PRIOR-INDEPENDENCE: A NEW LENS FOR MECHANISM DESIGN

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF COMPUTER SCIENCE
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Qiqi Yan
June 2012

Abstract

We propose to study revenue-maximizing auctions in the prior-independent analysis framework. The goal is to identify a single auction mechanism for all underlying valuation distributions, so that its expected revenue approximates that of the optimal mechanism tailored for the underlying distribution, under standard weak conditions on the distribution.

We use the prior-independent analysis framework to analyze natural and practical auction mechanisms such as welfare-maximization with reserve prices, limiting supply to induce artificial scarcity, sequentially posting prices, etc. Our framework allows us to argue that these simple mechanisms give near-optimal revenue guarantee in a very robust manner.

Acknowledgments

Pursuing a PhD degree is an amazing intellectual journey. This journey would not be successful or as fun without the help and support from the various people along the way.

First I'd like to thank my PhD advisor Tim Roughgarden. Tim taught me a lot on almost every important aspect of research. Tim taught me to judge whether a research problem is important based on examining its interpretations carefully, helped me form the belief that there is always a simple explanation for apparently complex stuff, which is a powerful tool in problem solving, and showed me what is the right way of presenting things either in writing papers or giving talks. I benefited tremendously from learning these things from Tim. I also thank Tim for always being encouraging to students, and his stress-free advising style.

I also want to thank Jason Hartline, Shaddin Dughmi, and Mukund Sundararajan not only for being brilliant collaborators, but also for having deeply shaped my way of thinking. In particular, I thank Jason for teaching me a whole lot about his intuition and methodology on mechanism design, plus his magic in connecting inequalities, Shaddin for his ridiculous amount of optimism on research and for sharing his knowledge on submodularity and algorithmic mechanism design, and Mukund for giving various advice, from research, to pingpong, to life.

I thank Inbal Talgam-Cohen and Peerapong Dhangwatnotai for being great collaborators and in particular for contributing many ideas to this thesis. I also thank Mohammad Mahdian, Nitish Korula, Vahab Mirrokni for the wonderful summer internship experiences, which were eye-openers to me. I thank my defense committee Yinyu (chair), Tim, Serge, Ashish, Ryan for giving me various feedback on the thesis.

I thank Stanford for providing such a great academic environment, and for surrounding me with an amazing group of smart people: Ananth, Aleksandra, Aneesh, Bahman, Damon, David A., Eric L., Florian, Hart, Ian, Kshipra, Martin, Michael K., Nikola, Pranav, Shayan, Shipra, Uri, Ying W., Ying X., among others. In particular I thank Kshipra Bhawalker for being a great officemate, who constantly tries to improve the research/living environment around her, helping making our office, the theory wing, and even the department a better place. I also thank my first-year buddies: Aditya, Ankur, Eric K., Quoc, Steven, former roommates Marc, Jose, Wen, and Janet, pingpong pals Calvin, Chenguang, Ran, Xin, and various other friends. I also want to thank Don Knuth for the signed free books, thank Jan Vondrak for being a valuable oracle on submodularity.

My interest in theory was ignited by a class on set theory and logic given by Prof. Enshao Shen at Shanghai Jiao Tong University, who later became my master thesis advisor. I thank Prof. Shen for guiding me into theory, and for being extremely kind while being stern on mathematical rigor.

Finally, last but not least, I thank my parents and my wife Shanshan for the unconditional support throughout.

Contents

List of Figures

Part I

Background and Overview

Chapter 1

Introduction

1.1 Roadmap

This thesis is about a new way of looking at auction mechanisms, namely the prior-independent way. In particular, we look at the average performance (often revenue) of a mechanism over an underlying prior distribution that is not available to the mechanism, yet we demand that the mechanism has good performance compared to an optimal mechanism tailored for the distribution. In particular we want the guarantee to be true for every possible underlying prior.

Our prior-independent way is different from the standard prior-dependent average-case way of economics, or the prior-free or worst-case way of computer science. As a middle ground approach between these two, our prior-independent way leads to new insights and new results, as described in [25, 76, 66, 46]. This thesis aims to be a coherent summary of all these results. (Differences in organization and treatment from the original papers are outlined in Section 1.9). The rest of the thesis is organized as follows.

In the rest of this chapter, we introduce some background on auction mechanism design, motivate the prior-independent analysis framework, and briefly summarize our main results.

In Chapter 2, we discuss Bayesian mechanism design, surveying some of the most important results, notably Myerson's optimal auction theory and the Bulow-Klemperer theorem, and at the same time, fix our terminology and notations.

In Chapter 3, we formally define the prior-independent analysis framework, study the digital goods case with two bidders to extract several important ideas, and show a general reduction to Bulow-Klemperer-style statements.

In Chapters 4, 5, and 6, we present three approaches to prior-independent mechanisms, which are based on VCG with sampled reserves, supply limiting, and sequential posted prices, respectively.

In Chapters 7 and 8, we propose to consider prior-free mechanisms that approximate the envy-free optimal revenue benchmark. Such mechanisms have prior-independent guarantees as well as good worst-case guarantees for every input. Notably, such results are achieved via a connection to a theory of optimal envy-free outcomes, which can be of independent interest.

Finally, we conclude with several open directions in Chapter 9.

1.2 Background: Auction Mechanism Design

Auctions Auctions have a long history, and have been commonly used for the selling of collectibles, automobiles, real estate, electricity, etc. Recently, with the surge of Internet applications, auctions have also been used extensively to allocate advertisement space of search engines. (see e.g., [72], for an introduction to Auctions)

We interpret the notion of auctions very broadly, allowing it to contain essentially all selling procedures where buying agents' valuations are private. Roughly speaking, an auction is a process of allocating goods or resources, where the buying agents, or bidders, who want the goods bid or report how much they value the goods, and the auctioneer, or the seller, will use these bids to determine who win the goods, and how much they pay. An auction mechanism refers to the allocation rule for determining who gets what items and the payment rule for determining how much to charge each winner. Note that agents may report fake valuations for their own sake, and this is one of the distinguishing features that set auctions apart from other allocation procedures.

Examples Two canonical examples of auctions include the well-known first-price auction and second-price auction. Consider the sealed-bid format, where every bidder submits a bid for an item in an envelope to the auctioneer. In both auction formats, the highest bidder wins. In a first-price auction, the highest bidder pays the amount she bids, while in a second-price auction, the highest bidder pays the amount the second highest bidder bids. We highlight an important difference between these two formats. In a first-price auction, if a bidder bids her true value, her utility (value minus payment) is at most 0, and she is incentivized to bid lower than her true value. On the other hand, in a second-price auction, it is not difficult to verify that a bidder maximizes her utility by bidding her true value, no matter what the other bidders do. In this sense we say that second-price auction is truthful, or ex post incentive compatible (or ex post IC or even simply IC in short).

Truthfulness Truthful mechanisms are desirable in many ways. From each bidding agent's perspective, the problem of deciding what to report is made very simple, as she is incentivized to simply report the truth. From the auctioneer's perspective, when agents are not exploring different bids to game the system, the auction's outcome is much more stable, predictable, and easier to optimize. In this thesis, we will focus on the design of truthful mechanisms, and their revenue guarantees. By default, a mechanism is truthful in this thesis.

On Non-Truthful Mechanisms When agents' values are drawn from distributions, it is also natural to study mechanisms under the weaker solution concept of Bayes-Nash equilibrium. As we will mention in Remark 2.13, Myerson [61] proved that such a relaxed solution concept does not give us additional power in (expected) revenue maximization. As our primary goal is revenue maximization, it is without loss of generality that we will confine ourselves to truthful, or ex post IC mechanisms in this thesis.

1.3 An Algorithmic View

In the past decade, mechanism design has attracted the attention from many computer scientists. One main reason is that with the surge of Internet applications, a lot of algorithmic problems now involve the participation of self-interested agents, and for an algorithm to behave properly, incentives must be taken into consideration. In other words, algorithms for the Internet are essentially mechanisms.

On the other hand, a somewhat technical but also important reason is that mechanism design is very similar to algorithm design.

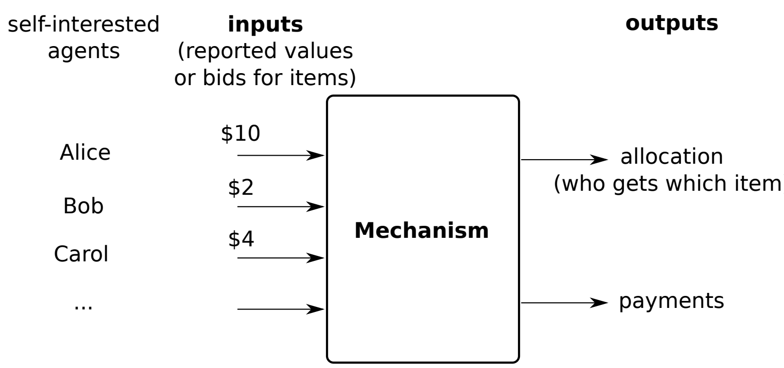

Figure 1.3.1: A Mechanism As An Algorithm

Mechanisms As Algorithms We depict the working of a (sealed-bid) mechanism in Figure 1.3.1. In words, a mechanism can be seen as an algorithm that transforms the input, which is bids from the agents, to output, which is allocation and payments. What makes a mechanism different from a standard algorithm is that a mechanism needs to obey incentive constraints.

Mechanism $\approx$ Monotone Algorithms Interestingly, for many auction environments, it can be proved that the incentive constraints translate into certain monotonicity conditions. For the class of single-dimensional domains, the monotonicity condition simply says that if you bid more, you win more (you are more likely to win) [61, 6]. For the class of multi-dimensional domains, certain monotonicity conditions also exist, albeit less intuitive (see e.g.,[67]). In other words, the design of truthful mechanism boils down to the design of monotone algorithms, which makes the problem even more algorithmic in nature.

Average-Case vs. Worst-Case Analysis Interestingly, for cultural and various reasons, economists traditionally analyzed mechanisms using average-case analysis, which is different from the common choice of worst-case analysis used by computer scientists to analyze algorithms. This difference appears to be huge, but seen from a different angle, it also suggests the possibility of cross-fertilization. After all, mechanisms are similar to algorithms, and economists and computer scientists simply have been looking at them from two different perspectives. Each of these two angles has its pros and cons, and there should be opportunities in achieving the pros of both, while alleviating the cons. Indeed, this is one of the main goals of this thesis.

1.4 Average-Case Analysis

We first look at average-case analysis, and understand why it is a natural choice for mechanism design.

Remember that we aim to study revenue-maximizing mechanisms. However, formalizing this goal can be tricky. As shown by the following example, there is no single mechanism that is optimal for revenue for every input.

Example 1.1. Consider the case of a single agent and a single item of good, where the agent has a value of v for the good. Any truthful mechanism corresponds to offering a fixed price p in this case, and we have a sale if and only if $v \geq p$. So for a fixed v, the optimal revenue is achieved by the mechanism that sets p to be exactly equal to v. Therefore, for different input values, the mechanisms that maximize revenue are in general different.

In other words, there is no single mechanism that is optimal for all inputs. Ideally, we want to quantify a mechanism's revenue performance using a single scalar number, so that it is easy to compare the performance of different mechanisms. A natural way to achieve this is then to specify weights or probabilities to different inputs, by imposing a probability distribution on top of them.

In average case analysis, as is commonly used in the economics literature, we assume that the input values are drawn from a probability distribution. Commonly, we assume that each agent's value is drawn from a distribution independently. Often these distributions are assumed to be identical too.

Example 1.2. Now given a distribution over inputs, every mechanism's revenue performance can be measured by a single scalar value, which is the expected revenue of the mechanism over the distribution. Importantly, this implies that an optimal mechanism exists. In fact, for single-dimensional domains, Myerson's seminal work characterizes the optimal mechanisms for revenue. The following states Myerson's result in the case of single-item auction and symmetric bidders. (The regularity condition will be defined in Section 2.3.)

Theorem 1.3 (Myerson'81)**.** *For a single-item auction over n bidders whose values are drawn i.i.d. from a regular distribution F, an optimal mechanism that maximizes expected revenue is the second-price auction with the monopoly reserve price $p^* = argmax_p p \cdot (1 - F(p))$.*

Note that to run the optimal mechanism, it is crucial here that we know about the distribution F beforehand, as it is needed for computing the monopoly price p^*. For more general settings, Myerson's optimal mechanisms correspond to so-called virtual surplus maximizers, which depend on the distribution information in much more complicated ways.

1.4.1 Problems with the Known Distribution Assumption

By the previous discussion, distribution information is crucial in allowing us to compare the revenue performance of different mechanisms. However, assuming that the distribution information is given to us beforehand can be problematic in various ways.

1. It is often difficult to estimate the distribution information accurately.

   This can be the case for various reasons. For example, we may not have enough access to the distribution information to form a meaningful estimate, or the distribution actually changes over time.

2. Estimating distribution information involves incentive issues.

> When one tries to estimate distribution information, one needs to be careful about incentive compatibility. For example, a common approach would be to run a market survey. But agents then might be tempted to lie in the market survey, hoping to benefit themselves in the later auction phase. We need to take such incentive issues into consideration.

Essentially, obtaining distribution information is something that every mechanism designer needs to do. In a sense it is in fact part of the mechanism design process. However, by assuming that the distribution is given, it is not possible for us to study this important part of the process in a formal way.

1.5 Worst-Case Analysis

In the past decade, computer scientists have also studied mechanism design, but mostly from a prior-free or worst-case point of view prior to the work described in this thesis. Many works have been done for problems such as digital goods auctions [37], which models the selling of software licenses or digital music downloads, with or without a supply limit on the total number of goods sold.

By the previous discussion, there is no single mechanism that is optimal for every input. To be able to conduct worst-case analysis, for every input profile $\mathbf{v}$ that contains the values of agents, we need to decide on a revenue benchmark $B(\mathbf{v})$ to compare to. It is in fact not clear what makes a meaningful revenue benchmark. For simple settings like digital goods, researchers have chosen such benchmarks in an ad hoc manner. For example, in [36], the "best single-price revenue" benchmark $\mathcal{F}(\mathbf{v})$ was proposed for digital goods auctions, which appears to be a natural and amenable choice. Unfortunately, this benchmark turns out to be too strong to work with, and people instead work with the slightly modified $\mathcal{F}^{(2)}$ benchmark, which is the best single price revenue from selling at least two items.

Given such a point-wise benchmark, the researchers' goal was then to look for a mechanism $\mathcal{M}$ such that for every input profile $\mathbf{v}$, the revenue of $\mathcal{M}$ over $\mathbf{v}$ is at least a ρ fraction of the benchmark $\mathcal{F}^{(2)}(\mathbf{v})$, where we want ρ to be a constant that is as large as possible. Unfortunately, due to the worst-case nature of the analysis, the approximation ratio ρ we get is often less than desirable. Prior to works in this thesis, the best known ratio was $\frac{1}{3.25}$ for the simple setting of digital goods auction with the $\mathcal{F}^{(2)}$ benchmark [43], and $\frac{1}{6.5}$ for multi-unit auctions. (The approximation ratio we achieve with our new analysis framework will be $\max\{\frac{1}{2}, 1-o(1)\}$ for these two cases.) Moreover, the mechanisms that are tailored to achieve best worst-case approximation ratio tend to be highly unnatural and impractical.

1.5.1 Hartline-Roughgarden's Revenue Benchmark

Hartline and Roughgarden [44] proposed to use the optimal revenue $\mathrm{OBO}(\mathbf{v})$ from a Bayesian optimal mechanism as a point-wise benchmark for input $\mathbf{v}$. There are several benefits of this new benchmark.

First for all single-dimensional auction environments, Bayesian optimal mechanisms are fully characterized. Therefore the benchmark is explicitly defined for a wide range of environments.

Second, if a mechanism can approximate such a benchmark point-wise for every input, it follows that if the input is drawn from a distribution satisfying weak assumptions, then the mechanism's expected revenue also approximates that of the optimal mechanism tailored for the distribution.

Overall, the Hartline-Roughgarden approach gives a systematic way of defining prior-free revenue benchmarks that is well-grounded in Bayesian optimal auction theory. However, it is still worst-case in nature. As an outcome, the approximation ratios proved based on this approach tend to be less than desirable (worse than $\frac{1}{10}$ for the specific problem studied in [44]).

1.6 Our Prior-Independent Analysis Framework

The Hartline-Roughgarden approach connects worst-case analysis to average-case analysis in the sense that a worst-case approximation to a suitable benchmark implies approximation guarantees in the average-case as well, with a constant factor loss in approximation ratio.

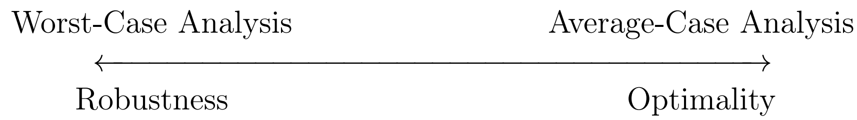

Figure 1.6.1: Trade-Off in Analysis Frameworks

We, on the other hand, propose the notion of parameterized prior-independent approximation, which contain both average-case and worst-case analyses as special cases with somewhat extreme choices of the parameter. We then pick the parameter carefully to identify a middle ground between these two to achieve desirable properties.

1.6.1 Optimality vs. Robustness

To generalize both average-case and worst-case analyses, we observe that informally, in choosing analysis frameworks, there seems to be an inherent trade-off between optimality and robustness, where average-case analysis is on the optimality side, and worst-case analysis is on the robustness side (see Figure 1.6.1). In particular, in average-case analysis, with given distribution information, we can achieve optimal revenue using Myerson's mechanism for a variety of settings, but the guarantee fails if the distribution turns out to be inaccurate, which is not very robust. On the other hand, in worst-case analysis, we can only achieve approximate optimality with a constant ratio, but the guarantee holds robustly no matter what the input is.

A natural question is then whether there a middle ground between these two analysis frameworks, where we can achieve a balanced trade-off between optimality and robustness. To identify such a middle ground, it is important that we first formalize the optimality versus robustness trade-off.

1.6.2 Formalizing the Optimality vs. Robustness Trade-Off

We formalize the trade-off between optimality and robustness using the notion of parameterized prior-independent approximation.

