, \dots, b_{\pi_{m-m'}}$. Note that $\sum_{k=1}^{m-m'} v_{b_{\pi_k}} \geq \sum_{k=m'+1}^{m} v_{b_k}$. Combining this with the sum from above leads to

$$\sum_{k=1}^{m'} v_{b'_k} + \sum_{k=1}^{m-m'} v_{b_{\pi_k}} \geq \sum_{k=1}^{m'} v_{b'_k} + \sum_{k=m'+1}^{m} v_{b_k} > \sum_{k=1}^{m'} v_{b_k} + \sum_{k=m'+1}^{m} v_{b_k} = \sum_{k=1}^{m} v_{b_k} \quad ,$$

which is a contradiction to the fact that $\{b_1, \dots, b_m\}$ is a maximum weight basis in the matroid given X truncated by r.

Having this, we can argue about the impact of adding i to X on $\{b_1, \dots, b_m\}$ and $\{b_1, \dots, b_{m'}\}$ respectively. Consider two parallel executions of the Greedy algorithm computing $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X, r)$ and $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid X \cup \{i\}, r)$. The first Greedy will compute $\{b_1, \dots, b_m\}$ whereas the second Greedy will choose exactly the same elements except for an element $b_{(i)}$ for which $\{b_1, \dots, b_{(i)}\} \cup \{i\}$ contains a circuit. Therefore, the difference on the left-hand side of Inequality (1) is equal to $v_{b_{(i)}}$.
Applying the same argument for the difference on the right-hand side of Inequality (1), there is an element $b_{(i)'}$ which is chosen in the first Greedy execution but not in the second one as $\{b_1, \dots, b_{(i)'}\} \cup \{i\}$ contains a circuit in the matroid contracted with X and truncated with r'. Therefore, the difference on the right-hand side is equal to $v_{b_{(i)'}}$.
We argue that $b_{(i)'}$ cannot be later than $b_{(i)}$ in the basis $\{b_1, \dots, b_m\}$ which allows us to conclude as elements in $b_1, \dots, b_m$ are sorted by weight in decreasing order.
We show the claim by contradiction, so assume that $b_{(i)'}$ is an element after $b_{(i)}$ and $b_{(i)'}$ is the first element such that $X \cup \{b_1, \dots, b_{(i)'}\} \cup \{i\}$ contains a circuit in the matroid truncated with r'. Now, $b_{(i)'}$ is later than $b_{(i)}$, so $X \cup \{b_1, \dots, b_{(i)'}\} \cup \{i\}$ is a superset of $Y := X \cup \{b_1, \dots, b_{(i)}\} \cup \{i\}$. Note that $|Y| \leq |X \cup \{b_1, \dots, b_{(i)'}\} \cup \{i\}| \leq r'$. By assumption on $b_{(i)}$, $X \cup \{b_1, \dots, b_{(i)}\} \cup \{i\}$ contains a circuit in the matroid truncated with r and hence also needs to contain a circuit in the matroid truncated with r', so either $b_{(i)'}$ is not the first element in $\{b_1, \dots, b_{m'}\}$ which leads to a circuit with i or $b_{(i)'}$ is before $b_{(i)}$ in the order of the basis. In the first case, apply the same argument again iteratively, in the second case, we showed the desired contradiction. Since there are only finitely many elements, the iterative application of the argument will terminate and hence, we proved the first inequality.

To see that Inequality (2) holds, we consider the matroid $\mathcal{M}$ truncated to rank r'. Denote this matroid by $\mathcal{M}_{r'}$. Expressed differently, this is the intersection of the matroid $\mathcal{M}$ with the r'-uniform matroid defined on the same ground set. Using Lemma 3 from Kleinberg and Weinberg [2012], the function $f_{r'}(Y) = \mathbf{v}(\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid Y, r'))$ is submodular in Y where now $\mathrm{OPT}_{\mathcal{B}}(\mathbf{v} \mid Y, r')$ is a maximum weight basis in the matroid $\mathcal{M}_{r'}$. This implies Inequality (2). □

Next, we consider a fixed buyer i. Note that by the order in which we approach the sellers, pricing buyer i is equivalent to choosing the cheapest current seller out of all available ones and compute the price with respect to the current $A_{\mathcal{B}}$ and r. In other words, as buyer i arrives, the price we offer is $\min_{j \in T} p_{i,j}(A_{\mathcal{B}}, r)$, where T denotes the set of available sellers.

Lemma 5.3.5. (Non-decreasing prices for buyers) *Fix any buyer i. Then for any $T' \subseteq T \subseteq \mathcal{S}$, $A'_{\mathcal{B}} \supseteq A_{\mathcal{B}}$ and $r' \leq r$, we have $\min_{j \in T} p_{i,j}(A_{\mathcal{B}}, r) \leq \min_{j \in T'} p_{i,j}(A'_{\mathcal{B}}, r')$.*

As a short remark, we never delete agents from the set $A_{\mathcal{B}}$ in our mechanism. Further, the number r never increases and sellers are only removed from M_{SELL} and never added. Therefore, in other words, Lemma 5.3.5 states that for any fixed buyer i, the prices are non-decreasing as the allocation process proceeds.

Proof. Observe that the minimum over T contains at least any possible seller $j \in T'$. Hence the minimum on the left is taken over a superset of T'. Therefore, the claim follows by applying Lemma 5.3.4, i.e. $p_{i,j}(A_{\mathcal{B}}, r)$ is non-decreasing with respect to adding agents to $A_{\mathcal{B}}$ and decreasing the number r. □

In addition to the properties shown above, we used the inequality

$$\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \text{OPT}_{\mathcal{B}}(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}))} \mathbf{E}_{\tilde{\mathbf{v}}}\left[p_i(A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v}), \tilde{\mathbf{v}})\right]\right] \leq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\text{OPT}_{\mathcal{B}}\left(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})\right)\right)\right]$$

in our proof. This inequality is a direct application of a proposition from Kleinberg and Weinberg [2012]. Adapted to our setting, we consider the matroid $\mathcal{M}_r$ which is the matroid over the set of buyers $\mathcal{M}_{\mathcal{B}}$ truncated to rank r (note that the intersection of $\mathcal{M}_{\mathcal{B}}$ with the r-uniform matroid over the same ground set is again a matroid). Denote by $\mathcal{I}_r$ the independent sets in $\mathcal{M}_r$. We apply Proposition 2 from Kleinberg and Weinberg [2012] to our setting.

Lemma 5.3.6. [adapted version of Kleinberg and Weinberg, 2012, Proposition 2] *Fix valuation profile $\tilde{\mathbf{v}}$ and r and let $A_{\mathcal{B}} \in \mathcal{I}_r$. For any from $A_{\mathcal{B}}$ disjoint set $V \in \mathcal{I}_r$ with $A_{\mathcal{B}} \cup V \in \mathcal{I}_r$, it holds $\sum_{i \in V} p_i(A_{\mathcal{B}}, r, \tilde{\mathbf{v}}) \leq \tilde{\mathbf{v}}\left(\text{OPT}_{\mathcal{B}}(\tilde{\mathbf{v}} \mid A_{\mathcal{B}}, r)\right)$.*

Setting $V = \text{OPT}_{\mathcal{B}}\left(\mathbf{v}' \mid A_{\mathcal{B}}(\mathbf{v}), r(\mathbf{v})\right)$ as well as $A_{\mathcal{B}} = A_{\mathcal{B}}(\mathbf{v})$, we get the desired inequality pointwise for any fixed $\mathbf{v}$ and $\mathbf{v}'$. Hence, we can conclude by taking the expectation on both sides, using linearity and the fact that $\tilde{\mathbf{v}}$ and $\mathbf{v}'$ are independent and identically distributed.

5.4 Matroid Double Auctions with Weak Budget Balance and Online Arrival

We consider the same setting as in Section 5.3, so the set of buyers who receive an item needs to be an independent set in $\mathcal{M}_{\mathcal{B}} = (\mathcal{B}, \mathcal{I}_{\mathcal{B}})$ and each sellers brings one identical item to the market. In contrast to Section 5.3, our mechanism can deal with buyers arriving online with an adversary specifying the order. The adversary may even adapt the choices depending on the set of already considered agents and their valuations.

Also in this section, we start by a description of the mechanism in Section 5.4.1, the pricing is described afterwards in Section 5.4.2. Properties of the mechanism and a proof of the approximation guarantee are discussed in Section 5.4.3.

5.4.1 The Mechanism

Let $A_\mathcal{B}$ be the set in which we will store all buyers who receive an item, hence $A_\mathcal{B} \in \mathcal{I}_\mathcal{B}$. Further, $A_\mathcal{S}$ denotes the set of sellers who decide to keep the item irrevocably. In addition, we define a set $A' = A'_\mathcal{B} \cup A'_\mathcal{S}$, which is the set of agents that we eventually use to set the prices. We first go through all the sellers asking whether seller j wants to irrevocably keep her item or try selling it knowing that she will receive at most an amount of p_j (in case we sell the item). Afterwards, we go through all buyers in any order one-by-one. When considering buyer i, we match i to an arbitrary seller who still tries selling her item (if available).

For this buyer-seller pair, we propose the following trade: Buyer i pays the specific price p_i but seller j only receives $\min\{p_i, T_j\}$, where T_j is the lowest price that we have ever offered to seller j up to this point. If seller j does not agree to trade, she irrevocably keeps the item; j is added to A but i is added to A'. Otherwise, if buyer i does not agree, she is irrevocably discarded. Seller j might get matched again but the price offered to her can only decrease.

Note that this mechanism does not require any specified order in which we process the agents — even further, the matching which we consider for possible trades can be arbitrary, even determined by an adversary. This is a sharp contrast to Section 5.3, where we consider buyer-seller pairs in a tailored way.

Algorithm 8: Mechanism for Matroid Double Auctions with Online Arrival

Result: Set $A = A_\mathcal{B} \cup A_\mathcal{S}$ of agents to get an item with $A_\mathcal{B} \subseteq \mathcal{B}$, $A_\mathcal{B} \in \mathcal{I}_\mathcal{B}$, $A_\mathcal{S} \subseteq \mathcal{S}$ and $|A| = |\mathcal{S}|$

$A \longleftarrow \emptyset$; $A' \longleftarrow \emptyset$; $M_{\text{SELL}} \longleftarrow \emptyset$; $T = (0, \ldots, 0)$ (T is vector of zeros of size $|\mathcal{S}|$)
for $j \in \mathcal{S}$ **do**
  if $v_j \geq p_j$ **then**
    $A \longleftarrow A \cup \{j\}$; $A' \longleftarrow A' \cup \{j\}$
  if $v_j < p_j$ **then**
    $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \cup \{j\}$; $T_j \longleftarrow p_j$
for $i \in \mathcal{B}$ **do**
  if $M_{\text{SELL}} \neq \emptyset$ **then**
    select $j \in M_{\text{SELL}}$ **arbitrarily**
    if $p_i \geq T_j$ **then**
      if $v_i \geq p_i$ **then**
        $A \longleftarrow A \cup \{i\}$; $A' \longleftarrow A' \cup \{i\}$; $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$
        buyer i pays p_i to the mechanism and seller j receives T_j
    if $p_i < T_j$ **then**
      if $v_j \geq p_i$ **then**
        $A \longleftarrow A \cup \{j\}$; $A' \longleftarrow A' \cup \{i\}$; $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$
      if $v_j < p_i$ **then**
        $T_j \longleftarrow p_i$
        if $v_i \geq p_i$ **then**
          $A \longleftarrow A \cup \{i\}$; $A' \longleftarrow A' \cup \{i\}$; $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$
          buyer i pays p_i to the mechanism and seller j receives T_j
return $A \cup M_{\text{SELL}}$

5.4.2 The Pricing

The key to setting the prices is the set $A' = A'_{\mathcal{B}} \cup A'_{\mathcal{S}}$ with $A'_{\mathcal{B}} \subseteq \mathcal{B}$ and $A'_{\mathcal{S}} \subseteq \mathcal{S}$, which is maintained in addition to the sets $A_{\mathcal{B}}$ and $A_{\mathcal{S}}$. The idea is that for agents in A' the respective agent-specific price can be charged to someone in our mechanism. For a buyer $i \in A'$, this can mean that the buyer herself received an item and paid for it or that the corresponding seller decided to keep the item. We calculate prices with respect to the set A' in the spirit of the pricing schemes of matroid Prophet Inequalities by Kleinberg and Weinberg [2019] and Dütting et al. [2020].

In more detail, concerning the buyers, we set prices to infinity if there are no items available anymore or if buyer i cannot be added to $A'_{\mathcal{B}}$, i.e. $p_i(A') = \infty$ if $A'_{\mathcal{B}} \cup \{i\} \notin \mathcal{I}_{\mathcal{B}}$ or $M_{\text{SELL}} = \emptyset$. In particular, this pricing only affects the buyers and will never occur as long as we go through the sellers (in the first for-loop) in our mechanism. Hence, for sellers, there will always be a finite seller-specific price.

Now, for any agent who can be feasibly added to A' (i.e. all sellers and all buyers in cases different to the ones mentioned above), we compute prices in the following way. Fix a valuation profile $\mathbf{v}$ and denote by $\text{OPT}(\mathbf{v} \mid X)$ the set of agents who receive an item in an optimal allocation given that we have already irrevocably allocated items to agents in X. In contrast to the mechanism in Section 5.3, the optimum is now computed over all agents, not only over the set of buyers. The value of this partial allocation is denoted by $\mathbf{v}\left(\text{OPT}(\mathbf{v} \mid X)\right)$, that is, the sum over the value v_i of all agents i who receive an item. Further, define $\text{OPT}(\mathbf{v}) = \text{OPT}(\mathbf{v} \mid \emptyset)$.

Denote by A_i and A'_i the state of set A and A' *after* processing agent i (also for $A_{\mathcal{B}}$ and $A_{\mathcal{S}}$) and let $p_i(A'_{i-1}, \mathbf{v}) = \mathbf{v}\left(\text{OPT}(\mathbf{v} \mid A'_{i-1})\right) - \mathbf{v}\left(\text{OPT}(\mathbf{v} \mid A'_{i-1} \cup \{i\})\right)$. For any seller and all buyers such that $A'_{\mathcal{B};i-1} \cup \{i\} \in \mathcal{I}_{\mathcal{B}}$, as long as there are items remaining, the price for agent i is computed as

$$\begin{aligned} p_i(A'_{i-1}) &= \frac{1}{2}\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[p_i(A'_{i-1}, \tilde{\mathbf{v}})\right] \\ &= \frac{1}{2}\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[\tilde{\mathbf{v}}\left(\text{OPT}(\tilde{\mathbf{v}} \mid A'_{i-1})\right) - \tilde{\mathbf{v}}\left(\text{OPT}(\tilde{\mathbf{v}} \mid A'_{i-1} \cup \{i\})\right)\right] \ . \end{aligned}$$

This way of setting prices also ensures feasibility with respect to the matroid constraint, meaning that $A_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}$. The reason is that $A'_{\mathcal{B}} \supseteq A_{\mathcal{B}}$ at all times as every time we add a buyer to $A_{\mathcal{B}}$, the buyer is also added to $A'_{\mathcal{B}}$. We even have $A'_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}$ because buyers have infinite prices as soon as they cannot be feasibly added to A'.

5.4.3 Properties of Our Mechanism

Now, we can state the following theorem.

Theorem 5.4.1. *The mechanism for matroid double auctions is DWBB, DSIC and IR for all buyers and sellers and $\frac{1}{2}$-competitive with respect to the optimal social welfare.*

Proof for DWBB, DSIC and IR. First, observe that DWBB is obtained via the price comparison of our mechanism: either a trade between some seller j and some buyer i happens at price p_i, or buyer i pays $p_i \geq T_j$ to the mechanism whereas seller j only receives T_j. The difference $p_i - T_j$ is extracted and never used again. Satisfying DSIC and IR for buyers can be seen easily as we only offer a trade to any buyer at most once. Also IR for sellers follows naturally. In order to obtain DSIC for sellers, the key observation is that the amount of money which we may pay to seller j is only non-increasing in the allocation process. Hence, as a seller, if she wants to sell her item, she wants to do so as early as possible which implies DSIC for sellers. □

Proof of approximation ratio via a reduction to matroid Prophet Inequality. Next, we prove the competitive ratio via a reduction to matroid Prophet Inequalities as in Kleinberg and Weinberg [2019] and Dütting et al. [2020]. For the sake of completeness, there is a self-contained proof in this thesis after the reduction-based one.

As we are optimizing over a set of agents which is partially (on the buyers' side) equipped with a matroid constraint, we start by extending this to an equivalent setting with a matroid over the whole set of agents. The ground set of this extended matroid is $\mathcal{B} \cup \mathcal{S}$. Afterwards, we show a correspondence of our setting to the one in Kleinberg and Weinberg [2019] and Dütting et al. [2020].

There is the matroid $\mathcal{M}_{\mathcal{B}} = (\mathcal{B}, \mathcal{I}_{\mathcal{B}})$ over the set of buyers. On the sellers' side we construct an artificial matroid by considering the $|\mathcal{S}|$-uniform matroid over the set of sellers, denoted by $\mathcal{M}_{\mathcal{S}} = (\mathcal{S}, \mathcal{I}_{\mathcal{S}})$. Afterwards, we consider the union of the two matroids $\widehat{\mathcal{M}} = (\mathcal{B} \cup \mathcal{S}, \mathcal{J})$, where a set $I = I_{\mathcal{B}} \cup I_{\mathcal{S}}$ is now independent, if $I_{\mathcal{B}} \in \mathcal{I}_{\mathcal{B}}$ and $I_{\mathcal{S}} \in \mathcal{I}_{\mathcal{S}}$. In order to mirror the feasibility constraint of having only $|\mathcal{S}|$ items, we intersect $\widehat{\mathcal{M}}$ with the $|\mathcal{S}|$-uniform matroid over $\mathcal{B} \cup \mathcal{S}$ and denote this matroid by $\mathcal{M}$. Observe that by construction, $\mathcal{M}$ is again a matroid. As a consequence, we can relate all feasible allocations with respect to $\mathcal{M}_{\mathcal{B}}$ to independent sets in the extended matroid $\mathcal{M}$.

Concerning our pricing scheme, first, observe that we calculated prices with respect to the set A' by setting $p_i = \infty$ if $A'_{\mathcal{B}} \cup \{i\} \notin \mathcal{I}_{\mathcal{B}}$ or if all items are irrevocably allocated. This corresponds to sets which are not independent in the extended matroid $\mathcal{M}$ over the ground set $\mathcal{B} \cup \mathcal{S}$. A finite price for i (in case i can feasibly be added to A') can also be interpreted in the extended matroid $\mathcal{M}$: If $A' \cup \{i\} \in \mathcal{I}$, the price for agent i is computed to be $p_i(A'_{i-1})$. Note that this is exactly the pricing scheme used in Kleinberg and Weinberg [2019] and Dütting et al. [2020]. In addition, we ensure that $\mathrm{OPT}(\mathbf{v}) \in \mathcal{I}$ and also $X \cup \mathrm{OPT}(\mathbf{v} \mid X) \in \mathcal{I}$ for any $X \in \mathcal{I}$. Note that in particular, the matroid $\mathcal{M}$ combines the feasibility constraints for buyers and the constraint of having $|\mathcal{S}|$ items. As a consequence, computing prices with respect to $\mathcal{M}$ is equivalent to our pricing strategy from Section 5.4.2.

It remains to argue why the competitive ratio of the matroid Prophet Inequality in Kleinberg and Weinberg [2019] and Dütting et al. [2020] also implies our approximation guarantee for matroid double auctions. To this end, let us first interpret sellers as buyers who keep their items if they exceed the offered price.

Base Value. Once an item is irrevocably allocated to agent i, denote the price for the offered trade by P_i. Summing over all agents who are allocated an item, we get

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in A(\mathbf{v})} P_i\right] &= \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in A'(\mathbf{v})} \frac{1}{2}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}}|A'_{i-1}(\mathbf{v}))\right) - \tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}}|A'_{i-1}(\mathbf{v}) \cup \{i\})\right)\right]\right] \\
&= \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right] - \frac{1}{2}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}}|A'(\mathbf{v}))\right)\right] \ .
\end{aligned}$$

As a consequence, the base value of our mechanism equals the revenue of the mechanism from Kleinberg and Weinberg [2019] and Dütting et al. [2020] in one-sided markets.

Surplus. Kleinberg and Weinberg [2019] use the following lower bound: consider agents in $\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)$ and the sum of their surpluses is a lower bound on the overall surplus.

Now, also in our two-sided environment, the only agents who cannot contribute to the overall surplus are the ones contained in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ as they do not receive an item. Note that by construction, any agent $i \in A'(\mathbf{v}) \setminus A(\mathbf{v})$ is a buyer. We observe that any other agent $i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})$ whose value v_i exceeds her corresponding price gets an item in our mechanism. Overall, any agent $i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})$ had the chance to obtain an item in our process if her value exceeded her price.

Next, we discuss why a buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$. We see that buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ if and only if her buyer-specific price p_i, the value v_j of seller j (the seller who is matched to i once she entered the market) and the lowest price T_j offered to seller j satisfy

$$p_i < T_j \qquad \text{and} \qquad p_i \leq v_j \ .$$

In particular, the decision whether buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ does not depend on her value v_i at all. So, for any deviation of buyer i to v_i', buyer i would end up in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ in the same cases as she would with valuation v_i. In addition, observe that $i \in \mathrm{OPT}(\mathbf{w} \mid A'(\mathbf{v}))$ implies $i \notin A'(\mathbf{v})$ for any valuation profile $\mathbf{w}$.

Combining this with the above observation, any agent $i \in \mathrm{OPT}(\mathbf{w} \mid A'(\mathbf{v}))$ has a considerable surplus if she exceeds her price. For an independently sampled valuation profile $\mathbf{v}' \sim \mathcal{D}$, we get

$$\begin{aligned}
\text{surplus}_i &\geq \left(v_i - p_i(A'_{i-1}(\mathbf{v}))\right)^+ \mathbb{1}_{i \notin A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))} \\
&\geq \left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \notin A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))} \\
&\geq \left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}'_{-i}) \mid A'((v_i', \mathbf{v}_{-i}))\right)} \ ,
\end{aligned}$$

where we use that prices are non-decreasing for any fixed agent in the second inequality (which follows from the pricing structure in the matroid Prophet Inequality from Kleinberg and Weinberg [2019]). Taking expectations on both sides and exploiting that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed, our lower bound for the surplus coincides with the one from the one-sided environment $\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v_i' - p_i(A'(\mathbf{v})))^+ \mathbb{1}_{i \in \mathrm{OPT}(\mathbf{v}' \mid A'(\mathbf{v}))}\right]$.

As also the welfare of the expected (offline) optimum is the same, we obtain the same guarantee as the competitive ratio of 1/2 in Kleinberg and Weinberg [2019]. □

For the reader not familiar with the literature on matroid Prophet Inequalities, we also provide a full proof of the approximation guarantee without reducing to Kleinberg and Weinberg [2019] and Dütting et al. [2020]. It is mainly based on the insights from the reduction-based proof.

Proof of Approximation Guarantee without Reduction. In order to give a full proof of the approximation guarantee without using a reduction to the Prophet Inequality proof of Kleinberg and Weinberg [2019] and Dütting et al. [2020], we start as follows. Denote by $A(\mathbf{v})$ and $A'(\mathbf{v})$ the sets A and A' under valuation profile $\mathbf{v}$. We want to compare $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(A(\mathbf{v}))\right]$ to $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right]$. To this end, again, we split the welfare of our algorithm into two parts, the base value and the surplus, and bound each quantity separately. The base value is defined as follows: let agent i receive an item in our mechanism, i.e. $i \in A$. Any buyer who gets an item has paid her agent-specific price for an item. Any seller who decided to keep her item was asked to keep it for her agent-specific price or for the buyer-specific price she was matched to. The part of any agent i's value which is below this price is denoted the base value. The surplus is the part of any agent i's value above this threshold if it exists, otherwise it is zero.

Base Value. Our base part of the social welfare is defined via the prices. Note that all agents who are irrevocably allocated an item (i.e. which are in A before adding the remaining sellers) have a value which exceeds her agent-specific price, except for sellers $j \in A \setminus A'$ whose value for an item exceeds the price of the corresponding buyer. Denote this final price by P_i for agent i. For any seller $j \in A \setminus A'$, the corresponding buyer is stored in $A' \setminus A$, so we can replace the seller and the buyer when summing the base value of all agents who get an item. Fixing a valuation profile $\mathbf{v}$ and summing over all agents in $A(\mathbf{v})$ in the order that they were added to A, we can compute by a telescopic sum argument that

$$\begin{aligned}\sum_{i\in A(\mathbf{v})} P_i &= \sum_{i\in A'(\mathbf{v})} p_i(A'_{i-1}(\mathbf{v})) = \frac{1}{2}\sum_{i\in A'(\mathbf{v})} \mathbf{E}_{\tilde{\mathbf{v}}}\left[p_i(A'_{i-1}(\mathbf{v}),\tilde{\mathbf{v}})\right] \\ &= \frac{1}{2}\sum_{i\in A'(\mathbf{v})} \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}} \mid A'_{i-1}(\mathbf{v}))\right) - \tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}} \mid A'_{i-1}(\mathbf{v}) \cup \{i\})\right)\right] \\ &= \frac{1}{2}\left(\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] - \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}}|A'(\mathbf{v}))\right)\right]\right) \quad .\end{aligned}$$

Taking the expectation over all valuation profiles $\mathbf{v}$, exploiting linearity of expectation and using that $\tilde{\mathbf{v}} \sim \mathcal{D}$ implies

$$\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in A(\mathbf{v})} P_i\right] = \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}\left(\mathrm{OPT}(\mathbf{v})\right]\right. - \frac{1}{2}\mathbf{E}_{\mathbf{v},\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}} \mid A'(\mathbf{v}))\right)\right] \quad .$$

Surplus. We start with two observations which will be helpful later. First of all, we will see in Lemma 5.4.2 that agent-specific prices are non-decreasing: For a fixed agent i, it holds that $p_i(A'_{i-1}) \leq p_i(A'_{i'})$ for any step $i' > i$, so in particular it holds $p_i(A'_{i-1}) \leq p_i(A')$.

Second of all, we will interrupt for a moment and focus on the set $A'(\mathbf{v}) \setminus A(\mathbf{v})$. This set contains all buyers whose agent-specific price was paid by a seller, i.e. the seller decided to keep the item for price p_i. Hence, any buyer $i \in A'(\mathbf{v}) \setminus A(\mathbf{v})$ does not get an item in the end, so their surplus is necessarily zero. Note that by construction, any agent $i \in A'(\mathbf{v}) \setminus A(\mathbf{v})$ is a buyer. We observe that any other agent $i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})$ whose value v_i exceeds her corresponding price gets an item in our mechanism. Additionally, there might be some sellers keeping their items in the end and some sellers keeping items for buyer-specific prices later in the process, but we will not take these contributions to the surplus into account. Overall, any agent $i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})$ had the chance to obtain an item in our process if her value exceeded her price.

Next, we want to observe why a buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$. Having a closer look at our algorithm, we see that buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ if and only if her buyer-specific price p_i, the value v_j of seller j (the seller who is matched to i once she entered the market) and the lowest price T_j offered to seller j satisfy

$$p_i < T_j \quad \text{and} \quad p_i \leq v_j \ .$$

In particular, the decision whether buyer i is in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ does not depend on her value v_i at all. So, for any deviation of buyer i to v_i', buyer i would end up in $A'(\mathbf{v}) \setminus A(\mathbf{v})$ in the same cases as she would with valuation v_i. Therefore, $i \in A'(\mathbf{v}) \setminus A(\mathbf{v})$ holds if and only if $i \in A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))$.

Further, any agent who can be feasibly added to a set $A'(\mathbf{v})$ cannot be contained in the set $A'(\mathbf{v})$. Hence, $i \in \mathrm{OPT}(\mathbf{w} \mid A'(\mathbf{v}))$ implies $i \notin A'(\mathbf{v})$ for any valuation profile $\mathbf{w}$ and in particular, $i \in \mathrm{OPT}(\mathbf{w} \mid A'(\mathbf{v}))$ implies $i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})$. Combining this with the above observation, any agent $i \in \mathrm{OPT}(\mathbf{w} \mid A'(\mathbf{v}))$ has a considerable surplus if she exceeds her price.

We can now consider the surplus of an agent i. Let $\mathbf{v}' \sim \mathcal{D}$ be an independently sampled valuation profile. Now, the price for agent i depends on $A'_{i-1}(\mathbf{v})$. But A'_{i-1} only depends on agents $1, \ldots, i-1$, so in particular we could replace v_i by v_i' and use that prices are non-decreasing for any fixed agent (which we show later in Lemma 5.4.2). Combining this with the above observations, we can bound the surplus of agent i from below via

$$\begin{aligned}
\text{surplus}_i &\geq \left(v_i - p_i(A'_{i-1}(\mathbf{v}))\right)^+ \mathbb{1}_{i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})} \\
&= \left(v_i - p_i(A'_{i-1}(\mathbf{v}))\right)^+ \mathbb{1}_{i \notin A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))} \\
&\geq \left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \notin A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))} \\
&\geq \left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \notin A'((v_i', \mathbf{v}_{-i})) \setminus A((v_i', \mathbf{v}_{-i}))} \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}'_{-i}) \mid A'((v_i', \mathbf{v}_{-i}))\right)} \\
&= \left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}'_{-i}) \mid A'((v_i', \mathbf{v}_{-i}))\right)} \ .
\end{aligned}$$

Taking expectations on both sides and exploiting that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed leads to

$$\begin{aligned}
&\mathbf{E}_{\mathbf{v}}\left[\left(v_i - p_i(A'_{i-1}(\mathbf{v}))\right)^+ \mathbb{1}_{i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})}\right] \\
&\geq \mathbf{E}_{\mathbf{v}, \mathbf{v}'}\left[\left(v_i - p_i(A'((v_i', \mathbf{v}_{-i})))\right)^+ \mathbb{1}_{i \in \mathrm{OPT}\left((v_i, \mathbf{v}'_{-i}) \mid A'((v_i', \mathbf{v}_{-i}))\right)}\right] \\
&= \mathbf{E}_{\mathbf{v}, \mathbf{v}'}\left[\left(v_i' - p_i(A'(\mathbf{v}))\right)^+ \mathbb{1}_{i \in \mathrm{OPT}(\mathbf{v}' \mid A'(\mathbf{v}))}\right] \ .
\end{aligned}$$

Summing over all agents (i.e. buyers and sellers), we can lower bound the overall surplus by

$$
\begin{aligned}
\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in A(\mathbf{v})} (v_i - P_i)^+\right] &\geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathcal{B} \cup \mathcal{S}} (v_i - P_i)^+ \mathbb{1}_{i \notin A'(\mathbf{v}) \setminus A(\mathbf{v})}\right] \\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} (v'_i - p_i(A'(\mathbf{v})))^+\right] \\
&\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} v'_i\right] - \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} p_i(A'(\mathbf{v}))\right] \\
&= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)\right)\right] - \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} p_i(A'(\mathbf{v}))\right] \\
&\geq \frac{1}{2}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)\right)\right] \ .
\end{aligned}
$$

The last inequality follows by bounding

$$
\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} p_i(A'(\mathbf{v}))\right] \leq \frac{1}{2}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)\right)\right]
$$

which we prove below in Lemma 5.4.3.