Definition 1.4 (Parameterized Prior-Independent Approximation)**.** Suppose we are given the following two parameters:

Robustness Parameter A distribution class $\mathcal{C}$.

Optimality Parameter An approximation ratio $\rho \in [0, 1]$.

- We say that a mechanism $\mathcal{M}$ gives a prior-independent ρ-approximation w.r.t. distribution class $\mathcal{C}$, if for all distribution $F \in \mathcal{C}$,

$$\mathrm{E}_{\mathbf{v}\sim F}[\text{revenue}(\mathcal{M}, \mathbf{v})] \geq \rho \cdot \mathrm{E}_{\mathbf{v}\sim F}[\text{revenue}(OPT_F, \mathbf{v})],$$

  where OPT_F denotes Myerson's optimal mechanism tailored for F.

- In particular, we require that $\mathcal{M}$ does not know the specific F except that it comes from $\mathcal{C}$.

In other words, we do still assume that the inputs are drawn from a distribution, but this distribution is *unknown* to the mechanism. Despite this handicap, we require that the performance of our mechanism is approximately as good as an optimal mechanism that is tailored for the distribution.

Intuitively, the robustness parameter controls how widely our guarantee should hold, and the optimality parameter controls how close we are to optimal.

As shown by the examples below, parameterized prior-independence generalizes both average-case analysis and worst-case analysis as somewhat extreme special cases.

Example 1.5 (Average-Case Analysis)**.** Average-case analysis with distribution F is equivalent to parameterized prior-independent approximation with $\mathcal{C} = \{F\}$, where ρ can be as good as 1.

Example 1.6 (Worst-Case Analysis)**.** Worst-case analysis is equivalent to parameterized prior-independent approximation where $\mathcal{C}$ contains all (degenerate) distributions over single-point supports. Here ρ cannot be 1, and cannot be too close to 1 in general.

Now that we have formalized the optimality vs. robustness trade-off, which is mostly controlled by the choice of distribution class $\mathcal{C}$, to find a balanced middle ground, we should look for a class of distributions that is large enough to capture most "natural" distributions, while still allowing ρ to be good.

Worst-Case Analysis	Prior-Independence	Average-Case Analysis
←		→
Robustness	**Balance**	Optimality

Figure 1.6.2: Prior-Independence As A Balance Point

1.6.3 Independent Regular Distributions

In this thesis, our default choice of distribution class $\mathcal{C}$ is the class of independent regular distributions. We will define the regularity condition in Chapter 2. For now it suffices to say that it is a standard condition in auction theory that contains most common distributions such as uniform, exponential, normal, and log-normal distributions.

As we show in Example 3.5, prior-independence is not possible with completely arbitrary distributions. We choose the class of regular distributions very carefully:

1. This class is sufficiently general so that (1) it contains most common distributions, and (2) for the most part, Myerson's theory uses regularity as the standard condition (although it can be extended to irregular distributions as well). So our prior-independent theory is no less general than much of Myerson's theory.

2. The class is not so general that we are forced to derive complicated mechanisms with poor approximation ratios. In fact, the mechanisms we end up proposing are simple and natural ones, with good ratios. On the other hand, in the chapters on prior-free mechanisms, we relax the regularity assumption to tail-regularity assumption, and for those cases simple mechanisms fail to work.

To summarize, our opinion is that the class of regular distributions is a sweet-spot that gives a good balance between optimality and robustness.

1.7 Prior-Independent Mechanisms

1.7.1 Mechanism with a Single Dial

For now we assume that agents are a priori indistinguishable, or in other words their valuation distributions are identical (and independent). We often adopt a three-step approach toward prior-independent mechanisms.

1. Identify a class of simple mechanisms, which are controlled by a single parameter[1] (such as price, supply limit, etc.).

   We study three classes of such mechanisms, which are:

   (a) Welfare-maximizing VCG with a reserve price.
   Here the reserve price is the parameter.

   (b) Welfare-maximizing VCG with a supply limit.
   Here the supply limit, or the number of items available, is the parameter.

   (c) Sequential posted-price mechanisms.
   Here the price we post for the agents is the parameter.

2. Prove that there is a choice of the parameter based on distribution information, so that the mechanism achieves approximately-optimal revenue. (This has already been done in previous work for some of the above mechanisms.)

3. Prove that we can set this parameter without knowledge about the prior, at a bounded loss compared to step 2.

We remark that the mechanisms we obtain at the end are based on simple modifications to revenue-maximizing strategies that naturally occur in practice. They are also in general simpler than Myerson's virtual surplus maximizers, which rely on complete distribution information to be able to compute so-called virtual values for allocation (see Section 2.4).

[1] In some general environments, we will allow multiple parameters, e.g., one reserve price for each bidder.

1.7.2 VCG with Sampled Reserve

The celebrated VCG mechanism [73, 21, 38] allocates to maximize welfare, or total values of the winners. In Chapter 4, we study our first type of prior-independent mechanism, which is based on VCG with a sampled reserve price. Loosely speaking, we set reserve prices for bidders based on other participants' bids, and then run the VCG mechanism with these reserve prices to maximize welfare. Such mechanisms are natural to bidding agents, as welfare-maximization and reserve pricing are among the most common things that auctions do in practice. In contrast, Myerson's optimal mechanism might need to calculate virtual values, which can be difficult to understand for agents.

Interestingly, we prove that if we use a single bidder's bid as the reserve price, we can achieve a good constant factor approximation. An interpretation of this result is that even a single sample from a distribution — some bidder's valuation — is sufficient information to obtain approximately-optimal revenue.

1.7.3 Supply-Limiting Mechanisms

In Chapter 5, we study our second type of mechanism, which is based on artificially limiting the supply. In such a mechanism, we do not explicitly set any prices. Instead, we make sure that only a limited fraction of total available items are allocated, so that agents have to compete for the items, which drives up the prices they end up paying.

It turns out that for a variety of settings, by simply halving the available supply, we can already achieve prior-independent constant factor approximation. Notably, this supply-limiting approach applies to a multi-dimensional matching problem. For such multi-dimensional domains, identifying the optimal mechanism is considered a difficult problem.

1.7.4 Sequential Posted-Price Mechanisms

In Chapter 6, we study our third type of mechanisms, which is based on Sequential Posted-price Mechanisms (SPMs). In such a mechanism, we make take-it-or-leave-it pre-determined price offers to agents according to a pre-determined fixed ordering. Such a mechanism resembles the Buy-It-Now mechanism available on eBay, and how we sell goods like bikes to friends.

Chawla et al. [19] proved that such simple mechanisms can achieve approximately-optimal revenue. We give a more structured proof of this fact using the notion of correlation gap and submodularity. Then based on the insights from the new proof, we show how to set the price in an SPM without knowing the prior distribution, while still achieving approximately-optimal revenue.

1.7.5 Reduction to Bulow-Klemperer-Style Statements

A Bulow-Klemperer-style statement says that instead of running the optimal mechanism over the original environment, we can achieve expected revenue that is (approximately) as good by first expanding the environment by drawing more bidders (from the same distribution), and then running the VCG mechanism. Such a statement is prior-independent in nature, except that we need more bidders that are not available in the original environment.

In Section 5.3, we present a general reduction, for both single-dimensional and multi-dimensional environments, showing that proving prior-independent approximation can be reduced to proving a corresponding Bulow-Klemperer-style statement. We apply this reduction to prove prior-independent approximation guarantees in Section 4.4 and in Chapter 5.

1.8 Prior-Free Mechanisms and Envy-Freeness

A prior-independent guarantee is a guarantee in expectation. One would often hope for a more robust guarantee that holds for every input profile point-wise. However, it is not clear what it means to have good guarantee for every input.

To this end, in Chapters 7 and 8, we propose to consider the concept of envy-freeness that is tightly related to truthfulness. Unlike truthfulness, for every input profile, there is an envy-free outcome that is optimal in revenue. We characterize the optimal revenue of an envy-free outcome as the (ironed) virtual surplus maximizer, which parallels Myerson's theory for optimal truthful mechanisms, and propose to use this optimal revenue as a point-wise revenue benchmark. We prove that this benchmark is well-motivated in the sense that if a mechanism can approximate this benchmark for every input, then prior-independent approximation follows automatically for a variety of settings. Correspondingly, we give mechanisms that approximate the envy-free optimal benchmark for general settings.

1.9 Differences From Original Papers

This thesis aims to be a coherent summary of our works related to prior-independence, as published in [25, 76, 66, 46]. What's new or different is summarized below:

- Chapter 2 revisits Bayesian mechanism design. We present the cleanest known proofs wherever possible.
- Chapter 3 contains the basic definitions of prior-independence, a case study for the digital goods case with two bidders, and a general reduction to Bulow-Klemperer-style theorems. In particular:
  - Section 3.1 formally defines prior-independence as a special case of parameterized prior-independence. The notion of parameterized prior-independence unifies average-case and worst-case analyses under the same umbrella, and this concept was not explicit in previous papers.
  - Section 3.3 studies the digital goods case with two bidders, which arguably gives the most insight into prior-independence. A new small result is that the Vickrey auction achieves the best prior-independent approximation ratio among all mechanisms.
  - Section 5.3 gives a general reduction of prior-independence to Bulow-Klemperer-style statements in the flavor of Bulow and Klemperer [15]. To prove the reduction, we prove that optimal revenue is fractionally subadditive, slightly strengthening the subadditivity lemma proved in [66].

- Chapter 4 is based on the technical contents of [25]. The reduction in Section 4.4 is based on duplicating and pairing agents and was not presented in previous papers.
- Chapter 5 is based on [66], with improved organization and technical presentation.
- Chapter 6 is based on [76], with a new section (Section 6.5) on how to achieve prior-independence with SPMs.
- Chapters 7 and 8 are based on [46], where we define the basic settings more explicitly for ease of understanding.

1.10 Related Work

We mention some works that are generally related to the concept of prior-independence. More works that are technically related can be found in the various chapters.

Economics Literature Most of the vast literature on revenue-maximizing auctions studies designs tailored to a known distribution over bidders' private information (see, e.g., [55]). In particular, Myerson characterized the optimal auction mechanism for a given distribution [61]. Here, we mention only the works related to approximation guarantees for prior-independent mechanisms. [62] considers single-item auctions with i.i.d. bidders, and quantifies the fraction of the optimal welfare extracted as revenue by the (prior-independent) Vickrey auction, as a function of the number of bidders. [69] and [11] prove asymptotic optimality results for certain prior-independent mechanisms when bidders are symmetric, goods are identical, and the number of bidders is large.

Prior-Freeness In prior-free auction design, distributions are not even used to evaluate the performance of an auction — the goal is to design an auction with good revenue for every valuation profile, rather than in expectation. [36] proposes a revenue benchmark approach, which has been applied successfully to a number of auction settings.

Approximation in this revenue benchmark framework is stronger than the prior-independent approximation goal pursued in most part of this thesis, at a bounded loss in approximation factor; this fact was made explicit in [44] and pursued further by [22, 45] for simple auction settings, where the goods are in unlimited supply and/or the bidders are symmetric; see [42] for a survey. Our work in [46] generalizes the revenue benchmark approach to much more general asymmetric settings.

For more discussion on the motivation of approximation in mechanism design, see the survey of Hartline [41].

Risk Aversion and Utility Objective In this thesis we assume that the auctioneer or seller's goal is to maximize expected revenue, which implicitly assumes that the seller is risk-neutral. Sundararajan and Yan [70] studies a setting where the seller is risk-averse, and is endowed with a monotone concave utility function on top of revenue. The objective is then to maximize expected utility. For such a setting, several mechanisms are proved to give utility-oblivious approximation guarantees, in the sense that the mechanism achieves a constant fraction of the optimal utility of a mechanism, even without knowledge about the utility function. This type of guarantee is very similar to prior-independence in spirit.

Prior-Independence in General In this thesis, we study prior-independence as a framework for revenue-maximizing auctions. In fact, prior-independence can be defined not only for auction problems, but also for algorithmic problems in general. Notably, prior-independence is tightly connected to the random permutation model that has been studied in several online algorithm problems. In particular, an i.i.d. distribution can be seen as a convex combination of random permutation sample spaces. It follows that the approximation guarantee of an algorithm for the random permutation model automatically extends to the i.i.d. distribution model, even if the distribution is unknown (cf. Chapter 7).

As an example, the classical online matching problem of Karp et al. [53] was studied in [33, 58, 52] in a model where the online nodes arrive in random order. The guarantees proved for this model directly imply guarantees for the setting proposed in [31] where the online nodes are drawn according to an i.i.d. distribution.

The random permutation model was also considered in online packing problems [3, 32], as well as in the literature on secretary problems and their matroid variants (see e.g., [8]).

Chapter 2

Preliminaries

In this chapter we review basic facts about Bayesian mechanism design. Readers who are familiar with the area are encouraged to skim but not skip this chapter.

Section 2.1 defines the single-dimensional auction environments considered in this thesis. In particular Section 2.1.1 defines matroid environments and list some of its properties. Section 2.2 defines mechanisms and truthfulness, along with several important examples of truthful mechanisms. Section 2.3 defines the various classes of distributions we work with. Section 2.4 reviews Myerson's optimal auction theory, and Section 2.5 reviews a classical result of Bulow and Klemperer.

2.1 Single-Dimensional Environments

We mainly study auction environments that are single-dimensional, in the sense that each bidder has a single private value for winning a service or good. Multi-dimensional environments will be studied in Chapter 5.

In our setting, the seller or auctioneer sells services (or goods) to a set of n bidders $N = \{1, \ldots, n\}$. Each bidder i has a private value v_i for winning the service, and 0 otherwise, where each v_i is drawn independently from a distribution F_i. For simplicity we assume that every distribution is over the support $[0, \infty)$, and has a positive smooth density function. It is only feasible for the seller to serve certain subsets of the bidders simultaneously, and we let $\mathcal{I} \subseteq 2^N$ with $\emptyset \in \mathcal{I}$ represent all the feasible subsets. We assume that the environment is downward-closed, in the sense that the subset of a feasible set is also feasible.

Auction environments can be classified by the set system $(N, \mathcal{I})$. We will be studying the following auction environments:

Single-item Auctions We can serve at most one bidder. Here $\mathcal{I}$ contains all sets of size at most 1.

Digital Goods Auctions Every set of bidders is feasible (as digital goods such as music downloads can be reproduced at no cost). Here $\mathcal{I}$ contains all bidder subsets.

k-Unit Auctions/Multi-Unit Auctions A bidder set is feasible if its size is at most k. Single-item auctions correspond to 1-unit auctions, and digital goods auctions correspond to n-unit auctions. Here $\mathcal{I}$ contains all sets of size at most k.

Matroid Environments $(N, \mathcal{I})$ forms a matroid (to be defined in Section 2.1.1). This contains multi-unit auctions as a special case. Examples of matroid environments that are not multi-unit auctions include partition matroid environments, and certain constrained matching markets (see Section 2.1.1).

p-Independent Environments $(N, \mathcal{I})$ forms a p-independent set system (to be defined in Section 6.4.3). This contains matroid environments as special cases.

Downward-closed Environments $(N, \mathcal{I})$ forms a downward-closed set system. This contains all above environments as special cases. A typical example is combinatorial auctions with single-minded bidders, where each bidder is interested in a specific bundle of items, and feasible sets correspond to sets of bidders desiring mutually disjoint bundles.

An alternative way to capture the allocation constraint is via the (weighted) rank function, which tells us how many bidders in a given bidder set can be served simultaneously.

Weighted Rank Functions For a set system $(N, \mathcal{I})$ with nonnegative weights $(w_i)_{i \in N}$ on the elements, we define the weighted rank function $w^*(S)$ for $S \subseteq N$ as the maximum of $\sum_{i \in T} w_i$ over all $T \subseteq S$ with $T \in \mathcal{I}$. The (unweighted) rank functions are defined with weights set to one.

2.1.1 Matroids

Matroid environments have a rich history in mechanism design, see e.g., [71], [12], and [45]. It is an abstraction of many allocation constraints modeled by set systems, and in some sense captures the structurally "nice" allocation constraints. See [64] for an in-depth treatment of matroid theory.

Matroid A set system $(N, \mathcal{I})$ (with $\emptyset \in \mathcal{I}$) is a matroid if (1) $S \subseteq T \in \mathcal{I}$ implies that $S \in \mathcal{I}$, and (2) if $S, T \in \mathcal{I}$ and $|S| > |T|$, then for some $e \in S \backslash T$, $T \cup \{e\} \in \mathcal{I}$. Here the second property is called exchange property.

Matroids are also closely related to the greedy algorithm.

Greedy Given a set system $(N, \mathcal{I})$ with nonnegative weights $(w_i)_{i \in N}$, and a subset S of N, the greedy algorithm over S starts with an empty solution set A, and for each bidder i in S in decreasing order of w_i, adds i into the solution set A whenever $A \cup \{i\}$ is in $\mathcal{I}$. Finally it outputs A. We let *greedy*(S) denote the final output of the greedy algorithm over S.

Note that the greedy algorithm is *ordinal* in the sense that only the relative order (but not the magnitude) of weights matters in determining the outcome.

In a matroid set system, the following are true:

1. For every set S, the total weight of the set *greedy*(S) equals to the weighted rank of set S.

2. The weighted rank function is monotone ($\forall S, T : w(S) \leq w(T)$ whenever $S \subseteq T$) and submodular ($\forall S, T : w(S) + w(T) \geq w(S \cup T) + w(S \cap T)$).

We give three important examples of matroids below.

k-Uniform Matroid In a k-uniform matroid, a set is feasible if and only if its size is at most k. This allows us to model the constraint of a k-unit auction.

Partition Matroid In a partition matroid, the ground set is partitioned into ℓ parts, where each part is associated with a limit on how many elements can be taken from it. A set is feasible if the number of elements it contains from each part is within the limit for the part. This allows us to model an auction setting where the bidder set is divided into sub-markets, and we have supply limit for each sub-market.

Transversal Matroid A transversal matroid over element set N is associated with a bipartite graph with node set N on one side, and M on the other. A subset S of N is feasible if there is a matching between N and M where every element of S is matched. This matroid allows us to model a constrained matching scenario.

2.2 Mechanisms

A (deterministic) mechanism $\mathcal{M}$ comprises an *allocation rule* $\mathbf{x}$ that maps every bid vector $\mathbf{b}$ to a characteristic vector of a feasible set (in $\{0,1\}^n$), and a *payment rule* $\mathbf{p}$ that maps every bid vector $\mathbf{b}$ to a non-negative payment vector in $[0,\infty)^n$. We assume that every bidder i aims to maximize its quasi-linear utility $u_i(\mathbf{b}) = v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$, where v_i is her private valuation for winning.

Measure-Theoretic Assumption As we will be studying expected revenue of a mechanism over a distribution, we shall also assume that x_i, p_i's are Lebesgue-measurable, so that expected revenue can be well-defined. We will skip measure-theoretic discussions regarding whether each expectation in this thesis is well-defined, to make our presentation clean.

Definition 2.1. We say a mechanism $\mathcal{M}$ is *truthful* or ex post incentive compatible if for every bidder i and fixed bids $\mathbf{b}_{-i}$ of the other bidders, bidder i maximizes its utility by setting its bid b_i to its private valuation v_i. Formally, for all $i, \mathbf{b}_{-i}, b$, we have $v_i \cdot x_i(\mathbf{b}_{-i}, v_i) - p_i(\mathbf{b}_{-i}, v_i) \geq v_i \cdot x_i(\mathbf{b}_{-i}, b) - p_i(\mathbf{b}_{-i}, b)$. By default, we also require individual rationality, which means that for all $i, \mathbf{b}_{-i}$, we have $v_i \cdot x_i(\mathbf{b}_{-i}, v_i) - p_i(\mathbf{b}_{-i}, v_i) \geq 0$.

Since we only consider truthful mechanisms in this thesis, in the rest of the thesis we use the terms of values and bids interchangeably. We will also use mechanisms to refer to truthful mechanisms for brevity.

A well-known characterization of truthful mechanisms in single-dimensional environments [61, 6] states that a mechanism $(\mathbf{x}, \mathbf{p})$ is truthful if and only if the allocation rule is monotone — that is, $x_i(b'_i, \mathbf{b}_{-i}) \geq x_i(\mathbf{b})$ for every i, $\mathbf{b}$, and $b'_i \geq b_i$ — and the payment rule is given by a certain formula involving the allocation rule. We often specify a truthful mechanism by its monotone allocation rule, with the understanding that it is supplemented with the unique payment rule that yields a truthful mechanism. For deterministic mechanisms like those studied in this thesis, the payment of a winning bidder is simply the smallest (or the infimum of) bid she needs to bid to remain a winner.

Yet another equivalent way of defining the truthfulness constraints is that for each bidder i, if we fix the valuations $\mathbf{v}_{-i}$ of the other bidders, bidder i faces a take-it-or-leave-it offer at a price $p_i(\mathbf{v}_{-i})$ that is independent of bidder i's own value v_i.

Occasionally, we consider randomized mechanisms, where a randomized mechanism is simply a distribution over deterministic truthful mechanisms.

2.2.1 Objectives

There are two commonly studied objectives in mechanism design, which are welfare and revenue. Here welfare (also called efficiency) refers to the total value of the winners, or formally $\sum_i v_i \cdot x_i(\mathbf{v})$, and revenue refers to the total payment the auctioneer receives from the bidders, or formally $\sum_i p_i(\mathbf{v}) \cdot x_i(\mathbf{v})$. By individual rationality, the revenue of a mechanism outcome is bounded above by its welfare.

The main objective of this thesis is revenue, although most of our theory does extend to objectives that are convex combinations of welfare and revenue. The revenue objective is more interesting and challenging than the welfare objective. In particular, the well-known VCG mechanism achieves the best possible welfare for every valuation profile $\mathbf{v}$, while on the other hand, there is no mechanism such that for every valuation profile, this mechanism is optimal in revenue compared to all other mechanisms.

2.2.2 Example: VCG and Variants

For both single-dimensional and multi-dimensional environments, the celebrated VCG mechanism [73, 21, 38] maximizes welfare, and is truthful. In particular, it chooses the feasible set $S \in \mathcal{I}$ that maximizes the welfare $\sum_{i \in S} v_i$, and charges each bidder i its externality, which is the maximum welfare of all other bidders when i is not in the system, minus the welfare of all other bidders in the chosen allocation. Ties are broken in an arbitrary way that does not depend on the bids.

A useful variant of the VCG mechanism is based on artificially setting a supply limit. For example for the k-unit auction setting, the VCG$^{\leq \frac{n}{2}}$ mechanism allocates to maximize welfare, subject to the constraint that at most half of the bidders can win. More elaborate forms of supply-limiting exist, even for multi-dimensional environments. In general such mechanisms are special cases of maximal-in-range mechanisms (see [63]), which are well-known to be truthful.

For single-dimensional environments, two variants of the VCG mechanism are also important. Let r_i be a reserve price for bidder i. The VCG mechanism with eager reserves $\mathbf{r}$ (VCG-E) works as follows, given bids $\mathbf{v}$: (1) delete all bidders i with $v_i < r_i$; (2) run the VCG mechanism on the remaining bidders to determine the winners; (3) charge each winning bidder i the larger of r_i and its VCG payment in step (2). In the VCG mechanism with lazy reserves $\mathbf{r}$ (VCG-L), steps (1) and (2) are reversed. Both of these mechanisms are feasible and truthful in every downward-closed environment. In simple environments, these two variants are equivalent (Corollary 4.4), and will be both called VCG with reserves. But they are different in general.

2.2.3 Example: Virtual Surplus Maximizer

For single-dimensional environments, a virtual surplus maximizer is specified by a non-decreasing virtual value function $\phi : [0, \infty) \to (-\infty, \infty)$. The virtual surplus maximizer mechanism allocates to maximize the total virtual value of the winners, or formally, $\sum_i x_i(\mathbf{v}) \cdot \phi(v_i)$, and charges the appropriate payments. Ties are broken in an arbitrary way that does not depend on the bids. It is straightforward to verify that such a mechanism is monotone and hence truthful. In general, we can also specify a virtual value function for each bidder respectively.

2.2.4 Example: Sequential Posted-Price Mechanisms

Given an ordering of bidders and a price p_i for each bidder i, a Sequential Posted-price Mechanism (SPM) first initializes the allocated set A to be $\emptyset$, and for all bidders i in the given order, do the following: if serving i is feasible, i.e., $A + i \in \mathcal{I}$, offer to serve bidder i at the pre-determined price p_i, and add i to A if bidder i accepts.

The ordering of bidders can be crucial in SPMs. The greedy ordering will be of particular interest to us, where we make sure that bidders according to our ordering have descending posted prices.

2.3 Distributions

In most part of this thesis, we will make the assumption that the each bidder's value is drawn independently from a distribution (although this distribution may or may not be given to us). For technical convenience, we make the following basic assumptions on all distributions we work with:

- It has a support of the form of either $[L, \infty)$ or $[L, H]$ for some $L, H \geq 0$.
- It has a positive and smooth density function with one allowed exception. The exception is that if the support is of the form of $[L, H]$, then a constant amount of probability mass can concentrate on the point H, and the density function can be not smooth at this point.
- It has a finite expectation.

The above-mentioned exception allows us to define distributions such as the equal revenue distributions (Example 2.9) directly.

For technical convenience, we often assume that $L = 0$, as a valid distribution with $L > 0$ can be adapted slightly to be a valid distribution over a support with $L = 0$, with a negligible effect on the analysis.

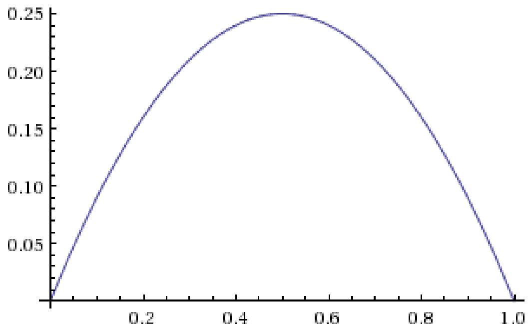

(a) Uniform Distribution over [0,1]

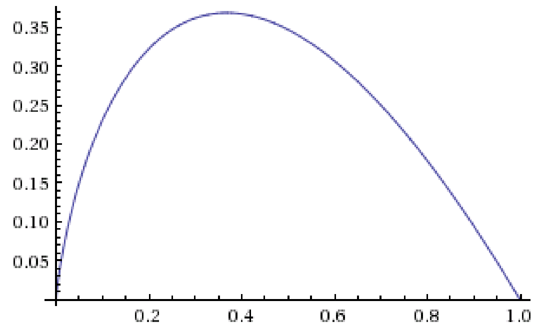

(b) Exponential Distribution with Rate 1

Figure 2.3.1: Revenue Curves of Uniform and Exponential Distributions

Distribution information is commonly captured by either the Cumulative Distribution Function (CDF in short) F, or the Probability Density Function (PDF in short) f. However, for the purpose of revenue maximization, a third way of capturing distribution information turns out to be crucial, which is based on the notion of the revenue curve function.

Definition 2.2 (Revenue Curve Function)**.** Given a distribution F, the revenue curve function is defined as $R_F(q) = q \cdot F^{-1}(1-q)$ for all probability $q \in (0,1]$, and $R_F(0) = 0$.

In the case that F is defined over a support $[L, H]$ where a constant amount q_H of probability mass is concentrated at H, we let $F^{-1}(1-q) = H$ for $q < q_H$.

To associate the revenue curve function with a semantic meaning, consider a single-bidder single-item auction with underlying distribution F. If we offer a price such that the probability of sale is q, then the price should be $F^{-1}(1-q)$, and then $R_F(q)$ is the expected revenue we achieve. We can define a distribution F by specifying an appropriate $R_F(\cdot)$ function.

Figures 2.3.1a and 2.3.1b depict the revenue curve function of the uniform distribution over [0,1], and the exponential distribution with rate 1, respectively. In both cases, the revenue curves are concave.

2.3.1 Regular Distributions

In most part of the thesis, we assume that the distributions satisfy a standard *regularity* condition, which we define as follows (cf. [61], [15]).

Definition 2.3 (Regular Distributions)**.** A distribution F is *regular* if the revenue function $R_F(\cdot)$ w.r.t. F is concave.

An equivalent way of defining regular distributions is based on the notion of virtual values.

Definition 2.4 (Virtual Value Function)**.** Given a distribution F, its virtual value function is $\phi_F(v) = v - \frac{1}{h_F(v)}$, where $h_F(v) = \frac{f(v)}{1-F(v)}$ is the hazard rate function w.r.t. F.

Lemma 2.5. *Given a distribution F and value p with $q = 1 - F(p)$. Then $\frac{dR_F(q)}{dq} = \phi_F(p)$.*

Proof. The lemma follows from the following equalities:

$$\begin{aligned}
\frac{dR_F(q)}{dq} &= \frac{d(qF^{-1}(1-q))}{dq} \\
&= F^{-1}(1-q) + q\frac{dF^{-1}(1-q)}{dq} \\
&= p - q \cdot \frac{1}{f(F^{-1}(1-q))} \\
&= p - \frac{1-F(p)}{f(p)} \\
&= \phi_F(p).
\end{aligned}$$

□

Corollary 2.6. *A distribution is regular if and only if its virtual value function is non-decreasing.*

Three extreme examples of regular distributions include the class of left-triangle distributions, the class of right-triangle distributions, and the class of equal-revenue distributions. We give the definitions below, and depict the corresponding revenue curves in Figure 2.3.2.

Example 2.7 (Left-Triangle Distribution)**.** A left-triangle distribution F with parameters H is given by its revenue curve function $R(q) = qH$ for $q \in [0, \frac{1}{H+1}]$, and $R(q) = 1 - q$ for $q \in [\frac{1}{H+1}, 1]$. More explicitly, its CDF is given by $F(v) = 1 - \frac{1}{1+z}$ for $[0, H)$ and $F(H) = 1$.

Example 2.8 (Right-Triangle Distribution)**.** A right-triangle distribution F with parameter H is essentially the single-point distribution that $Pr_{v \sim F}[v = H] = 1$. Its revenue curve function is $R(q) = qH$ for $q \in [0, 1]$.

Example 2.9 (Equal-Revenue Distribution)**.** An equal-revenue distribution capped at H has CDF F defined as $F(z) = 1 - \frac{1}{z}$ for $x \in [1, H)$ and $F(H) = 1$.[1] It is called equal-revenue distribution because every price from $[0, H]$ leads to the same expected revenue of 1.

As many important distributions are regular, regularity has been considered as a standard assumption. We do note that there exist distributions that are not regular, e.g., bimodal distributions. We refer the readers to [28] for a list of regular distributions as well as a list of irregular distributions. Yet another alternative definition of regularity is also given in [28].

2.3.2 Monotone Hazard Rate Distributions

An important (strict) subclass of regular distributions is the class of monotone hazard rate distributions.

Definition 2.10 (Monotone Hazard Rate Distributions)**.** A distribution F satisfies the *monotone hazard rate* condition (or simply F is m.h.r.), if $h_F(v) = \frac{f(v)}{1-F(v)}$ is nondecreasing in v.

Some examples of m.h.r. distributions are uniform, exponential, and normal distributions. The log-normal distributions are regular but not m.h.r. The three extreme regular distributions in Examples 2.7, 2.8, and 2.9 are also not m.h.r.

[1] Here the capping at H is essential. Without capping, the distribution would have an infinite expectation.

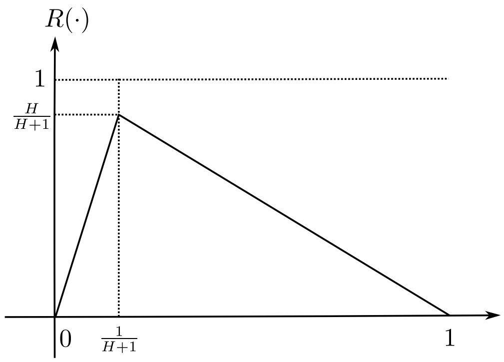

(a) Left-Triangle Distribution with Parameter H

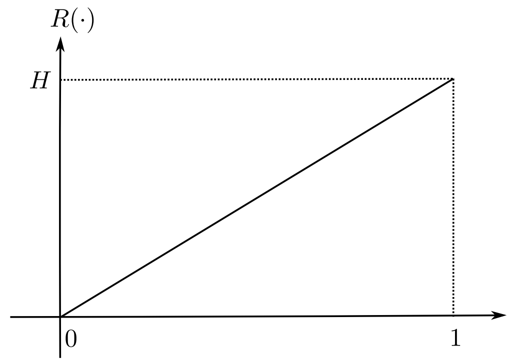

(b) Right-Triangle Distribution with Parameter H

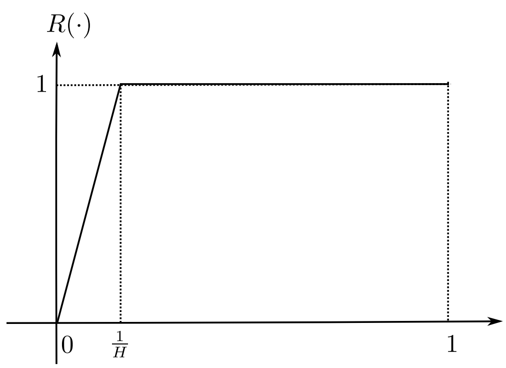

(c) Equal-Revenue Distribution Capped at H

Figure 2.3.2: Revenue Curves of Several Extreme Regular Distributions

2.4 Myerson's Optimal Auction Theory

If bidders' values are drawn from a joint distribution F, then there exists an optimal mechanism OPT_F that achieves as much expected revenue (over F) as any other mechanism. Myerson [61] characterizes the optimal mechanisms for single-dimensional environments when bidders' values are drawn from independent distributions. A key element in the characterization is the notion of virtual values (Definition 2.4).

To understand virtual values, consider the simple setting of a single bidder and a single item. The bidder's value for the item is drawn from a distribution F. Suppose that given a target probability of sale q, we offer the corresponding price $p = F^{-1}(1-q)$ to the bidder. The expected revenue as a function of q can be written as $R_F(q) = q \cdot F^{-1}(1-q)$. On the other hand, by Lemma 2.5, the virtual value of v is equal to $\frac{dR_F(q)}{dq}$ where $q = 1 - F(v)$, which we collect whenever $v \geq p$. Therefore the expected virtual value we achieve is $\int_p^\infty 1_{v \geq p} \phi(v) dv = \int_0^q \frac{dR_F(q)}{dq} dq$, which equals to $R_F(q)$.

Essentially, for the single-bidder case, no matter what price the bidder faces, expected revenue equals to expected virtual value.

Now suppose there are multiple bidders. The same discussion applies, because from a fixed bidder's point of view, the effect of having other bidders simply means that she now faces a threshold price that she has to outbid to win, where this threshold price is induced by other bidders. No matter what this threshold price is, the expected revenue we get from this fixed bidder equals to expected virtual value by the same argument as in the single-bidder case. Therefore we have the following lemma:

Lemma 2.11 (Myerson's Lemma)**.** *For every truthful mechanism* $(\mathbf{x}, \mathbf{p})$*, the expected payment of a bidder* i *with valuation distribution* F_i *satisfies*

$$E_{\mathbf{v}}[p_i(\mathbf{v})] = E_{\mathbf{v}}[\varphi_{F_i}(v_i) \cdot x_i(\mathbf{v})].$$

Moreover, this identity holds even after conditioning on the bids $\mathbf{v}_{-i}$ *of the bidders other than* i*.*

In other words, expected revenue maximization can be reduced to the maximization of virtual surplus (sum of total virtual values). Consider the mechanism that chooses a feasible set S that maximizes the virtual surplus $\sum_{i \in S} \varphi_{F_i}(v_i)$. Because distributions are regular, virtual value functions are monotone, and it follows that this mechanism is monotone, and hence truthful.

Theorem 2.12 (Myerson's Optimal Mechanism)**.** *When distributions are independent and regular, Myerson's mechanism that allocates to the feasible set that maximizes total virtual surplus also maximizes expected revenue.*

The role of regularity is to ensure that this allocation rule is indeed monotone; otherwise, additional ideas are needed [61].