Summing the base value and the surplus proves our claim as we can exploit that $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed.

In order to conclude, we need to prove two remaining facts: first, agent-specific prices are non-decreasing, second, we need to show that

$$
\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} p_i(A'(\mathbf{v}))\right] \leq \frac{1}{2}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)\right)\right] \ .
$$

As mentioned in the reduction-based proof before, we can interpret the problem in an extended matroid containing buyers and sellers. This allows to directly apply propositions from Kleinberg and Weinberg [2012] and Dütting et al. [2020] with respect to our pricing scheme.

Properties of Prices. Looking at our optimization problem and in particular on the prices from the viewpoint of the matroid $\mathcal{M}$ over all agents, we can use Lemma E.2 in Dütting et al. [2020] to show that agent-specific prices are non-decreasing for any fixed agent i.

Lemma 5.4.2. [Dütting et al., 2020, Lemma E.2] *Consider any independent sets $X, Y \in \mathcal{I}$ with $X \subseteq Y$. Then, for any agent i, we have $p_i(X) \leq p_i(Y)$.*

Having that agent-specific prices are only non-decreasing, it remains to show that $\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}' | A'(\mathbf{v}))} p_i(A'(\mathbf{v}))\right] \leq \frac{1}{2}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)\right)\right]$ in order to conclude. Again, we use the matroid $\mathcal{M}$ as constructed above and apply a proposition from Kleinberg and Weinberg [2012].

Lemma 5.4.3. [Kleinberg and Weinberg, 2012, Proposition 2] *Fix valuation profile* $\tilde{\mathbf{v}}$ *and let* $A' \in \mathcal{I}$. *For any disjoint set* $V \in \mathcal{I}$ *with* $A' \cup V \in \mathcal{I}$, *it holds* $\sum_{i \in V} p_i(A', \tilde{\mathbf{v}}) \leq \tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}} \mid A')\right)$.

Setting $V = \mathrm{OPT}\left(\mathbf{v}' \mid A'(\mathbf{v})\right)$ as well as $A' = A'(\mathbf{v})$, we get the desired inequality pointwise for any fixed $\mathbf{v}$ and $\mathbf{v}'$. Hence, we can conclude by taking the expectation and using that $\tilde{\mathbf{v}}$ and $\mathbf{v}'$ are independent and identically distributed. □

5.5 Combinatorial Double Auctions with Strong Budget Balance

In combinatorial double auctions, there is a set of n buyers $\mathcal{B}$, a set of k sellers $\mathcal{S}$ and a set of m possibly heterogeneous items M. Agents have valuation functions over item bundles, so in contrast to Section 5.3 and Section 5.4, buyers may be interested in multiple items. As before, our mechanism is described first in Section 5.5.1, the pricing is discussed afterwards in Section 5.5.2, followed by proofs for the properties of the mechanism in Section 5.5.3.

5.5.1 The Mechanism

Our mechanism which is stated in Algorithm 9 works as follows: Given static and anonymous item prices p_j, we first ask any seller l which of her items in I_l she would like to keep and which to sell at prices p_j. After this, we have a set of available items M_{SELL} which we try to sell to the buyers via a sequential posted pricing procedure. That is, buyers are considered online one-by-one. Once buyer i is considered, she buys a utility-maximizing bundle among the available items and pays the respective prices to any seller from whom she gets an item. After running the mechanism, all items which are unallocated are returned to their corresponding sellers. As a side remark, the mechanism is robust concerning the arrival order of buyers; the order can even be chosen by an adversary.

5.5.2 The pricing

We restrict the class of valuation functions for both, buyers and sellers, to valuations which can be represented by XOS functions. In order to compute suitable prices, we mimic the pricing approach from Feldman et al. [2015] and Dütting et al. [2020] and apply this to two-sided markets. As a side remark, one could also use the LP-based approach from Chapter 3 to argue about the existence of static item prices.

Let ALLOC be an algorithm which allocates all items among the agents. We assume that ALLOC can only allocate items to sellers which are in their initial bundle. Fix a valuation profile $\tilde{\mathbf{v}}$ and denote by $\mathrm{ALLOC}_i(\tilde{\mathbf{v}})$ the set of items allocated to agent i when running allocation algorithm ALLOC on valuation profile $\tilde{\mathbf{v}}$. Using the XOS property, we can denote the additive function with which buyer i would evaluate set $\mathrm{ALLOC}_i(\tilde{\mathbf{v}})$ by $w_i^{\mathrm{ALLOC}_i(\tilde{\mathbf{v}})}$.

Then, for any $j \in \mathrm{ALLOC}_i(\tilde{\mathbf{v}})$, we can interpret $w_i^{\mathrm{ALLOC}_i(\tilde{\mathbf{v}})}(\{j\})$ as the contribution of item j to the overall social welfare of allocation algorithm ALLOC given valuation profile $\tilde{\mathbf{v}}$. In other words, for fixed valuation profile $\tilde{\mathbf{v}}$, we consider the allocation of

Algorithm 9: Mechanism for Combinatorial Double Auctions

Result: Allocation $\mathbf{X} = (X_i)_{i \in \mathcal{B} \cup \mathcal{S}}$ of items to agents such that for any seller l we have $X_l \subseteq I_l$ and $\bigcup_{i \in \mathcal{B} \cup \mathcal{S}} X_i = \bigcup_{l \in \mathcal{S}} I_l = M$

$X_i \longleftarrow \emptyset$ for all $i \in \mathcal{B} \cup \mathcal{S}$; $M_{\text{SELL}} \longleftarrow \emptyset$
for $l \in \mathcal{S}$ **do**
  Show prices p_j for each item $j \in I_l$ to seller l
  Ask seller l which items she wants to keep or try selling
  $X_l \longleftarrow \{j \in I_l :$ seller l wants to keep item $j\}$
  $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \cup \{j \in I_l :$ seller l tries selling item $j\}$
for $i \in \mathcal{B}$ **do**
  Show prices p_j for each item $j \in M_{\text{SELL}}$ to buyer i
  Ask buyer i which items she wants to buy
  $X_i \longleftarrow \{j \in M_{\text{SELL}} :$ buyer i wants to buy item $j\}$
  $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus X_i$
  Buyer i pays $\sum_{j \in X_i} p_j$
  Any seller with $j \in I_l$ for some $j \in X_i$ receives p_j and item j is traded to buyer i
for $l \in \mathcal{S}$ **do**
  $X_l \longleftarrow X_l \cup (M_{\text{SELL}} \cap I_l)$
return X

ALLOC, the additive set function $w_i^{\text{ALLOC}_i(\tilde{\mathbf{v}})}$ which represents $\tilde{v}_i(\text{ALLOC}_i(\tilde{\mathbf{v}}))$ and evaluate this only for a single item $j \in \text{ALLOC}_i(\tilde{\mathbf{v}})$. Now, we compute the price for item j as

$$p_j = \frac{1}{2}\mathbf{E}_{\tilde{\mathbf{v}} \sim \mathcal{D}}\left[\sum_{i \in \mathcal{B} \cup \mathcal{S}} w_i^{\text{ALLOC}_i(\tilde{\mathbf{v}})}(\{j\})\mathbb{1}_{j \in \text{ALLOC}_i(\tilde{\mathbf{v}})}\right] .$$

Observe that these prices are static and anonymous item prices for any item $j \in M$. Further, note that for ALLOC we have multiple choices: if we do not care about computational issues, we could use an optimal algorithm which computes an optimal allocation with respect to social welfare.

5.5.3 Properties of Our Mechanism

We consider two different settings for the chosen classes of valuation functions for buyers and sellers. First, we restrict to the case of unit-supply sellers, i.e. each seller bringing one non-identical item to the market. For buyers, we assume XOS valuation functions. Note that the valuation functions for sellers can also be represented by fractionally subadditive functions and hence, in order to prove the competitive ratio, we can apply Lemma 5.5.3 which allows to state the following theorem.

Theorem 5.5.1. *The mechanism for combinatorial double auctions with unit-supply sellers and buyers having XOS valuation functions is DSBB, DSIC and IR for all buyers and sellers and 1/2-competitive with respect to the optimal social welfare for any online adversarial arrival order of agents.*

For our second result, we assume buyers and sellers to have additive valuation functions. Note that this allows sellers to bring more than one item to the market. Since any additive valuation function can trivially be represented by a fractionally subadditive one, we can again apply Lemma 5.5.3 and hence, state the following theorem.

Theorem 5.5.2. *The mechanism for combinatorial double auctions with buyers and sellers having additive valuation functions is DSBB, DSIC and IR for all buyers and sellers and 1/2-competitive with respect to the optimal social welfare for any online adversarial arrival order of agents.*

Concerning the proof of these theorems, observe that by construction, the mechanism consists of bilateral trades where an item is transferred from one seller to one buyer and in exchange, a static and anonymous item price is paid by this buyer to the corresponding seller. Hence, we satisfy DSBB. In addition, IR is also satisfied as any agent can withdraw by buying nothing (as a buyer) or keeping the initial bundle (as a seller). The mechanism is further DSIC for buyers as any buyer is asked once in our mechanism which bundle she wants to purchase. As a unit-supply seller l only has one item and we offer her a price of p_j for the item, the mechanism is also DSIC for unit-supply sellers.
In the case of additive valuations for buyers and sellers, the mechanism is still DSIC for sellers: By additivity, any seller has a value $v_l(\{j\})$ for any $j \in I_l$ and hence, we can rewrite the utility as $\sum_{j \in X_l} v_l(\{j\}) + \sum_{j \in I_l \setminus X_l} p_j$. Since all buyers also have additive valuations, some buyer i will buy an available item j if and only if $v_i(\{j\}) > p_j$. In the case that for all buyers $v_i(\{j\}) < p_j$, the item is returned to the seller anyway. Hence, for any seller it is a dominant strategy to try selling all items for which $v_l(\{j\}) \leq p_j$ and keeping the items with $v_l(\{j\}) > p_j$ in order to maximize utility.

Hence, what remains to show is the desired guarantee with respect to the social welfare obtained by our mechanism.

Lemma 5.5.3. *The mechanism for combinatorial double auctions is 1/2-competitive with respect to the social welfare of* ALLOC *for any (possibly adversarial) arrival order of buyers and sellers when buyers have valuation functions which can be represented by fractionally subadditive functions and sellers are unit-supply or both have additive valuation functions over item bundles.*

Concerning the competitive ratio, we give a reduction-based proof of Lemma 5.5.3 first, a self-contained proof can be found afterwards.

Proof of Approximation Ratio via Reduction to Prophet Inequalities. In order to show that our mechanism achieves the desired competitive ratio, we already argued that in any of the two settings, a seller l will initially keep an item j if and only if $v_l(\{j\}) \geq p_j$. As a consequence, the competitive ratio directly follows by an application of the results from Feldman et al. [2015] and Dütting et al. [2020]: Interpret the sellers as buyers which are considered first, offer each to keep any item in I_l and sell the remaining items afterwards to all buyers via a sequential posted-prices mechanism. □

For any reader not familiar with the literature on Prophet Inequalities for combinatorial auctions, there is a self-contained proof below.

Proof of Approximation Guarantee without Reduction. Instead of applying a reduction, we can also mimic the techniques from Feldman et al. [2015] and Dütting et al. [2020]. Therefore, we split the contribution to social welfare into base value and surplus and bound each quantity separately.
Before we start, note that our mechanism consists of three phases. First, we ask all sellers which items should be sold and which they would like to keep. Afterwards, in the second phase, we ask all buyers which items they would like to buy. In the last phase,

all unsold items are returned to the corresponding sellers. Note that the last phase only increases the welfare of our mechanism compared to a mechanism which would stop after the second phase and dispose all unallocated items. We do not consider the increase in welfare in the third phase and argue about the welfare which we achieved in the first and second phase. This is a lower bound on the overall social welfare of our mechanism.

Base Value. Note that any item which is irrevocably allocated in the first two phases of the mechanism is allocated to an agent who has an item-specific value at least as high as the item price. Denote by A the set of irrevocably allocated items after the second phase, i.e. all items which are allocated before the last for-loop in the mechanism where we return unallocated items to their corresponding sellers. Note that A depends on the valuation profile $\mathbf{v}$. We write $A(\mathbf{v})$ in order to specify this dependence. Therefore, we can state the base value as

$$\mathbf{E}_{\mathbf{v}}\left[\text{Base Value}(\mathbf{v})\right] = \sum_{j \in M} \mathbf{Pr}_{\mathbf{v}}\left[j \in A(\mathbf{v})\right] \cdot p_j \ .$$

Surplus. For the surplus, we split the set of agents in buyers and sellers and consider them separately.

Sellers: Fix seller l. If l is holding one item j initially, then seller l irrevocably keeps the item if $v_l(\{j\}) \geq p_j$. Therefore, seller l has a considerable surplus if $v_l(\{j\}) - p_j \geq 0$. The same argument extends to the case of additive valuation functions, as seller l will initially keep all items for which $v_l(\{j\}) \geq p_j$ in order to maximize utility. Counting the surplus of seller l only for items in $\text{ALLOC}_l\left((v_l, \mathbf{v}'_{-l})\right)$ is a feasible lower bound for the surplus of seller l in our mechanism. Choosing $\mathbf{v}' \sim \mathcal{D}$ to be an independent sample, we can bound the surplus via

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[\text{surplus}_l(\mathbf{v})\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j \in \text{ALLOC}_l\left((v_l,\mathbf{v}'_{-l})\right)} \left(w_l^{\text{ALLOC}_l\left((v_l,\mathbf{v}'_{-l})\right)}(\{j\}) - p_j\right)^+\right] \\ &= \mathbf{E}_{\mathbf{v}'}\left[\sum_{j \in \text{ALLOC}_l(\mathbf{v}')} \left(w_l^{\text{ALLOC}_l(\mathbf{v}')}(\{j\}) - p_j\right)^+\right] \ ,\end{aligned}$$

where we used that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed.

Buyers: Fix buyer i. Extending the notation from above, denote by $A_i(\mathbf{v})$ the set of irrevocably allocated items as agent i is considered in the mechanism. Note that the set A_i does not depend on v_i but only on the agents which were considered before i. Hence, $A_i(\mathbf{v}) = A_i\left((v'_i, \mathbf{v}_{-i})\right)$ for any other valuation v'_i of buyer i. Buyer i could purchase the set $\text{ALLOC}_i\left((v_i, \mathbf{v}'_{-i})\right) \setminus A_i\left((v'_i, \mathbf{v}_{-i})\right)$. As buyer i maximizes utility, the utility which buyer i obtains must be at least as high as the utility when purchasing items in $\text{ALLOC}_i\left((v_i, \mathbf{v}'_{-i})\right) \setminus A_i\left((v'_i, \mathbf{v}_{-i})\right)$ for which the marginal surplus is non-negative when evaluating with the additive function in the XOS support $w_i^{\text{ALLOC}_i\left((v_i,\mathbf{v}'_{-i})\right)}(\cdot)$. As the utility of buyer i is captured in the surplus, we can bound the surplus of buyer i via

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[\text{surplus}_i(\mathbf{v})\right] &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in \text{ALLOC}_i\left((v_i,\mathbf{v}'_{-i})\right)\setminus A_i\left((v'_i,\mathbf{v}_{-i})\right)}\left(w_i^{\text{ALLOC}_i\left((v_i,\mathbf{v}'_{-i})\right)}(\{j\})-p_j\right)^+\right]\\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{j\in \text{ALLOC}_i(\mathbf{v}')\setminus A_i(\mathbf{v})}\left(w_i^{\text{ALLOC}_i(\mathbf{v}')}(\{j\})-p_j\right)^+\right] .\end{aligned}$$

Combination. Next, we sum over all buyers and sellers. Further, we use that once an item is irrevocably allocated, it remains so until the end of the mechanism, hence $A_i(\mathbf{v}) \subseteq A(\mathbf{v})$ for any agent i and any valuation profile $\mathbf{v}$. In order to simplify notation, note that $A_l(\mathbf{v}) \cap I_l = \emptyset$ as we ask seller l which items she wants to keep or try selling. As $\text{ALLOC}_l(\mathbf{v}) \subseteq I_l$, we know that $\text{ALLOC}_l(\mathbf{v}) \setminus A_l(\mathbf{v}) = \text{ALLOC}_l(\mathbf{v})$ as seller l arrives. Hence, we get

$$\begin{aligned}&\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in\mathcal{B}\cup\mathcal{S}}\text{surplus}_i(\mathbf{v})\right]\\ &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\sum_{i\in\mathcal{B}\cup\mathcal{S}}\sum_{j\in M}\left(w_i^{\text{ALLOC}_i(\mathbf{v}')}(\{j\})-p_j\right)^+\cdot \mathbb{1}_{j\in\text{ALLOC}_i(\mathbf{v}')}\cdot\mathbb{1}_{j\notin A_i(\mathbf{v})}\right]\\ &\geq \sum_{j\in M}\sum_{i\in\mathcal{B}\cup\mathcal{S}}\mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(w_i^{\text{ALLOC}_i(\mathbf{v}')}(\{j\})-p_j\right)^+\cdot \mathbb{1}_{j\in\text{ALLOC}_i(\mathbf{v}')}\cdot\mathbb{1}_{j\notin A(\mathbf{v})}\right]\\ &\geq \sum_{j\in M}\mathbf{Pr}_{\mathbf{v}}\left[j\notin A\left(\mathbf{v}\right)\right]\cdot\mathbf{E}_{\mathbf{v}'}\left[\sum_{i\in\mathcal{B}\cup\mathcal{S}}\left(w_i^{\text{ALLOC}_i(\mathbf{v}')}(\{j\})-p_j\right)\cdot\mathbb{1}_{j\in\text{ALLOC}_i(\mathbf{v}')}\right]\\ &\geq \sum_{j\in M}\mathbf{Pr}_{\mathbf{v}}\left[j\notin A\left(\mathbf{v}\right)\right]\cdot p_j \ ,\end{aligned}$$

where the last inequality uses the definition of the prices as well as the fact that ALLOC can allocate any item at most once.

Finally, adding base value and surplus implies the desired guarantee using that $\sum_{j\in M} p_j = \frac{1}{2}\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in\mathcal{B}\cup\mathcal{S}} v_i(\text{ALLOC}_i(\mathbf{v}))\right]$.

□

Note that we bound the competitive ratio of our mechanism with respect to the social welfare of the algorithm ALLOC. As said, ALLOC can either be an optimal algorithm, leading to the desired $1/2$-competitive mechanism with respect to the optimal welfare. Still, it is known that for deterministic XOS valuation functions, computing the optimal allocation is NP-hard [Lehmann et al., 2001]. Hence, we can also set ALLOC to be an approximation algorithm for the optimal social welfare. In this case, an α-approximation algorithm ALLOC leads to an $\alpha/2$-competitive mechanism. For example, one could use the algorithm by Feige [2009] with an approximation ratio of $1-1/e$. As a side remark, Dobzinski et al. [2010] show that it is NP-hard to approximate the optimal allocation for XOS valuations within any constant factor better than $1-1/e$.

5.6 Knapsack Double Auctions with Strong Budget Balance and Online Customized Arrival

As in Sections 5.3 and 5.4, we assume buyers to be unit-demand and sellers to be unit-supply each bringing exactly one identical item to the market, hence $k = m$. In contrast to matroid double auctions, we now work in a setting with a knapsack constraint. That is, each of the n buyers has a publicly known weight $w_i \in [0, 1]$. The set of buyers $A_{\mathcal{B}}$ who are allocated an item after our mechanism needs to satisfy $\sum_{i \in A_{\mathcal{B}}} w_i \leq 1$. Notation is simplified by interpreting $\mathbf{v}$ as the $|\mathcal{B} \cup \mathcal{S}|$-dimensional vector with non-negative real entries in which each entry v_i corresponds to the value of an agent for being allocated an item. Also, we denote by j the seller as well as the corresponding item.

In this section, we start by a description of the mechanism and the pricing strategy in Section 5.6.1. Afterwards, we discuss properties as DSBB, DSIC, IR and the approximation guarantee in Section 5.6.2.

5.6.1 The Mechanism

We start by a restriction to the case of $w_i \leq \frac{1}{2}$ for all buyers $i \in \mathcal{B}$. The general case will be discussed later.

Algorithm 10: Mechanism for Knapsack Double Auctions with Strong Budget-Balance

Result: Set $A = A_{\mathcal{B}} \cup A_{\mathcal{S}}$ of agents to get an item with $A_{\mathcal{B}} \subseteq \mathcal{B}$, $\sum_{i \in A_{\mathcal{B}}} w_i \leq 1$, $A_{\mathcal{S}} \subseteq \mathcal{S}$ and $|A| = |\mathcal{S}|$

$A \longleftarrow \emptyset$; $W \longleftarrow 0$; $i \longleftarrow 1$; $j \longleftarrow 1$
while $i \leq n$ and $j \leq k$ **do**
  if $W + w_i^* > 1$ **then**
    $i \longleftarrow i + 1$
  if $W + w_i^* \leq 1$ **then**
    if $v_j \geq p_i$ **then**
      $A \longleftarrow A \cup \{j\}$; $W \longleftarrow W + w_i^*$; $j \longleftarrow j + 1$
    if $v_j < p_i$ **then**
      if $v_i \geq p_i$ **then**
        $A \longleftarrow A \cup \{i\}$; $W \longleftarrow W + w_i^*$; $j \longleftarrow j + 1$
        transfer item from seller j to buyer i for price p_i
      $i \longleftarrow i + 1$
return $A \cup \{j' \in \mathcal{S} : j \leq j' \leq k\}$

We state our mechanism in Algorithm 10 and give a quick description: We sort buyers in a way such that $w_1 \geq w_2 \geq \cdots \geq w_n$ and compute artificial weights w_i^* for any buyer via $w_i^* := \max\left(w_i; \frac{1}{k}\right)$. Now, let the buyer-specific price be

$$p_i := \frac{2}{9} \cdot w_i^* \cdot \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] \quad ,$$

where $\mathrm{OPT}(\tilde{\mathbf{v}})$ denotes the optimal allocation of all items among all agents such that the set of selected buyers satisfies the knapsack constraint. We choose $\tilde{\mathbf{v}}$ to be drawn independently from the same distribution as $\mathbf{v}$. Further, we initialize $W = 0$ which will be our variable controlling feasibility with respect to w_i^*. In particular, if for some buyer i we have $W + w_i^* > 1$, we will not consider buyer i for a trade. In the other case where buyer i's artificial weight w_i^* can feasibly be added to W, we first ask the current seller j if she wants to keep or try selling the item for price p_i. If she considers selling, we ask buyer i if she wants to purchase the item.

Feasibility Considerations

We need to compute a feasible allocation A. In other words, the set $A_{\mathcal{B}} = A \cap \mathcal{B}$ needs to be feasible with respect to the knapsack constraint. In our mechanism, we instead compute an allocation with respect to the artificial weights w_i^*. To see that this is also feasible with respect to the initial weights w_i, observe that we always ensure $W + w_i^* \leq 1$ for any buyer i to whom we propose a trade. Since $w_i \leq w_i^*$ for any buyer i and we add w_i^* to W any time an item is irrevocably allocated, we ensure $\sum_{i \in A_{\mathcal{B}}} w_i \leq \sum_{i \in A_{\mathcal{B}}} w_i^* \leq 1$. Further, every time an item is allocated, we add some w_i^* to W. Since any $w_i^* \geq \frac{1}{k}$, we do never allocate more than k items in total.

5.6.2 Properties of Our Mechanism

Theorem 5.6.1. *Let all buyers' weights be no larger than half of the total capacity. Then, Mechanism 10 for knapsack double auctions is DSBB, DSIC and IR for all buyers and sellers and $\frac{1}{9}$-competitive with respect to the optimal social welfare when ordering buyers such that $w_1 \geq w_2 \geq \cdots \geq w_n$.*

Using this theorem allows also to state a generalized version without restrictions on the weights.

Theorem 5.6.2. *There is a mechanism for knapsack double auctions which is DSBB, DSIC and IR for all buyers and sellers and $\frac{1}{12}$-competitive with respect to the optimal social welfare.*

We now give a proof of Theorem 5.6.1. The generalized version is discussed below.

Proof of Theorem 5.6.1. Again, the arguments for DSBB, IR and DSIC for buyers follow in similar ways as in the previous sections. Concerning DSIC for sellers, note that we sorted buyers by non-increasing weight. Hence, the price for trades is non-increasing in the ongoing process. As a consequence, as a seller, you would like to sell your item as early as possible (if you want to sell it at all). Therefore, reporting a lower valuation might end in a trade at some price lower than your actual value. On the other hand, reporting a higher valuation may block a possibly beneficial trade. Overall, misreporting does not increase the seller's utility compared to truth-telling.

Again, denote by $A(\mathbf{v})$ the set of agents who receive an item A under valuation profile $\mathbf{v}$. Also W depends on $\mathbf{v}$, so in the same way we denote by $W(\mathbf{v})$ the value of W under valuation profile $\mathbf{v}$. We want to compare $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(A(\mathbf{v}))\right]$ to $\mathbf{E}_{\mathbf{v}}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right]$. To this end, again, we split the welfare of our algorithm into two parts, the base value and the surplus, and bound each quantity separately. The base value is defined as follows: let agent i receive an item in our mechanism, i.e. $i \in A$. Any buyer who gets an item has paid her buyer-specific price for an item. Any seller who decided to keep her item was asked to keep it for some buyer-specific price. The part of any agent's value which is below this price is denoted the base value. The surplus is the part of any agent's value above this threshold if it exists, otherwise it is zero.

Base Value. Our base part of the social welfare is defined via the prices. Summing over all agents in $A(\mathbf{v})$, we can compute the following. In particular, we sum over all prices for which either a buyer purchased an item or a seller irrevocably kept it:

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in A(\mathbf{v})} p_i\right] &= \frac{2}{9}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right]\cdot\mathbf{E}_{\mathbf{v}}\left[W(\mathbf{v})\right] \\ &\geq \frac{2}{9}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right]\cdot\frac{1}{2}\mathbf{Pr}_{\mathbf{v}}\left[W(\mathbf{v})\geq\frac{1}{2}\right] \quad .\end{aligned}$$

Surplus. We consider buyers and sellers separately and combine their respective contributions to the surplus afterwards.

Sellers: We observe that a seller j might be matched to some buyer i in the mechanism. Denote by i_j the first buyer that seller j is matched to and let $i_j =\perp$ and $w^*_{i_j} = 0$ if seller j is never matched to a buyer. Note that this initial matching is independent of seller j's value. Further, prices are only non-increasing in the ongoing process. Thus, we use that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed combined with linearity of expectation to get

$$\begin{aligned}\mathbf{E}_{\mathbf{v}}\left[\mathrm{surplus}_j\right] &\geq \mathbf{E}_{\mathbf{v}}\left[\left(v_j - p_{i_j}\right)^+\right] \\ &\geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v_j - p_{i_j}\right)^+\cdot\mathbb{1}_{W\left((v'_j,\mathbf{v}_{-j})\right)\leq\frac{1}{2}}\cdot\mathbb{1}_{j\in\mathrm{OPT}\left((v_j,\mathbf{v}'_{-j})\right)}\right] \\ &= \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v'_j - p_{i_j}\right)^+\cdot\mathbb{1}_{W(\mathbf{v})\leq\frac{1}{2}}\cdot\mathbb{1}_{j\in\mathrm{OPT}(\mathbf{v}')}\right] \quad .\end{aligned}$$

Buyers: When considering the buyers, we first argue under which circumstances buyer i gets an item in our mechanism. First, buyer i's value needs to exceed her price p_i. Second, there needs to be a time step t such that $W_t + w^*_i \leq 1$, where W_t denotes the value of W at time step t. Third, there needs to exist a seller j such that $v_j \leq p_i$ as otherwise, there will be no item available for buyer i. We make use of the following observation: Buyer i is never asked to purchase an item until we either can offer her an item for price p_i or buyer i becomes infeasible with respect to W and w^*_i. Therefore, everything happening before this event is independent of buyer i's value. Hence, when considering the value of W on a hallucinated valuation profile v'_i drawn independently from the same distribution as v_i, we get the following: if $W\left((v'_i,\mathbf{v}_{-i})\right) \leq 1/2$, i.e. the value of W on valuation profile $(v'_i,\mathbf{v}_{-i})$ is at most $1/2$ after running the mechanism, then buyer i could be feasibly added at the end of the mechanism. As W is only non-decreasing, buyer i could have also been feasibly added at time t. Using this, we can bound the surplus of buyer i via

$$\mathrm{surplus}_i \geq (v_i - p_i)^+\cdot\mathbb{1}_{W\left((v'_i,\mathbf{v}_{-i})\right)\leq\frac{1}{2}} \geq (v_i - p_i)^+\cdot\mathbb{1}_{W\left((v'_i,\mathbf{v}_{-i})\right)\leq\frac{1}{2}}\cdot\mathbb{1}_{i\in\mathrm{OPT}\left((v_i,\mathbf{v}'_{-i})\right)} \quad .$$

Again, using linearity of expectation as well as exploiting that $\mathbf{v}'$ and $\mathbf{v}$ are independent and identically distributed, we get

$$\mathbf{E}_{\mathbf{v}}\left[\mathrm{surplus}_i\right] \geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v'_i - p_i)^+\cdot\mathbb{1}_{W(\mathbf{v})\leq\frac{1}{2}}\cdot\mathbb{1}_{i\in\mathrm{OPT}(\mathbf{v}')}\right] \quad .$$

Combination. Summing over all buyers and sellers, we can combine the two bounds:

$$\begin{aligned}
&\sum_{i \in \mathcal{B} \cup \mathcal{S}} \mathbf{E}_{\mathbf{v}}\left[\text{surplus}_i\right] \\
&\geq \sum_{i \in \mathcal{B}} \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v'_i - p_i)^+ \cdot \mathbb{1}_{W(\mathbf{v}) \leq \frac{1}{2}} \cdot \mathbb{1}_{i \in \mathrm{OPT}(\mathbf{v}')}\right] + \sum_{j \in \mathcal{S}} \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v'_j - p_{i_j}\right)^+ \cdot \mathbb{1}_{W(\mathbf{v}) \leq \frac{1}{2}} \cdot \mathbb{1}_{j \in \mathrm{OPT}(\mathbf{v}')}\right] \\
&= \mathbf{Pr}_{\mathbf{v}}\left[W(\mathbf{v}) \leq \frac{1}{2}\right] \cdot \left(\mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{B}} (v'_i - p_i)^+\right] + \mathbf{E}_{\mathbf{v}'}\left[\sum_{j \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{S}} \left(v'_j - p_{i_j}\right)^+\right]\right) \\
&\geq \mathbf{Pr}_{\mathbf{v}}\left[W(\mathbf{v}) \leq \frac{1}{2}\right] \cdot \left(\mathbf{E}_{\mathbf{v}'}\left[\mathbf{v}'\left(\mathrm{OPT}(\mathbf{v}')\right)\right] - \mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{B}} p_i\right] - \mathbf{E}_{\mathbf{v}'}\left[\sum_{j \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{S}} p_{i_j}\right]\right) \quad .
\end{aligned}$$

Again, for the first equality we use that $\mathbf{v}$ and $\mathbf{v}'$ are independent and the respective terms each only depend on one of the two. In order to bound the sums of prices, we note that also OPT is restricted with a total capacity of one as well as OPT can also not allocate more than k items, so

$$\sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{B}} w^*_i \leq \sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{B}} w_i + \sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{B}} \frac{1}{k} \leq 1 + 1 = 2 \quad .$$

In a similar way, we can bound the sum for the sellers $\sum_{i \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{S}} w^*_{i_j}$. First, it is crucial that we only go to the next seller in our mechanism once the previous seller has either sold or kept her item. Second, note that buyers are sorted such that weights are non-increasing. As a consequence, we can charge the weight of buyer i_j always to the previous seller, except for the first buyer-seller-pair. Still, the weight of the first buyer is upper bounded by half. Using that we only make an offer if the buyer is feasible with respect to the current status of W, we obtain $\sum_{j \in \mathrm{OPT}(\mathbf{v}') \cap \mathcal{S}} w^*_{i_j} \leq \sum_{j \in \mathcal{S}} w^*_{i_j} \leq 2$.

Therefore, the overall surplus fulfills

$$\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathcal{B} \cup \mathcal{S}} \text{surplus}_i\right] \geq \mathbf{Pr}_{\mathbf{v}}\left[W(\mathbf{v}) \leq \frac{1}{2}\right] \cdot \left(1 - 4 \cdot \frac{2}{9}\right) \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] \quad .$$

Summing the base value and the surplus proves our claim as we can exploit that $\mathbf{v}$, $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed.

□

In order to extend this to the general case in Theorem 5.6.2 when $w_i \in [0, 1]$ instead of $w_i \leq 1/2$, we split the set of buyers in high- and low-weighted ones and use our constructed mechanism on the latter — for the former, we use the DSBB-mechanism for matroids from Section 5.3. Estimating the better of the two mechanisms upfront and running it allows to formulate Theorem 5.6.2.

To be a bit more precise, high-weights buyers are the ones with $w_i > \frac{1}{2}$, low-weighted ones satisfy $w_i \leq \frac{1}{2}$. Observe that in an instance of high-weighted buyers, we can allocate at most one item which corresponds to a 1-uniform matroid constraint over the set of buyers. Concerning the use of the mechanism for matroid double auctions, note that we do not need to insist on a offline order of buyers now. We can rather fix the arrival sequence of buyers beforehand as we consider the 1-uniform matroid over all buyers. This implies that all buyers face the same take-it-or-leave-it buyer-specific contribution to the prices and hence allow easier arguments concerning the properties of the mechanism from Section 5.3.

5.7 Knapsack Double Auctions with Weak Budget Balance and Online Adversarial Arrival

Next, we would like to argue that when allowing the mechanism to be only weakly budget balanced, we can improve the approximation ratio while even allowing a worst-case trading order. The setting is equivalent to the one of Section 5.6. Each of the n unit-demand buyers has a weight $w_i \in [0,1]$ with the constraint that $\sum_{i \in A_{\mathcal{B}}} w_i \leq 1$. In contrast, our mechanism can handle online adversarial arrival orders of agents, even with an (adaptive) adversary specifying the order.

The Mechanism

As in Section 5.6, we first restrict weights to the case of $w_i \leq \frac{1}{2}$ for all buyers $i \in \mathcal{B}$. The generalization to weights in $w_i \in [0,1]$ is equivalent to the reasoning from Section 5.6.

Algorithm 11: Mechanism for Knapsack Double Auctions with Weak Budget Balance

Result: Set $A = A_{\mathcal{B}} \cup A_{\mathcal{S}}$ of agents to get an item with $A_{\mathcal{B}} \subseteq \mathcal{B}$, $\sum_{i \in A_{\mathcal{B}}} w_i \leq 1$, $A_{\mathcal{S}} \subseteq \mathcal{S}$ and $|A| = |\mathcal{S}|$

$A \longleftarrow \emptyset$; $M_{\text{SELL}} \longleftarrow \emptyset$
for $j \in \mathcal{S}$ **do**
  if $v_j \geq p_j$ **then**
    $A \longleftarrow A \cup \{j\}$
  if $v_j < p_j$ **then**
    $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \cup \{j\}$
for $i \in \mathcal{B}$ **do**
  if $\sum_{i' \in A} w^*_{i'} \leq 1 - w^*_i$ **then**
    if $v_i \geq p_i$ **then**
      $A \longleftarrow A \cup \{i\}$
      pick one $j \in M_{\text{SELL}}$, transfer item from j to i, i pays p_i, j receives p_j
      $M_{\text{SELL}} \longleftarrow M_{\text{SELL}} \setminus \{j\}$
$A \longleftarrow A \cup M_{\text{SELL}}$

Algorithm 11 works as follows: Compute artificial weights w^*_i for all buyers as $w^*_i := \max\left(w_i; \frac{1}{k}\right)$ and for all sellers as $w^*_j := \frac{1}{k}$. For any agent, we set the agent-specific price to be

$$p_i := \frac{2}{5} \cdot w^*_i \cdot \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] \quad ,$$

where $\mathrm{OPT}(\tilde{\mathbf{v}})$ denotes the optimal allocation of all items among all agents such that the set of selected buyers satisfies the knapsack constraint. Now, we first go through all sellers asking each if she wants to keep or try selling the item if we might pay an amount of p_j to her later-on. Afterwards, we go through all buyers, asking each of them if she wants to purchase an item for price p_i if buyer i can be feasibly added to the chosen set of agents with respect to the artificial weights w^*_i.

Feasibility Considerations

Arguing about the feasibility of our solution, we can proceed similar to Section 5.6, as we again compute a feasible allocation with respect to the artificial weights w_i^*. To see that this is also feasible with respect to the initial weights w_i, note that we ensure $1 \geq \sum_{i \in A} w_i^* \geq \sum_{i \in A_{\mathcal{B}}} w_i^* \geq \sum_{i \in A_{\mathcal{B}}} w_i$ throughout our mechanism. Further, we need to ensure that we do not allocate more than k items in total. This is mirrored by the fact that $w_i^* \geq \frac{1}{k}$ for any agent i and, as we do not allocate items if $\sum_{i \in A} w_i^* > 1$, we get that $|A| = k \cdot \sum_{i \in A} \frac{1}{k} \leq k \cdot \sum_{i \in A} w_i^* \leq k$.

Properties of Our Mechanism

Theorem 5.7.1. *The mechanism for knapsack double auctions is DWBB, DSIC and IR for all buyers and sellers and $\frac{1}{5}$-competitive with respect to the optimal social welfare if all buyers' weights are no larger than half of the total capacity.*

Proof. Our mechanism is DWBB, as by construction, the mechanism consists of bilateral trades where an item is traded from one seller to one buyer. For any seller j we have that $w_j^* \leq w_i^*$ for all buyers i, and hence $p_j \leq p_i$ for any buyer-seller pair (i, j). The buyer pays p_i to the mechanism and the seller receives p_j, so we get DWBB. IR follows naturally, DSIC from the fact that any agent is asked at most once in our mechanism.

The set of agents who receive an item A depends on $\mathbf{v}$, so we denote by $A(\mathbf{v})$ the set A under valuation profile $\mathbf{v}$. We want to compare $\mathbf{E_v}\left[\mathbf{v}(A(\mathbf{v}))\right]$ to $\mathbf{E_v}\left[\mathbf{v}(\mathrm{OPT}(\mathbf{v}))\right]$. To this end, again, we split the welfare of our algorithm into two parts, the base value and the surplus, and bound each quantity separately.

Base Value. Our base part of the social welfare is defined via the prices. Summing over all agents in $A(\mathbf{v})$, we can compute the following:

$$\begin{aligned}\mathbf{E_v}\left[\sum_{i \in A(\mathbf{v})} p_i\right] &= \frac{2}{5}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] \cdot \mathbf{E_v}\left[\sum_{i \in A(\mathbf{v})} w_i^*\right] \\ &\geq \frac{2}{5}\mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\mathrm{OPT}(\tilde{\mathbf{v}})\right)\right] \cdot \frac{1}{2}\mathbf{Pr_v}\left[\sum_{i \in A(\mathbf{v})} w_i^* \geq \frac{1}{2}\right] \quad .\end{aligned}$$

Surplus. We consider buyers and sellers separately and combine their respective contributions to the surplus afterwards.

Sellers: Note that any seller whose value exceeds her corresponding price can keep the item, so

$$\mathrm{surplus}_j \geq (v_j - p_j)^+ \geq (v_j - p_j)^+ \cdot \mathbb{1}_{\sum_{i' \in A\left((v_j', \mathbf{v}_{-j})\right)} w_{i'}^* \leq \frac{1}{2}} \cdot \mathbb{1}_{j \in \mathrm{OPT}\left((v_j, \mathbf{v}'_{-j})\right)} \quad .$$

Now, we use that $\mathbf{v}$ and $\mathbf{v}'$ are independent and identically distributed combined with linearity of expectation to get

$$\mathbf{E_v}\left[\mathrm{surplus}_j\right] \geq \mathbf{E_v}\left[(v_j - p_j)^+\right] \geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[\left(v_j' - p_j\right)^+ \cdot \mathbb{1}_{\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}} \cdot \mathbb{1}_{j \in \mathrm{OPT}(\mathbf{v}')}\right] \quad .$$

Buyers: Concerning the buyers, note that buyer i gets an item if buyer i's value exceeds her price and if the sum of the weights of agents in A does allow i to be added. That is, denote by A_{i-1} the set of accepted agents A after processing buyer $i-1$. Then, we ensure $\sum_{i' \in A_{i-1}(\mathbf{v})} w_{i'}^* \leq 1 - w_i^*$. Note that A_{i-1} does not depend on buyer i, so in particular $A_{i-1}(\mathbf{v}) = A_{i-1}\left((v_i', \mathbf{v}_{-i})\right)$. Further, an even stronger condition is that $\sum_{i' \in A\left((v_i', \mathbf{v}_{-i})\right)} w_{i'}^* \leq \frac{1}{2}$ as $w_i \leq \frac{1}{2}$ by assumption and we can assume that $k \geq 2$ (the case $k = 1$ can be easily understood via our results in the previous sections with improved approximation guarantees). Therefore, we can bound

$$\begin{aligned} \text{surplus}_i &\geq (v_i - p_i)^+ \cdot \mathbb{1}_{\sum_{i' \in A\left((v_i', \mathbf{v}_{-i})\right)} w_{i'}^* \leq \frac{1}{2}} \\ &\geq (v_i - p_i)^+ \cdot \mathbb{1}_{\sum_{i' \in A\left((v_i', \mathbf{v}_{-i})\right)} w_{i'}^* \leq \frac{1}{2}} \cdot \mathbb{1}_{i \in \text{OPT}\left((v_i, \mathbf{v}_{-i}')\right)} \quad . \end{aligned}$$

Again, using linearity of expectation as well as choosing v_i' and v_i to be independent and identically distributed, we get

$$\mathbf{E}_{\mathbf{v}}\left[\text{surplus}_i\right] \geq \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v_i' - p_i)^+ \cdot \mathbb{1}_{\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}} \cdot \mathbb{1}_{i \in \text{OPT}(\mathbf{v}')}\right] \quad .$$

Combination: Summing over all buyers and sellers, we can combine the two bounds:

$$\begin{aligned} \mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathcal{B} \cup \mathcal{S}} \text{surplus}_i\right] &\geq \sum_{i \in \mathcal{B} \cup \mathcal{S}} \mathbf{E}_{\mathbf{v},\mathbf{v}'}\left[(v_i' - p_i)^+ \cdot \mathbb{1}_{\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}} \cdot \mathbb{1}_{i \in \text{OPT}(\mathbf{v}')}\right] \\ &= \mathbf{Pr}_{\mathbf{v}}\left[\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}\right] \cdot \mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \text{OPT}(\mathbf{v}')} (v_i' - p_i)^+\right] \\ &\geq \mathbf{Pr}_{\mathbf{v}}\left[\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}\right] \cdot \left(\mathbf{E}_{\mathbf{v}'}\left[\mathbf{v}'\left(\text{OPT}(\mathbf{v}')\right)\right] - \mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \text{OPT}(\mathbf{v}')} p_i\right]\right) \quad . \end{aligned}$$

To get the equality, note that $\mathbf{v}$ and $\mathbf{v}'$ are independent and the respective terms each only depend on one of the two.

Now, in order to bound the sum of prices, we calculate

$$\mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \text{OPT}(\mathbf{v}')} p_i\right] = \frac{2}{5} \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\text{OPT}(\tilde{\mathbf{v}})\right)\right] \cdot \mathbf{E}_{\mathbf{v}'}\left[\sum_{i \in \text{OPT}(\mathbf{v}')} w_i^*\right] \quad .$$

We can bound $w_i^* = \max(w_i, \frac{1}{k}) \leq w_i + \frac{1}{k}$ on the buyers' side as well as $w_i^* = \frac{1}{k}$ for all sellers to get

$$\sum_{i \in \text{OPT}(\mathbf{v}')} w_i^* \leq \sum_{i \in \text{OPT}(\mathbf{v}') \cap \mathcal{B}} w_i + \sum_{i \in \text{OPT}(\mathbf{v}')} \frac{1}{k} \leq 1 + 1 = 2$$

as the sum over the weights of all buyers in any feasible allocation is upper bounded by one and further, we cannot allocate more than k items in any feasible allocation, so $|\text{OPT}(\mathbf{v}')| \leq k$.

Therefore, we can bound the overall surplus by

$$\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in \mathcal{B} \cup \mathcal{S}} \text{surplus}_i\right] \geq \mathbf{Pr}_{\mathbf{v}}\left[\sum_{i' \in A(\mathbf{v})} w_{i'}^* \leq \frac{1}{2}\right] \cdot \left(1 - \frac{4}{5}\right) \mathbf{E}_{\tilde{\mathbf{v}}}\left[\tilde{\mathbf{v}}\left(\text{OPT}(\tilde{\mathbf{v}})\right)\right] \quad .$$

Summing the base value and the surplus proves our claim as we can exploit that $\mathbf{v}$, $\mathbf{v}'$ and $\tilde{\mathbf{v}}$ are independent and identically distributed. □

In order to extend this to the general case, the only difference to Section 5.6 is to use the mechanism from Section 5.4 for the high-weighted items. Again, estimating the expected welfare of each of the two options and selecting the better one allows us to state the following theorem.

Theorem 5.7.2. *There is a mechanism for knapsack double auctions which is DWBB, DSIC and IR for all buyers and sellers and* $\frac{1}{7}$*-competitive with respect to the optimal social welfare for any adversarial online arrival order of agents.*

Chapter 6

Asymptotically Optimal Welfare of Posted Pricing with MHR Distributions

For the design of the algorithms in the previous chapters, a crucial assumption is that distributions are independent across buyers (or more general, for Chapter 5, buyers and sellers), but are not required to be identical. Still, Example 1.1.1 shows that for non-identical distributions, competitive ratios better than half are impossible. This naturally raises the question if insisting on stronger assumptions concerning the distributions allows to improve guarantees. In this chapter, we follow this path and shift our perspective to independent *and* identically distributed (i.i.d.) valuation functions.

As before, we assume that the valuation functions $v_1, \ldots, v_n$ are unknown a priori. Still, in this chapter, all valuation functions are drawn independently from the *same*, publicly known distribution. Our algorithms under consideration are pricing based: For every item, we can either set a static item price or change the prices dynamically over time. Buyers arrive one-by-one and each of them chooses the set of items that maximizes her utility given the current prices among the remaining items. Static prices have the advantage that they are easier to explain and thus give easier mechanisms. However, dynamic prices can yield both higher welfare and revenue because they can be adapted to the remaining supply and the remaining number of buyers to appear.

We would like to understand which fraction of the expected offline optimal welfare $\mathbf{E}[\mathrm{OPT}]$ can be guaranteed by posted-pricing mechanisms when distributions are identical. The case of a single item is well understood: For a static price and a single item, the best competitive ratio is $\zeta = 1 - \frac{1}{e} \approx 0.63$ [Correa et al., 2017, Ehsani et al., 2018]; for dynamic pricing and a single item, it is $\zeta \approx 0.745$ [Abolhassani et al., 2017, Correa et al., 2017]. There are a number of extensions of these results to multiple items, many of which are $O(1)$-competitive. The competitive ratios of $\zeta = 1 - \frac{1}{e} \approx 0.63$ and $\zeta \approx 0.745$ are optimal in the sense that there are distributions and choices of n such that no better guarantee can be obtained. Importantly, they are still tight when imposing a lower bound on n. That is, even for large n, there is a distribution such that if all values are drawn from this distribution the respective bound cannot be beaten.

Distributions with Monotone Hazard Rate. Faced by these tight guarantees, in this chapter, we strengthen previous results by restricting the class of distributions to ones with monotone hazard rate. The single-item case is defined as follows. Consider a probability distribution on the reals with probability density function (PDF) f and cumulative distribution function (CDF) F, its hazard rate h is defined by $h(x) = \frac{f(x)}{1-F(x)}$ for x with $F(x) < 1$. It has a *monotone hazard rate* (MHR) — more precisely, increasing hazard rate — if h is a non-decreasing function. It has become a common and well-studied approach to model buyer preferences by MHR distributions. One of the reasons is that many standard distributions exhibit a monotone hazard rate such as, for example, uniform, normal, exponential and logistic distributions[1]. Furthermore, the monotone hazard rate of distributions is also preserved under certain operations; for example, order statistics of MHR distributions also have an MHR distribution. Additionally, every MHR distribution is regular in the sense that its virtual value function [Chawla et al., 2007, Myerson, 1981] is increasing.

Multiple Items with MHR Distributions. Our assumption for multiple items will be as follows: We consider *independent item valuations*, i.e. $v_{i,j} \sim \mathcal{D}^{(j)}$ is an independent draw from a distribution $\mathcal{D}^{(j)}$. In other words, the value of item j is independent of the value of item j' and both values are drawn from (possibly different) MHR distributions as defined above.

As we will see, asymptotically optimal welfare can be guaranteed. That is, if n grows large, the social welfare when suitably choosing prices is within a $1 - o(1)$ factor of the optimum, where the $o(1)$ term is independent of the distribution as long as its marginals satisfy the MHR property. Stated differently, there is a sequence $(\zeta_n)_{n\in\mathbb{N}}$ with $\zeta_n \to 1$ for $n \to \infty$ such that for every number of buyers n there exists a posted-prices mechanism that takes any distribution with MHR marginals as input and guarantees $\mathbf{E}[\text{ALG}] \geq \zeta_n \cdot \mathbf{E}[\text{OPT}]$.

The technically most interesting contribution is the one on dynamic pricing when values are independent across items. The idea is to set prices so that the offline optimum is mimicked. If item j is allocated in the optimal allocation with probability q_j, then we would like it to be sold in every step with an ex-ante probability of $\frac{q_j}{n}$. However, analyzing such a selling process is still difficult because items are incomparable and bounds for MHR distributions cannot be applied directly to draws from multiple distributions, which are not necessarily identical. To bypass this problem, we introduce a reduction that allows us to view item valuations not only as independent but also as identically distributed. To this end, we compare the selling process of our mechanism to a hypothetical setting, in which buyers do not make their decisions based on the actual utility but in quantile space. Only afterwards, we can apply a concentration bound due to the MHR property.

Additional Related Work on MHR Distributions

There are surprisingly few results on pricing and Prophet Inequalities that derive better guarantees by imposing additional assumptions on the distribution. Babaioff et al. [2017] consider the problem of maximizing revenue when selling a single item to one of n buyers drawn i.i.d. from an *unknown* MHR distribution with bounded support $[1, h]$. If n is large enough compared to h, they get a constant-factor approximation to the optimal revenue using dynamic posted prices. Note that in contrast, in this chapter, we assume to know the underlying distributions perfectly. Giannakopoulos and Zhu [2018] consider revenue

[1] For a much more extensive list, see Rinne [2014].

maximization in the single-item setting with valuations drawn independently from the same MHR distribution. They show that achieving asymptotically optimal revenue is possible by offering the item for the same static price to all bidders. More precisely, one of their main results is that one gets within a factor of $1 - O\left(\frac{\ln \ln n}{\ln n}\right)$. While they claim this result is "essentially tight", we will see that the best factor is indeed $1 - \Theta\left(\frac{\ln \ln \ln n}{\ln n}\right)$ because it is a special case of the results in Section 6.3 and Section 6.4. Jin et al. [2019] also consider revenue maximization in the single-item setting with identical and independent MHR values but in a non-asymptotic sense, providing a bound for every n.

Chapter Organization and Remarks

This chapter is based on *Asymptotically Optimal Welfare of Posted Pricing for Multiple Items with MHR Distributions* [Braun et al., 2021], which is joint work with Matthias Buttkus and Thomas Kesselheim. More detailed bibliographic notes can be found in Section 1.5.

In Section 6.1, we give some preliminaries as well as useful properties of MHR distributions which we will use in the technical sections. Afterwards, we discuss our algorithms for independent item values in Section 6.2, followed by providing impossibility results in Section 6.3. Also, an extension of the techniques to subadditive valuations is possible, as we will see in Section 6.4.

6.1 Preliminaries

For the main results in this chapter, we consider a setting of n buyers $[n]$ and a set M of m items. Every buyer has a unit-demand valuation function. The functions $v_1, \dots, v_n$ are unknown a priori but all drawn independently from the same, publicly known distribution $\mathcal{D}$. In this chapter, we abuse notation and denote by v_i the valuation function as well as the vector $v_i = (v_{i,1}, \dots, v_{i,m})$ with the corresponding value per item as entries.

Let $\mathcal{D}^{(j)}$ be the marginal distribution of $v_{i,j}$, which is the value of a buyer for being allocated item j. We assume that $\mathcal{D}^{(j)}$ is a continuous, real, non-negative distribution with monotone hazard rate. That is, let F_j be the cumulative distribution function of $\mathcal{D}^{(j)}$ and f_j its probability density function. The distribution's hazard rate is defined as $h_j(x) = f_j(x)/(1 - F_j(x))$ for all x such that $F_j(x) < 1$. We assume a *monotone hazard rate*, which means that h_j is a non-decreasing function. Equivalently, we can require $x \mapsto \log(1 - F_j(x))$ to be a concave function.

We work in an online setting: Buyers arrive one-by-one and have the choice between all items which have not been allocated so far. Let $M^{(i)}$ denote the set of available items as buyer i arrives. The mechanism presents buyer i a menu of prices $p_j^{(i)}$ for all items $j \in M^{(i)}$. The buyer then picks the item $j_i \in M^{(i)}$ which maximizes her *utility* $v_{i,j_i} - p_{j_i}^{(i)}$ if positive. Buyer i and item j_i are matched immediately and irrevocably. If buyer i has negative utility for all items $j \in M^{(i)}$, then buyer i does not buy any item and remains unmatched.

We denote by j_i the item allocated to buyer i (set $j_i = \perp$ if i remains unmatched in the algorithm). Our goal is to maximize the *expected social welfare* of the mechanism, given by $\mathbf{E}\left[\sum_{i=1}^n v_{i,j_i}\right] =: \mathbf{E}[\text{ALG}]$. We compare this quantity to the expected offline optimum, denoted by $\mathbf{E}\left[\sum_{i=1}^{n} v_{i,j_i^*}\right] =: \mathbf{E}[\mathrm{OPT}]$ and aim for a desirable competitive ratio.

6.1.1 Useful Properties of MHR Distributions

Before we start with the main technical part, we first state a few useful properties of MHR distributions. The first important property of MHR distributions is a useful lemma from Babaioff et al. [2017]. It allows to compare the expectation of the maximum of n and $n' \leq n$ draws from independent and identically distributed random variables, if the distribution has a monotone hazard rate.

Lemma 6.1.1 (Lemma 5.3 in Babaioff et al. [2017])**.** *Consider a collection $(Y_i)_i$ of independent and identically distributed random variables, where their distribution has a monotone hazard rate. Then, for any $n' \leq n$, we have*

$$\frac{\mathbf{E}\left[\max_{i\in[n']} Y_i\right]}{\mathbf{E}\left[\max_{i\in[n]} Y_i\right]} \geq \frac{H_{n'}}{H_n} \geq \frac{\log n'}{\log n} \ .$$

In addition, we can reformulate a useful lemma from Giannakopoulos and Zhu [2018] which nicely complements the previous lemma. To this end, given any sequence of real numbers $(w_1, \dots, w_n)$, we let $w_{(k)}$ denote the k-th highest order statistic. That is, $w_{(k)}$ is the largest x such that there are at least k entries in $(w_1, \dots, w_n)$ whose value is at least x. For random numbers, denote its expectation by $\mathbf{E}\left[w_{(k)}\right] = \mu_k$. As a side remark, for MHR distributions, also order statistics $w_{(k)}$ are distributed according to an MHR distribution [Rinne, 2014].

Lemma 6.1.2. *For the expected k-th order statistic μ_k of n i.i.d. draws from an MHR distribution, we have that for all q and k such that $\exp(H_{k-1} - H_n) \leq q \leq 1$, it holds*

$$F^{-1}(1-q) \geq -\frac{\log(q)}{H_n - H_{k-1}}\mu_k \ .$$

Proof. Use Lemma 3 in Giannakopoulos and Zhu [2018] with $c = -\frac{\log(q)}{H_n - H_{j-1}}$ and apply the quantile function on both sides to prove the result. □

In addition to these two lemmas, we provide another useful bound for MHR distributions.

Lemma 6.1.3. *Let $z \in [0,1]$ and $\Gamma \in \mathbb{N}$. Further, let $\mathcal{D}$ be a distribution with monotone hazard rate with CDF F, let $X, (Y_i)_i \sim \mathcal{D}$ be independent and identically distributed. For $\max\left(1, \frac{1+\log(1/z)}{H_\Gamma}\right) \leq \alpha \leq \frac{1}{\Gamma \cdot z}$,we have*

$$\mathbf{E}\left[X \,\middle|\, X \geq F^{-1}(1-z)\right] \leq \alpha \cdot \mathbf{E}\left[\max_{i\in[\Gamma]} Y_i\right] \ .$$

Proof. First, we define by $g(y) = \frac{1-y}{f(F^{-1}(y))}$ the inverse of the hazard rate at point $F^{-1}(y)$ for $y \in [0,1]$. In other words, for the hazard rate $h(x) = \frac{f(x)}{1-F(x)}$, we have $h\left(F^{-1}(y)\right) = \frac{f(F^{-1}(y))}{1-F(F^{-1}(y))} = \frac{f(F^{-1}(y))}{1-y} = \frac{1}{g(y)}$. Observe that by the MHR property, $h(x)$ is non-decreasing and hence $g(y)$ is non-increasing.

Second, we aim for a suitable expression of $\mathbf{E}\left[X \mid X \geq F^{-1}(1-z)\right]$ which we can compute as follows:

$$\begin{aligned}\mathbf{E}\left[X \,\middle|\, X \geq F^{-1}(1-z)\right] - F^{-1}(1-z) &= \frac{1}{z}\int_{F^{-1}(1-z)}^{\infty} 1 - F(x)dx \\ &= \frac{1}{z}\int_{F^{-1}(1-z)}^{\infty} \frac{1-F(x)}{f(x)} f(x)dx \\ &= \frac{1}{z}\int_{F^{-1}(1-z)}^{\infty} \frac{1}{h(x)} f(x)dx \\ &= \frac{1}{z}\int_{1-z}^{1} g(y)dy \ .\end{aligned}$$

In addition, observe that we can write $F^{-1}(1-z) = \int_0^{1-z} \frac{g(y)}{1-y}dy$.