Remark 2.13. We have constrained ourselves to truthfulness, or ex post incentive compatibility. Myerson also studied the weaker solution concepts of Bayesian Incentive Compatibility (Bayes-IC) and Bayesian Nash Equilibrium (Bayes-NE). In such a Bayesian setting, each bidder knows her own value, and the distributions (but not values) of the other bidders. In a Bayes-IC mechanism, every bidder maximizes her expected utility (in expectation over distributions of the other bidders) by bidding truthfully, given that other bidders also bid truthfully. Myerson's optimal mechanism in fact maximizes expected revenue among all Bayes-IC mechanisms. An application of the revelation principle even allows this guarantee to hold for the even larger class of Bayes-NE mechanisms. In other words, relaxing to these weaker solution concepts do not give us additional power in revenue maximization, and it is without loss of generality that we focus on truthfulness, or ex post incentive compatibility in this thesis.

2.5 Bulow-Klemperer Theorem and Its Generalizations

The classical Bulow-Klemperer theorem says that instead of running Myerson's optimal mechanism for selling a single item to n bidders with values drawn i.i.d. from a regular distribution, the auctioneer is better off running the Vickrey auction with one additional bidder. The following is a straightforward generalization of the theorem to k-unit auctions.

Theorem 2.14 (Generalized Bulow-Klemperer Theorem)**.** *[15, 26] For every k-unit environment with i.i.d. regular bidders, the expected revenue of VCG with k additional bidders is at least as high as the optimal expected revenue.*

The following proof is based on generalizing the proof of [54] in a straightforward way.

Proof. Starting with the optimal mechanism for k-unit auctions over n bidders, we construct an intermediate mechanism for k-unit auctions over $n+k$ bidders, which first runs the optimal mechanism over the first n out of $n+k$ bidders w.r.t. a fixed bidder ordering, and then give the left-over unallocated items to the last k bidders for free. It is easy to verify that this is a truthful mechanism. The theorem now follows from the following two simple observations.

(1) The intermediate mechanism has the same revenue as the optimal mechanism. This is because the last k bidders never pay any amount.

(2) Among all mechanisms that always allocates all k items (which includes the intermediate mechanism in particular), the VCG mechanism maximizes expected revenue. To see this, among all mechanisms that allocate all k items, the one that maximizes expected revenue allocates to the k bidders with the highest virtual values by Myerson's Lemma. By the regularity assumption, the virtual values are monotone in values. It follows that the VCG mechanism, which maximizes total value, also maximizes total virtual value, and hence expected revenue as well. □

2.6 Properties of the Optimal Revenue Function

Given an auction environment (single-dimensional or multi-dimensional), we abuse notation to let $\mathrm{OPT}(S)$ denote both the optimal mechanism and its expected revenue for the sub-environment induced by bidder set S. In this section we study the properties satisfied by this set function.

We mention a few standard definitions regarding set functions. For a finite ground set N, a nonnegative set function $f(S) : 2^N \to [0, \infty)$ is:

- monotone, if $f(S) \le f(T)$ whenever $S \subseteq T$

- submodular if $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ for all S, T
- fractionally subadditive if for every bidder set S and a fractional cover of S, i.e., T_j for all j and coefficients $0 \leq \alpha_j \leq 1$ for all j such that for every $i \in S$, $\sum_{j|i \in T_j} \alpha_j \geq 1$, we have $f(S) \leq \sum_i \alpha_j f(T_j)$.
- subadditive if $f(S) + f(T) \geq f(S \cup T)$ for all S, T.

Assuming that f is monotone, submodularity implies fractional subadditivity, which then implies subadditivity. Readers are referred to Feige [30] for more discussions on these classes of set functions.

All single- or multi- dimensional environments studied in this thesis satisfy the following condition: every feasible allocation for a smaller bidder set S is also a feasible allocation for a bigger bidder set T that contains S. This condition allows us to argue that $\mathrm{OPT}(\cdot)$ is monotone. This is because an optimal mechanism for S is trivially a mechanism for T (where bidders in $T \backslash S$ never wins anything), and the optimal mechanism for T can only be better.

Dughmi et al. [26] prove that for single-dimensional matroid environments, $\mathrm{OPT}(\cdot)$ is monotone and submodular. In Chapter 5 we prove for both single-dimensional and multi-dimensional environments, that $\mathrm{OPT}(\cdot)$ is monotone and fractionally subadditive. The statement of the lemma is as follows.

Lemma 2.15 (Fractional Subadditivity of Optimal Expected Revenue)**.** *For every (single-dimensional or multi-dimensional) auction environment,* $\mathrm{OPT}(\cdot)$ *is fractionally subadditive.*

Proof. By revenue monotonicity, it is sufficient to prove the claim for the case that $T_j \subseteq S$ for all j.

For every j, the optimal mechanism $\mathrm{OPT}(S)$ induces a randomized mechanism $\mathcal{M}_j$ for set T_j. In particular, this mechanism randomly draws values for bidders from $S \backslash T_j$ according to their distributions, and simulates $\mathrm{OPT}(S)$ on T_j, cancelling outcome for fake bidders of $S \backslash T_j$. The expected revenue of $\mathcal{M}_j$ (also denoted by $\mathcal{M}_j$) is bidder-wise the same as the expected revenue of $\mathrm{OPT}(S)$ for all bidders in T_j. By the fractional covering assumption, $\sum_j \alpha_j \mathcal{M}_j \geq \mathrm{OPT}(S)$. Our lemma follows from the fact that $\mathrm{OPT}(T_j) \geq \mathcal{M}_j$ for every j. $\square$

This lemma allows us to easily bound the loss in removing bidders from the system. For example, if we randomly remove one of the n bidders, then the optimal revenue of the remaining environment is at least a $(1 - \frac{1}{n})$-fraction of the original environment.

Part II

Prior-Independent Mechanisms

Chapter 3

Prior-Independence: Definition, Example, and Reduction

In this chapter, we formally define the prior-independent analysis framework, study how to achieve prior-independent approximation in a simple illuminating setting, and prove a general reduction of prior-independent approximation to Bulow-Klemperer-style statements.

3.1 Prior-Independent Approximation

We repeat the definition of parameterized prior-independence from Chapter 1.

Definition 3.1 (Parameterized Prior-Independence). Given an auction environment with at least two bidders, we say that a mechanism $\mathcal{M}$ gives a prior-independent ρ-approximation w.r.t. distribution class $\mathcal{C}$ if:

$$\mathrm{E}_{\mathbf{v}\sim F}[\text{revenue of } \mathcal{M}(\mathbf{v})] \geq \rho \cdot \mathrm{E}_{\mathbf{v}\sim F}[\text{revenue of } \mathrm{OPT}_F(\mathbf{v})]$$

for all $F \in \mathcal{C}$.

Here the mechanism $\mathcal{M}$ can dependent on $\mathcal{C}$ but not the specific F.

Informally, we require that the expected revenue of our mechanism over every distribution from the clas is approximately as good as the optimal expected revenue of a mechanism tailored for the distribution. Moreover, our mechanism should have no specific knowledge about the distribution, except it comes from certain distribution class.

Remark 3.2. It should be obvious that there need to be at least two bidders for prior-independence to make sense. This is because for single-bidder case, every truthful mechanism is essentially a fixed price. It is impossible to set a good price without any idea about the scale of the distribution. See also [36].

3.2 Distribution Classes

What we left open in the definition is the choice of distribution class. We will explore several different options to the choice of $\mathcal{C}$, which are listed below in order of increasing generality.

$$\begin{array}{cc} & \text{monotone hazard rate distributions} \\ \subseteq & \text{regular distributions} \\ \subseteq & \text{tail-regular distributions} \\ \subseteq & \text{arbitrary distributions} \end{array}$$

We assume that distributions of different agents are independent.

Our main focus will be on prior-independence w.r.t. independent regular distributions. Monotone hazard rate distributions will be studied in Chapter 4, and tail-regular distributions will be studied in Chapters 7 and 8. Example 3.5 shows that if we allow arbitrary distributions, prior-independence is not achievable.

3.3 Case Study: Digital Goods with Two Bidders

A complete understanding of a simple special case will be illuminating, revealing the core insight that will lead us to general results. In this section, we convey a thorough study of digital goods auctions with two i.i.d. regular bidders, which will give us the most insight on prior-independence. In particular, we will learn that:

1. A single sample can be sufficient information about the distribution to achieve approximately optimal revenue.

2. Supply limiting can be a good strategy for revenue maximization.

3. Proving a prior-independent approximation is connected to proving Bulow-Klemperer-style statements in the flavor of Bulow and Klemperer [15].

Many results in Chapters 4 and 5 can be seen as generalizations of results in this section in various ways.

3.3.1 Characterizing All Truthful Mechanisms

Although in general the space of truthful mechanisms can be too rich to understand, in the case of digital goods auctions with two bidders, we can give a very explicit description of the space of truthful mechanisms [36].

Let the two bidders be named Alice and Bob. Their values v_A and v_B respectively are drawn i.i.d. from a regular distribution F. Every (deterministic) truthful mechanism can be essentially characterized by two "threshold" functions $t_A, t_B : [0, \infty) \to [0, \infty)$, where Alice wins if and only if $v_A \geq t_A(v_B)$ (or $v_A > t_A(v_B)$) and Bob wins if and only if $v_B \geq t_B(v_A)$ (or $v_B > t_B(v_A)$). Recall that we assume that these two threshold functions are Lebesgue-measurable, so that expected revenue of the mechanism is well-defined.

3.3.2 Revenue of the Vickrey Auction

Without knowledge about the prior distribution, perhaps the most natural thing to do is to set $t_A(v_B) = v_B$ for all v_B, and $t_B(v_A) = v_A$ for all v_A. In fact, this is what happens in the Vickrey auction (up to what happens in the case of a tie), where Alice wins if and only if she outbids Bob, and vice versa.

It turns out that Vickrey auction gives a prior-independent $\frac{1}{2}$-approximation. The following proof is from [25].

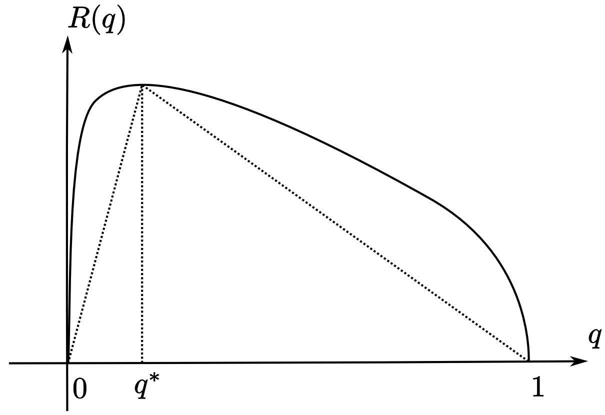

Figure 3.3.1: Random vs. Optimal Price: Revenue Comparison

Theorem 3.3. *For digital goods auctions with two i.i.d. regular bidders, the Vickrey auction gives a prior-independent $\frac{1}{2}$-approximation to optimal expected revenue.*

Proof. First assume that the underlying distribution is defined over $[0, \infty)$.

In the Vickrey auction for two bidders, from each bidder's point of view, she/he faces a random price drawn from distribution F. On the other hand, in the optimal mechanism for F, each bidder is offered the optimal price p^*. It suffices to quantify how much we lose by using a random price instead of the optimal price.

We claim that the expected revenue from using a random price is at least half of that from using the optimal price, for regular distributions. A random price $p \sim F$ has a selling probability $q = 1 - F(p)$ that is distributed uniformly over $[0, 1]$. It follows that the expected revenue from offering a random price can be written as $\mathrm{E}_{q \sim U[0,1]}[R(q)]$. Pictorially, as depicted in Figure 3.3.1, this equals to the area under the curve defined by R, and above the q-axis. On the other hand, the expected revenue from offering the optimal price equals to $R(q^*)$, which equals to the height of the curve. By concavity, the area under the curve is at least the area of the triangle in Figure 3.3.1, which equals to half of the height.

For the case that the support of F starts at $L > 0$, it is easy to slightly modify the distribution so that it becomes a regular distribution defined with support starting from $L = 0$, with a negligible effect on the analysis.

Finally we deal with the slightly tricky case where the support of F is over $[0, H]$, where a constant amount of probability mass q_H is concentrated on H. Assume randomized tie-breaking of Vickrey. I.e., there is a sale with probability $\frac{1}{2}$ if the value exactly equals the price. Conditioning on that the random price turns out to be H, which happens with probability q_H, our expected revenue is $\frac{1}{2}q_H H$. So our total expected revenue from the case that the random price equals H is $\frac{1}{2}q_H^2 H$. Recall we defined $R(q) = qH$ for $q < q_H$ in this case, and hence the corresponding "area" $\int_0^{q_H} R(q)dq$ also equals $\frac{1}{2}q_H^2 H$, and our geometric argument for previous case extends. □

3.3.3 Interpretation I: A Single Sample is Near-Optimal

In the proof of Theorem 3.3, we reduced the claim that Vickrey gives a prior-independent $\frac{1}{2}$-approximation to the claim that a random price is half-optimal. Essentially, we use a single sample from the distribution, and this gives sufficient information to achieve approximately optimal revenue, at least in digital goods auctions with two bidders.

In fact, this holds much more generally. In Chapter 4, we prove that for a variety of settings, a random reserve price (taken from a participating agent's bid) appropriately combined with VCG mechanism can give good prior-independent guarantees.

3.3.4 Interpretation II: Supply-Limiting is Good for Revenue

In a digital goods auction with two bidders, we have two items available for two bidders. On the other hand, the Vickrey auction commits to selling exactly one of the two items, effectively limiting the supply from two to one. This makes the two bidders compete against each other, whereas competition drives up revenue.

In fact, in Chapter 5, we prove that supply-limiting is a good revenue-maximizing strategy in general, which gives a prior-independent guarantees for a variety of settings, including a multi-dimensional matching environment.

3.3.5 Interpretation III: Market Expansion is Good for Revenue

Now suppose we only had one item and one regular bidder. The optimal revenue in this case is exactly half of that of the two-bidder case we just studied. Theorem 3.3 can then be reinterpreted as a Bulow-Klemperer-style statement.

- Compared to running the optimal auction for one bidder, with an extra bidder (with the same value distribution), the Vickrey auction gives at least as much expected revenue.

In fact, this is the same as the Bulow-Klemperer theorem (cf. Theorem 2.14) for the one bidder case.

Essentially, a Bulow-Klemperer-style statement is very close to a prior-independent guarantee, as the Vickrey auction does not reference distribution information. The only difference is that the Vickrey auction runs over an expanded set of bidders.

In fact, prior-independent approximation and Bulow-Klemperer-style statements are tightly related. In Section 5.3, we show that prior-independent approximations can be reduced to Bulow-Klemperer-style statements for a variety of environments, and toward a prior-independent approximation, what we commonly do is to first prove a corresponding version of a Bulow-Klemperer-style statement.

3.3.6 Vickrey is the Best Prior-Independent Mechanism

In this subsection, for digital auctions with two bidders, we study whether there exists a mechanism that achieves a better prior-independent approximation ratio than Vickrey auction.

Proposition 3.4. *For digital goods auctions with two i.i.d. regular bidders, every truthful mechanism achieves a prior-independent approximation ratio of at most* $\frac{1}{2}$.

Proof. Suppose for contradiction that some mechanism with threshold functions $t_A(\cdot)$ and $t_B(\cdot)$ achieves a prior-independent approximation ratio of strictly larger than $\frac{1}{2}$.

We first argue that $t_A(v) \leq v$ and $t_B(v) \leq v$. Suppose for contradiction that for some v, $t_A(v) > v$. Let the distribution be the single-point distribution over v (i.e., the right-triangle distribution of Example 2.8 with parameter v). Then we get zero revenue from Alice, and we can get at best a $\frac{1}{2}$-approximation to optimal revenue, contradiction.

Next consider the left-triangle distribution (see Example 2.7) with parameters H for sufficiently large H. For this distribution, a higher price within the range of $[0, H)$ gives higher expected revenue. Therefore among all mechanisms with $t_A(v) \leq v$ and $t_B(v) \leq v$ for $v \leq H$, the revenue-optimal mechanism satisfies $t_A(v) = t_B(v) = v$ for all $v < H$. For such a mechanism, by an area vs. height argument as in proof of Theorem 3.3, the approximation ratio is $\frac{1}{2}$, but not strictly larger than $\frac{1}{2}$. What happens when $v = H$ will not matter, as a random value is in this range with negligible probability for sufficiently large H. □

The following example shows that if $\mathcal{C}$ contains arbitrary independent distributions, then no constant factor approximation is possible, even if randomized mechanisms are allowed.

Example 3.5. Consider a digital goods auction with two i.i.d. bidders. Let mechanism $\mathcal{M}$ have $t_A(0) = h_A$, and $t_B(0) = h_B$. Let H be a much larger number than h_A and h_B, and consider a distribution where a random draw takes value H with probability $\frac{1}{H}$, and 0 otherwise. The optimal mechanism is to offer a price H to each bidder, achieving an expected revenue of 1 from each bidder. On the other hand, to calculate the expected revenue of $\mathcal{M}$ from Alice, with probability $1 - \frac{1}{H}$, Bob's value is 0, and $\mathcal{M}$ collects expected revenue of $\frac{1}{H}h_A$ from Alice, and with probability $\frac{1}{H}$, Bob's value is H, and no matter what price Alice is offered, the expected revenue from Alice in this case is upper-bounded by 1. In total the expected revenue from Alice (and Bob, symmetrically) is negligible compared to 1.

If the mechanism is a randomized one, the threshold functions t_A and t_B map real values to random real values. However, the argument extends directly if we replace h_A and h_B by the expected values of $t_A(0)$ and $t_B(0)$, respectively.

Chapter 4

VCG with Sampled Reserves

In this chapter, we propose mechanisms based on the VCG mechanism with random reserve prices, and prove that they give prior-independent approximation. This chapter is mainly based on [25].

4.1 Introduction

4.1.1 Settings

We consider attribute-based single-dimensional downward-closed environments. In such an environment, each bidder has an observable attribute, and we assume that the valuations of bidders with a common attribute are drawn i.i.d. from a distribution that is unknown to the seller. Bidders with different attributes can have valuations drawn (independently) from completely different distributions. For example, based on (publicly observable) eBay bidding history, one might classify bidders into "bargain-hunters", "typical", and "aggressive", with the expectation that bidders in the same class are likely to have similar valuations, without necessarily knowing how their valuations for a given item are distributed. We assume that the environment is *non-singular*, meaning that there is no bidder with a unique attribute.[1]

[1] No prior-independent auction has a non-trivial approximation guarantee when there is a bidder with a unique attribute. The reasoning is similar to that above for arbitrary valuation distributions; see Section 3.5 and also [36].

4.1.2 Approach

Toward prior-independence, we instantiate the three-step approach described in Section 1.7.1.

1. Identify a class of mechanisms with certain parameters.

   In particular we consider both VCG-E and VCG-L, i.e., welfare-maximizing VCG with both eager and lazy reserves (see Section 2.2.2 for definitions), which are parameterized by the reserves set for bidders.

2. Given distribution information, prove that for appropriate reserve prices VCG with these reserves is approximately optimal.

3. Without distribution information, set these reserve prices, at a bounded loss in revenue.

We mention how we achieve 2 and 3 for regular matroid environments in the following. Similar results holds for m.h.r. downward-closed environments.

Toward 2, if bidders are i.i.d., VCG (either VCG-E or VCG-L) with monopoly reserve happens to be the same as Myerson's optimal mechanism. For the non-i.i.d. case, Hartline and Roughgarden [45] prove that the VCG-E mechanism with monopoly reserves are approximately-optimal. We prove that VCG-L with monopoly reserves are also approximately-optimal.

To achieve 3, we study how to learn the monopoly reserve by taking samples from the distributions. We propose the prior-independent *Single Sample* mechanism. This mechanism is essentially the Vickrey-Clarke-Groves (VCG) mechanism, supplemented with reserve prices chosen at random from participants' bids. We show that the expected revenue of the Single Sample mechanism is close to that of the VCG-L mechanism with monopoly reserves. Since the Single Sample mechanism uses random reserves and the VCG-L mechanism uses monopoly reserves, this is essentially a generalization of the argument in Section 3.3.2.

4.1.3 Main Results

We prove that under reasonably general assumptions, the Single Sample mechanism gives prior-independent approximation guarantee. Conceptually, our analysis shows that even a single sample from a distribution — some bidder's valuation — is sufficient information to obtain near-optimal expected revenue.

Our first main result considers *matroid* environments, where bidders satisfy a type of "generalized substitutes" condition (Section 2.1). Examples of such environments include k-unit auctions and certain matching markets. Here, we prove that the Single Sample mechanism gives a prior-independent approximation of a factor of $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ when there are at least $\kappa \geq 2$ bidders of every present attribute. When all bidders have a common attribute and thus have i.i.d. valuations, we improve the approximation factor to $\frac{1}{2}$ for every $\kappa \geq 2$.

Our second main result is that, for every downward-closed environment in which every valuation distribution has a monotone hazard rate (as defined in Section 2.4), the expected revenue of the Single Sample mechanism is at least a constant fraction of the expected optimal welfare (and hence revenue) in that environment. The approximation factor is $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$, and our analysis of our mechanism is tight (for a worst-case distribution) for each κ. This factor is $\frac{1}{8}$ when $\kappa = 2$ and quickly approaches $\frac{1}{4}$ as κ grows. This gives, as an example special case, the first revenue guarantee for combinatorial auctions with single-minded bidders outside of the standard Bayesian setup with known distributions [56, 45].

4.1.4 Other Results

An Alternative Approach In fact, there is an alternative approach to achieve results in this chapter. This approach is based on a subtle form of reduction to Bulow-Klemperer-style theorems for duplicate environments of Hartline and Roughgarden [45]. The resulting mechanisms are complicated with worse approximation ratio, but the approach is general: once we have a new version of Bulow-Klemperer-style theorem for duplicate environments, this reduction gives us a prior-independent mechanism. We present this reduction in Section 4.4.

Many Samples As one can expect, taking more samples helps. We modify the Single Sample mechanism to give better bounds as κ tends to infinity. A weak version of this result, which does not give quantitative bounds on the number of bidders required, can be derived from the Law of Large Numbers. To prove our distribution-independent bound on the number of bidders required, we show that there exists a set of "quantiles" that is simultaneously small enough that concentration bounds can be usefully applied, and rich enough to guarantee a good approximation for every regular valuation distribution.

4.2 Revenue Guarantees with a Single Sample

In this section, we design a mechanism that gives a prior-independent constant-factor approximation to the optimal expected revenue in every non-singular regular matroid environment, and in every non-singular m.h.r. downward-closed environment. Section 4.2.1 defines our mechanism. Section 4.2.2 introduces some of our main analysis techniques in the simpler setting of i.i.d. matroid environments — here, we also obtain better approximation bounds. Section 4.2.3 gives an overview of our general proof approach. Sections 4.2.4 and 4.2.5 prove our approximation guarantees for regular matroid and m.h.r. downward-closed environments, respectively. Section 4.2.6 shows that there is no common generalization of these two results, in that the Single Sample mechanism does not have a constant-factor approximation guarantee in regular downward-closed environments. Section 4.2.7 discusses computationally efficient variants of our mechanism.

4.2.1 The Single Sample Mechanism

Recall that each bidder has a publicly observable attribute that belongs to a known set A. We assume that each bidder with attribute a has a private valuation for winning that is an independent draw from a distribution F_a.

We propose and analyze the *Single Sample* mechanism: we randomly pick one bidder of each attribute to set a reserve price for the other bidders with that attribute, and then run the VCG-L mechanism (Section 2.2.2) on the remaining bidders.

Definition 4.1 (Single Sample)**.** Given a non-singular downward-closed environment, the *Single Sample* mechanism is the following:

(1) For each represented attribute a, pick a *reserve bidder* i_a with attribute a uniformly at random from all such bidders.

(2) Run the VCG mechanism on the sub-environment induced by the non-reserve bidders to obtain a preliminary winning set P.

(3) For each bidder $i \in P$ with attribute a, place i in the final winning set W if and only if $v_i \geq v_{i_a}$. Charge every winner $i \in W$ with attribute a the maximum of its VCG payment computed in step (2) and the reserve price v_{i_a}.

The Single Sample mechanism is clearly prior-independent — that is, it is defined independently of the F_a's — and it is easy to verify that it is truthful. Section 4.3 shows how to use multiple samples to obtain better approximation factors there are more than two bidders with each represented attribute.

4.2.2 Warm-Up: I.I.D. Matroid Environments

To introduce some of our primary analysis techniques in a relatively simple setting, we first consider matroid environments (recall Section 2.1) in which all bidders have the same attribute (i.e., have i.i.d. valuations).

Theorem 4.2 (I.I.D. Matroid Environments)**.** *For every i.i.d. regular matroid environment with at least $n \geq 2$ bidders, the Single Sample mechanism gives a prior-independent $\frac{1}{2} \cdot \frac{n-1}{n}$-approximation to optimal expected revenue.*

The factor of $(n-1)/n$ can be removed with a minor tweak to the mechanism (Remark 4.6).

What's so special about i.i.d. regular matroid environments? Recall that a *monopoly reserve price* of a valuation distribution F is a price in $\mathrm{argmax}_p[p \cdot (1 - F(p))]$.

The following proposition follows immediately from Myerson's Lemma, the fact that the greedy algorithm maximizes welfare in matroid environments, and the fact that the virtual valuation function is order-preserving when valuations are drawn i.i.d. from a regular distribution (see Section 2.1.1).

Proposition 4.3. *In every i.i.d. regular matroid environment, the VCG-E mechanism with monopoly reserves is a revenue-maximizing mechanism.*

The matroid assumption also allows us to pass from eager to lazy reserves.

Corollary 4.4. *In every i.i.d. regular matroid environment, the VCG-L mechanism with monopoly reserves is a revenue-maximizing mechanism.*

Proof. The VCG mechanism can be implemented in a matroid environment via the greedy algorithm: bidders are considered in non-increasing order of valuations, and a bidder is added to the winner set if and only if doing so preserves feasibility, given the previous selections. With a common reserve price r, it makes no difference whether bidders with valuations below r are thrown out before or after running the greedy algorithm. Thus in matroid environments, the VCG-L and VCG-E mechanisms with an anonymous reserve price are equivalent. □

Proving an approximate revenue-maximization guarantee for the Single Sample mechanism thus boils down to understanding the two ways in which it differs from the VCG-L mechanism with monopoly reserves — it throws away a random bidder, and it uses a random reserve rather than a monopoly reserve. The damage from the first difference is easy to control. Applying Lemma 2.15, in expectation over the choice of the reserve bidder, the expected revenue of an optimal mechanism for the environment induced by the non-reserve bidders is at least an $\frac{n-1}{n}$ fraction of the expected revenue of an optimal mechanism for the original environment.

The crux of the proof of Theorem 4.2 is to show that a random reserve price serves as a sufficiently good approximation of a monopoly reserve price. The next key lemma formalizes this goal for the case of a single bidder, which is a generalization of the argument of Theorem 3.3. For a distribution F, define the *revenue function* by $\widehat{R}(p) = p(1 - F(p))$, the expected revenue earned by posting a price of p on a good with a single bidder with valuation drawn from F.

Lemma 4.5. *Let F be a regular distribution with monopoly price r^* and revenue function $\widehat{R}$. Let v denote a random valuation from F. For every nonnegative number $t \geq 0$,*

$$\mathbf{E}_v[\widehat{R}(\max\{t, v\})] \geq \frac{1}{2} \cdot \widehat{R}(\max\{t, r^*\}). \tag{4.2.1}$$

Proof. For simplicity, assume the distribution is over support of $[0, \infty)$. Other cases can be handled in a way similar to the proof of Theorem 3.3. We can rewrite the claim using the revenue curve function. Essentially, for every $q_t \in [0, 1]$, and concave R with $q^* = argmax_q R(q)$, we want to prove that:

$$E_{q \in U[0,1]}[R(\min\{q_t, q\})] \geq \frac{1}{2} \cdot R(\min\{q_t, q^*\}).$$

First suppose that $q_t = 1$, so that the claim is equivalent to the assertion that the expectation of $R(q)$ is at least half of $R(q^*)$, which was proved in Section 3.3.2.

If $q_t \geq q^*$, then the right-hand side is unchanged. The left-hand side can only be higher, as $R(q)$ is decreasing in $[q^*, q]$, and so $q > q_t$ implies that $R(q_t) \geq R(q)$. Finally, if $q_t < q^*$, then the right-hand side is $R(q_t)$; and the left-hand side is a convex combination of $R(q_t)$ (when $q > q_t$) and the expected value of $R(q)$ when q is drawn uniformly from $[0, q_t]$, which by argument similar to 3.3, leveraging concavity, is at least $R(q_t)/2$. □

We prove Theorem 4.2 by extending the approximation bound in Lemma 4.5 from a single bidder to all bidders.

Proof. (of Theorem 4.2) Condition on the choice of the reserve bidder j. Fix a non-reserve bidder i and condition on all valuations except those of i and j. Recall that j, as a reserve bidder, does not participate in the VCG computation in step (2) of the Single Sample mechanism. Thus, there is a "threshold" $t(\mathbf{v}_{-i})$ for bidder i such that i wins if and only if its valuation is at least $t(\mathbf{v}_{-i})$, in which case its payment is $t(\mathbf{v}_{-i})$.

With this conditioning, we can analyze bidder i as in a single-bidder auction, with an extra external reserve price of $t(\mathbf{v}_{-i})$. Let r^* and $\widehat{R}$ denote the monopoly price and revenue function for the underlying regular distribution F, respectively.

The conditional expected revenue that i contributes to the revenue-maximizing solution in the sub-environment of non-reserve bidders is $\widehat{R}(\max\{t(\mathbf{v}_{-i}), r^*\})$. The conditional expected revenue that i contributes to the Single Sample mechanism is $\mathbf{E}_{v_j}[\widehat{R}(\max\{t(\mathbf{v}_{-i}), v_j\})]$. Since v_i, v_j are independent samples from the regular distribution F, Lemma 4.5 implies that the latter conditional expectation is at least 50% of the former. Taking expectations over the previously fixed valuations of bidders other than i and j, summing over the non-reserve bidders i and applying linearity of expectation, and finally taking the expectation over the choice of the reserve bidder j proves the theorem. □

Remark 4.6 (Optimized Version of Theorem 4.2). We can improve the approximation guarantee in Theorem 4.2 from $\frac{1}{2} \cdot \frac{n-1}{n}$ to $\frac{1}{2}$. Instead of discarding the reserve bidder j, we include it in the VCG computation in step (2) of the Single Sample mechanism. An arbitrary other bidder h is used to set a reserve price v_h for the reserve bidder j. Like the other bidders, the reserve bidder is included in the final winning set W if and only if it is chosen by the VCG mechanism in step (2) and also has a valuation above its reserve price ($v_j \geq v_h$). Its payment is then the maximum of its VCG payment and v_h.

The key observation is that, for every choice of a reserve bidder j, a non-reserve bidder i, and valuations $\mathbf{v}$, bidder i wins with bidder j included in the VCG computation in step (2) if and only if it wins with bidder j excluded from the computation. Like Corollary 4.4, this observation can be derived from the fact that the VCG mechanism can be implemented via a greedy algorithm in i.i.d. regular matroid environments. If $v_i \leq v_j$, then i cannot win in either case (it fails to clear the reserve); and if $v_i > v_j$, then the greedy algorithm considers bidder i before j even if the latter is included in the VCG computation.

Thus, the expected revenue from non-reserve bidders is the same in both versions of the Single Sample mechanism. In the modified version, an application of Lemma 4.5 implies that the reserve bidder also contributes, in expectation, a $\frac{1}{2} \cdot \frac{1}{n}$ fraction of the expected revenue of an optimal mechanism. Combining the contributions of the reserve and non-reserve bidders yields an approximation guarantee of $\frac{1}{2}$ for the modified mechanism. This analysis, and hence also the bound in Lemma 4.5, is tight in the worst case — even in a digital goods auction with two bidders, and a left-triangle distribution F for large H.

4.2.3 Proof Framework

Relaxing the matroid or i.i.d. assumptions of Section 4.2.2 introduces new challenges in the analysis of the Single Sample mechanism. The expected revenue-maximizing mechanism becomes complicated — nothing as simple as the VCG mechanism with reserve prices. In addition, eager and lazy reserve prices are not equivalent.

Our proof framework hinges on the VCG-L mechanism with monopoly reserves, which we use as a proxy for the optimal mechanism. The analysis proceeds in two steps: (cf. Section 1.7.1)

1. Prove that the expected revenue of the VCG-L mechanism with monopoly reserves is close to that of an optimal mechanism.

2. Prove that the expected revenue of the Single Sample mechanism is close to that of the VCG-L mechanism with monopoly reserves in the sub-environment induced by the non-reserve bidders.

Given two such approximation guarantees, we can show that the expected revenue of the Single Sample mechanism is a constant fraction of that of the optimal mechanism. Section 4.2.2 implemented this plan for the special case of i.i.d. regular matroid environments, where the VCG-L mechanism with monopoly reserves is optimal.

The arguments in Section 4.2.2 essentially accomplish the second step of the proof framework, with an approximation factor of 2, for all regular downward-closed non-singular environments. The harder part is the first step. The next two sections establish such approximation guarantees under two incomparable sets of assumptions, via two different arguments: regular matroid environments, and m.h.r. downward-closed environments.

For regular matroid environments, we prove that the expected revenue of the VCG-L mechanism with monopoly reserves is at least half that of an optimal mechanism (Theorem 4.7), which in turn implies an approximation guarantee of $\frac{1}{4}\frac{\kappa-1}{\kappa}$ for the Single Sample mechanism (Theorem 4.8).

For m.h.r. downward-closed environments, we prove that the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $1/e$ fraction of the optimal welfare (Theorem 4.9). This implies that the expected revenue of the Single Sample mechanism is at least a $\frac{1}{2e} \cdot \frac{\kappa-1}{\kappa}$ fraction of that of an optimal mechanism when there are at least $\kappa \geq 2$ bidders of every present attribute (Theorem 4.10). Via an optimized analysis, we also prove an approximation factor of $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ (Theorem 4.12). This factor is $\frac{1}{8}$ when $\kappa = 2$ and quickly approaches $\frac{1}{4}$ as κ grows.

4.2.4 Regular Matroid Environments

This section proves an approximation guarantee for the Single Sample mechanism for regular matroid environments. We follow the proof framework outlined in Section 4.2.3, step 1 of which involves proving an approximation bound for the VCG-L mechanism with monopoly reserves.

[45] proved that the expected revenue of the VCG-E mechanism with monopoly reserves (Section 2.2.2) is at least half that of an optimal mechanism in regular matroid environments. The VCG-E and VCG-L mechanisms do not coincide in matroid environments unless all bidders face a common reserve price (cf., Corollary 4.4), and the results of [45] have no obvious implications for the VCG-L mechanism with monopoly reserves in matroid environments with non-i.i.d. bidders. We next supplement the arguments in [45] with some new ideas to prove an approximation guarantee for this mechanism.

Theorem 4.7 (VCG-L With Monopoly Reserves)**.** *For every regular matroid environment, the expected revenue of the VCG-L mechanism with monopoly reserves is at least half of that of an optimal mechanism.*

Proof. Consider a regular matroid environment. For a valuation profile $\mathbf{v}$, let $W(\mathbf{v})$ and $W'(\mathbf{v})$ denote the winning bidders in the VCG-L mechanism with monopoly reserves and in the optimal mechanism, respectively. We claim that

$$\mathbf{E}_{\mathbf{v}}\left[\sum_{i \in W(\mathbf{v}) \setminus W'(\mathbf{v})} \varphi_i(v_i)\right] \geq 0 \tag{4.2.2}$$

and

$$\mathbf{E}_{\mathbf{v}}\left[\sum_{i\in W(\mathbf{v})} p_i(\mathbf{v})\right] \geq \mathbf{E}_{\mathbf{v}}\left[\sum_{i\in W'(\mathbf{v})\setminus W(\mathbf{v})} \varphi_i(v_i)\right]. \tag{4.2.3}$$

Given these two claims, Lemma 3.9 in [45] immediately implies the theorem.[2]

Inequality (4.2.2) holds because the VCG-L mechanism with monopoly reserves allocates to a bidder only if its valuation is above its monopoly reserve. Since a monopoly reserve r_i^* satisfies $\varphi_i(r_i^*) = 0$ and the virtual valuation function is nondecreasing (by regularity), this occurs only when a bidder has a non-negative virtual valuation.