Third, note that we can calculate

$$\begin{aligned}\mathbf{E}\left[\max_{i\in[\Gamma]} Y_i\right] &= \int_0^{\infty} 1 - F^{\Gamma}(x)dx = \int_0^{\infty} \frac{1-F^{\Gamma}(x)}{f(x)} f(x)dx \\ &= \int_0^{\infty} \frac{1}{h(x)} \frac{1-F^{\Gamma}(x)}{1-F(x)} f(x)dx \\ &= \int_0^1 g(y)\frac{1-y^{\Gamma}}{1-y}dy \ .\end{aligned}$$

As a consequence, the claim of our lemma holds if and only if

$$\alpha\int_0^1 g(y)\frac{1-y^{\Gamma}}{1-y}dy - \int_0^{1-z}\frac{g(y)}{1-y}dy - \int_{1-z}^1 \frac{g(y)}{z}dy \geq 0 \ .$$

We can split the left-hand side in the following (possibly empty) integrals: First, split the first integral into two ranges from 0 to $1-z$ and the remainder starting from $1-z$. Then, combine the respective integrals over equal ranges and define a threshold $y^* = \min\left\{\sqrt[\Gamma]{1-\frac{1}{\alpha}}, 1-z\right\}$ as the point at which the sign of $\alpha\left(1-y^{\Gamma}\right)-1$ changes from positive to negative. This allows to rewrite the integrals of the left-hand side as

$$\int_0^{y^*} \frac{g(y)}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{y^*}^{1-z} \frac{g(y)}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{1-z}^1 g(y)\left(\alpha\sum_{i=0}^{\Gamma-1} y^i - \frac{1}{z}\right)dy \ .$$

Observe that $g(y)$ is non-increasing by the MHR property. Further, note that $\alpha\sum_{i=0}^{\Gamma-1} y^i - \frac{1}{z} \leq 0$ as $\alpha\Gamma \leq \frac{1}{z}$. Setting $c = g(y^*)$, we can compute

$$\begin{aligned}&\int_0^{y^*} \frac{g(y)}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{y^*}^{1-z} \frac{g(y)}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{1-z}^1 g(y)\left(\alpha\sum_{i=0}^{\Gamma-1} y^i - \frac{1}{z}\right)dy \\ \geq &\int_0^{y^*} \frac{c}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{y^*}^{1-z} \frac{c}{1-y}\left(\alpha\left(1-y^{\Gamma}\right)-1\right)dy + \int_{1-z}^1 c\left(\alpha\sum_{i=0}^{\Gamma-1} y^i - \frac{1}{z}\right)dy \\ = &\ c\alpha\int_0^1\sum_{i=0}^{\Gamma-1} y^i dy - c\int_0^{1-z}\frac{1}{1-y}dy - c\int_{1-z}^1 \frac{1}{z}dy = c\alpha H_{\Gamma} - c\left(1+\ln\left(\frac{1}{z}\right)\right) \ .\end{aligned}$$

By our choice of $\alpha \geq \frac{1+\ln\left(\frac{1}{z}\right)}{H_{\Gamma}}$, observe that $\alpha H_{\Gamma} - \left(1+\ln\left(\frac{1}{z}\right)\right) \geq 0$. So the integral is non-negative for any $c \geq 0$. As a consequence, the claim holds. □

These will be the key lemmas for exploiting the MHR property of the distributions in the next sections.

6.2 Asymptotically Tight Bounds for Independent Item Values

In this section, we show how to derive bounds if the buyers' values are independent across items. That is, each $v_{i,j} \sim \mathcal{D}^{(j)}$ is drawn independently from a distribution with monotone hazard rate. This is a standard assumption when considering multiple items [Chawla et al., 2007, 2010]. As a consequence, the distribution over item values is a product distribution $v_i = (v_{i,1}, \ldots, v_{i,m}) \sim \mathcal{D} = \times_{j=1}^{m} \mathcal{D}^{(j)}$ for any $i \in [n]$ and every marginal $\mathcal{D}^{(j)}$ satisfies the MHR condition.

6.2.1 Dynamic prices

We first consider the case of dynamic pricing mechanisms. Without loss of generality, we can assume that $m \geq n$. If we have less items than buyers, i.e. $m < n$, we can add dummy items with value 0 to ensure $m = n$. Matching i to one of these dummy items in the mechanism then corresponds to leaving i unmatched. Observe that technically a point mass on 0 is not a MHR distribution. However, all relevant statements still apply.

Our mechanism is based on a pricing rule which balances the probability of selling a specific item. Let $M^{(i)}$ be the set of remaining items as buyer i arrives. We determine dynamic prices such that one item is allocated for sure in every step. Therefore, always $|M^{(i)}| = m - i + 1$. We can now define $q_j^{(i)}$ to be the probability that item j is allocated in the "remaining" offline optimum on $M^{(i)}$ and $n - i + 1$ buyers if $j \in M^{(i)}$ and 0 else. In other words, if $j \in M^{(i)}$, $q_j^{(i)}$ is the probability that item j is allocated in the offline optimum constrained to buyers $1, \ldots, i-1$ receiving the items from $M \setminus M^{(i)}$. The prices $(p_j^{(i)})_{j \in M^{(i)}}$ are now chosen such that buyer i buys item j with probability $\frac{q_j^{(i)}}{n-i+1}$ and one item is allocated for sure[2]. This allows us to state the following theorem.

Theorem 6.2.1. *The posted-prices mechanism with dynamic prices and independent item-valuations is* $1 - O\left(\frac{1}{\log n}\right)$*-competitive with respect to the expected offline optimal social welfare.*

The remainder of Section 6.2.1 is dedicated to the proof of Theorem 6.2.1. The roadmap will be as follows: We first introduce a quantile allocation rule in which we allocate items with respect to quantiles rather than allocating the item which maximizes utility. Then, in Lemma 6.2.3, we argue about the distribution of the value of buyer i for item j conditioned on allocating j to i in the quantile allocation rule. Using two observations on the probability of allocating item j in the remaining offline optimum allows to bound the expected contribution of item j to the social welfare. More precisely, we can guarantee to achieve a reasonable fraction of the contribution to the ex-ante relaxation of the offline optimum.

Proof. In order to bound the social welfare obtained by the posted-prices mechanism, we consider the following *quantile allocation rule*. For any $j \in M^{(i)}$ with $q_j^{(i)} > 0$, compute

[2] To see that such prices exist, we can use Theorem 4 in Banihashem et al. [2024] and argue in a similar way as we did in the proof of Theorem 4.0.1. In particular, we can use our Algorithm 12 introduced below as input to Theorem 4 of Banihashem et al. [2024] and get an algorithm as output which uses dynamic prices with our described properties.

the weighted quantile

$$R_j^{(i)} := F_j(v_{i,j})^{\frac{1}{q_j^{(i)}}}$$

and allocate buyer i the item j which maximizes $R_j^{(i)}$, formally stated in Algorithm 12.

Algorithm 12: Quantile Allocation Rule for Independent Item Values

1. $M^{(1)} \longleftarrow M$
2. **for** $i \in [n]$ **do**
3. Compute $F_j(v_{i,j})^{1/q_j^{(i)}}$ for any $j \in M^{(i)}$ with $q_j^{(i)} > 0$
4. Set j_i such that it attains $\max_{j \in M^{(i)}} F_j(v_{i,j})^{1/q_j^{(i)}}$ and remove j_i from $M^{(i+1)}$

Observe that by this definition for any i, any j and any $t \in [0,1]$,

$$\begin{aligned} \mathbf{Pr}\left[R_j^{(i)} \leq t\right] &= \mathbf{Pr}\left[F_j(v_{i,j}) \leq t^{q_j^{(i)}}\right] = \mathbf{Pr}\left[v_{i,j} \leq F_j^{-1}\left(t^{q_j^{(i)}}\right)\right] \\ &= F_j\left(F_j^{-1}\left(t^{q_j^{(i)}}\right)\right) = t^{q_j^{(i)}} \quad . \end{aligned} \tag{6.1}$$

To get some intuition, note that for $q_j^{(i)} = 1$, this is exactly the CDF of a random variable drawn from Unif$[0,1]$.

Next, we define indicator variables $X_{i,j}$ which are one if buyer i is allocated item j in the quantile allocation rule. Then, we can observe the following for the conditional probability of allocating item j to buyer i.

Observation 6.2.2. *It holds*

$$\mathbf{Pr}\left[X_{i,j} = 1 \;\middle|\; M^{(i)}\right] = \frac{q_j^{(i)}}{n-i+1} \quad .$$

Note that by this, conditioned on the set of available items $M^{(i)}$, the probability of allocating item j in step i via the quantile allocation rule is $q_j^{(i)}/n-i+1$, exactly as in the posted-prices mechanism.

Proof of Observation 6.2.2. We allocate item j in the quantile allocation rule if $R_j^{(i)} \geq R_{j'}^{(i)}$ for any $j' \in M^{(i)}$. For fixed $M^{(i)}$, also the values of $q_j^{(i)}$ are fixed. We can use this as well as the independence of the $v_{i,j}$ variables to compute:

$$\begin{aligned} \mathbf{Pr}\left[X_{i,j} = 1 \;\middle|\; M^{(i)}\right] &= \mathbf{Pr}\left[\max_{j' \neq j} R_{j'}^{(i)} \leq R_j^{(i)} \;\middle|\; M^{(i)}\right] \\ &= \int_0^1 \mathbf{Pr}\left[\max_{j' \neq j} R_{j'}^{(i)} \leq t \;\middle|\; M^{(i)}, R_j^{(i)} = t\right] q_j^{(i)} t^{q_j^{(i)}-1} dt \\ &= \int_0^1 \prod_{j' \neq j} \mathbf{Pr}\left[R_{j'}^{(i)} \leq t \;\middle|\; M^{(i)}, R_j^{(i)} = t\right] q_j^{(i)} t^{q_j^{(i)}-1} dt \\ &= \int_0^1 \left(\prod_{j' \neq j} t^{q_{j'}^{(i)}}\right) q_j^{(i)} t^{q_j^{(i)}-1} dt \\ &= q_j^{(i)} \int_0^1 t^{(n-i+1)-1} dt = \frac{q_j^{(i)}}{n-i+1} \quad , \end{aligned}$$

where we use that $\sum_{j \in M^{(i)}} q_j^{(i)} = n - i + 1$ for any value of i. □

Now, the crucial observation is that the expected contribution of any buyer to the social welfare in the posted-prices mechanism is at least as large as under the quantile allocation rule. To see this, fix buyer i and split buyer i's contribution to the social welfare into revenue and utility. Concerning revenue, note that in both cases the probability of selling any item j to buyer i is equal to $q_j^{(i)}/n-i+1$ and we allocate one item for sure. So, the expected revenue is identical. Further, since we maximize utility in the posted-prices mechanism, the achieved utility is always at least as large as the utility of the quantile allocation rule. So, overall, we get $\mathbf{E}[\mathrm{ALG}] \geq \mathbf{E}[\mathrm{ALG}_{\text{quantile}}]$.

Next, we aim to control the distribution of $v_{i,j}$ given that $X_{i,j} = 1$ in order to get access to the value that agent i has when being allocated item j in the quantile allocation rule. To this end, we use the following lemma.

Lemma 6.2.3. *For all i, j and $M^{(i)}$, we have*

$$\mathbf{Pr}\left[v_{i,j} \leq t \,\middle|\, X_{i,j} = 1, M^{(i)}\right] = F_j(t)^{\frac{n-i+1}{q_j^{(i)}}} \quad .$$

Proof. Fix a set of available items $M^{(i)}$ and let us assume for simplicity that all items $j' \in M^{(i)}$ satisfy $q_{j'}^{(i)} > 0$ (e.g. by removing all items from the set $M^{(i)}$ for which $q_{j'}^{(i)} = 0$). Observe that in the vector $\left(R_j^{(i)}\right)_{j \in M^{(i)}}$, we choose j to maximize $R_j^{(i)}$. As a consequence, conditioned on picking item $j \in M^{(i)}$, we have $R_j^{(i)} = \max_{j' \in M^{(i)}}\left(R_{j'}^{(i)}\right)$.

Recall that $R_{j'}^{(i)}$ are random variables with support $[0,1]$ and cumulative distribution function $t^{q_{j'}^{(i)}}$ by Equation (6.1). By independence across items and using the fact that $\sum_{j' \in M^{(i)}} q_{j'}^{(i)} = n - i + 1$, the random variable $\max_{j' \in M^{(i)}}\left(R_{j'}^{(i)}\right)$ has cumulative distribution function t^{n-i+1}.

As a consequence, the conditional cumulative distribution function of $R_j^{(i)}$ conditioned on $X_{i,j} = 1$ becomes t^{n-i+1}. Using that we can rewrite

$$v_{i,j} = F_j^{-1}\left(\left(R_j^{(i)}\right)^{q_j^{(i)}}\right) \quad ,$$

we get

$$\begin{aligned}\mathbf{Pr}\left[v_{i,j} \leq t \,\middle|\, X_{i,j} = 1, M^{(i)}\right] &= \mathbf{Pr}\left[F_j^{-1}\left(\left(R_j^{(i)}\right)^{q_j^{(i)}}\right) \leq t \,\middle|\, X_{i,j} = 1, M^{(i)}\right] \\ &= \mathbf{Pr}\left[R_j^{(i)} \leq F_j(t)^{1/q_j^{(i)}} \,\middle|\, X_{i,j} = 1, M^{(i)}\right] \\ &= F_j(t)^{\frac{n-i+1}{q_j^{(i)}}} \quad .\end{aligned}$$

□

For integral values of $n-i+1/q_j^{(i)}$ (in particular, if e.g. $q_j^{(i)} = 1$), observe that this is exactly the CDF of the maximum of $n-i+1/q_j^{(i)}$ independent draws from distribution F_j. This mirrors the following intuition: Say all $q_j^{(i)} = 1$. If we assign item j in the quantile allocation rule, item j's quantile needs to be the largest of $n - i + 1$ independent draws in quantile space. Transforming all quantiles back into the value space of item j, the value of buyer i for item j is the maximum of $n-i+1$ i.i.d. draws from distribution $\mathcal{D}^{(j)}$.

Having this, our overall goal is to show the following bound.

Proposition 6.2.4. *Let* OPT_j *denote the random variable indicating the contribution of item* j *to the social welfare of the optimal offline solution. Then, we have*

$$\mathbf{E}\left[\sum_{i=1}^{n} v_{i,j} X_{i,j}\right] \geq \left(1 - O\left(\frac{1}{\log n}\right)\right) \mathbf{E}\left[\mathrm{OPT}_j\right] \quad .$$

Note that showing this proposition proves the claim by taking a sum over all $j \in M$. To this end, we argue on the random variables $q_j^{(i)}$. As a first remark, the variables $q_j^{(i)}$ are independent of the values $v_{i,j}$ as we define $q_j^{(i)}$ without any knowledge on the $v_{i,j}$. The following observation shows that $\mathbf{E}\left[q_j^{(i)}/n-i+1\right]$ is exactly $q_j^{(1)}/n$. Note that $q_j^{(1)}$ is deterministic because it is the a priori probability that item j is allocated in the offline optimum.

Observation 6.2.5. *For all* i*, we have* $\mathbf{E}\left[\frac{q_j^{(i)}}{n-i+1}\right] = \frac{q_j^{(1)}}{n}$.

Observe that the left-hand side is the expected probability that we allocate item j in step i of the generalized quantile allocation rule.

Proof. Let $M_*^{(i)}$ be the set of items not allocated buyers $1, \ldots, i-1$ by the optimal offline solution. Note that $M^{(i)}$ and $M_*^{(i)}$ are identically distributed. Therefore, the optimal offline solution assigns item j to one of the buyers $i, \ldots, n$ with a probability of $\mathbf{E}\left[q_j^{(i)}\right]$ a priori. By symmetry across buyers, each buyer is assigned item j in the optimal offline solution with the same probability, namely $\frac{1}{n-i+1}\mathbf{E}\left[q_j^{(i)}\right] = \frac{1}{n} q_j^{(1)}$. □

In other words, Observation 6.2.5 states that $\mathbf{E}\left[q_j^{(i)}\right] = \frac{n-i+1}{n} \cdot q_j^{(1)}$, i.e. in expectation, the probability of assigning item j in steps $i, \ldots, n$ is exactly a $n-i+1/n$-fraction of the a priori probability. Complementing this, we can also argue that the probabilities $q_j^{(i)}$ are always upper bounded by the a priori probability that item j is allocated in the offline optimum.

Observation 6.2.6. *We have* $q_j^{(i)} \leq q_j^{(1)}$ *for all* i *and* j*.*

Proof. If $j \notin M^{(i)}$, we have $q_j^{(i)} = 0$, so the statement is clear. Otherwise, recall that $q_j^{(i)}$ is the probability that item j is allocated in the offline optimum constrained to buyers $1, \ldots, i-1$ receiving the items from $M \setminus M^{(i)}$. Next, note that we always allocate exactly one item to every buyer, so the offline optimum which is constrained to the previous allocation only has the flexibility to allocate items to the remaining $n-i+1$ buyers.

We aim to show that $q_j^{(i)} \geq q_j^{(i+1)}$ for every buyer i which allows us to conclude. To this end, consider the edge which we add to our allocation in step i, denoted by (i, j'). In other words, the optimum matchings which contribute to $q_j^{(i+1)}$ are restricted in the way that buyer i and item j' are not available.

Next, consider item j for which we would like to bound the probability of allocating in the remaining offline optimum starting from step i, i.e. $q_j^{(i)}$. We can distinguish two cases: First, we consider all valuation profiles for which item j' is not allocated in the offline optimum among $i, \dots, n$. When comparing these to the offline optima which are restricted to match (i, j'), we observe that from the perspective of item j, there is one additional buyer available when not matching i to j'. Hence, as buyers are symmetric, the probability of matching j in the offline optimum is at least as large as in the restricted case.

Second, for all valuation profiles in which item j' is allocated to some buyer in the remaining offline optimum among buyers $i, \dots, n$, we can use symmetry across buyers again. In particular, by symmetry, we end up in the same situation as when adding the edge (i, j') to the matching and hence, the probability of allocating j in the remaining offline optimum is the same.

Combining these two cases, the probability that item j is allocated in the remaining offline optimum among buyers $i, \dots, n$ is at least as large as allocating item j in the remaining offline optimum among buyers $i+1, \dots, n$. Hence, the claim follows. □

Having these observations, we consider the contribution of buyer i and item j conditioned on $M^{(i)}$. First, observe that conditioned on $M^{(i)}$, $q_j^{(i)}$ is not random anymore. Hence, by the use of Observation 6.2.2, we get that

$$\begin{aligned}\mathbf{E}\left[v_{i,j} X_{i,j} \,\middle|\, M^{(i)}\right] &= \mathbf{Pr}\left[X_{i,j} = 1 \,\middle|\, M^{(i)}\right] \cdot \mathbf{E}\left[v_{i,j} \,\middle|\, X_{i,j} = 1, M^{(i)}\right] \\ &= \frac{q_j^{(i)}}{n-i+1} \cdot \mathbf{E}\left[Y_{i,j} \,\middle|\, M^{(i)}\right] \quad ,\end{aligned}$$

where we use Lemma 6.2.3 to introduce $Y_{i,j}$, a random variable with CDF $F_j(t)^{\frac{n-i+1}{q_j^{(i)}}}$. By Observation 6.2.6, we can bound $\mathbf{E}\left[Y_{i,j} \,\middle|\, M^{(i)}\right] \geq \mathbf{E}\left[Y'_{i,j}\right]$ where $Y'_{i,j}$ is a random variable with CDF $F_j(t)^{n-i+1/q_j^{(1)}}$. Note that the latter CDF does not depend on $q_j^{(i)}$ anymore, but only on $q_j^{(1)}$ which is deterministic. As a consequence, we can bound

$$\mathbf{E}\left[v_{i,j} X_{i,j} \,\middle|\, M^{(i)}\right] \geq \frac{q_j^{(i)}}{n-i+1} \cdot \mathbf{E}\left[Y'_{i,j}\right] \quad .$$

Taking the expectation over all possible sets $M^{(i)}$, we get

$$\mathbf{E}\left[v_{i,j} X_{i,j}\right] \geq \mathbf{E}\left[\frac{q_j^{(i)}}{n-i+1} \cdot \mathbf{E}\left[Y'_{i,j}\right]\right] = \frac{q_j^{(1)}}{n} \cdot \mathbf{E}\left[Y'_{i,j}\right] \quad .$$

Now, we can sum over all buyers i to bound the contribution of one item to the quantile allocation rule. Recall that $Y'_{i,j}$ has CDF $F_j(t)^{n-i+1/q_j^{(1)}}$. When rounding the value of $n-i+1/q_j^{(1)}$ to the next smaller integer, we get a random variable which is the maximum of $\left\lfloor \frac{n-i+1}{q_j^{(1)}} \right\rfloor$ draws from F_j. Hence, we get

$$\mathbf{E}\left[\sum_{i=1}^{n} v_{i,j} X_{i,j}\right] \geq \sum_{i=1}^{n} \frac{q_j^{(1)}}{n} \cdot \mathbf{E}\left[Y'_{i,j}\right] \geq \sum_{i=1}^{n} \frac{q_j^{(1)}}{n} \cdot \mathbf{E}\left[\max_{i' \in \left[\left\lfloor n-i+1/q_j^{(1)} \right\rfloor\right]} \{v_{i',j}\}\right] \quad .$$

Now, define $\Gamma := \left\lceil n/2q_j^{(1)} \right\rceil$ and $\Gamma_i := \left\lfloor n-i+1/q_j^{(1)} \right\rfloor$. We make a case distinction if $\Gamma_i \geq \Gamma$ or not. Denote by i^* the last i for which $\Gamma_i \geq \Gamma$.

Case 1: $\Gamma_i \geq \Gamma$. Note that for all i with $\Gamma_i = \left\lfloor n-i+1/q_j^{(1)} \right\rfloor \geq \Gamma$, i.e. $i = 1, \dots, i^*$, we get

$$\mathbf{E}\left[\max_{i' \in [\Gamma_i]} \{v_{i',j}\}\right] \geq \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \geq \frac{\log(n-i)}{\log(n)} \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \quad ,$$

where the last inequality trivially multiplies the non-negative expectation with a number smaller than one.

Case 2: $\Gamma_i < \Gamma$. For all i such that $\left\lfloor n-i+1/q_j^{(1)} \right\rfloor < \Gamma$, we can exploit the MHR property via an application of Lemma 6.1.1 in order to bound

$$\begin{aligned}\mathbf{E}\left[\max_{i' \in [\Gamma_i]} \{v_{i',j}\}\right] &\geq \frac{\log\left(\left\lfloor n-i+1/q_j^{(1)} \right\rfloor\right)}{\log(\Gamma)} \cdot \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \\ &\geq \frac{\log(n-i) - \log\left(q_j^{(1)}\right)}{\log(n+2) - \log\left(2q_j^{(1)}\right)} \cdot \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \quad .\end{aligned}$$

For the last inequality, we use that

$$\begin{aligned}\log(\Gamma) = \log\left(\left\lceil n/2q_j^{(1)} \right\rceil\right) &\leq \log\left(\frac{n}{2q_j^{(1)}} + 1\right) \\ &= \log\left(\frac{n + 2q_j^{(1)}}{2q_j^{(1)}}\right) = \log\left(n + 2q_j^{(1)}\right) - \log\left(2q_j^{(1)}\right) \\ &\leq \log(n+2) - \log\left(2q_j^{(1)}\right) \quad .\end{aligned}$$

By similar calculations, we get that $\log\left(\left\lfloor n-i+1/q_j^{(1)} \right\rfloor\right) \geq \log(n-i) - \log\left(q_j^{(1)}\right)$. Observe that for $i = n$, the latter expression is not well defined. We will not consider these variables and only take the sum until $n-1$ into account.

Combination: Now, we aim to apply Lemma 6.1.3 for suitably chosen values of α, Γ and z. We set $z = \frac{q_j^{(1)}}{n}$ and $\Gamma = \left\lceil \frac{1}{2z} \right\rceil = \left\lceil \frac{n}{2 \cdot q_j^{(1)}} \right\rceil$ as above and $\alpha = \frac{1+\ln(n)}{H_{n/2}}$. Let us denote by v_j a draw from distribution $\mathcal{D}^{(j)}$. Then, we can compute

$$\begin{aligned}\mathbf{E}\left[\sum_{i=1}^{n} v_{i,j} X_{i,j}\right] &\geq \sum_{i=1}^{i^*} \frac{q_j^{(1)}}{n} \cdot \frac{\log(n-i)}{\log(n)} \cdot \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \\ &\quad + \sum_{i=i^*+1}^{n-1} \frac{q_j^{(1)}}{n} \cdot \frac{\log(n-i) - \log\left(q_j^{(1)}\right)}{\log(n+2) - \log\left(2q_j^{(1)}\right)} \cdot \mathbf{E}\left[\max_{i' \in [\Gamma]} \{v_{i',j}\}\right] \\ &\geq \sum_{i=1}^{i^*} \frac{q_j^{(1)}}{n} \cdot \frac{1}{\alpha} \cdot \frac{\log(n-i)}{\log(n)} \cdot \mathbf{E}\left[v_j \,\middle|\, v_j \geq F_j^{-1}\left(1 - \frac{q_j^{(1)}}{n}\right)\right] \\ &\quad + \sum_{i=i^*+1}^{n-1} \frac{q_j^{(1)}}{n} \cdot \frac{1}{\alpha} \cdot \frac{\log(n-i) - \log\left(q_j^{(1)}\right)}{\log(n+2) - \log\left(2q_j^{(1)}\right)} \cdot \mathbf{E}\left[v_j \,\middle|\, v_j \geq F_j^{-1}\left(1 - \frac{q_j^{(1)}}{n}\right)\right] \quad .\end{aligned}$$

Now, observe that $q_j^{(1)} \mathbf{E}\left[v_j \,\middle|\, v_j \geq F_j^{-1}\left(1 - \frac{q_j^{(1)}}{n}\right)\right] \geq \mathbf{E}\left[\mathrm{OPT}_j\right]$ because the former is exactly the contribution of item j to the ex-ante relaxation and hence it is an upper bound for the expected contribution of item j to the offline optimum. In addition, we can use the integral estimation $\sum_{i=1}^{n} \log(i) = \sum_{i=1}^{n} \log(n-i+1) \geq n \log n - n + 1$. Therefore,

$$\begin{aligned}
\mathbf{E}\left[\sum_{i=1}^{n} v_{i,j} X_{i,j}\right] &\geq \mathbf{E}\left[\mathrm{OPT}_j\right] \cdot \frac{1}{\alpha} \cdot \left(\sum_{i=1}^{i^*} \frac{1}{n} \cdot \frac{\log(n-i)}{\log(n)} + \sum_{i=i^*+1}^{n-1} \frac{1}{n} \cdot \frac{\log\left(n-i\right) - \log\left(q_j^{(1)}\right)}{\log\left(n+2\right) - \log\left(2q_j^{(1)}\right)} \right) \\
&\geq \mathbf{E}\left[\mathrm{OPT}_j\right] \cdot \frac{1}{\alpha} \cdot \frac{1}{n \cdot \log(n)} \left(\sum_{i=1}^{n-1} \log(n-i) \right) \\
&\geq \left(1 - O\left(\frac{1}{\log n}\right)\right) \mathbf{E}\left[\mathrm{OPT}_j\right] \quad ,
\end{aligned}$$

which allows us to conclude the proof of Theorem 6.2.7. □

6.2.2 Static prices

Next, we would like to demonstrate how to use static prices. We consider the case that the number of items m is upper bounded by $\frac{n}{(\log \log n)^2}$. We set the price for item j to

$$p_j = F_j^{-1}\left(1 - q\right) \text{ , where } q = \frac{\log \log n}{n} \text{ ,}$$

which allows us to state the following theorem.

Theorem 6.2.7. *For $m \leq \frac{n}{(\log \log n)^2}$, the posted-prices mechanism with static prices and independent item-valuations is $1 - O\left(\frac{\log \log \log n}{\log n}\right)$-competitive with respect to the expected offline optimum social welfare.*

The proof will work as follows: First, observe that we can bound the probability of selling item j to buyer i by the probability of the event that buyer i has only non-negative utility for this item. This implies a bound on the probability of selling item j in our algorithm. Finally, we combine this with a lower bound on the price p_j and hence are able to bound the revenue (and thus the welfare) obtained by our algorithm. Observe that our guarantee only applies if the number of items m is bounded by $\frac{n}{(\log \log n)^2}$. We leave the extension to the general case as an open problem. As a first step, one could try to derive a suitable bound on the utility of agents in order to extend the result.

Proof of Theorem 6.2.7. We start by considering the probability of selling a fixed item j. We can lower bound the probability that buyer i buys the item by the event that buyer i only has positive utility for item j, i.e.

$$\begin{aligned}
\mathbf{Pr}\left[i \text{ buys item } j \,\middle|\, M^{(i)} \right] &\geq \left(1 - F_j(p_j)\right) \prod_{j' \in M^{(i)} \setminus \{j\}} F_{j'}(p_{j'}) \geq \frac{\log \log n}{n} \left(1 - \frac{\log \log n}{n}\right)^m \\
&\geq \frac{\log \log n}{n} \left(1 - m \cdot \frac{\log \log n}{n}\right) \geq \frac{\log \log n}{n} \left(1 - \frac{1}{\log \log n}\right) \quad ,
\end{aligned}$$

where the third inequality is an application of Bernoulli's inequality $(1+x)^r \geq 1+xr$ for any $x \geq -1$ and integer r and the last inequality follows as $m \leq \frac{n}{(\log\log n)^2}$. Taking the expectation over all possible sets $M^{(i)}$, the probability of the counter event, namely that buyer i does not buy item j, is upper bounded by

$$\mathbf{Pr}\left[i \text{ does not buy item } j\right] \leq 1 - \frac{\log\log n - 1}{n} \quad .$$

As a consequence, the probability that item j is not sold during the process is upper bounded by

$$\mathbf{Pr}\left[j \text{ unsold}\right] \leq \left(1 - \frac{\log\log n - 1}{n}\right)^n \leq \exp\left(1 - \log\log n\right) = \frac{\mathrm{e}}{\log n} \quad .$$

Therefore, the probability of selling item j is lower bounded by $\mathbf{Pr}\left[j \text{ sold}\right] \geq 1 - \frac{\mathrm{e}}{\log n}$. Additionally, we can bound the price of item j by Lemma 6.1.2 via

$$\begin{aligned} p_j = F_j^{-1}\left(1 - \frac{\log\log n}{n}\right) &\geq \frac{\log n - \log\log\log n}{H_n} \cdot \mathbf{E}\left[\max_{i\in[n]} v_{i,j}\right] \\ &\geq \left(1 - \frac{\log\log\log n + 1}{\log n}\right) \cdot \mathbf{E}\left[\max_{i\in[n]} v_{i,j}\right] \quad . \end{aligned}$$

Having this, we can conclude by the some fundamental calculus to get

$$\begin{aligned} \mathbf{E}\left[\text{ALG}\right] &\geq \mathbf{E}\left[\text{revenue}_{\text{pp}}\right] \geq \sum_{j=1}^{m} \mathbf{Pr}\left[j \text{ sold}\right] \cdot p_j \\ &\geq \sum_{j=1}^{m} \left(1 - \frac{\mathrm{e}}{\log n}\right)\left(1 - \frac{\log\log\log n + 1}{\log n}\right) \mathbf{E}\left[\max_{i\in[n]} v_{i,j}\right] \\ &\geq \left(1 - O\left(\frac{\log\log\log n}{\log n}\right)\right) \sum_{j=1}^{m} \mathbf{E}\left[\max_{i\in[n]} v_{i,j}\right] \geq \left(1 - O\left(\frac{\log\log\log n}{\log n}\right)\right) \mathbf{E}\left[\text{OPT}\right] \quad . \end{aligned}$$

□

6.3 Asymptotic Upper Bounds on the Competitive Ratios

The competitive ratios from Section 6.2 are asymptotically tight. To see this, we provide matching upper bounds in this section showing asymptotic optimality. We consider the case of selling a single item with static and dynamic prices respectively. In any of the two cases, we can achieve asymptotic upper bounds on the competitive ratio of posted prices mechanisms which match our results from the previous section. In particular, we prove that these bounds are the best possible ones for *any* choice of pricing strategy.

6.3.1 Dynamic prices

We consider the guarantee of our dynamic-pricing algorithm first. Even with a single item and types drawn from an exponential distribution, the best competitive ratio is $1 - \Omega\left(\frac{1}{\log n}\right)$. We simplify notation by omitting indices when possible.

Proposition 6.3.1. *Let $v_1, \dots, v_n \in \mathbb{R}_{\geq 0}$ be random variables where each v_i is drawn i.i.d. from the exponential distribution with rate one, so $v_1, \dots, v_n \sim \mathrm{Exp}(1)$. For all dynamic prices, the competitive ratio of the mechanism picking the first v_i with $v_i \geq p^{(i)}$ is upper bounded by $1 - \Omega\left(\frac{1}{\log n}\right)$.*

In order to prove Proposition 6.3.1, we use that the expected value of the optimal offline solution (the best value in hindsight) is given by $\mathbf{E}\left[\max_{i \in [n]} v_i\right] = H_n$ [Arnold et al., 2008]. Therefore, it suffices to show that the expected value of any dynamic pricing rule is upper bounded by $H_n - c$ for some constant $c > 0$.

To upper-bound the expected social welfare of any dynamic pricing rule, we use the fact that this problem corresponds to a Markov decision process and the optimal dynamic prices[3] are given by

$$p^{(n)} = 0 \quad \text{and} \quad p^{(i)} = \mathbf{E}\left[\max\{v_{i+1}, p^{(i+1)}\}\right] \quad \text{for } i < n \ .$$

Furthermore, $p^{(0)}$ is exactly the expected social welfare of this mechanism. Therefore, the following lemma with $k = n$ directly proves our claim.

Lemma 6.3.2. *Let $v_1, \dots, v_n \in \mathbb{R}_{\geq 0}$ be random variables where each v_i is drawn i.i.d. from the exponential distribution $\mathrm{Exp}(1)$. Moreover, let $p^{(n)} = 0$ and $p^{(i)} = \mathbf{E}\left[\max\{v_{i+1}, p^{(i+1)}\}\right]$ for $i < n$. Then, we have $p^{(n-k)} \leq H_k - \frac{1}{8}$ for all $2 \leq k \leq n$.*

Proof. We are going to prove the statement by induction on index variable k.

First, consider the induction base $k = 2$. By definition of the thresholds we know that $p^{(n)} = 0$ and $p^{(n-1)} = \mathbf{E}\left[\max\{v_n, p^{(n)}\}\right] = \mathbf{E}\left[\max\{v_n, 0\}\right] = \mathbf{E}\left[v_n\right] = 1$. Next, consider the threshold $p^{(n-2)}$ defined by

$$\begin{aligned} p^{(n-2)} &= \mathbf{E}\left[\max\{v_{n-1}, p^{(n-1)}\}\right] = \mathbf{E}\left[\max\{v_{n-1}, 1\}\right] \\ &= \int_0^\infty \mathbf{Pr}\left[\max\{v_{n-1}, 1\} \geq x\right] \mathrm{d}x = \int_0^1 1 \mathrm{d}x + \int_1^\infty \mathrm{e}^{-x} \mathrm{d}x = 1 + \frac{1}{\mathrm{e}} \approx 1.368\,. \end{aligned}$$

Thus, we can easily verify that $p^{(n-2)} \leq 1.375 = H_2 - \frac{1}{8}$.

For the inductive step, we move from k to $k+1$. By the induction hypothesis, we have

$$p^{(n-(k+1))} = \mathbf{E}\left[\max\{v_{n-k}, p^{(n-k)}\}\right] \leq \mathbf{E}\left[\max\left\{v_{n-k}, H_k - \frac{1}{8}\right\}\right] \ .$$

Furthermore,

$$\begin{aligned} \mathbf{E}\left[\max\left\{v_{n-k}, H_k - \frac{1}{8}\right\}\right] &= \int_0^\infty \mathbf{Pr}\left[\max\left\{v_{n-k}, H_k - \frac{1}{8}\right\} \geq x\right] \mathrm{d}x \\ &= \int_0^{H_k - \frac{1}{8}} 1 \mathrm{d}x + \int_{H_k - \frac{1}{8}}^\infty \mathrm{e}^{-x} \mathrm{d}x \\ &= H_k - \frac{1}{8} + \mathrm{e}^{-H_k + \frac{1}{8}} \ . \end{aligned}$$

[3] See e.g. Chow et al. [1971].

We now use the fact that the k-th harmonic number H_k for $k \geq 2$ is bounded from below by $H_k \geq \log k + \gamma$, in which $\gamma \approx 0.577$ denotes the Euler-Mascheroni constant. So for $k \geq 2$

$$\mathrm{e}^{-H_k + \frac{1}{8}} \leq \mathrm{e}^{-(\log k + \gamma) + \frac{1}{8}} = \frac{\mathrm{e}^{\frac{1}{8} - \gamma}}{k} \leq \frac{\mathrm{e}^{\frac{1}{8} - 0.57}}{k} \leq \frac{1}{\frac{3}{2} \cdot k} \leq \frac{1}{k+1} \quad .$$

In combination, this gives us

$$p^{(n-(k+1))} \leq H_k - \frac{1}{8} + \frac{1}{k+1} = H_{k+1} - \frac{1}{8} \quad .$$

□

6.3.2 Static prices

For static pricing rules, we show that any mechanism is $1 - \Omega\left(\frac{\log\log\log n}{\log n}\right)$-competitive. Again, this bound even holds for a single item and the valuations being drawn from an exponential distribution.

Proposition 6.3.3. *Let $v_1, \ldots, v_n \in \mathbb{R}_{\geq 0}$ be random variables where each v_i is drawn i.i.d. from the exponential distribution with rate one, so $v_1, \ldots, v_n \sim \mathrm{Exp}(1)$. For all static prices $p \in \mathbb{R}_{\geq 0}$ the competitive ratio of the mechanism picking the first v_i with $v_i \geq p$ is upper bounded by $1 - \Omega\left(\frac{\log\log\log n}{\log n}\right)$.*

The idea of the proof is as follows. The expected welfare obtained by a static pricing mechanism using price p is given by $\mathbf{E}\left[\mathrm{ALG}\right] = \mathbf{E}\left[v \mid v \geq p\right] \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] = (p+1) \cdot (1 - (1 - \mathrm{e}^{-p})^n)$. This has to be compared to the expected value of the optimal offline solution $\mathbf{E}\left[\max_{i \in [n]} v_i\right] = H_n$ [Arnold et al., 2008].

Proof of Proposition 6.3.3. Consider $n \geq \exp(\exp(\exp(4)))$.

Observe that always $\mathbf{E}\left[\mathrm{OPT}\right] = \mathbf{E}\left[\max_{i \in [n]} v_i\right] = H_n$ [Arnold et al., 2008]. Now, we will bound $\mathbf{E}\left[\mathrm{ALG}\right]$ for any choice of a static price $p \in \mathbb{R}_{\geq 0}$. Regardless of p, we have

$$\mathbf{E}\left[\mathrm{ALG}\right] = \mathbf{E}\left[v \mid v \geq p\right] \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] = (p+1) \cdot \left(1 - (1 - \mathrm{e}^{-p})^n\right) \quad .$$

We will show that for any choice of a static price p,

$$(p+1) \cdot \left(1 - (1 - \mathrm{e}^{-p})^n\right) \leq H_n - c \log\log\log n$$

for some constant c, which then immediately proves the claim as $H_n = \Theta(\log n)$.

To this end, we will consider three cases for the choice of p.

Case 1: $0 \leq p < \log n - \frac{1}{2}\log\log\log n$: We use the trivial upper bound of 1 for the probability term in $\mathbf{E}\left[\mathrm{ALG}\right]$, so

$$\begin{aligned} (p+1) \cdot \left(1 - (1 - \mathrm{e}^{-p})^n\right) &\leq (p+1) < \log n - \frac{1}{2}\log\log\log n + 1 \\ &\leq \log n - \frac{1}{2}\log\log\log n + \frac{1}{4}\log\log\log n \\ &= \log n - \frac{1}{4}\log\log\log n \end{aligned}$$

as we assumed that $n \geq \exp(\exp(\exp(4)))$.

Case 2: $\log n - \frac{1}{2}\log\log\log n \le p \le H_n - 1$: Again, observe that the expected value of the algorithm can be upper bounded by

$$\begin{aligned}&(p+1)\cdot\left(1-(1-\mathrm{e}^{-p})^n\right) \le H_n\cdot\left(1-(1-\mathrm{e}^{-p})^n\right)\\ &\le H_n\cdot\left(1-\left(1-\mathrm{e}^{-\log n+\frac{1}{2}\log\log\log n}\right)^n\right) = H_n\cdot\left(1-\left(1-\frac{\sqrt{\log\log n}}{n}\right)^n\right).\end{aligned}$$

Next, we want to lower bound $\left(1-\frac{\sqrt{\log\log n}}{n}\right)^n$ in order to get the desired upper bound on $\mathbf{E}\left[\mathrm{ALG}\right]$. For this purpose, we use the following inequality. For $n > 1$ and $x \in \mathbb{R}$ with $|x| \le n$, we have $(1+\frac{x}{n})^n \ge \mathrm{e}^x\cdot\left(1-\frac{x^2}{n}\right)$. This way, we get

$$\begin{aligned}\left(1-\frac{\sqrt{\log\log n}}{n}\right)^n &\ge \mathrm{e}^{-\sqrt{\log\log n}}\cdot\left(1-\frac{(-\sqrt{\log\log n})^2}{n}\right) = \mathrm{e}^{-\sqrt{\log\log n}}\cdot\underbrace{\left(1-\frac{\log\log n}{n}\right)}_{\ge\frac{1}{2},\ \forall n}\\ &\ge \frac{1}{2}\mathrm{e}^{-\sqrt{\log\log n}}.\end{aligned}$$

Note that if $\log\log n \ge 4$, then also $\sqrt{\log\log n} \le \frac{1}{2}\log\log n$, and hence, we get that $\mathrm{e}^{-\sqrt{\log\log n}} \ge \mathrm{e}^{-\frac{1}{2}\log\log n} = \frac{1}{\sqrt{\log n}}$. This gives us

$$\begin{aligned}(p+1)\cdot\left(1-(1-\mathrm{e}^{-p})^n\right) &\le H_n\cdot\left(1-\left(1-\frac{\sqrt{\log\log n}}{n}\right)^n\right) \le H_n\cdot\left(1-\frac{1}{2}\mathrm{e}^{-\sqrt{\log\log n}}\right)\\ &\le H_n\cdot\left(1-\frac{1}{2\sqrt{\log n}}\right) \le H_n\cdot\left(1-\frac{1}{2}\cdot\frac{\log\log\log n}{\log n}\right),\end{aligned}$$

where in the last step we use that $\sqrt{\log n} \ge \log\log\log n$ and therefore $\sqrt{\log n} \le \frac{\log n}{\log\log\log n}$.

Case 3: $p > H_n - 1$: In this case we use the fact that Giannakopoulos and Zhu [2018] showed that the revenue function $p \mapsto p\cdot(1-(1-\mathrm{e}^{-p})^n)$ is non-increasing on $[H_n - 1, \infty)$.

This implies that

$$\begin{aligned}(p+1)\cdot\left(1-(1-\mathrm{e}^{-p})^n\right) &\le p\cdot\left(1-(1-\mathrm{e}^{-p})^n\right)+1 \le H_n\cdot\left(1-\left(1-\mathrm{e}^{-(H_n-1)}\right)^n\right)+1\\ &\le H_n\cdot\left(1-\left(1-\frac{\mathrm{e}}{n}\right)^n\right)+1\\ &\le \frac{99}{100}\cdot H_n + 1 = H_n\cdot\left(\frac{99}{100}+\frac{1}{H_n}\right),\end{aligned}$$

where the last inequality holds for $n \ge 4$. Now, we can use that for n large enough, we have $\frac{99}{100}+\frac{1}{H_n} \le 1 - \frac{c\log\log\log n}{H_n}$ and hence, we get the desired bound. □

6.4 Extensions to Subadditive Buyers and Revenue Considerations

In this section, we illustrate that the same style of mechanisms used for unit-demand buyers can be extended to buyers with subadditive valuations. To generalize the MHR property, we assume that the subadditive valuation functions are drawn from distributions with MHR marginals. That is, $v_i \sim \mathcal{D}$ and we assume that $v_i(\{j\})$ has a marginal distribution with monotone hazard rate. Buyers arrive online one-by-one and purchase the bundle of items which maximizes the buyer's utility.

We construct a dynamic-pricing mechanism which is $1 - O\left(\frac{1+\log m}{\log n}\right)$-competitive in Section 6.4.1. In the static pricing environment, our mechanism in Section 6.4.2 is $1 - O\left(\frac{\log\log\log n}{\log n} + \frac{\log m}{\log n}\right)$-competitive for subadditive buyers. Note that the guarantees now depend on the number of items m. To make them meaningful, we need $m = o(n)$. This makes them significantly worse than the ones we obtain for unit-demand functions in Section 6.2 with a much more careful treatment. However, they are stronger in one aspect, namely that in both cases we will bound the revenue of the mechanism in terms of the optimal social welfare. In particular, this means that they are also approximations of the optimal revenue.

Distributions with MHR Marginals

Concerning subadditivity, we first start with a straightforward extension of the definitions from Section 6.1 to the case of buyers' valuation functions being subadditive. As before, we assume that the functions $v_1, \ldots, v_n$ are drawn independently from a publicly known distribution $\mathcal{D}$. Now, the distribution has *MHR marginals*. We define MHR marginals as follows. Let $\mathcal{D}^{(j)}$ be the marginal distribution of $v_i(\{j\})$, which is the value of a buyer for being allocated only item j. We assume that $\mathcal{D}^{(j)}$ is a continuous, real, non-negative distribution with monotone hazard rate. Note that this allows arbitrary correlation between items.

Some Additional Remarks for Subadditive Valuations

We quickly recall that in posted-pricing mechanisms, we assume that buyer i picks the bundle of items $S \subseteq M^{(i)}$ which maximizes her *utility* $v_i(S) - \sum_{j \in S} p_j^{(i)}$ if positive. Also, remember that the *expected social welfare* of the mechanism is given by $\mathbf{E}\left[\sum_{i=1}^n v_i(S_i)\right] =: \mathbf{E}\left[\text{ALG}\right]$. Its *expected revenue* is given by $\mathbf{E}\left[\sum_{i=1}^n \sum_{j \in S_i} p_j^{(i)}\right] =: \mathbf{E}\left[\text{revenue}_{\text{pp}}\right]$.

Our benchmark is the expected welfare of the offline optimum allocation, denoted by $\mathbf{E}\left[\text{OPT}\right] := \mathbf{E}\left[\sum_{i=1}^n v_i(S_i^*)\right]$. Using the subadditivity of buyers' valuations, we can upper-bound the expected optimal social welfare by

$$\mathbf{E}\left[\sum_{i=1}^n v_i(S_i^*)\right] \leq \mathbf{E}\left[\sum_{i=1}^n \sum_{j \in S_i^*} v_i(\{j\})\right] \leq \sum_{j=1}^m \mathbf{E}\left[\max_{i \in [n]} v_i(\{j\})\right] .$$

Furthermore, let $\mathbf{E}\left[\text{revenue}_{\text{opt}}\right]$ denote the maximum expected revenue of any individually rational mechanism. Due to individual rationality, we have $\mathbf{E}\left[\text{revenue}_{\text{opt}}\right] \leq \mathbf{E}\left[\text{OPT}\right]$ and $\mathbf{E}\left[\text{revenue}_{\text{pp}}\right] \leq \mathbf{E}\left[\text{ALG}\right]$.

6.4.1 Dynamic Pricing for Subadditive Valuations with MHR marginals

We first consider dynamic prices. That is, buyer i faces prices depending on the set of available items. Our strategy is to sell specific items only to a subgroup of buyers in order to gain control over the selling process. We can implement this by imposing the following item prices which decrease as the selling process proceeds.

We split the set of buyers in groups of size $\lfloor n/m \rfloor =: n'$ and sell item j only to the group of buyers

$$\{(j-1)\cdot n'+1, ..., (j-1)\cdot n'+n'\} =: N_j \ .$$

For the k-th buyer in N_j, we set the price for item j to

$$p_j^{((j-1)n'+k)} = F_j^{-1}\left(1 - \frac{1}{\lfloor \frac{n}{m} \rfloor - k + 1}\right) \tag{6.2}$$

and the prices for all other unsold items to infinity. This choice of prices ensures that the first item is sold among the first $\lfloor \frac{n}{m} \rfloor$ buyers, the second item among the second $\lfloor \frac{n}{m} \rfloor$ buyers, and so on. As a consequence, all items are sold in our process.

Theorem 6.4.1. *The posted-prices mechanism with subadditive buyers and dynamic prices is* $1 - O\left(\frac{1+\log m}{\log n}\right)$*-competitive with respect to the expected optimal offline social welfare.*

Proof. We start by considering the case of selling one item among $n' = \lfloor \frac{n}{m} \rfloor$ buyers where the prices for the item are as in Equation (6.2). Note that we simplify notation in this context and omit the index of the item. Let X_i be a random variable which is equal to one if buyer i buys the item and zero otherwise. Note that by our choice of prices, the item is sold in step i with probability $\frac{1}{n'-i+1}$ what leads to

$$\mathbf{E}\left[X_i\right] = \mathbf{Pr}\left[X_i = 1\right] = \frac{1}{n'-i+1} \cdot \prod_{i'=1}^{i-1}\left(1 - \frac{1}{n'-i'+1}\right) = \frac{1}{n'} \ .$$

Further, buyer i only buys the item if v_i exceeds the price. Using Lemma 6.1.2 allows to calculate

$$p^{(i)} = F^{-1}\left(1 - \frac{1}{n'-i+1}\right) \geq \frac{\log\left(n'-i+1\right)}{H_{n'}} \cdot \mathbf{E}\left[\max_{i\in[n']} v_i\right] \ .$$

Note that this bound is deterministic. We will next make use of an application of the integral estimation $\sum_{i=1}^{n'} \log\left(n'-i+1\right) \geq n' \log n' - n' + 1$ as well as bound the harmonic number $H_{n'} \leq \log n' + 1$. This leads to a lower bound for the expected social welfare via the expected revenue of

$$\begin{aligned}\sum_{i=1}^{n'} \mathbf{E}\left[p^{(i)} X_i\right] &\geq \sum_{i=1}^{n'} \frac{\log\left(n'-i+1\right)}{n' H_{n'}} \mathbf{E}\left[\max_{i\in[n']} v_i\right] \geq \frac{n' \log n' - n' + 1}{n' H_{n'}} \mathbf{E}\left[\max_{i\in[n']} v_i\right] \\ &\geq \frac{\log n' - 1 + \frac{1}{n'}}{\log n' + 1} \mathbf{E}\left[\max_{i\in[n']} v_i\right] \geq \left(1 - \frac{2}{\log n'}\right) \mathbf{E}\left[\max_{i\in[n']} v_i\right] \ .\end{aligned}$$

Now, we apply Lemma 6.1.1 which states that the quotient of the expectation of the maximum of n' and n i.i.d. random draws from an MHR distribution is lower bounded by $\log n' / \log n$ for $n' \leq n$. This leads to the bound

$$\sum_{i=1}^{n'} \mathbf{E}\left[p^{(i)} X_i\right] \geq \left(1 - \frac{2}{\log n'}\right) \mathbf{E}\left[\max_{i\in[n']} v_i\right] \geq \left(1 - \frac{2}{\log n'}\right) \frac{\log n'}{\log n} \mathbf{E}\left[\max_{i\in[n]} v_i\right] \ .$$

In order to generalize this to the case of m items, our pricing strategy ensures that we can apply the received bound for every item separately. To this end, note that only buyers in N_j will consider buying item j. Further, also by our prices, every buyer will buy at most one item. By the introduction of indicator random variables X_{ij} indicating if buyer i buys item j, we can conclude

$$\begin{aligned}\mathbf{E}\left[\text{revenue}_{\text{pp}}\right] &= \sum_{j=1}^{m}\sum_{i=1}^{n}\mathbf{E}\left[p_j^{(i)}X_{i,j}\right] = \sum_{j=1}^{m}\sum_{i\in N_j}\mathbf{E}\left[p_j^{(i)}X_{i,j}\right] \\ &\geq \left(1-\frac{2}{\log n'}\right)\sum_{j=1}^{m}\mathbf{E}\left[\max_{i\in[n']} v_i(\{j\})\right] \\ &\geq \left(1-O\left(\frac{1+\log m}{\log n}\right)\right)\sum_{j=1}^{m}\mathbf{E}\left[\max_{i\in[n]} v_i(\{j\})\right] \\ &\geq \left(1-O\left(\frac{1+\log m}{\log n}\right)\right)\mathbf{E}\left[\text{OPT}\right] \ .\end{aligned}$$

□

Corollary 6.4.2. *The expected revenue of the posted-prices mechanism with subadditive buyers and dynamic prices is a* $1-O\left(\frac{1+\log m}{\log n}\right)$*-fraction of the expected optimal revenue.*

Note that the assumption of buyers' valuations being identically distributed is actually a too strong requirement for these results. For the proofs in this chapter it would be sufficient to consider buyers having identical marginals on single item sets, but correlations between items might be buyer-specific.

6.4.2 Static Pricing for Subadditive Valuations with MHR marginals

For the case of static prices, we give a $1-O\left(\frac{\log\log\log n}{\log n}+\frac{\log m}{\log n}\right)$-competitive mechanism.

The general design idea for our mechanism is as follows. Setting fairly low prices will put high probability on the event of selling all items. Although we cannot control which buyer will buy which bundle of items, we can extract all social welfare of the posted prices mechanism as revenue. Therefore, having prices which still ensure that the revenue can be lower bounded by a suitable fraction of the optimal social welfare will lead to the desired bound.

For any item, we set the price of item j to

$$p_j = F_j^{-1}\left(1-q\right), \text{ where } q = \frac{m\log\log n}{n}$$

and F_j denotes the marginal distribution of $v_i(\{j\})$. Observe the similarity to the pricing structure in Section 6.2.2. This allows us to prove the following theorem.

Theorem 6.4.3. *The posted-prices mechanism with subadditive buyers and static prices is* $1-O\left(\frac{\log\log\log n}{\log n}+\frac{\log m}{\log n}\right)$*-competitive with respect to the expected optimal offline social welfare.*

Proof. Lower bounding the expected revenue by a suitable fraction of the optimal social welfare will allow us to prove the theorem.

We start with an application of Lemma 6.1.2 to get a bound on p_j exploiting the MHR property:

$$\begin{aligned} p_j = F_j^{-1}(1-q) &\geq \frac{-\log q}{H_n} \mathbf{E}\left[\max_{i \in [n]} v_i(\{j\})\right] \\ &= \frac{\log n - \log\log\log n - \log m}{H_n} \mathbf{E}\left[\max_{i \in [n]} v_i(\{j\})\right] \\ &\geq \left(1 - \frac{\log\log\log n + \log m + 1}{\log n}\right) \mathbf{E}\left[\max_{i \in [n]} v_i(\{j\})\right] \enspace . \end{aligned}$$

Now, we aim for a lower bound on the probability that all items are sold in our mechanism. To this end, let $M^{(i)}$ denote the (random) set of items that are still unsold as buyer i arrives. Observe that buyer i will buy at least one item if $v_i(\{j\}) > p_j$ for some $j \in M^{(i)}$. We defined the prices such that $\mathbf{Pr}\left[v_i(\{j\}) > p_j\right] = q$. Consequently, $\mathbf{Pr}\left[\text{buyer } i \text{ buys an item} \;\middle|\; M^{(i)} \neq \emptyset\right] \geq q$ for all i.

To bound the probability of selling all items, consider the following thought experiment: For every buyer i, we toss a coin which shows head with probability q. Denote by Z the random variable counting the number of occurring heads in n coin tosses. By the above considerations, the probability for tossing head in our thought experiment is a lower bound on the probability that buyer i buys at least one item as long as there are items remaining. As a consequence, the probability for the event of seeing at least m times head is a lower bound on the probability of selling all items in our mechanism.

Using that

$$\mathbf{E}\left[Z\right] = nq = m \log\log n \enspace ,$$

a Chernoff bound with $\delta = 1 - \frac{1}{\log\log n}$ yields

$$\begin{aligned} \mathbf{Pr}\left[Z < m\right] = \mathbf{Pr}\left[Z < (1-\delta)\mathbf{E}\left[Z\right]\right] &\leq \exp\left(-\frac{1}{2}\delta^2 \mathbf{E}\left[Z\right]\right) \\ &= \exp\left(-\frac{1}{2}\frac{(\log\log n - 1)^2}{\log\log n} m\right) \overset{(\diamond)}{\leq} \exp\left(-\log\log n + 2\right) = \frac{\mathrm{e}^2}{\log n} \enspace , \end{aligned}$$

where in $(\diamond)$ we assumed that $m \geq 2$. Observe that the case of $m = 1$ is covered by our results in Section 6.2.

Combining these, we can lower bound the expected social welfare of the posted prices mechanism by

$$\begin{aligned} \mathbf{E}\left[\text{ALG}\right] &\geq \mathbf{E}\left[\text{revenue}_{\text{pp}}\right] \geq \mathbf{Pr}\left[\text{all items are sold}\right] \left(\sum_{j=1}^{m} p_j\right) \\ &\geq \mathbf{Pr}\left[Z \geq m\right] \sum_{j=1}^{m} F_j^{-1}(1-q) \\ &\geq \left(1 - \frac{\mathrm{e}^2}{\log n}\right)\left(1 - \frac{\log\log\log n + \log m + 1}{\log n}\right) \sum_{j=1}^{m} \mathbf{E}\left[\max_{i \in [n]} v_i(\{j\})\right] \\ &\geq \left(1 - O\left(\frac{\log\log\log n}{\log n} + \frac{\log m}{\log n}\right)\right) \mathbf{E}\left[\text{OPT}\right] \enspace . \end{aligned}$$

□

Observe that the proof of Theorem 6.4.3 only requires to bound the expected revenue of our mechanism. Bounding the expected optimal revenue by the expected optimal social welfare, we can state the following corollary.

Corollary 6.4.4. *The expected revenue achieved by the posted prices mechanism with subadditive buyers and static prices yields a* $1 - O\left(\frac{\log\log\log n}{\log n} + \frac{\log m}{\log n}\right)$*-fraction of the expected optimal revenue.*

Chapter 7

The Secretary Problem with Predicted Additive Gap

In Chapter 3 to Chapter 6, we made the crucial assumption that values of buyers for items are always drawn from known probability distributions. Now, we change the model and enter the terrain with adversarial weights which are revealed in random order, a.k.a. the Secretary problem.

In this chapter, we assume that there is a single item for sale, every buyer has a weight w_i for being allocated the item which is fixed by an adversary, and buyers arrive in random order. To fix notation, we assume that $w_1 \geq w_2 \geq \cdots \geq w_n$. Still, the assumption from Section 1.1.2 on the Secretary problem of having no prior information on the weights at all is highly pessimistic in the modern era. Usually, there is a huge amount of data and past information which we can use to predict the weights in (more or less) accurate ways.

To capture this access to information, we study the Secretary problem in the recently very popular model of algorithms with predictions as introduced by Lykouris and Vassilvitskii [2021] and Purohit et al. [2018]. Antoniadis et al. [2020] and Dütting et al. [2021b] already studied the Secretary problem with a prediction of the largest weight in the sequence, and resolve this setting with an algorithm which yields a nice robustness-consistency trade-off. Fujii and Yoshida [2023] consider the Secretary problem with an even stronger prediction: a prediction for every weight in the sequence. Given the prior work, we ask the meta question:

What is the weakest piece of information we can predict that still allows us to break the $1/e$ barrier from Theorem 1.1.3?

Stated another way, is there a different parameter we can predict, one that does not require us to learn the best value, but is still strong enough to beat $1/e$? This brings us to the idea of predicting the *gap* between the highest and k-th highest weight, or in other words, predicting how valuable w_k is with respect to w_1. To this end, we consider predicting an *additive gap* $w_1 - w_k$.

Predicting the additive gap versus w_1 can be motivated when considering markets displaying some sort of translation invariance. For example, imagine we are trying to sell cars and buyers are arriving each day with an offer. Different car models will typically have very different highest offers. However, the additive gap between the highest and second highest offer may be a more stable parameter to learn. From a more theoretical perspective, the additive gap between w_1 and w_k can also be viewed as interpolating between two previously studied setups: when $w_1 - w_k$ gets small, we get closer towards the k-best Secretary problem (see e.g. Gilbert and Mosteller [1966], Buchbinder et al. [2014]), and when $w_1 - w_k$ is very large, the additive gap acts as a surrogate prediction of w_1, the prediction setting in Antoniadis et al. [2020] and Dütting et al. [2021b]. As we will see, even though the additive gap is much weaker than a direct prediction for w_1, it strikes the perfect middle ground: it is strong enough to beat $1/e$ by a constant for any possible value of the gap $w_1 - w_k$ (and even if we do not know what k is upfront).