Proving inequality (4.2.3) requires a more involved argument. Fix a valuation profile $\mathbf{v}$ and let $W''(\mathbf{v})$ denote the winning bidders under the VCG mechanism with zero reserve prices. Recall the definition of a matroid environment in terms of the "exchange property" (Section 2.1). This property implies that all maximal feasible sets have equal size and, since $W''(\mathbf{v})$ must be maximal, that we can choose a subset $S \subseteq W''(\mathbf{v})\setminus W'(\mathbf{v})$ such that $S \cup W'(\mathbf{v})$ and $W''(\mathbf{v})$ have the same size.

We next use a non-obvious but well-known property of matroids (see e.g., Schrijver [68, Corollary 39.12a]): given two feasible sets of equal size, such as $W''(\mathbf{v})$ and $S \cup W'(\mathbf{v})$, there is a bijection f from $(S \cup W'(\mathbf{v}))\setminus W''(\mathbf{v})$ to $W''(\mathbf{v})\setminus(S \cup W'(\mathbf{v}))$ such that, for every bidder i in the domain, $W''(\mathbf{v})\setminus\{f(i)\}\cup\{i\}$ is a feasible set. Since $S \subseteq W''(\mathbf{v})$, the domain of f is simply $W'(\mathbf{v})\setminus W''(\mathbf{v})$. Since the VCG mechanism chooses a welfare-maximizing set, the threshold bid (and hence the payment) of a winning bidder $f(i)$ in the range of the function f is at least v_i. Summing over all bidders in $W'(\mathbf{v})\setminus W''(\mathbf{v})$ and using that f is a bijection, the revenue of the VCG

[2]The proof goes as follows. First, using (4.2.2), the virtual welfare of the optimal mechanism from bidders in $W(\mathbf{v}) \cap W'(\mathbf{v})$ is at most that of the virtual welfare of the VCG-L mechanism with monopoly reserves. Second, using (4.2.3), the virtual welfare of the optimal mechanism from bidders in $W'(\mathbf{v})\setminus W(\mathbf{v})$ is at most the revenue of the VCG-L mechanism with monopoly reserves. Finally, applying Myerson's Lemma (Lemma 2.11) completes the proof.

The inequalities (4.2.2) and (4.2.3) almost correspond to the definition of "commensurate" in Hartline and Roughgarden [45, Definition 3.8], but our second inequality is weaker. Nonetheless, the proof of Lemma 3.9 in [45] carries over unchanged.

mechanism is at least

$$\sum_{i \in W'(\mathbf{v}) \setminus W''(\mathbf{v})} v_i \geq \sum_{i \in W'(\mathbf{v}) \setminus W''(\mathbf{v})} \varphi_i(v_i), \tag{4.2.4}$$

where the inequality follows from the definition of a virtual valuation.

Finally, we use in two ways the fact that the allocation rule of the VCG-L mechanism with monopoly reserves differs from that of the VCG mechanism only via the removal of bidders with negative virtual valuations. First, since $W'(\mathbf{v})$ includes no bidders with negative virtual valuations, the right-hand side of (4.2.4) equals

$$\sum_{i \in W'(\mathbf{v}) \setminus W(\mathbf{v})} \varphi_i(v_i),$$

and, by Myerson's Lemma (Lemma 2.11), the expected revenue of the VCG mechanism is at least the expected value of this quantity. Second, again using Lemma 2.11, the expected virtual welfare and hence the expected revenue of the VCG-L mechanism with monopoly reserves is at least that of the VCG mechanism. This completes the proof of inequality (4.2.3) and the theorem. □

Theorem 4.7 establishes step 1 of our main technique. The arguments in Section 4.2.2 now imply that the expected revenue of the Single Sample mechanism is almost half that of the VCG-L mechanism with monopoly reserves (step 2). Precisely, mimicking the proof of Theorem 4.2, gives the following result.

Theorem 4.8 (Single Sample Guarantee)**.** *For every regular matroid environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ fraction of that of an optimal mechanism for the environment.*

4.2.5 M.H.R. Downward-Closed Environments

We now implement the proof framework outlined in Section 4.2.3 for m.h.r. downward-closed environments. We carry out the arguments for expected welfare, rather than expected revenue, because this gives a stronger result. (Such a result is not achievable for regular matroid environments.) We first observe that by applying Lemma 2.15, the expected optimal welfare in the sub-environment induced by non-reserve bidders is at least a $(\kappa - 1)/\kappa$ fraction of that in the original environment.

Analogous to Lemma 4.5, we require a technical lemma about the single-bidder case to establish step 1 of our proof framework.

Lemma. *Let F be an m.h.r. distribution with monopoly price r^* and revenue function $\widehat{R}$. Let $V(t)$ denote the expected welfare of a single-item auction with a posted price of t and a single bidder with valuation drawn from F. For every nonnegative number $t \geq 0$,*

$$\widehat{R}(\max\{t, r^*\}) \geq \frac{1}{e} \cdot V(t). \tag{4.2.5}$$

Proof. Let s denote $\max\{t, r^*\}$. Recall that, by the definition of the hazard rate function, $1 - F(x) = e^{-H(x)}$ for every $x \geq 0$, where $H(x)$ denotes $\int_0^x h(z)dz$. Note that since $h(z)$ is non-negative and nondecreasing, $H(x)$ is nondecreasing and convex. We can write the left-hand side of (4.2.5) as $s \cdot (1 - F(s)) = s \cdot e^{-H(s)}$ and, for a random sample v from F,

$$\begin{aligned} V(t) &= \Pr[v \geq t] \cdot \mathbf{E}[v \mid v \geq t] \\ &= e^{-H(t)} \cdot \left[t + \int_t^\infty e^{-(H(v) - H(t))} dv \right]. \end{aligned} \tag{4.2.6}$$

By convexity of the function H, we can lower bound its value using a first-order approximation at s:

$$H(v) \geq H(s) + H'(s)(v - s) = H(s) + h(s)(v - s) \tag{4.2.7}$$

for every $v \geq 0$. There are now two cases. If $t \leq r^* = s$, then $h(s) = 1/s$ since r^* is a monopoly price.[3] Starting from (4.2.6) and using that H is nondecreasing, and then

[3] One proof of this follows from the first-order condition for the revenue function $p(1 - F(p))$; alternatively, applying Myerson's Lemma to the single-bidder case shows that $r^* = \varphi_F^{-1}(0)$ and hence $r^* - 1/h(r^*) = \varphi_F(r^*) = 0$.

substituting (4.2.7) yields

$$\begin{aligned} V(t) &\leq \int_0^\infty e^{-H(v)} dv \\ &\leq \int_0^\infty e^{-(H(s)+\frac{v-s}{s})} dv \\ &= e \cdot s \cdot e^{-H(s)}. \end{aligned}$$

If $r^* \leq t = s$, then the m.h.r. assumption implies that $h(s) \geq 1/s$ and (4.2.7) implies that $H(v) \geq H(t) + (v-t)/t$ for all $v \geq t$. Substituting into (4.2.6) gives

$$\begin{aligned} V(t) &\leq e^{-H(t)} \cdot \left[t + \int_t^\infty e^{-(H(t)+\frac{v-t}{t}-H(t))} dv \right] \\ &\leq e^{-H(t)} \cdot \int_0^\infty e^{-\frac{v-t}{t}} dv \\ &= e \cdot s \cdot e^{-H(s)}, \end{aligned}$$

where in the second inequality we use that $e^{-(v-t)/t} \geq 1$ for every $v \leq t$. □

Lemma 4.2.5 implies that the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $\frac{1}{e}$ fraction of the expected optimal welfare in every downward-closed environment with m.h.r. valuation distributions.

Theorem 4.9 (VCG-L With Monopoly Reserves)**.** *For every m.h.r. downward-closed environment, the expected revenue of the VCG-L mechanism with monopoly reserves is at least a $\frac{1}{e}$ fraction of the expected efficiency of the VCG mechanism.*

Proof. Fix a bidder i and valuations $\mathbf{v}_{-i}$. This determines a winning threshold t for bidder i under the VCG mechanism (with no reserves). Lemma 4.2.5 implies that the conditional expected revenue obtained from i in the VCG-L mechanism with monopoly reserves is at least a $1/e$ fraction of the conditional expected welfare obtained from i in the VCG mechanism (with no reserves). Taking expectations over $\mathbf{v}_{-i}$ and then summing over all the bidders proves the theorem.

Considering a single bidder with an exponentially distributed valuation shows that the bounds in Lemma 4.2.5 and Theorem 4.9 are tight in the worst case. □

Theorem 4.9 establishes step 1 of our main technique. The arguments in Section 4.2.2 now imply that the expected revenue of the Single Sample mechanism is almost half that of the VCG-L mechanism with monopoly reserves (step 2). Precisely, mimicking the proof of Theorem 4.2, gives the following result.

Theorem 4.10 (Single Sample Guarantee #1)**.** *For every m.h.r. downward-closed environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{2e} \cdot \frac{\kappa-1}{\kappa}$ fraction of the expected optimal welfare in the environment.*

We can improve the guarantee in Theorem 4.10 by optimizing jointly the two single-bidder guarantees in Lemmas 4.2.5 (step 1) and 4.5 (step 2). This is done in the next lemma.

Lemma 4.11. *Let F be an m.h.r. distribution with monopoly price r^* and revenue function $\widehat{R}$, and define $V(t)$ as in Lemma 4.2.5. For every nonnegative number $t \geq 0$,*

$$\mathbf{E}_v[\widehat{R}(\max\{t, v\})] \geq \frac{1}{4} \cdot V(t). \tag{4.2.8}$$

Proof. Define the function H as in the proof of Lemma 4.2.5, and recall from that proof that $V(t)$ can be written as in (4.2.6). We show that the left-hand side of (4.2.8) is at least 25% of that quantity.

Consider two i.i.d. samples v_1, v_2 from F. We interpret v_2 as the random reserve price v in (4.2.8) and v_1 as the valuation of the single bidder. The left-hand side of (4.2.8) is equivalent to the expectation of a random variable that is equal to t if $v_2 \leq t \leq v_1$, which occurs with probability $F(t) \cdot (1 - F(t))$; equal to v_2 if $t \leq v_2 \leq v_1$, which occurs with probability $\frac{1}{2}(1-F(t))^2$; and equal to zero, otherwise. Hence,

$$
\begin{aligned}
\mathbf{E}_v[\widehat{R}(\max\{t,v\})] \geq \frac{1}{2}\,&(F(t)\cdot(1-F(t))\cdot t \\
&+(1-F(t))^2\cdot \mathbf{E}[\min\{v_1,v_2\}\mid \min\{v_1,v_2\}\geq t]\big) \\
&= \frac{1}{2}(1-F(t))\cdot\Big(t\cdot F(t)+(1-F(t))\cdot \\
&\qquad\qquad \Big[t+e^{2H(t)}\int_t^\infty e^{-2H(v)}dv\Big]\Big) \\
&\geq \frac{1}{2}(1-F(t))\cdot\left[t+e^{H(t)}\int_t^\infty e^{-H(2v)}dv\right] && (4.2.9) \\
&= \frac{1}{4}(1-F(t))\cdot\left[2t+\int_{2t}^\infty e^{-(H(v)-H(t))}dv\right] \\
&\geq \frac{1}{4}(1-F(t))\cdot\left[t+\int_{t}^\infty e^{-(H(v)-H(t))}dv\right], && (4.2.10)
\end{aligned}
$$

where in (4.2.9) and (4.2.10) we are using that H is non-negative, nondecreasing, and convex. Comparing (4.2.6) and (4.2.10) proves the lemma. □

We then obtain the following optimized version of Theorem 4.10.

Theorem 4.12 (Single Sample Guarantee #2)**.** *For every m.h.r. downward-closed environment with at least $\kappa \geq 2$ bidders of every present attribute, the expected revenue of the Single Sample mechanism is at least a $\frac{1}{4}\cdot\frac{\kappa-1}{\kappa}$ fraction of the expected optimal welfare in the environment.*

The proof of Theorem 4.12 is the same as that of Theorem 4.2, with the following substitutions: the welfare of the VCG mechanism (with no reserves) plays the previous role of the revenue of the VCG-L mechanism with monopoly reserves; Lemma 4.11 replaces Lemma 4.5.

Remark 4.13 (Theorem 4.12 Is Tight). Our analysis of the Single Sample mechanism is tight for all values of $\kappa \geq 2$, as shown by a digital goods environment with κ bidders with valuations drawn i.i.d. from an exponential distribution with rate 1: the expected optimal welfare is κ, and a calculation shows that the expected revenue of Single Sample is $(\kappa-1)/4$.

Since the revenue of every mechanism is bounded above by its welfare, we have the following corollary.

Corollary 4.14. *For every m.h.r. environment with at least $\kappa \geq 2$ bidders of every present attribute, the Single Sample mechanism gives a prior-independent $\frac{1}{4} \cdot \frac{\kappa-1}{\kappa}$ approximation to the optimal expected revenue.*

4.2.6 Counterexample for Regular Downward-Closed Environments

We now sketch an example showing that a restriction to m.h.r. valuation distributions (as in Section 4.2.5) or to matroid environments (as in Section 4.2.4) is necessary for the VCG-L mechanism with monopoly reserves and the Single Sample mechanism to have constant-factor approximation guarantees. The following example is adapted from Hartline and Roughgarden [45, Example 3.4].

For n sufficiently large, consider two "big" bidders and n "small" bidders $1, 2, \ldots, n$. The feasible subsets are precisely those that do not contain both a big bidder and a small bidder. Fix an arbitrarily large constant H. Each big bidder's valuation is deterministically $\frac{1}{2}n\sqrt{\ln H}$, so the expected revenue of an optimal mechanism is clearly at least $n\sqrt{\ln H}$. The small bidders' valuations are i.i.d. draws from the left-triangle distribution with parameter H: $F(z) = 1 - \frac{1}{z+1}$ on $[0, H)$ and $F(H) = 1$. For n sufficiently large, the sum of the small bidders' valuations is tightly concentrated around $n \ln H$.

We complete the sketch for the VCG-L mechanism with monopoly reserves; the argument for the Single Sample mechanism is almost identical. The VCG mechanism almost surely chooses all small bidders as its preliminary winner set, with a threshold bid of zero for each. The expected revenue extracted from each small winner, via its monopoly reserve H, is at most 1.[4] Thus, the expected revenue of the VCG-L

[4] A subtle point is that each small bidder's valuation is now drawn at random from F, conditioned on the event that the VCG mechanism chose all of the small bidders. But since the small bidders are chosen with overwhelming probability (for large n and H), the probability that a given small bidder is pivotal is vanishingly small, so it still contributes at most 1 to the expected revenue of the mechanism.

mechanism with monopoly reserves is not much more than n, which is arbitrarily smaller than the maximum-possible as $H \to \infty$.

4.2.7 Computationally Efficient Variants

In the second step of the Single Sample mechanism, a different mechanism can be swapped in for the VCG mechanism. One motivation for using a different mechanism is computational efficiency (although this is not a first-order goal in this thesis). For example, for combinatorial auctions with single-minded bidders — where feasible sets of bidders correspond to those desiring mutually disjoint bundles of goods — implementing the VCG mechanism requires the solution of a packing problem that is NP-hard, even to approximate.

For example, the proof of Theorem 4.12 evidently implies the following more general statement: if step (2) of the Single Sample mechanism uses a truthful mechanism guaranteed to produce a solution with at least a $1/c$ fraction of the maximum welfare, then the expected revenue of the corresponding Single Sample mechanism is at least a $\frac{1}{4c}\frac{\kappa-1}{\kappa}$ fraction of the expected optimal welfare (whatever the underlying m.h.r. downward-closed environment). For example, for knapsack auctions — where each bidder has a public size and feasible sets of bidders are those with total size at most a publicly known budget — we can substitute the polynomial-time, $(1+\epsilon)$-approximation algorithm by [14]. For combinatorial auctions with single-minded bidders, we can use the algorithm of [57] to obtain an $O(\sqrt{m})$-approximation in polynomial time, where m is the number of goods. This factor is essentially optimal for polynomial-time approximation, under appropriate computational complexity assumptions [57].

4.3 Revenue Guarantees with Multiple Samples

This section modifies the Single Sample mechanism to achieve improved guarantees via an increased number of samples from the underlying valuation distributions, and provides quantitative and distribution-independent polynomial bounds on the number of samples required to achieve a given approximation factor.

4.3.1 Estimating Monopoly Reserve Prices

Improving the revenue guarantees of Section 4.2 via multiple samples requires thoroughly understanding the following simpler problem: Given an accuracy parameter ϵ and a regular distribution F, how many samples m from F are needed to compute a reserve price r that is $(1-\epsilon)$-*optimal*, meaning that $\widehat{R}(r) \geq (1-\epsilon) \cdot \widehat{R}(r^*)$ for a monopoly reserve price r^* for F? Recall from Section 4.2.2 that $\widehat{R}(p)$ denotes $p \cdot (1-F(p))$. We pursue bounds on m that depend *only* on ϵ and not on the distribution F — such bounds do not follow from the Law of Large Numbers and must make use of the regularity assumption.

Given m samples from F, renamed so that $v_1 \geq v_2 \geq \cdots \geq v_m$, an obvious idea is to use the reserve price that is optimal for the corresponding empirical distribution, which we call the *empirical reserve*:

$$\text{argmax}_{i \geq 1} \; i \cdot v_i. \tag{4.3.1}$$

Interestingly, this naive approach does *not* in general give distribution-independent polynomial sample complexity bounds. Intuitively, with a heavy-tailed distribution F, there is a constant probability that a few large outliers cause the empirical reserve to be overly large, while a small reserve price has much better expected revenue for F.

Our solution is to forbid the largest samples from acting as reserve prices, leading to a quantity we call the *guarded empirical reserve* (with respect to an accuracy parameter ϵ):

$$\text{argmax}_{i \geq \epsilon m} \, i \cdot v_i. \tag{4.3.2}$$

We use the guarded empirical reserve to prove distribution-independent polynomial bounds on the sample complexity needed to estimate the monopoly reserve of a regular distribution.

Lemma 4.15 (Estimating the Monopoly Reserve). *For every regular distribution F and sufficiently small $\epsilon, \delta > 0$, the following statement holds: with probability at least $1 - \delta$, the guarded empirical reserve (4.3.2) of $m \geq c(\epsilon^{-3}(\ln \epsilon^{-1} + \ln \delta^{-1}))$ samples from F is a $(1-\epsilon)$-optimal reserve, where c is a constant that is independent of F.*

Proof. Set $\gamma = \epsilon/11$ and consider m samples $v_1 \geq v_2 \geq \cdots \geq v_m$ from F. Define "q-values" via by $q_t = 1 - F(v_t)$ and $q^* = 1 - F(r^*)$, where r^* is a monopoly price for F. Since the q's are i.i.d. samples from the uniform distribution on $[0, 1]$, the expected value of the quantile q_t is $t/(m+1)$, which we estimate by t/m for simplicity. An obvious approach is to use Chernoff bounds to argue that each q_t is close to this expectation, followed by a union bound. Two issues are: for small t's, the probability that t/m is a very good estimate of q_t is small; and applying the union bound to such a large number of events leads to poor probability bounds. In the following, we restrict attention to a carefully chosen small subset of quantiles, and take advantage of the properties of the revenue functions of regular distributions to get around these issues.

First we choose an integer index sequence $0 = t_0 < t_1 < \cdots < t_L = m$ in the following way. Let $t_0 = 0$ and $t_1 = \lfloor \gamma m \rfloor$. Inductively, if t_i is defined for $i \geq 1$ and $t_i < m$, define t_{i+1} to be the largest integer in $\{1, \ldots, m\}$ such that $t_i < t_{i+1} \leq (1+\gamma)t_i$. If $m = \Omega(\gamma^{-2})$, then $t_i + 1 \leq (1+\gamma)t_i$ for every $t_i \geq \gamma m$ and hence such a t_{i+1} exists. Observe that $L \approx \log_{1+\gamma} \frac{1}{\gamma} = O(\gamma^{-2})$ and $t_{i+1} - t_i \leq \gamma m$ for every $i \in \{0, \ldots, L-1\}$.

We claim that, with probability 1, a sampled quantile q_t with $t \geq \gamma m$ differs from t/m by more than a $(1 \pm 3\gamma)$ factor only if some quantile q_{t_i} with $i \in \{1, 2, \ldots, L\}$ differs from t_i/m by more than a $(1 \pm \gamma)$ factor. For example, suppose that $q_t > \frac{(1+3\gamma)t}{m}$ with $t \geq \gamma m$; the other case is symmetric. Let $i \in \{1, 2, \ldots, L\}$ be such that $t_i \leq t \leq t_{i+1}$. Then

$$q_{t_{i+1}} \geq q_t > \frac{(1+3\gamma)t}{m} \geq \frac{(1+3\gamma)t_i}{m} \geq \frac{(1+3\gamma)t_{i+1}}{(1+\gamma)m} \geq \frac{(1+\gamma)t_{i+1}}{m},$$

as claimed.

We next claim that the probability that q_{t_i} differs from t_i/m by more than a $(1 \pm \gamma)$ factor for some $i \in \{1, 2, \ldots, L\}$ is at most $2Le^{-\gamma^3 m/4}$. Fix $i \in \{1, 2, \ldots, L\}$. Note that $q_{t_i} > (1+\gamma)\frac{t_i}{m}$ only if less than t_i samples have q-values at most $(1+\gamma)\frac{t_i}{m}$. Since the expected number of such samples is $(1+\gamma)t_i$, Chernoff bounds (e.g., [60]) imply that the probability that $q_{t_i} > (1+\gamma)\frac{t_i}{m}$ is at most

$$\exp\{-\gamma^2 t_i/3(1+\gamma)\} \leq \exp\{-\gamma^2 t_i/4\} \leq \exp\{-\gamma^3 m/4\},$$

where the inequalities use that γ is at most a sufficiently small constant and that $t_i \geq \gamma m$ for $i \geq 1$. A similar argument shows that the probability that $q_{t_i} < (1-\gamma)\frac{t_i}{m}$ is at most $\exp\{-\gamma^3 m/4\}$, and a union bound completes the proof of the claim. If $m = \Omega(\gamma^{-3}(\log L + \log\delta^{-1})) = \Omega(\epsilon^{-3}(\log\epsilon^{-1} + \log\delta^{-1}))$, then this probability is at most δ.

Now condition on the event that every quantile q_{t_i} with $i \in \{1, 2, \ldots, L\}$ differs from t_i/m by at most a $(1 \pm \gamma)$ factor, and hence every quantile q_t with $t \geq \gamma m$ differs from t/m by at most a $(1 \pm 3\gamma)$ factor. We next show that there is a candidate for the guarded empirical reserve (4.3.2) which, if chosen, has good expected revenue. Choose $i \in \{0, 1, \ldots, L-1\}$ so that $t_i/m \leq q^* \leq t_{i+1}/m$. Define t^* as t_i if $q^* \geq 1/2$ and t_{i+1} otherwise. Assume for the moment that $q^* \leq 1/2$. By the concavity of revenue function in probability space $R(q)$ — recall Section 4.2.2 — $R(q_{t_{i+1}})$ lies above the line segment between $R(q^*)$ and $R(1)$. Since $R(1) = 0$, this translates to

$$R(q_{t^*}) \geq R(q^*) \cdot \frac{1 - q_{t_{i+1}}}{1 - q^*} \geq R(q^*) \cdot \frac{1 - (1+3\gamma)\left(\frac{t_i}{m} + \gamma\right)}{1 - \frac{t_i}{m}} \geq (1 - 5\gamma) \cdot R(q^*),$$

where in the final inequality we use that $\frac{t_i}{m} \leq \frac{1}{2}$ and γ is sufficiently small. For the case when $q^* \geq \frac{1}{2}$, a symmetric argument (using $R(0)$ instead of $R(1)$ and q_{t_i} instead of $q_{t_{i+1}}$) proves that $R(q_{t^*}) \geq (1 - 5\gamma) \cdot R(q^*)$.

Finally, we show that the guarded empirical reserve also has good expected revenue. Let the maximum in (4.3.2) correspond to the index $\hat{t}$. Since $\hat{t}$ was chosen over t^*, we have $\hat{t} \cdot v_{\hat{t}} \geq t^* \cdot v_{t^*}$. Using that each of $q_{\hat{t}}, q_{t^*}$ is approximated up to a $(1 \pm 3\gamma)$ factor by $\hat{t}/m, t^*/m$ yields

$$R(q_t) = q_{\hat{t}} v_{\hat{t}} \geq \frac{(1-3\gamma)\hat{t}}{m} v_{\hat{t}} \geq \frac{(1-3\gamma)t^*}{m} v_{t^*} \geq \frac{1-3\gamma}{1+3\gamma} q_{t^*} v_{t^*} = \frac{1-3\gamma}{1+3\gamma} R(q_{t^*})$$

and hence

$$R(q_t) \geq \frac{(1-5\gamma)(1-3\gamma)}{1+3\gamma} R(q^*) \geq (1-11\gamma)R(q^*).$$

Since $\gamma = \epsilon/11$, the proof is complete. □

Remark 4.16 (Optimization for M.H.R.Distributions). There is a simpler and stronger version of Lemma 4.15 for m.h.r. distributions. We use a simple fact, first noted in Hartline et al. [47, Lemma 4.1], that the selling probability q^* at the monopoly reserve r^* for an m.h.r. distribution is at least $1/e$. Because of this, we can take the parameter t_1 in the proof of Lemma 4.15 to be $\lfloor m/e \rfloor$ instead of $\lfloor \gamma m \rfloor$ without affecting the rest of the proof. This saves a γ factor in the exponent of the bound on the probability that some q_{t_i} is not well approximated by t_i/m, which translates to a new sample complexity bound of $m \geq c(\epsilon^{-2}(\ln \epsilon^{-1} + \ln \delta^{-1}))$, where c is some constant that is independent of the underlying distribution. Also, this bound remains valid even for the empirical reserve (4.3.1) — the guarded version in (4.3.2) is not necessary.

4.3.2 The Many Samples Mechanism

In the following *Many Samples mechanism*, we assume that an accuracy parameter ϵ is given, and use m to denote the sample complexity bound of Lemma 4.15 (for regular valuation distributions) or of Remark 4.16 (for m.h.r. distributions) corresponding to the accuracy parameter $\frac{\epsilon}{3}$ and failure probability $\frac{\epsilon}{3}$. The mechanism is only defined if every present attribute is shared by more than m bidders.

(1) For each represented attribute a, pick a subset S_a of m *reserve bidders* with attribute a uniformly at random from all such bidders.

(2) Run the VCG mechanism on the sub-environment induced by the non-reserve bidders to obtain a preliminary winning set P.

(3) For each bidder $i \in P$ with attribute a, place i in the final winning set W if and only if v_i is at least the guarded empirical reserve r_a of the samples in S_a.

Charge every winner $i \in W$ with attribute a the maximum of its VCG payment computed in step (2) and the reserve price r_a.

We prove the following guarantees for this mechanism.

Theorem 4.17 (Guarantees for Many Samples)**.** *The expected revenue of the Many Samples mechanism is at least:*

(a) a $(1-\epsilon)$ fraction of that of an optimal mechanism in every i.i.d. regular matroid environment with at least $n \geq 3m/\epsilon = \Theta(\epsilon^{-4}\log \epsilon^{-1})$ bidders;

(b) a $\frac{1}{2}(1-\epsilon)$ fraction of that of an optimal mechanism in every regular matroid environment with at least $n \geq 3m/\epsilon = \Theta(\epsilon^{-4}\log \epsilon^{-1})$ bidders;

(c) a $\frac{1}{e}(1-\epsilon)$ fraction of the optimal expected welfare in every downward-closed m.h.r. environment with at least $\kappa \geq 3m/\epsilon = \Theta(\epsilon^{-3}\log \epsilon^{-1})$ bidders of every present attribute.

Bidders with i.i.d. and exponentially distributed valuations show that part (c) of the theorem is asymptotically optimal (as is part (a), obviously).

Proof. The lower bound $\kappa \geq 3m/\epsilon$ on the number of bidders of each attribute implies that at most an $\epsilon/3$ fraction of all bidders are designated as reserve bidders. Lemma 2.15 implies that the expectation, over the choice of reserve bidders, of the expected revenue of an optimal mechanism for and the expected optimal welfare of the sub-environment induced by the non-reserve bidders are at least a $(1-\frac{\epsilon}{3})$ fraction of those in the full environment.

Now condition on the choice of reserve bidders, but not on their valuations. Fix a non-reserve bidder i, and condition on the valuations of all other non-reserve bidders. Let t denote the induced threshold bid for i and r^* a monopoly price for the valuation distribution F of i. The conditional expected revenue obtained from i using the price $\max\{r^*, t\}$ is precisely that obtained by the VCG-L mechanism with monopoly reserves for the sub-environment induced by the non-reserve bidders.

The Many Samples mechanism, on the other hand, uses the price $\max\{r, t\}$, where r is the guarded empirical reserve of the reserve bidders that share i's attribute.

By Lemma 4.15 and our choice of m, r is $(1 - \frac{\epsilon}{3})$-optimal for F with probability at least $1 - \frac{\epsilon}{3}$. Concavity of the revenue function (cf., Figure 3.3.1) and an easy case analysis shows that, whenever r is $(1 - \frac{\epsilon}{3})$-optimal, the conditional expected revenue from i with the price $\max\{r, t\}$ is at least a $(1 - \frac{\epsilon}{3})$ fraction of that with the price $\max\{r^*, t\}$, for every value of t. Since valuations are independent of each other and the choice of the reserve bidders, the expected revenue from i in the Many Samples mechanism, conditioned on the choice of reserve bidders and on the valuations of the other non-reserve bidders, is at least a $(1 - \frac{\epsilon}{3})^2 \geq (1 - \frac{2}{3}\epsilon)$ fraction of that of the VCG-L mechanism with monopoly reserves. Removing the conditioning on the valuations of other non-reserve bidders; summing over the non-reserve bidders; and removing the conditioning on the choice of reserve bidders shows that the expected revenue of the Many Samples mechanism is at least a $(1 - \frac{2}{3}\epsilon)$ fraction of that of the VCG-L mechanism with monopoly reserves on the sub-environment induced by the non-reserve bidders. The three parts of the theorem now follow from Corollary 4.4, Theorem 4.7, and Theorem 4.9, respectively.

□

Remark 4.18 (Case Study: Digital Goods Auctions). Our results in this section have interesting implications even in the special case of digital goods auctions. We note that there is no interference between different bidders in such an auction, so the general case of multiple attributes reduces to the single-attribute i.i.d. case (each attribute can be treated separately).

The Deterministic Optimal Price (DOP) digital goods auction offers each bidder i a take-it-or-leave offer equal to the empirical reserve of the other $n - 1$ bidders. The expected revenue of the DOP auction converges to that of an optimal auction as the number n of bidders goes to infinity, provided valuations are i.i.d. samples from a distribution with bounded support [36] or from a regular distribution [69]. However, the number of samples required in these works to achieve a given degree of approximation depends on the underlying distribution F.

As an alternative, consider the variant of DOP that instead uses the guarded empirical reserve (4.3.2) of the other $n - 1$ bidders to formulate a take-it-or-leave-it offer for each bidder. Our Lemma 4.15 implies a *distribution-independent* bound for this auction: provided the number of bidders is $\Omega(\epsilon^{-3}\log\epsilon^{-1})$, its expected revenue is at least a $(1-\epsilon)$ times that of the optimal auction.

4.4 Prior-Independent Mechanisms Via Reductions

This section describes a subtle form of reduction to Bulow-Klemperer-style theorems for duplicate environments of Hartline and Roughgarden [45]. The resulting pairing mechanisms are more complicated than the single sample mechanism, and the approximation ratio is worse. However, the approach is very general: once we have a new version of Bulow-Klemperer-style theorem for duplicate environments, this reduction directly gives us a prior-independent approximation mechanism, which could be a good starting point for searching for simpler mechanisms with better guarantees.

Single-Dimensional Duplicate Environment

As discussed in Section 5.3, prior-independent approximation is tightly related to Bulow-Klemperer-style theorems. For single-dimensional downward-closed environments, a Bulow-Klemperer-style theorem was proved in Hartline-Roughgarden [45], which is based on duplicating bidders.

Formally, given a single-dimensional downward-closed environment, the duplicate environment is constructed as follows. For every original bidder i with distribution F_i, introduce an extra bidder i' into the system, called the duplicate of i, whose value is independently drawn from F_i. A bidder set S in the duplicate environment is feasible if (1) at most one from each bidder and its duplicate is in S, and (2) the set we obtain from replacing every duplicate bidder in S by its original bidder is feasible in the original environment.

Hartline and Roughgarden proved the following:

Theorem. *[45] Given a downward-closed environment, the expected revenue of the VCG mechanism over the duplicate environment is at least ρ fraction of that of the optimal mechanism over the original environment, where:*

- $\rho = \frac{1}{2}$ *for regular matroid environments*

- $\rho = \frac{1}{3}$ *for m.h.r. downward-closed environments.*

This theorem is also prior-independent in nature, except that VCG is run over a different set of bidders than the original ones. However, there does exist a subtle form of reduction, showing that such a theorem implies a prior-independent approximation, at the cost of additional constant loss in ratio.

A Reduction Based on A Pairing Mechanism

Definition 4.19 (Pairing Mechanism)**.** Given a non-singular downward-closed environment, the pairing mechanism does the following:

1. For every attribute a, group the bidders with attribute a into pairs arbitrarily. The bidders that are left out due to odd parity are removed from the system. Assign one bidder of each pair into E_1, and the other into E_2.

2. Choose p from 1,2 uniformly at random, and call E_p be the main bidders. (E_{3-p} will never win)

3. Construct the duplicate environment of E_p by drawing the paired bidders from E_{3-p} as the duplicates.

4. Run VCG over the duplicate environment to obtain a preliminary winning set W.

5. Let $W \cap E_p$ be the final winning set.

Note that the pairing mechanism is a truthful mechanism, because a bidder is either in E_{3-p} and never wins, or is part of E_p, participating in a VCG mechanism.

Theorem 4.20. *Given a downward-closed environment , if the expected revenue of VCG over its duplicate environment is at least ρ fraction of that of an optimal mechanism over the original environment, then the pairing mechanism gives a prior-independent $\frac{\rho}{6}$-approximation to optimal expected revenue.*

Proof. The expected revenue of the pairing mechanism is exactly half of the expected revenue of the VCG mechanism over the duplicate environment it constructed in the process, which by assumption is at least ρ fraction of the expected revenue of an optimal mechanism of the sub-environment induced by E_p. Removing the bidders E_{3-p} in worse case loses another factor of $\frac{1}{3}$ (when there are three bidders in each group). All in all we achieve a prior-independent $\frac{\rho}{6}$-approximation. $\square$

Chapter 5

Supply-Limiting Mechanisms

In this chapter, we propose mechanisms based on the welfare-maximization with supply-limiting, and prove that they give prior-independent approximations. This chapter is mainly based on [66].

5.1 Introduction

5.1.1 A Matching Problem

Consider the problem of matching agents to a set of non-identical items for sale, with the goal of maximizing the seller's revenue. For example, a travel website selling hotel accommodation would like to match agents to a set of hotel rooms. Each agent is unit-demand in the sense that she is interested in one item. For example each agent only needs one hotel room. Each agent has a different private value for each type of room. Agents' values are drawn from prior distributions, with one distribution per item (one distribution for a suite at the Ritz, another for a room at Best Western, and so on). The seller wishes to maximize her expected revenue. A far-seeing seller might also want to (approximately) maximize social welfare as well.

Maximizing expected revenue in the matching problem above is difficult even when the value distributions are known. The difficulty stems from the problem's *multi-dimensional* nature. The theory of optimal auction design stops short of solving

Algorithm 5.1 A Generic Prior-Independent Mechanism

1. Set a supply limit equal to half of the number of bidders.
2. Run the VCG mechanism subject to this supply limit.

settings in which the description of an agent's preferences requires multiple parameters. Recent breakthrough work of Chawla et al. [19] circumvents this limitation by introducing *approximately* optimal mechanisms. These mechanisms make use of a priori knowledge of the value distributions, and are somewhat complicated.