Algorithms Inspired from Classical Secretary

The algorithms in this chapter are inspired by Algorithm 2 for classical Secretary, but additionally incorporate the gap: Wait for some time to get a flavor for the weights in the sequence, set a threshold based on the past observations and the gap, afterwards pick the first element exceeding the threshold.

At first glance, this might not sound promising: In cases when the gap is small, incorporating the gap in the threshold does not really affect the best-so-far term. Hence, it may seem that beating $1/\mathrm{e}$ is still hard. However, in these cases, even though the threshold will be dominated by the best-so-far term most of the time, the gap reveals the information that the best value and all other values up to w_k are not too far off. That is, accepting any weight which is at least w_k ensures a sufficient contribution.

Our analyses use this fact in a foundational way: Either the gap is large in which case we do not consider many elements in the sequence for acceptance at all. Or the gap is small which implies that accepting one of the k highest weights is reasonably good. For each of the cases we derive lower bounds on the weight achieved by the algorithm.

Since we do not know upfront which case the instance belongs to, the waiting time cannot be tailored to the respective case but rather needs to be able to deal with both cases simultaneously. This introduces some sort of tension: For instances which have a large gap, we would like the waiting time to be small. By this, we could minimize the loss which we incur by waiting instead of accepting high weighted elements at the beginning of the sequence. For instances which have a small gap, the contribution of the gap to the algorithm's threshold can be negligible. This results in the need of a longer waiting time at the beginning to learn the magnitude of weights reasonably well. We solve this issue by using a waiting time which balances between these two extremes: It is (for most cases) shorter than the waiting time of $1/\mathrm{e}$ from the classical Secretary algorithm. Still, it is large enough to gain some information on the instance with reasonable probability.

As a corollary of the main theorem, we will see that we can beat the competitive ratio of $1/\mathrm{e}$ even if we only know the gap $w_1 - w_k$ but do not get to know the index k. In particular, this proves that even an information like "there is a gap of c in the instance" is helpful to beat $1/\mathrm{e}$, no matter which weights are in the sequence and which value c attains.

Additional Related Work on Algorithms with Predictions for the Secretary Problem

We already scratched the surface of the field on algorithms with predictions. Here, the algorithm has access to some machine learned advice upfront and may use this information to adapt decisions. Initiated by the work of Lykouris and Vassilvitskii [2021] and Purohit et al. [2018], there have been many new and interesting results in many different problems within the last years, including ski rental [Wei and Zhang, 2020], online bipartite matching [Lavastida et al., 2021], load balancing [Ahmadian et al., 2023], and many more (see e.g. Im et al. [2021], Zeynali et al. [2021], Almanza et al. [2021]). Since this area is developing very fast, we refer the reader to the website Algorithms-with-Predictions for references of literature.

As mentioned before, also the Secretary problem itself has been studied in this framework. Antoniadis et al. [2020] consider the Secretary problem when the machine learned advice predicts the weight of the largest element w_1. Their algorithm's performance depends on the error of the prediction as well as some confidence parameter by how much the decision maker trusts the advice. In complementary work, Dütting et al. [2021b] give a bigger picture for Secretary problems with machine learned advice. Their approach is LP based and can capture a variety of settings. They assume that the prediction is one variable for each weight (e.g. a 0/1-variable indicating if the current element is the best overall or not). Fujii and Yoshida [2023] assume an even stronger prediction: Their algorithm has access to a prediction for every weight in the sequence. In contrast, we go into the opposite direction and deal with a less informative piece of information in this chapter.

Chapter Organization and Remarks

This chapter is based on *The Secretary Problem with Predicted Additive Gap* [Braun and Sarkar, 2023], which is joint work with Sherry Sarkar. More detailed bibliographic notes can be found in Section 1.5.

In Section 7.2, we show how to beat $1/e$ when having access to the additive gap exactly. As this might be too much to hope for, we complement this result by studying the robustness-consistency trade-off in Section 7.3. It is crucial to mention that the error in the prediction might be unbounded. To complement this, when having an upper bound on the error of the prediction, we can get an improved guarantee, as shown in Section 7.4. In addition to the theoretical results, simulations can be found in Section 7.5 to support our theoretical findings.

7.1 Preliminaries

Recall that in the classical Secretary problem as introduced in Section 1.1.2, an adversary fixes n non-negative, real-valued weights, denoted $w_1 \geq w_2 \geq \cdots \geq w_n$. For each, there is an *arrival time* $t_i \overset{\text{iid}}{\sim} \text{Unif}[0,1]$. Weight w_i is revealed at time t_i and we immediately and irrevocably need to decide if we want to accept or reject this element.

Our goal is to maximize the selected weight and we measure the quality of an algorithm by the competitive ratio of $\text{OPT} = \max_i w_i$ compared to $\mathbf{E}[\text{ALG}]$, where the expectation is taken over the random arrival times of elements (and possible internal randomness of the algorithm).

In addition to the random arrival order, we assume to have access to a single prediction $\hat{c}_k$ for one *additive gap* together with its index k. The additive gap for some index $2 \leq k \leq n$ is $c_k := w_1 - w_k$. We say that an algorithm has access to an *exact* or *accurate* gap if $\hat{c}_k = c_k$ (as in Section 7.2).

When the algorithm gets a predicted additive gap $\hat{c}_k$ which might not be accurate (as in Section 7.3 or Section 7.4), we say that $\hat{c}_k$ has *error* $\eta = |\hat{c}_k - c_k|$. We call an algorithm *ρ-robust* if the algorithm is ρ-competitive regardless of error η and we say the algorithm is *ψ-consistent* if the algorithm is ψ-competitive when $\eta = 0$.

When clear from the context, we drop the index k at the gap and only call the gap c or $\hat{c}$ respectively.

7.2 Knowing an Exact Gap

We start with the setting of getting an accurate prediction for the gap, before diving into the cases where the predicted gap may be inaccurate in Section 7.3 and Section 7.4. Crucially for this section, we are given the exact gap $c_k = w_1 - w_k$ for some $2 \leq k \leq n$. We assume that we get to know the index k as well as the value of c_k, but neither w_1 nor w_k.

Our algorithm takes as input the gap c as well as a waiting time τ. This gives us the freedom to potentially choose τ independent of k if required. As a consequence, we could make the algorithm oblivious to the index k of the element to which the gap is revealed. We will use this in Corollary 7.2.4.

Algorithm 13: Secretary with Exact Additive Gap

```
Input: Additive gap c, time τ ∈ [0, 1]
Before time τ:
    Observe weights w_i
At time τ:
    Compute BSF(τ) = max_{i:t_i ≤ τ} w_i
After time τ:
    Accept first element with w_i ≥ max(BSF(τ), c)
```

This algorithm beats the prevalent competitive ratio of $1/e \approx 0.368$ by a constant.

Theorem 7.2.1. *Given any additive gap $c_k = w_1 - w_k$, for $\tau = 1 - (1/k+1)^{1/k}$, Algorithm 13 achieves a competitive ratio of*

$$\max\left(0.4, \frac{1}{2}\left(\frac{1}{k+1}\right)^{1/k}\right) .$$

Note that as k tends towards n and both become large, the competitive ratio approaches $1/2$.

We split the proof of Theorem 7.2.1 in the following two lemmas. Each of them gives a suitable bound on the performance of our algorithm for general waiting times τ in settings when w_k is small or large. We will only plug in our choice of τ in the conclusion later.

The first lemma gives a lower bound in cases when w_k is small in comparison to w_1.

Lemma 7.2.2. *If $w_k < \frac{1}{2}w_1$, then*

$$\mathbf{E}\left[\mathrm{ALG}\right] \geq (1-\tau)\left(\frac{1}{2} + \frac{1}{2(k-1)}\right) \cdot w_1 .$$

The second lemma will be used to give a bound when w_k is large compared to w_1. More precisely, we will give two bounds which both capture the weight achieved by the algorithm. The first bound is tighter for larger k, the second bound is independent of k.

Lemma 7.2.3. *If $w_k \geq \frac{1}{2}w_1$, then the following two bounds hold:*

(i) $\mathbf{E}[\text{ALG}] \geq \frac{k+1}{2k}\left(1-\tau-(1-\tau)^{k+1}\right)\cdot w_1$ *and*

(ii) $\mathbf{E}[\text{ALG}] \geq \left(\frac{3}{2}\tau\ln\left(\frac{1}{\tau}\right)-\frac{1}{2}\tau(1-\tau)\right)\cdot w_1$.

As a consequence, $\mathbf{E}[\text{ALG}]$ *is also at least as large as the maximum of the two bounds.*

We start with a proof of Lemma 7.2.2 when w_k is small compared to w_1.

Proof of Lemma 7.2.2. Let $w_k < \frac{1}{2}w_1$. Observe that in this case, the gap is quite large, as

$$c_k = w_1 - w_k > w_1 - \frac{1}{2}w_1 = \frac{1}{2}w_1 > w_k \geq \cdots \geq w_n \ .$$

In particular, the gap is large enough such that the algorithm either selects nothing (if w_1 arrives before τ) or some element among $w_1, \ldots, w_l$ for some $1 \leq l \leq k-1$. To see this, first observe that the gap never overshoots w_1 as we are always ensured that $w_1 \geq c_k$. In addition, we only accept elements strictly larger than w_k in this case. Let l be the index of the element for which $w_l \geq c_k > w_{l+1}$, i.e. the last element which is not excluded from a possible acceptance by c_k.

Hence, we can bound

$$\begin{aligned}
\mathbf{E}[\text{ALG}] &= \mathbf{Pr}[w_1 \text{ arrives after } \tau]\cdot\mathbf{E}[\text{ALG} \mid w_1 \text{ arrives after } \tau] \\
&\geq (1-\tau)\left(\frac{1}{2}w_1\frac{l-1}{l}+\frac{1}{l}w_1\right) \\
&\geq w_1(1-\tau)\left(\frac{1}{2}+\frac{1}{2(k-1)}\right) \ .
\end{aligned}$$

To see why the first inequality holds, note that the probability of an element to arrive after τ is precisely $1-\tau$. In addition, conditioned on the best element arriving after τ, we are ensured to accept some element among the first l elements. By the random arrival times, we accept the best element w_1 in at least a $1/l$-fraction of scenarios. In the remainder, we will select an element of weight at least $\frac{1}{2}w_1$. The second inequality uses $l \leq k-1$. □

In this case, we actually accept w_1 with much higher probability than $1/l$. In particular, observe that we exclude $w_{l+1}, \ldots, w_n$ by the gap in the threshold, so the problem boils down to solving a Secretary instance with l elements. However, the bound presented in the proof of Lemma 7.2.2 is sufficient for our purposes.

Next, we turn to the regime when w_k is large and prove the two bounds from Lemma 7.2.3.

Proof of Lemma 7.2.3 (i). Let $w_k \geq \frac{1}{2}w_1$. In this case, the gap can be quite small. Still, we are guaranteed that selecting any element among $w_2, \ldots, w_k$ achieves at least a weight of $\frac{1}{2}w_1$.

We condition on the event $w_i = \text{BSF}(\tau)$ for $i \in \{2, \ldots, n\}$. For any $i \leq k+1$, having element w_i as $\text{BSF}(\tau)$, we will always accept the first element among $w_1, \ldots, w_{i-1}$ arriving after τ. To see this, note that none of these elements will be excluded by the gap c_k in the threshold of the algorithm as

$$c_k = w_1 - w_k \leq w_1 - \frac{1}{2}w_1 = \frac{1}{2}w_1 \leq w_k \ .$$

We first give a bound on the expected weight achieved by the algorithm conditioned on seeing w_i as the BSF(τ) for some $2 \le i \le k+1$. When seeing w_i as the BSF(τ), we select w_1 in a $\frac{1}{i-1}$-fraction of scenarios. In addition, any element $w_2, \dots, w_{i-1}$ is ensured to have a weight at least $\frac{1}{2}w_1$. As a consequence, using that we only consider $i \le k+1$,

$$\mathbf{E}\left[\text{ALG} \mid w_i \text{ is BSF}(\tau)\right] \ge \frac{1}{i-1}w_1 + \frac{i-2}{i-1} \cdot \frac{1}{2}w_1 \ge \frac{1}{2}\left(1+\frac{1}{k}\right) \cdot w_1 \ . \tag{7.1}$$

Using this, we can derive the following lower bound. Note that if w_1 is BSF(τ), we will select nothing:

$$\begin{aligned}
\mathbf{E}\left[\text{ALG}\right] &\ge \sum_{i=2}^{k+1} \mathbf{Pr}\left[w_i \text{ is BSF}(\tau)\right] \cdot \mathbf{E}\left[\text{ALG} \mid w_i \text{ is BSF}(\tau)\right] \\
&= \sum_{i=2}^{k+1} \tau(1-\tau)^{i-1} \cdot \mathbf{E}\left[\text{ALG} \mid w_i \text{ is BSF}(\tau)\right] \\
&\overset{(7.1)}{\ge} w_1 \cdot \frac{1}{2}\left(1+\frac{1}{k}\right)\tau \sum_{i=2}^{k+1}(1-\tau)^{i-1} \\
&= w_1 \cdot \frac{1}{2}\left(1+\frac{1}{k}\right)\tau\left(\frac{1-(1-\tau)^{k+1}}{\tau} - 1\right) \\
&= w_1 \cdot \frac{1}{2}\left(1+\frac{1}{k}\right)\left(1-\tau-(1-\tau)^{k+1}\right) \ .
\end{aligned}$$

The second equality uses a geometric sum to simplify the expression. □

In addition, we can use an alternate analysis, which is tighter for small k.

Proof of Lemma 7.2.3 (ii). Let $w_k \ge \frac{1}{2}w_1$. We only consider the probability of selecting the best or second best element. Observe that the second best element satisfies $w_2 \ge w_k \ge \frac{1}{2}w_1$ by the case distinction, no matter for which k we observe the gap. Our goal will be to lower bound the acceptance probabilities of w_1 and w_2 with the ones from the classical Secretary problem. To this end, we first observe that for any element w_i, if $w_i \ge c_k$, then

$$\mathbf{Pr}\left[\text{Algorithm 13 selects } w_i\right] \ge \mathbf{Pr}\left[\text{An algorithm with threshold BSF}(\tau) \text{ selects } w_i\right] \ . \tag{7.2}$$

To see why this inequality holds, let w_i arrive after time τ. If w_i is the first element to surpass BSF(τ), by the hypothesis that $w_i \ge c_k$, w_i is also the first to surpass the threshold used by Algorithm 13.

Using this, we can lower bound the selection probabilities of w_1 and w_2 by the ones of an algorithm which is just using BSF(τ) as a threshold. The next few lines of calculations are very similar to folklore techniques for the classical Secretary problem, and we have seen those arguments in the proof of Theorem 1.1.3.

To bound the probability of selecting w_1 by an algorithm which uses BSF(τ) as a threshold, observe that if w_1 arrives before τ, it will be rejected. If it arrives at some time $x \in (\tau, 1]$, there are two cases in which we accept w_1: Either no other element arrived before (by the i.i.d. arrival times, this happens with probability $(1-x)^{n-1}$) or the best element before did arrive before τ. The latter happens with probability $\frac{\tau}{x}$. As a consequence, we get

$$\begin{aligned}
\mathbf{Pr}\left[\text{Algorithm 13 selects } w_1\right] &\overset{(7.2)}{\geq} \mathbf{Pr}\left[\text{An algorithm with threshold BSF}(\tau) \text{ selects } w_1\right] \\
&= \int_{\tau}^{1} (1-x)^{n-1} + \left(1-(1-x)^{n-1}\right)\frac{\tau}{x}\, dx \\
&\geq \int_{\tau}^{1} \frac{\tau}{x}\, dx = \tau \ln\left(\frac{1}{\tau}\right).
\end{aligned}$$

Similarly, for w_2 we can compute the same integral after conditioning on w_1 arriving after time x. Observe that the probability of w_1 arriving after time x is precisely $1-x$.

$$\begin{aligned}
\mathbf{Pr}\left[\text{Algorithm 13 selects } w_2\right] &\overset{(7.2)}{\geq} \mathbf{Pr}\left[\text{An algorithm with threshold BSF}(\tau) \text{ selects } w_2\right] \\
&= \int_{\tau}^{1} (1-x)\left((1-x)^{n-2} + \left(1-(1-x)^{n-2}\right)\frac{\tau}{x}\right) dx \\
&\geq \int_{\tau}^{1} (1-x)\frac{\tau}{x}\, dx = \tau \ln\left(\frac{1}{\tau}\right) - \tau(1-\tau)\ .
\end{aligned}$$

Using that $w_2 \geq \frac{1}{2}w_1$, we obtain

$$\begin{aligned}
\mathbf{E}\left[\text{ALG}\right] &\geq w_1 \tau \ln\left(\frac{1}{\tau}\right) + w_2\left(\tau \ln\left(\frac{1}{\tau}\right) - \tau(1-\tau)\right) \\
&\geq w_1\left(\tau \ln\left(\frac{1}{\tau}\right) + \frac{1}{2}\left(\tau \ln\left(\frac{1}{\tau}\right) - \tau(1-\tau)\right)\right) \\
&= w_1\left(\frac{3}{2}\tau \ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)\right)\ .
\end{aligned}$$

□

Having these two lemmas, we can now conclude the proof of the main theorem.

Proof of Theorem 7.2.1. We use the lower bound obtained by Lemma 7.2.2. From Lemma 7.2.3, we take the maximum of the two bounds into consideration. Since we do not know upfront to which case our instance belongs, we can only obtain the minimum of the two as a general lower bound on the weight achieved by the algorithm.

As a consequence, we obtain $\mathbf{E}\left[\text{ALG}\right] \geq \alpha \cdot w_1$ for

$$\alpha := \min\left(\frac{(1-\tau)k}{2(k-1)}; \max\left(\frac{k+1}{2k}\left(1-\tau-(1-\tau)^{k+1}\right); \frac{3}{2}\tau\ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)\right)\right) \tag{7.3}$$

which depends on the waiting time τ. Now, we plug in $\tau = 1 - \left(\frac{1}{k+1}\right)^{1/k}$. First, note that we can bound the maximum in α with the first of the two terms. Factoring out a $1-\tau$, we get

$$\begin{aligned}
\alpha &\geq (1-\tau) \cdot \min\left(\frac{k}{2(k-1)}; \frac{k+1}{2k}\left(1-(1-\tau)^k\right)\right) \\
&= \left(\frac{1}{k+1}\right)^{1/k} \cdot \min\left(\frac{k}{2(k-1)}; \frac{k+1}{2k} \cdot \frac{k}{k+1}\right) \\
&= \frac{1}{2}\left(\frac{1}{k+1}\right)^{1/k}\ .
\end{aligned}$$

To compensate for the poor performance of this bound for small k, we can use basic calculus to state the following.

After plugging in our choice of τ into Expression (7.3), the first term is minimized for $k = 7$ for a value of at least 0.43. For $2 \leq k \leq 11$, the last term is always at least 0.404 and for any $k \geq 12$, the second term exceeds 0.403. Hence, we always ensure that $\alpha \geq 0.4$. □

As a corollary of the proof of Theorem 7.2.1, we also get a lower bound on the weight achieved by the algorithm if we are only given the gap, but not the element which obtains this gap. That is, we are given c_k but not the index k.

Corollary 7.2.4. *If the algorithm only has access to c_k, but not k, setting $\tau = 0.2$ achieves $\mathbf{E}[\text{ALG}] \geq 0.4 \cdot w_1$.*

The proof mainly relies on the fact that the lower bound we obtained in the proof of Theorem 7.2.1 holds for any choice $\tau \in [0, 1]$. Also, the algorithm itself only uses the gap to contribute to the threshold. The index k is only used to compute τ. As a consequence, when choosing $\tau = 0.2$ independent of k, the algorithm is oblivious to the exact value of k, but only depends on the gap c_k. For this choice of τ, we can show that $\alpha \geq 0.4$.

Proof of Corollary 7.2.4. Note that the lower bound in Expression (7.3) holds for any choice of $\tau \in [0, 1]$. For $\tau = 0.2$, the value of α in Expression (7.3) satisfies

$$\begin{aligned}\alpha &= \min\left(0.8 \cdot \frac{k}{2(k-1)}; \max\left(\frac{k+1}{2k}\left(0.8 - 0.8^{k+1}\right); 0.3\ln(5) - 0.08\right)\right)\\ &\geq \min\left(0.8 \cdot \frac{1}{2}; \max\left(\frac{k+1}{2k}\left(0.8 - 0.8^{k+1}\right); 0.4\right)\right)\\ &\geq 0.4\end{aligned}$$

and hence, we get a competitive ratio of at least 0.4. □

As a consequence, very surprisingly, even if we only get to know *some* additive gap c_k and not even the index k, we can outperform the prevalent bound of $1/e$. Also, observe that this is independent of the exact value that c_k attains and holds for any small or large gaps.

As mentioned before, Algorithm 13 is required to get the exact gap as input. In particular, once the gap we use in the algorithm is a tiny bit larger than the actual gap c_k, we might end up selecting no element at all. In order to see this, we consider the following example.

Example 7.2.5. *We get to know the gap to the smallest weight $c_n = w_1 - w_n$ and the smallest weight w_n in the sequence satisfies $w_n = 0$. Let the gap which we use in Algorithm 13 be only some tiny $\delta > 0$ too large. In other words, we use $c_n + \delta$ as a gap in the algorithm. Still, this implies that our threshold $\max(\text{BSF}(\tau), c_n + \delta)$ after the waiting time satisfies*

$$\max(\text{BSF}(\tau), c_n + \delta) \geq c_n + \delta = w_1 + \delta > w_1 \geq w_2 \geq \cdots \geq w_n \ .$$

As a consequence, we end up selecting no weight at all and have $\mathbf{E}[\text{ALG}] = 0$.

This naturally motivates the need to introduce more robust algorithms in this setting. To this end, the results in Section 7.3 will show that a slight modification in the algorithm and its analysis allows to obtain robustness to errors in the predictions while still outperforming $1/e$ for accurate gaps.

7.3 Robustness-Consistency Trade-off

Motivated by the weakness of Algorithm 13 revealed in Example 7.2.5, we now shift our perspective towards the following question: how well can an algorithm perform when equipped with a potentially erroneous prediction?

We show how to modify our algorithm in order to still beat $1/e$ when getting the correct gap as input, but still be constant competitive in case the predicted gap is inaccurate. The modification leads to Algorithm 14 and works as follows: Initially, we run the same algorithm as before. After a certain amount of time $1-\gamma$, if the algorithm has not terminated yet, we will lower our threshold in order to hedge against an incorrect prediction.

Algorithm 14: Robust-Consistent Algorithm

Input: Predicted gap $\hat{c}$, times $\tau \in [0,1)$, $\gamma \in [0, 1-\tau)$
Before time τ:
    Observe weights w_i
At time τ:
    Compute $\mathrm{BSF}(\tau) = \max_{i:t_i \leq \tau} w_i$
Between time τ and time $1-\gamma$:
    Accept first element with $w_i \geq \max(\mathrm{BSF}(\tau), \hat{c})$
After time $1-\gamma$:
    Accept first element with $w_i \geq \mathrm{BSF}(\tau)$

Note that by $\gamma \in [0, 1-\tau)$, we ensure that $\tau < 1-\gamma$, i.e. the waiting time τ is not after time $1-\gamma$ and hence, the algorithm is well-defined. Now, we can state the following theorem which gives guarantees on the consistency and the robustness of Algorithm 14. We will discuss afterwards how to choose τ and γ in order to outperform the classical bound of $1/e$ by a constant for accurate predictions while satisfying constant robustness simultaneously.

Theorem 7.3.1. *Given a prediction $\hat{c}_k$ for the additive gap c_k, define*

- $\alpha_1 := \frac{(1-\tau-\gamma)\cdot k}{2(k-1)}$,
- $\alpha_2 := \frac{k+1}{2k}\left(1-\tau-(1-\tau)^{k+1}\right)$ *and*
- $\alpha_3 := \frac{3}{2}\tau \ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)$.

Then, Algorithm 14 is

(i) $\min\left(\alpha_1, \max\left(\alpha_2, \alpha_3\right)\right)$*-consistent and*

(ii) $\min\left(\tau\gamma, \tau \ln\left(1/\tau\right)\right)$*-robust.*

We split the proof of Theorem 7.3.1 into two parts: first, we argue about the consistency of our algorithm, second, we show that it is also robust.

Proof of Theorem 7.3.1 (i). For consistency, we assume that our prediction error is zero, hence our predicted gap $\hat{c}_k$ equals the actual gap c_k. We can perform the same case distinction as we did in the proof of Theorem 7.2.1.

In the first case, let $w_k < \frac{1}{2}w_1$, hence,

$$\hat{c}_k = c_k = w_1 - w_k > \frac{1}{2}w_1 \geq w_k \geq w_{k+1} \geq \cdots \geq w_n \ .$$

Now, we lower bound the expected weight obtained by Algorithm 14 via the expected weight if w_1 arrives between time τ and $1-\gamma$.

$$\begin{aligned}\mathbf{E}[\text{ALG}] &\geq \mathbf{Pr}[w_1 \text{ in } [\tau, 1-\gamma]] \cdot \mathbf{E}[\text{ALG} \mid w_1 \text{ in } [\tau, 1-\gamma]] \\ &= (1-\tau-\gamma) \cdot \mathbf{E}[\text{ALG} \mid w_1 \text{ in } [\tau, 1-\gamma]] \\ &\geq w_1 \left(\frac{(1-\tau-\gamma)\cdot k}{2(k-1)} \right) \ .\end{aligned}$$

Here, the first inequality lower bounds the expected weight of our algorithm obtained when w_1 arrives after $1-\gamma$ by zero. In particular, we only considers contributions which are made if w_1 arrives in $[\tau, 1-\gamma]$. The second inequality uses the same reasoning as in the proof of Lemma 7.2.2. So, for this case, we have $\mathbf{E}[\text{ALG}] \geq \alpha_1 \cdot w_1$.

In the second case, let $w_k \geq \frac{1}{2}w_1$, hence,

$$\hat{c}_k = c_k = w_1 - w_k \leq \frac{1}{2}w_1 \leq w_k \ .$$

Interestingly, the analysis of Algorithm 13 directly carries over in this case. Recall in this case, we only relied on the gap not excluding high weight elements (with indices $1, \ldots, k$), and dropping the gap as a threshold after time $1-\gamma$ all together preserves this property. When following the proof of Lemma 7.2.3 step by step, we can use exactly the same arguments also for Algorithm 14. Hence, we get the same bounds as in Lemma 7.2.3 (i) and (ii). Therefore, defining $\alpha_2 := \frac{k+1}{2k}\left(1-\tau-(1-\tau)^{k+1}\right)$ and $\alpha_3 := \frac{3}{2}\tau \ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)$, we obtain that

$$\mathbf{E}[\text{ALG}] \geq \alpha \cdot w_1$$

for $\alpha = \min(\alpha_1, \max(\alpha_2, \alpha_3))$. □

With the desired consistency guarantees in mind, we can now shift our perspective towards robustness.

Proof of Theorem 7.3.1 (ii). For robustness, we need to protect our algorithm against inaccurate gaps – no matter how bad the predicted gap is. To this end, we consider two cases.

When the predicted gap is larger than the actual gap (and perhaps even overshoots w_1), note that Algorithm 14 will correctly pick w_1 if w_2 appears before time τ and w_1 appears after time $1-\gamma$. Therefore, Algorithm 14 will pick w_1 with probability at least $\tau \cdot \gamma$ and hence, is at least $(\tau \cdot \gamma)$-competitive in this case.

When the predicted gap is smaller than the actual gap (and in particular, does not overshoot w_1 and therefore is still a valid threshold), we will select w_1 with probability at least $\tau \ln(1/\tau)$, as detailed in the proof of Lemma 7.2.3 (ii). Taking the minimum of these two cases yields the desired robustness. □

For example, when using a waiting time $\tau = 0.2$ as in Corollary 7.2.4 independent of the index k and a value of $\gamma = 0.05$, we get the following: Algorithm 14 is at least 0.375-consistent and 0.01-robust. In particular, we can outperform the prevalent bound of $1/e$ by a constant if the predicted gap is accurate while ensuring to be constant competitive even if our predicted gap is far off. Of course, when being more risk averse, one could also increase the robustness guarantee for the cost of decreasing the competitive ratio in the consistent case. Still, we note that the analysis for robustness as well as consistency is not tight. Finding a tight trade-off remains an open problem.

As a side remark, we highlight that these guarantees as well as Theorem 7.3.1 hold independent of any bounds on the error of the predicted gap. However, eventually, it might be reasonable to assume that the gap is bounded. We turn towards this setting in Section 7.4 and show that we can achieve much better competitive ratios when we know a range for the error.

7.4 Improved Guarantees for Bounded Errors

Complementing the previous sections where we had either access to the exact gap (Section 7.2) or no information on a possible error of the prediction (Section 7.3), we now assume that the error is bounded[1]. That is, we get to know some $\tilde{c}_k \in [c_k - \epsilon; c_k + \epsilon]$ which is ensured to be at most an ϵ off. Also, the bound ϵ on the error is revealed to us. Still, the true gap c_k remains unknown.

Our algorithm follows the template which we discussed before. Still, we slightly perturb $\tilde{c}_k$ to ensure that the threshold is not exceeding w_1.

Algorithm 15: Secretary with Bounded Prediction Error

```
Input: Approximate gap c̃, time τ ∈ [0,1], error bound ϵ
Before time τ:
    Observe weights w_i
At time τ:
    Compute BSF(τ) = max_{i:t_i ≤ τ} w_i
After time τ:
    Accept first element with w_i ≥ max(BSF(τ), c̃ − ϵ)
```

This algorithm allows to state an approximate version of Theorem 7.2.1. As a matter of fact, we can obtain the same lower bounds of α as in the exact gap case. Still, we suffer an additive loss depending on the range of the error ϵ.

Theorem 7.4.1. *Given any prediction of the gap $\tilde{c}_k \in [c_k - \epsilon; c_k + \epsilon]$, where $c_k = w_1 - w_k$, Algorithm 15 satisfies $\mathbf{E}[\mathrm{ALG}] \geq \alpha \cdot w_1 - 2\epsilon$.*

[1] In order to distinguish a bounded error from a possibly unbounded one, we use $\tilde{c}$ instead of $\hat{c}$ for the predicted gap in this section.

As in Section 7.2, for $\tau = 1 - (1/k+1)^{1/k}$ we get $\alpha \geq \max\left(0.4, \frac{1}{2}(1/k+1)^{1/k}\right)$ and for waiting time $\tau = 0.2$, we are still guaranteed $\alpha \geq 0.4$. Hence, also the results when not knowing the index k carry over. In particular, this nicely complements the robustness result from Theorem 7.3.1 as follows: Once we can bound the error in a reasonable range, even not knowing the gap exactly does not cause too much of an issue.

In the proof, we perform a case distinction whether w_k is at least $\frac{1}{2}w_1 - 2\epsilon$ or not. Afterwards, we derive bounds for the two cases separately and combine them to obtain the competitive ratio, similar to the proof of Theorem 7.2.1.

Proof of Theorem 7.4.1. First, we argue that the threshold in the algorithm is never too high to avoid acceptance of w_1 if w_1 arrives after τ. To see this, note that $\tilde{c}_k - \epsilon \leq (c_k + \epsilon) - \epsilon = c_k = w_1 - w_k \leq w_1$.

For the case distinction, we consider the cases that w_k is small or large with respect to w_1. Still, we need to incorporate the fact that we know a prediction with bounded error and not the exact gap.

Case 1. $w_k < \frac{1}{2}w_1 - 2\epsilon$.

Observe that in this case, the term $\tilde{c}_k - \epsilon$ in the threshold is quite large. In particular,

$$\tilde{c}_k - \epsilon \geq (c_k - \epsilon) - \epsilon = w_1 - w_k - 2\epsilon > w_1 - \left(\frac{1}{2}w_1 - 2\epsilon\right) - 2\epsilon = \frac{1}{2}w_1 \ . \tag{7.4}$$

As mentioned before, $\tilde{c}_k - \epsilon$ never exceeds w_1. Also, observe that $\frac{1}{2}w_1 > \frac{1}{2}w_1 - 2\epsilon > w_k \geq \cdots \geq w_n$ by our case distinction.

As a consequence, in this case, the algorithm either selects nothing or some element among $w_1, \ldots, w_l$ for some $1 \leq l \leq k-1$. As in the proof in Section 7.2, we can define l to be the index of element with $w_l \geq \tilde{c}_k - \epsilon > w_{l+1}$. Also, any element which is selected has a weight of at least $\frac{1}{2}w_1$ by Inequality (7.4).

Hence, we achieve the same bound as in the exact gap scenario of

$$\begin{aligned}
\mathbf{E}\left[\mathrm{ALG}\right] &\geq \mathbf{Pr}\left[w_1 \text{ arrives after } \tau\right] \cdot \mathbf{E}\left[\mathrm{ALG} \mid w_1 \text{ arrives after } \tau\right] \\
&\geq (1-\tau)\left(\frac{1}{2}w_1\frac{l-1}{l} + \frac{1}{l}w_1\right) \\
&\geq w_1(1-\tau)\left(\frac{1}{2} + \frac{1}{2(k-1)}\right) \ .
\end{aligned}$$

Case 2. $w_k \geq \frac{1}{2}w_1 - 2\epsilon$.

Observe that selecting any element among $w_2, \ldots, w_k$ achieves at least a weight of $\frac{1}{2}w_1 - 2\epsilon$. We now bound the expected weight of the algorithm in a similar way as in Section 7.2 by deriving two lower bounds.

Bound (i):

We condition on seeing elements $w_2, \ldots, w_{k+1}$ as $\mathrm{BSF}(\tau)$. As before,

$$\begin{aligned}
\mathbf{E}[\text{ALG}] &\geq \sum_{i=2}^{k+1} \mathbf{Pr}[w_i \text{ is BSF}(\tau)] \cdot \mathbf{E}[\text{ALG} \mid w_i \text{ is BSF}(\tau)] \\
&= \sum_{i=2}^{k+1} \tau(1-\tau)^{i-1} \cdot \mathbf{E}[\text{ALG} \mid w_i \text{ is BSF}(\tau)] \\
&\geq \sum_{i=2}^{k+1} \tau(1-\tau)^{i-1} \cdot \left(\frac{1}{i-1} w_1 + \frac{i-2}{i-1} \cdot \left(\frac{1}{2} w_1 - 2\epsilon\right)\right) \\
&\geq \left(w_1 \cdot \frac{1}{2}\left(1+\frac{1}{k}\right) - 2\epsilon\right) \tau \sum_{i=2}^{k+1} (1-\tau)^{i-1} \\
&\geq w_1 \cdot \frac{1}{2}\left(1+\frac{1}{k}\right)\left(1-\tau-(1-\tau)^{k+1}\right) - 2\epsilon \ .
\end{aligned}$$

The only difference to Section 7.2 is the lower bound for $w_2, \ldots, w_k$ which are only guaranteed to be at least $\frac{1}{2}w_1 - 2\epsilon$.

Bound (ii):

As before, in order to compensate for the weak lower bound in the small k regime, we only consider the probability of selecting the best or second best element. Observe that Inequality (7.2) also holds for Algorithm 15 when replacing the condition on the weights with $w_i \geq \tilde{c}_k - \epsilon$.

Still, for the following reason, we need to argue in a slightly different way than in the exact gap case. Setting the contribution to the threshold to $\tilde{c}_k - \epsilon$ ensures that the threshold never overshoots w_1. Still, the weight w_2 can now fall in two ranges: (a) $w_2 \geq \tilde{c}_k - \epsilon$ and hence, w_2 is not affected by the gap in the threshold or (b) $w_2 < \tilde{c}_k - \epsilon$ in which case the algorithm does not select w_2 as the threshold is too high. The latter case might occur as $w_2 \geq w_k \geq \frac{1}{2}w_1 - 2\epsilon$, but $\tilde{c}_k - \epsilon \leq c_k + \epsilon - \epsilon = w_1 - w_k \leq \frac{1}{2}w_1 + 2\epsilon$. Still, this will not introduce any problems.

Concerning (a), we argue as before. The second best element satisfies $w_2 \geq w_k \geq \frac{1}{2}w_1 - 2\epsilon$ by the case distinction. Similar to the proof for the exact gap, we bound the probabilities of selecting w_1 or w_2 and use the lower bound on w_2. This implies that

$$\begin{aligned}
\mathbf{E}[\text{ALG}] &\geq w_1 \tau \ln\left(\frac{1}{\tau}\right) + w_2\left(\tau \ln\left(\frac{1}{\tau}\right) - \tau(1-\tau)\right) \\
&\geq w_1 \tau \ln\left(\frac{1}{\tau}\right) + \left(\frac{1}{2}w_1 - 2\epsilon\right)\left(\tau \ln\left(\frac{1}{\tau}\right) - \tau(1-\tau)\right) \\
&\geq w_1\left(\frac{3}{2}\tau \ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)\right) - 2\epsilon \ .
\end{aligned}$$

Concerning (b), note that if w_2 is excluded from acceptance by the threshold, so is any $w_i \neq w_1$. Hence, the algorithm will always either select nothing (if w_1 appears before τ) or accept w_1 (if w_1 appears after τ). Here it is important that the contribution of $\tilde{c}_k - \epsilon$ never exceeds w_1. As a consequence, in this case,

$$\mathbf{E}[\text{ALG}] = w_1(1-\tau) \geq w_1\left(\frac{3}{2}\tau \ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)\right) - 2\epsilon \ .$$

Combination. When combining everything, we obtain $\mathbf{E}[\mathrm{ALG}] \geq \alpha \cdot w_1 - 2\epsilon$ for

$$\alpha := \min\left((1-\tau)\frac{k}{2(k-1)}; \max\left(\frac{k+1}{2k}\left(1-\tau-(1-\tau)^{k+1}\right); \frac{3}{2}\tau\ln\left(\frac{1}{\tau}\right) - \frac{1}{2}\tau(1-\tau)\right)\right) ,$$

which is the same bound as in Expression (7.3) with an additional additive loss of 2ϵ. □

7.5 Simulations

In order to gain a more fine-grained understanding of the underlying habits, we run experiments[2] with simulated weights and compare our algorithms among each other and to the classical Secretary algorithm.

In Section 7.5.1, we compare our Algorithm 13 to the classical Secretary algorithm stated in Algorithm 2. To this end, we draw weights i.i.d. from distributions and execute our algorithm and the classical one. As we will see, instances which are hard in the normal Secretary setting (without knowing any additive gap) become significantly easier with additive gap; we can select the best candidate with a much higher probability. We also demonstrate that for some instances, knowing the gap has a smaller impact, though Algorithm 13 still outperforms the classical one.

Second, in Section 7.5.2, we focus on inaccurate gaps and compare Algorithm 13 developed in Section 7.2 to the robust and consistent Algorithm 14 from Section 7.3. As a matter of fact, we will see that underestimating the exact gap is not as much of an issue as an overestimation. In particular, underestimating the gap implies a smooth decay in the competitive ratio while overestimating can immediately lead to a huge drop. Note that these findings align with Example 7.2.5, in which overestimating turned out to ruin all guarantees.

7.5.1 The Impact of Knowing the Gap

We compare our algorithm with additive gap to the classical Secretary algorithm (stated in Algorithm 2, see e.g. [Dynkin, 1963]) with a waiting time of $1/e$.

Experimental Setup

We run the comparison on three different classes of instances:

(i) *Pareto*: We first draw some $\theta \sim \mathrm{Pareto}(5/n, 1)$. Afterwards, each weight w_i is determined as follows: Draw $Y_i \sim \mathrm{Unif}[0, \theta]$ i.i.d. and set $w_i = Y_i^{\left(n^{1.5}\right)}$ (for more details on Pareto distributions and Secretary problems, see e.g. Ferguson [1989]).

(ii) *Exponential*: Here, all $w_i \sim \mathrm{Exp}(1)$.

(iii) *Chi-Squared*: Draw $w_i \sim \chi^2(10)$. That is, each w_i is drawn from a chi-squared distribution which sums over ten squared i.i.d. standard normal random variables.

For each class of instances, we average over 5000 iterations. In each iteration, we draw $n = 200$ weights i.i.d. from the respective distribution together with 200 arrival times which are drawn i.i.d. from $\mathrm{Unif}[0, 1]$.

[2] All experiments were implemented in Python 3.9 and executed on a machine with Apple M1 and 8 GB Memory.

The benchmark is Algorithm 2 with a waiting time of $\tau = 1/e$: Set the largest weight up to time τ as a threshold and accepts the first element afterwards exceeding this threshold. Algorithm 13 is executed with waiting times $\tau = 0.2$ as well as $\tau = 1 - (1/k+1)^{1/k}$.

Experimental Results

When weights are sampled based on the procedure explained in (i), we observe an interesting phenomenon (see Figure 7.1). For the classical Secretary algorithm, we achieve approximately the tight guarantee of $1/e \approx 0.368$. Our algorithm, however, achieves a competitive ratio of approximately 0.8 for $\tau = 0.2$. When having a waiting time depending on k, we improve the competitive ratio for large k while suffering a worse ratio for small k.

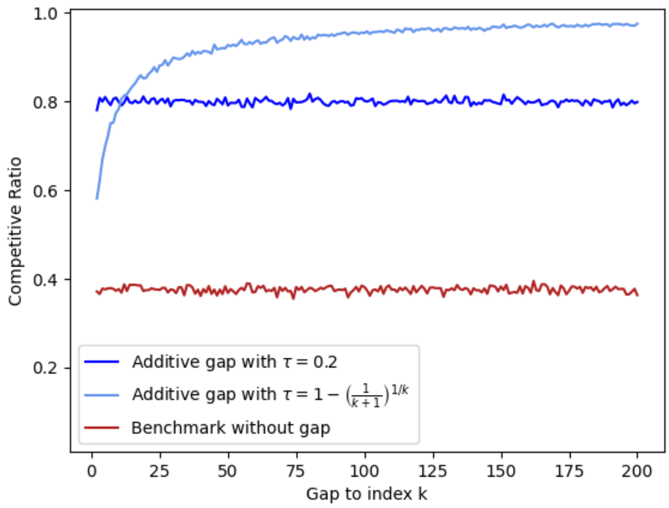

Figure 7.1: Competitive ratios for weights based on (i). On the x-axis, we have the index k from 2 to n. The y-axis shows the competitive ratios.

This can be explained as follows. Weights which are distributed according to (i) almost always have a very large gap between the highest and second highest weight. Hence, no matter which gap we observe, it will always be sufficiently large to exclude all elements except the best one. Therefore, we only incur a loss if we do not accept anything (which happens if and only if the best element arrives before the waiting time). As a consequence, for $\tau = 0.2$, we observe the ratio of 0.8 (which is the probability of the highest weight arriving after time τ). For the waiting times depending on k, the waiting time turns out to be larger for smaller k and vice versa. The improvement in the competitive ratio for large k is due to the reduced waiting time and hence a smaller probability of facing an arrival of w_1 during the waiting period.

Interestingly, this shows that there are instances for which Algorithm 2 almost obtains its tight guarantee of $1/e$ while these instances become easy when knowing an additive gap.

Given these observations, one might wonder if it is always true that the index k does not play a key role when using a constant waiting time $\tau = 0.2$. For exponentially distributed weights as in (ii), one can see that even with a static waiting time $\tau = 0.2$, larger indices (and hence automatically larger gaps) are helpful to boost the competitive ratio (see Figure 7.2).

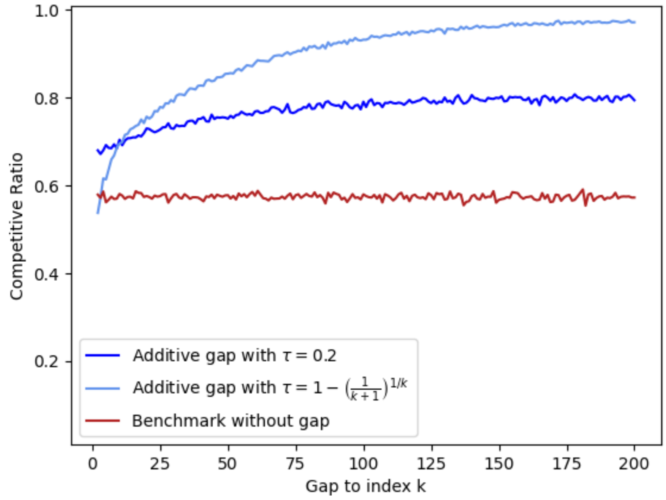

Figure 7.2: Competitive ratios for weights based on (ii)

In addition, for $k = 2$, the waiting time which depends on k does even worse than the classical algorithm. This phenomenon can also be observed for the weights produces by procedure (iii) (see Figure 7.3). It seems that the waiting time for $k = 2$ of $1 - 1/\sqrt{3} \approx 0.423$ is simply too large and suffers from losing too much during the exploration phase. Still, also for weights from a Chi-Squared distribution, we can observe that, first, knowing a gap helps most of the time, and second, larger gaps outperform smaller ones.

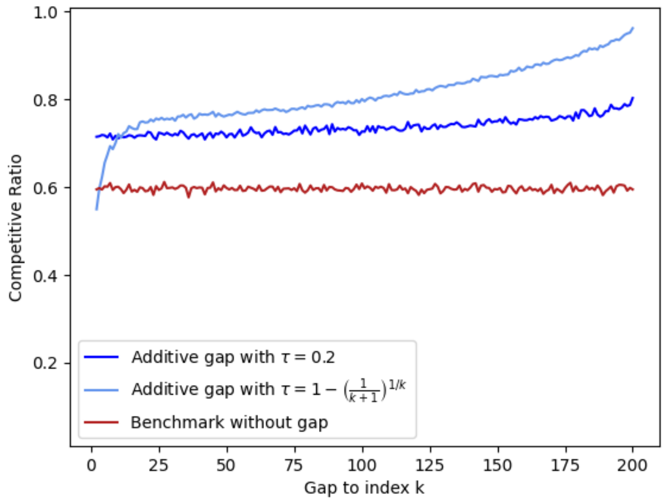

Figure 7.3: Competitive ratios for weights based on (iii)

Summing up, empowering the algorithm with the additional information of an additive gap helps to improve the competitive ratio. As a downside, it turned out that in a few cases it seems that using the index k to compute the waiting time is not beneficial.

Still, as a matter of fact, the waiting time depending on k is an artifact from our analysis. In particular, the waiting time which we used in the simulations was introduced to give provable guarantees. Hence, in the settings studied above, it is beneficial to use a waiting time of e.g.

$$\tau = \min\left(0.2; 1 - \left(\frac{1}{k+1}\right)^{1/k}\right)$$

when having access to the index k in order to avoid the waiting time to be too long.

7.5.2 Dealing with Inaccurate Gaps

In order to get a better understanding concerning inaccuracies in the gap, we run a simulation with different errors.

Experimental Setup

Again, we average over 5000 iterations. In each iteration, we set $n = 200$, draw arrival times as before and weights as follows:

(iv) *Exponential*: Here, all $w_i \sim \text{Exp}(1)$.

(v) *Exponential with superstar*: Here, $w_i \sim \text{Exp}(1)$ for $n-1$ weights and we add a superstar element with weight $100 \cdot \max_i w_i$.

We compare Algorithm 13 to Algorithm 14 both with waiting time $\tau = 0.2$. In addition, Algorithm 14 will drop the gap from the threshold after a time of $1 - \gamma = 0.95$, in other words $\gamma = 0.05$.

The comparison is done for three different gaps: A small one where $k = 2$, i.e. the gap between largest and second largest element, $k = n/2$ and $k = n$, i.e. the gap to the smallest element.

Given a multiplication factor σ for the error, we feed our algorithm with a predicted gap $\hat{c}_k = \sigma \cdot c_k$ for σ going from zero to three in step size of 0.1. In other words, for $\sigma = 1$, we get an accurate gap, for $\sigma < 1$, we underestimate the gap, for $\sigma > 1$ we overestimate the gap and for $\sigma = 0$, the algorithms are equivalent to the classical Secretary algorithms with waiting time τ.

Experimental Results

For exponentially distributed weights (see Figure 7.4), we can observe that underestimating the gap does not cause too many issues. In particular, when highly underestimating the gap (say, $\sigma < 0.5$), both algorithms achieve a competitive ratio of approximately 0.65, similar to an algorithm not knowing any gap. For an accurate gap, $\sigma = 1$, larger gaps are more helpful as they block more elements from being considered. Still, $\sigma > 1$ introduces a transition. For $\sigma > 1$ and gaps between the best and a small element (e.g. $k = 100$ or $k = 200$), overestimating the gap reduces the selection probability of *any* weight of Algorithm 13 to zero: The predicted gap is too large and even exceeds w_1. Still, Algorithm 14 is robust in a sense that we still achieve a competitive ratio of approximately 0.15. This constant depends on our choice of γ. As mentioned before, there is the natural trade-off: Increasing γ for an improved robustness and suffer a decrease in the competitive ratio for $\sigma = 1$.

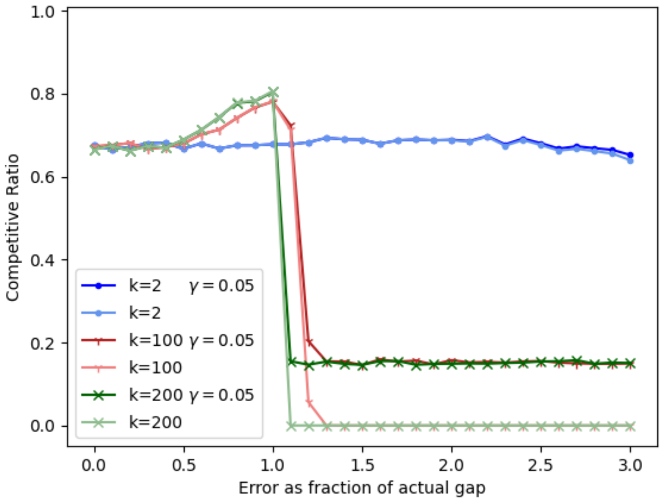

Figure 7.4: Competitive ratios for weights based on (iv). The x-axis shows σ, where the predicted gap $\hat{c}_k$ used by the algorithms satisfies $\hat{c}_k = \sigma \cdot c_k$ for $\sigma \in [0, 3]$.

Interestingly, for the gap between the best and second best element, both algorithms are much more robust. This can be explained as the gap is small in this case anyway, so overestimating by a factor of three does not cause any severe issues yet. One would require to overestimate by a much larger factor here to see a significant difference in the performance of both algorithms.

Next, we show what happens when shifting our perspective towards the more adversarial setting of exponential weights with one additional superstar as listed in (v). Note that in this setting, any algorithm can only achieve a reasonable competitive ratio by selecting the superstar. As illustrated in Figure 7.5, no matter if we consider the gap to $k = 2$, $k = 100$ or $k = 200$, the gap is always large enough to exclude mainly all elements.

In addition, even underestimating the gap by a lot (with $\sigma = 0.1$) does not cause any problems. On the other hand, once we overestimate only by a tiny bit, we mainly lose all guarantees and Algorithm 13 becomes not competitive. Our more robust variant in Algorithm 14 achieves a constant competitive ratio which could be increased when choosing larger values of γ. Again, this would lead to a decrease in the competitive ratio at $\sigma = 1$.

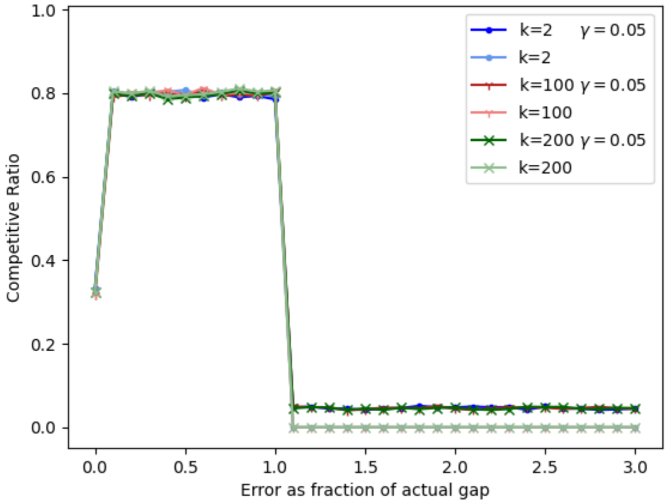

Figure 7.5: Competitive ratios for weights based on (v). The x-axis shows σ, where the predicted gap $\hat{c}_k = \sigma \cdot c_k$ for $\sigma \in [0, 3]$.

Chapter 8

Conclusion and Further Directions

In this thesis, we have seen new results and techniques which are related to three different areas in the intersection of algorithm theory, game theory and economics. In Bayesian online selection, we derived simplified proofs for the existence of prices for pricing-based mechanisms in Chapter 3. When changing the benchmark from the offline optimum to the (computationally unbounded) online optimum, we could give improved guarantees for valuation functions with bounded demand in Chapter 4. For a restricted class of distributions, pricing-based algorithms even allow to extract the optimal social welfare asymptotically, as illustrated in Chapter 6.

Once we have sellers which bring items to the market, we end up in strategically more complex scenarios as in Chapter 5. Still, also for these, we could design pricing-based algorithm with state-of-the-art approximation guarantees.

In a complementing direction where weights are chosen adversarially and not sampled from a distribution, we have seen in Chapter 7 that a single piece of information can help to beat prevalent bounds.

That said, there are a lot of directions for future research and open problems which are either related or inspired by the results of this thesis. As a concluding remark, we discuss some of them in more detail.

8.1 Future Directions in Bayesian Online Selection

The simplified proofs from Chapter 3 avoid to use arguments on specific valuations but only talk about the presence or absence of items. Still, solving the LP in Section 3.2.1 may require exponential time. When applying the proof steps from Section 3.3 constructively to the LP, one can solve for item prices. These prices are exactly the ones obtained by Feldman et al. [2015] and Dütting et al. [2020] for combinatorial auctions. Still, it remains an open problem if the general LP can be solved efficiently, for example by the help of demand oracles.

Open Question 8.1.1. *Can the LP from Section 3.2.1 be solved in polynomial time? Or is there another LP formulation which also leads to simplified proofs for the existence of prices and which can be solved efficiently?*

Complementing this, it would also be interesting to see the power of the approach via LP duality beyond combinatorial auctions and matroids. As such, one could consider arbitrary downward-closed feasibility constraints [Rubinstein, 2016] or other structures. Also, unifying several previous approaches would be a very desirable goal.

This LP-based approach is inspired by a paper on posted-pricing mechanisms for subadditive combinatorial auctions [Dütting et al., 2020]. They derive a Prophet Inequality whose competitive ratio depends on the number of items. Hence, one of the main open questions in this area still remains as follows.

Open Question 8.1.2. *Are there static and anonymous item prices such that a posted-pricing algorithm is constant-factor competitive for buyers with subadditive valuation functions?*

As a first step in this direction, Correa and Cristi [2023] showed that one can be within a constant-factor for some online policy compared to the expected offline optimum. Still, their algorithm is not pricing-based, and also using the reduction by Banihashem et al. [2024] only leads to dynamic prices rather than static ones.

One of the fundamental assumptions in the proofs for Prophet Inequalities in Section 3.2, Feldman et al. [2015] or Dütting et al. [2020] for combinatorial auctions is that buyers are utility-maximizers. That is, every buyer buys the bundle among the available items which maximizes $v_i(S) - \sum_{j \in S} p_j$ for prices p_j. What is happening if we replace this with another model. For example, buyers could maximize their return-of-investment (ROI) meaning that they maximize their value $v_i(S)$ subject to the set S leading to non-negative utility.

Open Question 8.1.3. *Can we also achieve a 1/2-competitive Prophet Inequality for combinatorial auctions with static item prices when buyers maximize their return-of-investment rather than quasi linear utilities?*

Deng et al. [2022b] give a 1/4-competitive Prophet Inequality by using previous techniques similar to Feldman et al. [2015] or Dütting et al. [2020]. Still, getting the tight guarantee could be interesting as it probably requires modified techniques.

Talking about combinatorial auctions, the results from Chapter 4 use the assumption that buyers sample a multi-demand valuation function. A natural open question is to find tight guarantees for these valuation functions.

Open Question 8.1.4. *Can the guarantee of* $0.5 + \kappa$ *for* $\kappa = 0.0115$ *from Algorithm 4 be improved? Maybe by using a stronger LP relaxation?*

Another crucial aspect of multi-demand valuation functions is the property that the contribution of one item is either zero or some fixed value. An inspiring extension would be to consider more general valuation functions, for example gross-substitutes, submodular or XOS functions.

Open Question 8.1.5. *Can we find better-than-half approximation algorithms for the optimal online policy for combinatorial auctions beyond multi-demand valuation functions?*

In other words, can we beat the approximation guarantee of half for more general classes of valuations than the ones in Chapter 4? To achieve this goal, one is probably required to use a different LP as well as a different approach than the one from Chapter 4 for several reasons: The current LP is heavily tailored to the multi-demand functions, so is the algorithm. For example, when considering submodular valuations, it is by far not clear how to bound the contribution of one item in a two proposal algorithm. Is an approach related to the one of Braverman et al. [2022] maybe more promising? Solving this problem could lead to new insights and technical novelties which might have a much broader impact.

In Chapter 6, we have seen that asymptotically optimal welfare can be achieved via pricing-based algorithms. As a crucial assumption, we used the MHR property of the distribution from which buyers' valuations are drawn.

Open Question 8.1.6. *Which other classes of distributions allow to extract the optimal welfare asymptotically via a pricing-based algorithm?*

To this end, the quantile allocation rule introduced in Section 6.2 seems very powerful when combined with Theorem 4 in Banihashem et al. [2024] to turn it into a dynamic pricing mechanism. Maybe, it can be used as a starting point to tackle this question.

8.2 Future Directions in Two-Sided Mechanism Design

As we have seen in Chapter 5, using the ideas from one-sided markets related to the ones in Chapter 3 allows to design truthful mechanisms in two-sided environments as well. Even though the mechanisms in Chapter 5 are DSIC and IR and fulfill budget balance constraints, tight approximation guarantees remain open.

Open Question 8.2.1. *What are the tight approximation guarantees for matroid, combinatorial and knapsack double auctions when restricting to DISC, IR and budget balanced mechanisms?*

To tackle this question, a useful starting point might be to extend the pricing techniques from Ehsani et al. [2018] or Correa et al. [2017] to two-sided markets in the spirit of our extension of Prophet Inequalities to two-sided markets in Chapter 5. As a first step, it could be even interesting to see which other techniques from Prophet Inequalities can be applied to the special case of bilateral trade instances.

Open Question 8.2.2. *Can we get improved guarantees for bilateral trade instances using techniques from Prophet Inequalities (e.g. free-order Prophet or Prophet Secretary)?*

When talking about combinatorial double auctions, in order to obtain DSIC and IR mechanisms for buyers and sellers, we were required to work with restrictive classes of valuations: Either, all agents are additive across items or sellers only bring a single item to the market. What is happening beyond? Can we also get reasonably good approximation guarantees when buyers and sellers have more complex valuation functions?

Open Question 8.2.3. *Can we design mechanisms which are DSIC and IR for buyers and sellers with more general (e.g. submodular or XOS) valuation functions, budget balanced and have a "reasonably good" approximation ratio?*

As a first step in this direction, Colini-Baldeschi et al. [2020] already developed a mechanism which is Bayesian incentive-compatible when combining additive sellers and buyers with XOS valuations. But still, the question remains open if there is a mechanism which is DSIC.

8.3 Future Directions on Secretary Problems with Gaps

When tackling the single-selection problem in the Secretary model, we have seen in Chapter 7 that a single, simple piece of information of the form "there is a gap of c in the instance" helps to improve the competitive ratio. In addition, we can also obtain guarantees with respect to robustness and consistency when having a prediction for the additive gap. Still, the guarantees seem to be not tight.

Open Question 8.3.1. *Can we achieve a better competitive ratio for any gap for the algorithm in Section 7.2? Or is there a matching hardness result?*

Maybe one is also required to use another algorithm in order to get a tight competitive ratio for the Secretary problem when knowing an exact additive gap.

Discussing the trade-off between robustness and consistency, our guarantees are also not tight.

Open Question 8.3.2. *What is the tight trade-off between robustness and consistency for Secretary problems with predicted additive gaps?*

Recall that we considered gaps of the form $w_1 - w_k$ for some k. Still, it might be too much to hope for a prediction of the gap to the largest weight. Hence, can we maybe do something once we get a gap between two arbitrary elements?

Open Question 8.3.3. *Can we also beat the competitive ratio of* $1/\mathrm{e}$ *by a constant when obtaining a gap* $w_i - w_j$ *for some* $1 \leq i < j \leq n$*?*

As a side remark, it seems unlikely that we can beat $1/\mathrm{e}$ for any gaps. To see this, consider knowing that $w_{n-1} - w_n = 0$, or in other words, the smallest element appears twice in the sequence. Still, there are also gaps for which the information is helpful, such as $w_2 - w_3 = 0$. In the latter case, we can indeed improve the guarantee and select the largest weight with a probability larger than $1/\mathrm{e}$.

One of the main motivation to study additive gaps as a prediction model was to see the power of less informative predictions than getting an estimate for the best weight in the sequence. It is also interesting to see the power of less informative predictions beyond the single-selection case, for example in multi-selection or knapsack environments.

Open Question 8.3.4. *For the multi-selection Secretary problem: Can we obtain improved competitive ratios when having access to an additive gap? And if yes, can the algorithm be made robust to inaccurate predictions?*

Also, understanding weaker notions of predictions in combinatorial auctions with random arrival would be a desirable goal. To this end, one potential starting point is to find a reasonable prediction model for unit-demand combinatorial auctions with random arrivals (a.k.a. matching) and maybe extend this later to submodular valuation functions, in a similar spirit as Kesselheim et al. [2013] do this for the classical setting.

Open Question 8.3.5. *Which prediction models for combinatorial auctions (with unit-demand or submodular valuations) in the random order model allow to break the* $1/\mathrm{e}$*-barrier?*

As mentioned, finding prediction models which do not give an advice for the maximum weight matching but rather give some indirect prediction could lead to very interesting technical challenges.

Appendix A

Deferred Proofs

A.1 Proof of Theorem 1.1.2

To familiarize the reader with fundamental techniques, we recall a standard proof for Theorem 1.1.2 from the literature below. As mentioned, the ideas will be a fundamental cornerstone for Chapter 3 and Chapter 5.

Proof of Theorem 1.1.2. First, we reformulate the expected value selected by Algorithm 1 via

$$\begin{aligned}\mathbf{E}\left[\mathrm{ALG}\right] &= \mathbf{E}\left[\sum_i v_i \mathbb{1}_{i \text{ selected by ALG}}\right] = \mathbf{E}\left[\sum_i (v_i - p + p)\mathbb{1}_{i \text{ selected by ALG}}\right] \\ &= \mathbf{E}\left[\sum_i (v_i - p)\mathbb{1}_{i \text{ selected by ALG}}\right] + \mathbf{E}\left[\sum_i p\mathbb{1}_{i \text{ selected by ALG}}\right] \\ &= \mathbf{E}\left[\sum_i (v_i - p)\mathbb{1}_{i \text{ selected by ALG}}\right] + p \cdot \mathbf{Pr}\left[\mathrm{ALG \text{ selects someone}}\right] \quad .\end{aligned}$$

Let us pause here for a moment to define some notation which will be useful later as well. The second summand can be interpreted as the *revenue* or *base value* obtained by running ALG which uses a price of p. We denote this revenue by rev. On the other hand, the term $(v_i - p)\mathbb{1}_{i \text{ selected by ALG}}$ can be thought of as the *utility* or *surplus* of buyer i who is selected by the algorithm, denoted by u_i. Note that for any buyer who is not selected, the utility u_i is simply zero.

Having this, we bound each of the two quantities separately.

Revenue. Observe that for the revenue, we have

$$\mathbf{E}\left[\mathrm{rev}\right] = p \cdot \mathbf{Pr}\left[\mathrm{ALG \text{ selects someone}}\right] = p \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] \quad .$$

Utility. We first bound the utility of a single buyer i and use linearity of expectation afterwards. To this end, note that

$$\mathbf{E}\left[u_i\right] = \mathbf{E}\left[(v_i - p)\mathbb{1}_{i \text{ selected by ALG}}\right] = \mathbf{E}\left[(v_i - p)\mathbb{1}_{v_i \geq p \text{ and } v_1, \dots, v_{i-1} < p}\right] \quad .$$

Now, we can exploit the independence across buyers, i.e. the value of buyer i is independent of the value of previously arrived buyers. In addition, we use $(\cdot)^+$ to denote $\max(\cdot, 0)$. Hence, we get

$$\mathbf{E}\left[u_i\right] = \mathbf{E}\left[(v_i - p)^+\right] \cdot \mathbf{Pr}\left[v_1, \dots, v_{i-1} < p\right] \geq \mathbf{E}\left[(v_i - p)^+\right] \cdot \mathbf{Pr}\left[\forall i' : v_{i'} < p\right] \quad .$$

Summing over all buyers i, we get

$$\begin{aligned}\mathbf{E}\left[\sum_i u_i\right] &\geq \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \mathbf{E}\left[\sum_i (v_i - p)^+\right] \geq \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \mathbf{E}\left[\max_i (v_i - p)^+\right] \\ &\geq \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \mathbf{E}\left[\max_i (v_i - p)\right] = \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \left(\mathbf{E}\left[\max_i v_i\right] - p\right) \quad .\end{aligned}$$

Combining Revenue and Utility. Overall, we end up with

$$\mathbf{E}\left[\mathrm{ALG}(\mathbf{v})\right] \geq p \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] + \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \left(\mathbf{E}\left[\max_i v_i\right] - p \right) \quad . \tag{A.1}$$

Observe that $\mathbf{Pr}\left[\exists i : v_i \geq p\right] + \mathbf{Pr}\left[\forall i : v_i < p\right] = 1$. In addition, note that we still did not define how to set the price p yet. This will turn out to be crucial in order to prove Theorem 1.1.2. In the literature, there are different choices for the price p, two of which we present in the following.

Option 1: Set $p = 1/2 \cdot \mathbf{E}\left[\max_i v_i\right]$. This option was introduced by Kleinberg and Weinberg [2012]. Plugging this choice for p into the right-hand side of Equation (A.1), we obtain

$$\begin{aligned} \mathbf{E}\left[\mathrm{ALG}(\mathbf{v})\right] &\geq \frac{1}{2}\mathbf{E}\left[\max_i v_i\right] \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] + \mathbf{Pr}\left[\forall i : v_i < p\right] \cdot \frac{1}{2}\mathbf{E}\left[\max_i v_i\right] \\ &= \frac{1}{2}\mathbf{E}\left[\max_i v_i\right] \quad . \end{aligned}$$

Option 2: Set p such that $\mathbf{Pr}\left[\exists i : v_i \geq p\right] = 1/2$. Samuel-Cahn [1984] introduced this variant already in the 1980s. Again, using Equation (A.1), we get

$$\begin{aligned} \mathbf{E}\left[\mathrm{ALG}(\mathbf{v})\right] &\geq p \cdot \mathbf{Pr}\left[\exists i : v_i \geq p\right] + \left(1 - \mathbf{Pr}\left[\exists i : v_i \geq p\right]\right) \cdot \left(\mathbf{E}\left[\max_i v_i\right] - p \right) \\ &= p \cdot \frac{1}{2} + \frac{1}{2}\left(\mathbf{E}\left[\max_i v_i\right] - p \right) = \frac{1}{2}\mathbf{E}\left[\max_i v_i\right] \quad . \end{aligned}$$

It is easy to see that such a p exists for continuous distributions without point masses. In the case of distributions with point masses, one can introduce a randomized tie-breaking rule to achieve this property. □

A.2 Proof of Theorem 1.1.3

The proof for Theorem 1.1.3 is a folklore result and well-known in the existing literature, but will also be useful to understand the results and proofs in Chapter 7.

Proof of Theorem 1.1.3. Note that as weights are chosen adversarially, we cannot control by how much w_1 exceeds any other weight w_i. Hence, we bound

$$\mathbf{E}\left[\text{ALG}\right] \geq w_1 \cdot \mathbf{Pr}\left[\text{select } w_1\right] \ ,$$

and continue by finding a suitable expression for $\mathbf{Pr}\left[\text{select } w_1\right]$. To this end, fix an arrival time $x \in [0,1]$ for weight w_1. Observe that we select w_1 if x is later than τ and one of the following two cases happens. Either no arrival took place before x which happens with probability $(1-x)^{n-1}$ due to the independence across arrival times. Or the highest weight which arrived before x did arrive before time τ. This event happens with probability τ/x. As a consequence, we get

$$\mathbf{Pr}\left[\text{select } w_1\right] = \int_\tau^1 (1-x)^{n-1} + \left(1-(1-x)^{n-1}\right)\frac{\tau}{x}\, dx \geq \int_\tau^1 \frac{\tau}{x}\, dx = \tau \ln\left(\frac{1}{\tau}\right) .$$

Plugging in our choice of $\tau = 1/\mathrm{e}$ proves the claim. □

Bibliography

M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. D. Kleinberg, and B. Lucier. Beating 1-1/e for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. Pages 61–71. 2017.

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)*. Pages 3–16. 2012.

M. Adamczyk and M. Włodarczyk. Random order contention resolution schemes. In *Proceedings of the 59th Symposium on Foundations of Computer Science (FOCS)*. Pages 790–801. 2018.

S. Ahmadian, H. Esfandiari, V. Mirrokni, and B. Peng. Robust load balancing with machine learned advice. *J. Mach. Learn. Res.*, 24:44:1–44:46, 2023.

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

S. Alaei, M. Hajiaghayi, and V. Liaghat. The online stochastic generalized assignment problem. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop (APPROX), and 17th International Workshop (RANDOM)*. Pages 11–25. 2013.

Algorithms-with-Predictions. Website with list of papers on algorithms with predictions. URL: https://algorithms-with-predictions.github.io. (Last access: April 7, 2024).

M. Almanza, F. Chierichetti, S. Lattanzi, A. Panconesi, and G. Re. Online facility location with multiple advice. In *Advances in Neural Information Processing Systems (NeurIPS)*. Pages 4661–4673. 2021.

N. Anari, R. Niazadeh, A. Saberi, and A. Shameli. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. Pages 91–92. 2019.

A. Antoniadis, T. Gouleakis, P. Kleer, and P. Kolev. Secretary and online matching problems with machine learned advice. In *Advances in Neural Information Processing Systems (NeurIPS)*. Pages 7933–7944. 2020.

C. Argue, A. Gupta, M. Molinaro, and S. Singla. Robust secretary and prophet algorithms for packing integer programs. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1273–1297. 2022.

B. C. Arnold, N. Balakrishnan, and H. N. Nagaraja. *A First Course in Order Statistics (Classics in Applied Mathematics)*. Society for Industrial and Applied Mathematics, 2008. ISBN 9780898716481.

M. Arsenis, O. Drosis, and R. Kleinberg. Constrained-order prophet inequalities. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 2034–2046. 2021.

S. Assadi and S. Singla. Improved truthful mechanisms for combinatorial auctions with submodular bidders. In *60th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. Pages 233–248. 2019.

S. Assadi, T. Kesselheim, and S. Singla. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 653–661. 2021.

V. Avadhanula, A. Celli, R. Colini-Baldeschi, S. Leonardi, and M. Russo. Fully dynamic online selection through online contention resolution schemes. In *Proceedings of the 37th Conference on Artificial Intelligence and 35th Conference on Innovative Applications of Artificial Intelligence and 13th Symposium on Educational Advances in Artificial Intelligence (AAAI)*. Pages 6693–6700. 2023.

P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1358–1377. 2014.

Y. Azar, A. Chiplunkar, and H. Kaplan. Prophet secretary: Surpassing the 1-1/e barrier. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 303–318. 2018.

M. Babaioff and N. Nisan. Concurrent auctions across the supply chain. *J. Artif. Int. Res.*, 21(1):595–629, May 2004.

M. Babaioff and W. E. Walsh. Incentive-compatible, budget-balanced, yet highly efficient auctions for supply chain formation. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*. Pages 64–75. 2003.

M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 10th International Workshop (APPROX), and 11th International Workshop (RANDOM)*. Pages 16–28. 2007a.

M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 434–443. 2007b.

M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer. Posting prices with unknown distributions. *ACM Trans. Econ. Comput.*, 5(2), Mar. 2017.

M. Babaioff, Y. Cai, Y. A. Gonczarowski, and M. Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Page 373. 2018a.

M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. Matroid secretary problems. *J. ACM*, 65(6), Nov 2018b.

M. Babaioff, K. Goldner, and Y. A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 2452–2471. 2020.

J. Baek and W. Ma. Prophet inequalities on the intersection of a matroid and a graph. *CoRR*, abs/1906.04899. 2019.

K. Banihashem, M. Hajiaghayi, D. R. Kowalski, P. Krysta, and J. Olkowski. Power of posted-price mechanisms for prophet inequalities. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 4580–4604. 2024.

M. Bateni, M. T. Hajiaghayi, and M. Zadimoghaddam. Submodular secretary problem and extensions. *ACM Trans. Algorithms*, 9(4):32:1–32:23, 2013.

B. Berger, A. Eden, and M. Feldman. On the Power and Limits of Dynamic Pricing in Combinatorial Markets. In *16th International Conference on Web and Internet Economics (WINE)*. Pages 206–219. 2020.

L. Blumrosen and S. Dobzinski. Reallocation mechanisms. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC)*. Page 617. 2014.

L. Blumrosen and S. Dobzinski. (Almost) efficient mechanisms for bilateral trading. *Games and Economic Behavior*, 130:369–383, 2021.

L. Blumrosen and Y. Mizrahi. Approximating gains-from-trade in bilateral trading. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE)*. Pages 400–413. 2016.

D. Bradac, A. Gupta, S. Singla, and G. Zuzic. Robust Algorithms for the Secretary Problem. In *11th Innovations in Theoretical Computer Science Conference (ITCS)*. Pages 32:1–32:26. 2020.

A. Braun and T. Kesselheim. Truthful mechanisms for two-sided markets via prophet inequalities. In *The 22nd ACM Conference on Economics and Computation (EC)*. Pages 202–203. 2021.

A. Braun and T. Kesselheim. Simplified prophet inequalities for combinatorial auctions. In *Symposium on Simplicity in Algorithms (SOSA)*. Pages 381–389. 2023a.

A. Braun and T. Kesselheim. Truthful mechanisms for two-sided markets via prophet inequalities. *Mathematics of Operations Research*, 48(4):1959–1986, 2023b.

A. Braun and S. Sarkar. The secretary problem with predicted additive gap. Under submission. 2023.

A. Braun, M. Buttkus, and T. Kesselheim. Asymptotically Optimal Welfare of Posted Pricing for Multiple Items with MHR Distributions. In *29th Annual European Symposium on Algorithms (ESA)*. Pages 22:1–22:16. 2021.

A. Braun, T. Kesselheim, T. Pollner, and A. Saberi. Approximating optimum online for capacitated resource allocation. Under submission. 2024.

M. Braverman, M. Derakhshan, and A. M. Lovett. Max-weight online stochastic matching: Improved approximations against the online benchmark. In *The 23rd ACM Conference on Economics and Computation (EC)*. Pages 967–985. 2022.

J. Brustle, Y. Cai, F. Wu, and M. Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 589–590. 2017.

N. Buchbinder, K. Jain, and M. Singh. Secretary problems via linear programming. *Math. Oper. Res.*, 39(1):190–206, 2014.

J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–94, 1996.

Y. Cai and M. Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. Pages 170–183. 2017.

Y. Cai and M. Zhao. Simple mechanisms for profit maximization in multi-item auctions. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 217–236. 2019.

Y. Cai, K. Goldner, S. Ma, and M. Zhao. On multi-dimensional gains from trade maximization. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1079–1098. 2021.

S. Chawla, J. D. Hartline, and R. D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings 8th ACM Conference on Electronic Commerce (EC)*. Pages 243–251. 2007.

S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*. Pages 311–320. 2010.

S. Chawla, N. R. Devanur, A. E. Holroyd, A. R. Karlin, J. B. Martin, and B. Sivan. Stability of service under time-of-use pricing. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. Pages 184–197. 2017.

S. Chawla, K. Goldner, A. R. Karlin, and J. B. Miller. Non-adaptive matroid prophet inequalities. *CoRR*, abs/2011.09406. 2020.

W. Chen, S. Teng, and H. Zhang. Capturing complementarity in set functions by going beyond submodularity/subadditivity. In *10th Innovations in Theoretical Computer Science Conference (ITCS)*. Pages 24:1–24:20. 2019.

Y. Chow, H. Robbins, H. Robbins, and D. Siegmund. *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin, 1971. ISBN 9780395053140.

E. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

R. Colini-Baldeschi, B. de Keijzer, S. Leonardi, and S. Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1424–1443. 2016.

R. Colini-Baldeschi, P. Goldberg, B. de Keijzer, S. Leonardi, and S. Turchetta. Fixed price approximability of the optimal gain from trade. In *13th International Conference on Web and Internet Economics (WINE)*. Pages 146–160. 2017.

R. Colini-Baldeschi, P. W. Goldberg, B. d. Keijzer, S. Leonardi, T. Roughgarden, and S. Turchetta. Approximately efficient two-sided combinatorial auctions. *ACM Trans. Econ. Comput.*, 8(1), Mar. 2020.

J. Correa and A. Cristi. A constant factor prophet inequality for online combinatorial auctions. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC)*. Pages 686–697. 2023.

J. Correa, P. Dütting, F. Fischer, and K. Schewior. Prophet inequalities for i.i.d. random variables from an unknown distribution. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 3–17. 2019a.

J. Correa, A. Cristi, B. Epstein, and J. Soto. Sample-driven optimal stopping: From the secretary problem to the i.i.d. prophet inequality. *Mathematics of Operations Research*, 49, Apr. 2023.

J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 169–186. 2017.

J. R. Correa, R. Saona, and B. Ziliotto. Prophet secretary through blind strategies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1946–1961. 2019b.

J. R. Correa, A. Cristi, L. Feuilloley, T. Oosterwijk, and A. Tsigonias-Dimitriadis. The secretary problem with independent sampling. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 2047–2058. 2021.

J. R. Correa, A. Cristi, A. Fielbaum, T. Pollner, and S. M. Weinberg. Optimal item pricing in online combinatorial auctions. In *23rd International Conference on Integer Programming and Combinatorial Optimization (IPCO)*. Pages 126–139. 2022.

Y. Deng, J. Mao, B. Sivan, and K. Wang. Approximately efficient bilateral trade. In *54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. Pages 718–721. 2022a.

Y. Deng, V. Mirrokni, and H. Zhang. Posted pricing and dynamic prior-independent mechanisms with value maximizers. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems (NeurIPS)*. Pages 24158–24169. 2022b.

S. Dobzinski, N. Nisan, and M. Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010.

S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. *J. Comput. Syst. Sci.*, 78(1):15–25, 2012.

S. Dughmi, T. Roughgarden, and Q. Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*. Pages 149–158. 2011.

P. Dütting and T. Kesselheim. Posted pricing and prophet inequalities with inaccurate priors. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*. Pages 111–129. 2019.

P. Dütting and R. Kleinberg. Polymatroid prophet inequalities. In *23rd Annual European Symposium on Algorithms (ESA)*. Pages 437–449. 2015.

P. Dütting, T. Roughgarden, and I. Talgam-Cohen. Modularity and greed in double auctions. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation (EC)*. Pages 241–258. 2014.

P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing (SICOMP)*, 49(3):540–582, 2020a.

P. Dütting, T. Kesselheim, and B. Lucier. An $O(\log\log m)$ prophet inequality for subadditive combinatorial auctions. In *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. Pages 306–317. 2020b.

P. Dütting, F. Fusco, P. Lazos, S. Leonardi, and R. Reiffenhäuser. Efficient two-sided markets with limited information. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. Pages 1452–1465. 2021a.

P. Dütting, S. Lattanzi, R. Paes Leme, and S. Vassilvitskii. Secretaries with advice. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. Pages 409–429. 2021b.

P. Dütting, E. Gergatsouli, R. Rezvan, Y. Teng, and A. Tsigonias-Dimitriadis. Prophet secretary against the online optimal. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*. Pages 561–581. 2023.

E. B. Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Math. Dokl*, 4:627–629, 1963.

S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 700–714. 2018.

H. Esfandiari, M. Hajiaghayi, V. Liaghat, and M. Monemizadeh. Prophet secretary. *SIAM J. Discret. Math.*, 31(3):1685–1701, 2017.

T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang. Online stochastic max-weight matching: prophet inequality for vertex and edge arrival models. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*. Pages 769–787. 2020.

U. Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.*, 39(1):122–142, 2009.

U. Feige and R. Izsak. Welfare maximization and the supermodular degree. In *Innovations in Theoretical Computer Science Conference (ITCS)*. Pages 247–256. 2013.

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 Conference on Artificial Intelligence (AAAI)*. Pages 872–878. 2015.

M. Feldman and R. Gonen. Removal and threshold pricing: Truthful two-sided markets with multi-dimensional participants. In *11th International Symposium on Algorithmic Game Theory (SAGT)*. Pages 163–175. 2018.

M. Feldman, J. Naor, and R. Schwartz. Improved competitive ratios for submodular secretary problems (extended abstract). In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop (APPROX), and 15th International Workshop (RANDOM)*. Pages 218–229. 2011.

M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 123–135. 2015.

M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1014–1033. 2016.

M. Feldman, O. Svensson, and R. Zenklusen. A simple $O(\log\log(\text{rank}))$-competitive algorithm for the matroid secretary problem. *Math. Oper. Res.*, 43(2):638–650, 2018.

T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.

P. R. Freeman. The secretary problem and its extensions: A review. *International Statistical Review / Revue Internationale de Statistique*, 51(2):189–206, 1983.

H. Fu, P. Lu, Z. G. Tang, A. Turkieltaub, H. Wu, J. Wu, and Q. Zhang. Oblivious online contention resolution schemes. In *Symposium on Simplicity in Algorithms (SOSA)*. Pages 268–278. 2022.

K. Fujii and Y. Yoshida. The secretary problem with predictions. *CoRR*, abs/2306.08340. 2023.

R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM (JACM)*, 53(3): 324–360, 2006.

M. Gerstgrasser, P. W. Goldberg, B. de Keijzer, P. Lazos, and A. Skopalik. Multi-unit bilateral trade. In *Proceedings of the Thirty-Third Conference on Artificial Intelligence (AAAI)*. Pages 1973–1980. 2019.

Y. Giannakopoulos and K. Zhu. Optimal pricing for MHR distributions. In *14th International Conference on Web and Internet Economics (WINE)*. Pages 154–167. 2018.

J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, 1966.

O. Göbel, M. Hoefer, T. Kesselheim, T. Schleiden, and B. Vöcking. Online independent set beyond the worst-case: Secretaries, prophets, and periods. In *41st International Colloquium on Automata, Languages, and Programming (ICALP)*. Pages 508–519. 2014.

N. Gravin and H. Wang. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. Pages 93–109. 2019.

T. Groves. Incentives in teams. *Econometrica*, 41(4):617–31, 1973.

A. Gupta and S. Singla. Random-order models. In *Beyond the Worst-Case Analysis of Algorithms* by Tim Roughgarden. Pages 234–258. 2020.

M. T. Hajiaghayi, R. D. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *Proceedings 5th ACM Conference on Electronic Commerce (EC)*. Pages 71–80. 2004.

M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd Conference on Artificial Intelligence (AAAI)*. Pages 58–65. 2007.

T. P. Hill, R. P. Kertz, et al. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability*, 10(2):336–345, 1982.

S. Im, R. Kumar, M. Montazer Qaem, and M. Purohit. Online knapsack with frequency predictions. In *Advances in Neural Information Processing Systems (NeurIPS)*. Pages 2733–2743. 2021.

N. Immorlica, S. Singla, and B. Waggoner. Prophet inequalities with linear correlations and augmentations. In *21st ACM Conference on Economics and Computation (EC)*. Pages 159–185. 2020.

J. Jiang, W. Ma, and J. Zhang. Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1221–1246. 2022.

Y. Jin, W. Li, and Q. Qi. On the approximability of simple mechanisms for mhr distributions. In *International Conference on Web and Internet Economics (WINE)*. Pages 228–240. 2019.

Z. Y. Kang and J. Vondrák. Strategy-proof approximations of optimal efficiency in bilateral trade. 2018.

H. Kaplan, D. Naori, and D. Raz. Competitive Analysis with a Sample and the Secretary Problem. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 2082–2095. 2020.

T. Kesselheim and M. Molinaro. Knapsack Secretary with Bursty Adversary. In *47th International Colloquium on Automata, Languages, and Programming (ICALP)*. Pages 72:1–72:15. 2020.

T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *21st Annual European Symposium on Algorithms (ESA)*. Pages 589–600. 2013.

T. Kesselheim, R. D. Kleinberg, and R. Niazadeh. Secretary problems with non-uniform arrival order. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing (STOC)*. Pages 879–888. 2015.

T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. Primal beats dual on online packing lps in the random-order model. *SIAM Journal on Computing*, 47(5):1939–1964, 2018.

R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 630–631. 2005.

R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing (STOC)*. Pages 123–136. 2012.

R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games and Economic Behavior*, 113:97–115, 2019.

N. Korula and M. Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Proceedings of the 36th Internatilonal Collogquium on Automata, Languages and Programming: Part II (ICALP)*. Pages 508–520. 2009.

U. Krengel and L. Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc*, 83 (4), 1977.

U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

O. Lachish. O(log log rank) competitive ratio for the matroid secretary problem. In *55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. Pages 326–335. 2014.

T. Lavastida, B. Moseley, R. Ravi, and C. Xu. Learnable and Instance-Robust Predictions for Online Matching, Flows and Load Balancing. In *Proceedings of the 29th Annual European Symposium on Algorithms (ESA)*. Pages 59:1–59:17. 2021.

E. Lee and S. Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*. Pages 57:1–57:14. 2018.

B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. In *Proceedings 3rd ACM Conference on Electronic Commerce (EC)*. Pages 18–28. 2001.

D. V. Lindley. Dynamic programming and decision theory. *Journal of The Royal Statistical Society Series C-applied Statistics*, 10:39–51, 1961.

B. Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1): 24–47, 2017.

T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. *J. ACM*, 68(4):24:1–24:25, 2021.

C. MacRury, W. Ma, and N. Grammel. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 1995–2014. 2023.

M. Mahdian and Q. Yan. Online bipartite matching with random arrivals: An approach based on strongly factor-revealing lps. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing (STOC)*. Pages 597–606. 2011.

R. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56(2): 434 – 450, 1992.

R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, Apr. 1983.

R. B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, 1981.

J. Naor, A. Srinivasan, and D. Wajc. Online dependent rounding schemes. *CoRR*, abs/2301.08680. 2023.

R. Niazadeh, Y. Yuan, and R. Kleinberg. Simple and near-optimal mechanisms for market intermediation. In *10th International Conference on Web and Internet Economics (WINE)*. Pages 386–399. 2014.

C. Papadimitriou, T. Pollner, A. Saberi, and D. Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*. Pages 763–764. 2021.

K. Pashkovich and A. Sayutina. Non-adaptive matroid prophet inequalities. *CoRR*, abs/2301.01700. 2023.

T. Pollner, M. Roghani, A. Saberi, and D. Wajc. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*. Pages 321–322. 2022.

M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems (NeurIPS)*. Pages 9684–9693. 2018.

H. Rinne. The Hazard rate: Theory and inference. Justus-Liebig-Universität. URL: http://geb.uni-giessen.de/geb/volltexte/2014/10793. 2014.

Y. Rinott and E. Samuel-Cahn. Comparisons of optimal stopping values and prophet inequalities for negatively dependent random variables. *The Annals of Statistics*, 15, Dec. 1987.

A. Rubinstein. Beyond matroids: secretary problem and prophet inequality with general constraints. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing (STOC)*. Pages 324–332. 2016.

A. Rubinstein, J. Z. Wang, and S. M. Weinberg. Optimal single-choice prophet inequalities from samples. In *11th Innovations in Theoretical Computer Science Conference (ITCS)*. Pages 60:1–60:10. 2020.

A. Saberi and D. Wajc. The greedy algorithm is not optimal for online edge coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP)*. Pages 109:1–109:18. 2021.

E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984.

A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

E. Segal-Halevi, A. Hassidim, and Y. Aumann. SBBA: A strongly-budget-balanced double-auction mechanism. In *9th International Symposium on Algorithmic Game Theory (SAGT)*. Pages 260–272. 2016.

J. A. Soto, A. Turkieltaub, and V. Verdugo. Strong algorithms for the ordinal matroid secretary problem. *Math. Oper. Res.*, 46(2):642–673, 2021.

A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *Proceedings of the 42nd Symposium on Foundations of Computer Science (FOCS)*. Pages 588–597. 2001.

A. Torrico and A. Toriello. Dynamic relaxations for online bipartite matching. *INFORMS Journal on Computing*, 2022.

W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

A. Wei and F. Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NeurIPS)*. Pages 8042–8053. 2020.

Q. Yan. Mechanism design via correlation gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Pages 710–719. 2011.

A. C. Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *18th Annual Symposium on Foundations of Computer Science (FOCS)*. Pages 222–227. 1977.

A. Zeynali, B. Sun, M. Hajiesmaili, and A. Wierman. Data-driven competitive algorithms for online knapsack and set cover. In *Proceedings of the 35th Conference on Artificial Intelligence (AAAI)*. Pages 10833–10841. 2021.

H. Zhang. Improved Prophet Inequalities for Combinatorial Welfare Maximization with (Approximately) Subadditive Agents. In *28th Annual European Symposium on Algorithms (ESA)*. Pages 82:1–82:17. 2020.