5.1.2 A Supply-Limiting Mechanism

Our mechanisms are extremely simple, and are based on the natural idea of artificially *limiting the supply* to increase bidder competition for the items.

Previous prior-independent mechanisms are based largely on some form of random sampling to estimate the prior distributions. Our mechanisms in this chapter are the first known prior-independent mechanisms for nontrivial multi-dimensional environments.[1] In addition, our mechanism also guarantees approximately-optimal social welfare, even though our weak distributional assumptions allow the revenue and welfare of other mechanism to be very far from each other.

This paper shows that under minimal regularity assumptions, the simple, prior-independent mechanism above and its revenue guarantee generalize to significantly more complex settings. In other words, we identify settings in which the prior-independent VCG mechanism with limited supply is guaranteed to have approximately-optimal expected revenue. For the matching problem discussed above, we prove the following theorem.

Theorem 5.1 (Prior-Independent Mechanism for Matching (Informal))**.** *For every matching environment, the supply-limiting mechanism in Algorithm 5.1 gives a prior-independent constant factor approximation to optimal expected revenue.*

[1] Simultaneously and independently, another group obtained similar results using different mechanisms; see Section 5.1.5 for a detailed discussion.

The constant fractions we achieve are quite good in many cases, e.g., we achieve a fraction of $\frac{1}{4}$ when the number of bidders exceeds the number of items.

5.1.3 Technical Approach: Reduction to Proving Bulow-Klemperer-Style Theorems

Our technical approach to establishing approximation properties of supply-limiting mechanisms is based on the general reduction to proving Bulow-Klemperer-style theorems. We describe this general reduction in Section 5.3.

To instantiate the reduction in various single-dimensional environments, we can use generalizations of the original Bulow-Klemperer result to matroid environments [26] and to non-i.i.d. bidders [45]. For the matching problem, we need to prove the first generalization of the original Bulow-Klemperer theorem to a nontrivial multi-dimensional environment.

Theorem 5.2 (Bulow-Klemperer-Style Theorem for Matching (Informal))**.** *For every matching environment with n bidders and m items, the expected revenue of the VCG mechanism with either (1) m additional bidders or (2) $O(n)$ additional bidders and a supply limit n, is at least a constant fraction of the optimal expected revenue in the original environment.*

5.1.4 Our Results and Organization

Our main result is a collection of approximately-optimal supply-limiting mechanisms for different auction environments with i.i.d. regular bidders, as detailed in Table 5.1.1. For single-dimensional environments, we show supply-limiting mechanisms for multi-unit, parallel multi-unit, and matroid environments. For multi-dimensional environments, we show three versions of a supply-limiting mechanism for matching environments, each with a different approximation guarantee. The choice among these versions should be according to the parameters of the environment at hand. Our results generalize to multi-unit matching environments as well.

Environment		Supply Limit / Augmentation	Ratios to OPT
k-Unit Auctions	SL	$\frac{n}{2}$	$\frac{1}{2}$ (Thm 5.5)
	BK	add k bidders	1 (Thm 2.14)
Matroid	SL	$\lfloor \frac{r}{2} \rfloor$	$\frac{1}{4}$ (Thm 5.8)
	BK	add a disjoint base	1 (Thm 5.9)
Parallel Multi-Unit	SL	$\frac{n}{2}$ globally, $\frac{k_j}{2}$ for item j	$\frac{1}{4}$ (Thm 5.10)
	BK	add k_j bidders for item j	$\frac{1}{2}$ (Thm 5.11)
Matching with $m \leq \frac{n}{2}$	SL	$\lfloor \frac{n}{2} \rfloor$	$\frac{1}{2}$ (Thm 5.14)
	BK	add m bidders	1 (Thm 5.17)
Matching with $m \geq \frac{n}{2}$	SL	$\frac{n}{2}$	$\frac{n}{4m}$ (Thm 5.15)
	BK	add n bidders	$\frac{n}{2m}$ (Thm 5.18)
	SL	$\frac{n}{3}$	$\frac{1}{27}$ (Thm 5.16)
	BK	add $O(n)$ bidders	$\frac{1}{9}$ (Thm 5.19)
Multi-Unit Matching	SL	$\frac{n}{3}$ globally, $\frac{k_j}{2}$ for item j	$\frac{1}{27}$ (Thm 5.35)
	BK	add $2n$ bidders	$\frac{1}{9}$ (Thm 5.36)

Here n is the number of bidders, assumed to be multiple of 2 or 3 if needed, m is the total number of units of items, r is the rank of the matroid, k_j is the number of units available for item j. SL=Supply Limiting. BK=Bulow-Klemperer-Style Theorem.

Table 5.1.1: Summary of Main Results

For simplicity of notation we will assume that $n/2$ (or $n/3$ where appropriate) is an integer if needed. If this is not the case, consider removing a minimum amount of bidders before running the mechanism so that this is true. Lemma 5.4 implies that we suffer a vanishingly small loss in factors.

5.1.5 Related Work

Most related to our results are the following. Dughmi et al. [26] investigate conditions under which VCG yields near-optimal revenue, and use a generalized Bulow-Klemperer result to show this is guaranteed in matroid environments with sufficient competition in the form of disjoint bases. Our reduction encompasses and generalizes this result. Hartline and Roughgarden [45]also study conditions for when simple Vickrey-based mechanisms achieve near-optimal revenue and in particular they derive an anonymous-reserve mechanism from one of their Bulow-Klemperer-style results [45, Theorem 5.1]. This mechanism however is not prior-independent and is inherently limited to single items [45, Example 5.4]. Chawla et al. [19] use posted-price mechanisms that rely on prior distributions — i.e., mechanisms that are *not* prior-independent — to achieve a $\frac{1}{6.75}$-approximation for the matching setting with multiple units and non-i.i.d. bidders, and also a $\frac{3}{32}$-approximation for a more general environment than what we consider, namely a graphical matroid with unit-demand bidders.

In terms of techniques, our limited-supply mechanisms are special cases of maximal-in-range mechanisms (see, e.g., [63]), which are well-known to be truthful. We apply a reduction of Chawla et al. that relates single- and multi-dimensional environments ([19], see Section 5.2). Some of our techniques are inspired by Chawla et al. [20]'s analysis of the performance of a VCG mechanism in a job scheduling context.

Finally, the paper of Devanur et al. [23] is very closely related to this work. Both considered essentially the same set of problems and gave similar results, though using different mechanisms. The mechanisms in [23] are arguably quite complicated, certainly more complex and less natural than the supply-limiting mechanisms studied here. On the other hand, the mechanisms in [23] seem to be a bit easier to analyze than supply-limiting mechanisms, and they also achieve better constant factors in the matching problem for the case when m is large. Both analyses share some common preliminary steps, but at their core are quite different, reflecting the different mechanisms studied. Finally, in the matching problems studied in the original version of the present work, each item was assumed to have unit supply; we were inspired by [23] to pursue the more general multi-unit results presented here.

5.2 Preliminaries

Items v.s. Units We distinguish between *items* — which are of different kinds, and *units* — which are different copies of the same item. While bidders have the same value for different units of the same item, their values for different items are independent from one another (although possibly drawn from the same distribution). We use the following notation: item j means the j-th kind of item sold in the auction,

k_j denotes the number of units available of item j, and $m = \sum_j k_j$ is the total number of units available of all items.

Multi-Dimensional Matching Environments

A (single-unit) *matching* environment is a multi-dimensional environment with n bidders and m different items for sale. We only have one unit of each item available, and a multi-unit version of matching environment will be studied in Section 5.10. Bidders are unit-demand, in the sense that each bidder can only win at most one item. Feasible allocations are all matchings of items to bidders, such that each bidder wins at most one item and each item is assigned to at most one bidder. We can also impose an additional supply limit of $\ell \leq m$, restricting the matching to have size at most ℓ. Bidder i has a private value $v_{i,j}$ for winning item j, which is drawn independently at random from a distribution $F_{i,j}$. We say the bidders are i.i.d. if $F_{i,j}$ does not depend on i, which we can simply denote by F_j. Our supply-limiting results apply to i.i.d. matching environments in which the bidders are i.i.d.

In this chapter, for matching problems, we assume that the distribution has a positive smooth density function all over the support. This allows us to say that the maximum weighted matching is unique with probability 1, so that tie-breaking issues can be ignored.

Representative Environments for Upper-Bounding Optimal Multi-Dimensional Revenue

Characterizing the optimal mechanisms (in the ex post IC sense) for multi-dimensional matching environments is currently an open question. Chawla et al. [19] introduced the concept of *representative environments* in order to upper-bound the optimal expected revenue of a mechanism in i.i.d. matching environments.

Given a matching environment Env with n i.i.d. bidders, m items, and value distributions $\{F_j\}_j$, the representative environment Env^{rep} has nm single-dimensional bidders, one for every pair of original bidder and item (i, j). The m bidders in Env^{rep} that corresponds to original bidder i are called i's representatives. Like bidder i's value for item j in Env, representative (i, j)'s value $v_{i,j}$ for winning in Env^{rep} is drawn independently at random from F_j. Note that every subset S of representatives in Env^{rep} corresponds to a (not necessarily feasible) allocation in Env — if representative (i, j) is in S then item j is allocated to bidder i in Env. Feasible allocations in Env^{rep} are subsets of representatives such that the corresponding allocation in Env is feasible. In particular, since every bidder i in Env is unit-demand, only one of its representatives in Env^{rep} can win at a time.

Given a truthful mechanism $\mathcal{M}$ for Env, its allocation rule can be used to construct a truthful mechanism $\mathcal{M}^{\text{rep}}$ for Env^{rep}. The following lemma of Chawla et al. [19] relates the expected revenue of the two mechanisms. Intuitively, Env^{rep} involves more competition than Env since representatives of the same bidder compete with one another, and so the expected revenue of $\mathcal{M}^{\text{rep}}$ is higher.

Lemma 5.3. *[19] For every matching environment, and a mechanism $\mathcal{M}$, the expected revenue of $\mathcal{M}^{\text{rep}}$ for Env^{rep} is at least the expected revenue of $\mathcal{M}$ for Env.*

5.3 Reduction to Bulow-Klemperer-Style Statements

5.3.1 General Reduction to Bulow-Klemperer-Style Statements

As suggested in Section 3.3.5, for digital goods auctions with two bidders, prior-independent approximation guarantee of a supply-limiting mechanism can follow from the Bulow-Klemperer theorem. In fact, there is a much more general reduction from proving a prior-independent approximation guarantee to proving a Bulow-Klemperer-style statement.

Essentially, we can re-interpret a supply-limiting mechanism as doing the following two steps.

1. **Restriction:** Restrict the auction environment by dropping bidders along with supply for them.

Guarantee: Optimal expected revenue of the restricted environment approximates that of the original environment.

Proof: By a general fractional subadditivity property of optimal expected revenue in the bidder set. (Lemma 5.4)

2. **Augmentation+VCG:** Augment the restricted environment by adding bidders without changing the supply.

   Guarantee: Expected revenue of VCG over augmented environment approximates optimal expected revenue of original environment.

   Proof: By a suitable Bulow-Klemperer-style theorem.

A technical requirement is that if we first apply the restriction operation, and then apply the expansion operation, the resulting environment is a sub-environment of the original environment. This is important as our prior-independent mechanism is essentially to run the VCG mechanism over this sub-environment.

5.3.2 Examples of Reductions

Let n be even. For digital goods auctions with n bidders, the restriction operation corresponds to removing half of the bidders (with total supply changed to $\frac{n}{2}$), and the expansion operation corresponds to adding back half of the bidders, but still maintaining the supply limit of $\frac{n}{2}$. The corresponding Bulow-Klemperer-style statement then says that instead of running the optimal mechanism for digital goods auction, we are better off first doubling the number of bidders, and then running the VCG mechanism. This is effectively the generalized Bulow-Klemperer theorem for n-unit auctions (Theorem 2.14). It follows that the VCG mechanism with half of the supply gives a prior-independent $\frac{1}{2}$-approximation.

In Chapter 5, we study various reductions where restriction corresponds to supply-limiting. In Section 4.4, we also describe a subtle form of reduction for single-dimensional environments where bidders are not i.i.d.

5.3.3 Fractional Subadditivity of Optimal Revenue

Given an auction environment, for a bidder set S, let $\mathrm{OPT}(S)$ be the optimal expected revenue of a mechanism over the sub-environment induced by S. The following property of $\mathrm{OPT}(S)$ is crucial in allowing our reduction to work in both single-dimensional and multi-dimensional environments.

Lemma 5.4 (Fractional Subadditivity of Optimal Expected Revenue)**.** *For every (single-dimensional or multi-dimensional) auction environment,* $\mathrm{OPT}(\cdot)$ *is fractionally subadditive.*

Proof. By revenue monotonicity, it is sufficient to prove the claim for the case that $T_j \subseteq S$ for all j.

For every j, the optimal mechanism $\mathrm{OPT}(S)$ induces a randomized mechanism $\mathcal{M}_j$ for set T_j. In particular, this mechanism randomly draws values for bidders from $S \backslash T_j$ according to their distributions, and simulates $\mathrm{OPT}(S)$ on T_j, cancelling outcome for fake bidders of $S \backslash T_j$. The expected revenue of $\mathcal{M}_j$ (also denoted by $\mathcal{M}_j$) is bidder-wise the same as the expected revenue of $\mathrm{OPT}(S)$ for all bidders in T_j. By the fractional covering assumption, $\sum_j \alpha_j \mathcal{M}_j \geq \mathrm{OPT}(S)$. Our lemma follows from the fact that $\mathrm{OPT}(T_j) \geq \mathcal{M}_j$ for every j. □

5.4 Supply-Limiting for I.I.D. k-Unit Auctions

In this section we formally prove the following theorem, in order to illustrate our general approach in a simple single-dimensional environment. To simplify notation we assume that the number of bidders n is even. This assumption is essentially without loss of generality since if n is odd, one can first remove an arbitrary bidder from the environment, losing at most a $1/n$-fraction of the optimal expected revenue by Lemma 2.15. Let $\mathrm{VCG}^{\leq n/2}$ be the mechanism that allocate to maximize welfare, subject to the supply constraint that at most $\frac{n}{2}$ bidders can win.

Theorem 5.5 (Supply-Limiting Mechanism for I.I.D. k-Unit Auctions)**.** *For every k-unit auction with $n \geq 2$ i.i.d. regular bidders, the supply-limiting mechanism* $\mathrm{VCG}^{\leq n/2}$

gives a prior-independent $\max\{\frac{1}{2}, 1 - \frac{k}{n}\}$*-approximation to the optimal expected revenue.*

The proof of Theorem 5.5 using the reduction requires Theorem 2.14, the statement of which we repeat in the following.

Theorem 5.6 (Generalized Bulow-Klemperer Theorem for Multi-Unit Auctions)**.** *For every k-unit environment with i.i.d. regular bidders, the expected revenue of VCG with k additional bidders is at least as high as the optimal expected revenue.*

Now we instantiate our general reduction to use Theorem 2.14 to prove Theorem 5.5.

Proof. (of Theorem 5.5) We reinterpret the supply-limiting mechanism as follows:

1. **Restriction**: Remove $\min\{\frac{n}{2}, k\}$ bidders from the environment, along with the relevant supply.
   - This results in a $\min\{\frac{n}{2}, k\}$-unit auction over $n - \min\{\frac{n}{2}, k\} = \max\{\frac{n}{2}, n-k\}$ bidders.
2. **Augmentation**: Add back $\min\{\frac{n}{2}, k\}$ bidders, maintaining the supply.
   - This results in a $\min\{\frac{n}{2}, k\}$-unit auction over n bidders.
3. **VCG**: Run the VCG mechanism.

By Lemma 2.15, the optimal expected revenue of the restricted environment is at least $\max\{\frac{1}{2}, 1 - \frac{k}{n}\}$ fraction of the original environment. By applying the generalized Bulow-Klemperer theorem (Theorem 2.14) to the restricted environment, the expected revenue of VCG over the augmented environment is at least as high as the optimal expected revenue for the restricted environment. Together, it follows that the supply-limiting mechanism gives a prior-independent $\max\{\frac{1}{2}, 1 - \frac{k}{n}\}$ approximation. □

The approximation factor in the above theorem is asymptotically tight. In fact, no supply limiting mechanism based on limiting the supply by a fixed ratio can achieve a better approximation factor.

Proposition 5.7 (Asymptotic Tightness)**.** *For every $0 \leq \rho \leq 1$, there exists an n-unit environment with n i.i.d. bidders whose values are drawn from a regular distribution F such that $\mathrm{VCG}^{\leq \rho n}$ gives in expectation at most $(\frac{1}{2} + o(1))$-fraction of the optimal expected revenue.*

Proof. First suppose that $\rho < 1/2$. Let the distribution F be uniform over $[1, 1+\epsilon]$ for a negligible positive ϵ. The optimal expected revenue is then roughly n, while VCG with supply limit ρn can extract at most $\rho n(1+\epsilon) < n/2$ in expectation.

Now suppose $\rho > 1/2$. Let the distribution be the left-triangle distribution with parameter H (Example 2.7) for sufficiently large H. I.e., $F(z) = \frac{z}{1+z}$ for $[0, H)$ and $F(H) = 1$. The optimal mechanism can offer price H to every bidder, achieving an expected revenue of $H(1 - \frac{H}{1+H}) = \frac{H}{1+H} \approx 1$ from every bidder. In VCG with supply limit ρn, with very high probability the $(\rho n+1)$-st highest bid is concentrated around $z = \frac{1-\rho}{\rho}$, the value of z such that $F(z) = 1-\rho$, and so we achieve an expected revenue of roughly $\frac{1-\rho}{\rho} \cdot \rho n = (1-\rho)n \leq \frac{n}{2}$. One can also show that all but negligible amount of expected revenue comes from this case. □

5.4.1 A Supply-Limiting Mechanism for I.I.D. Matroid Environments

A *matroid environment* is a single-dimensional environment in which the set system $(N = \{1, \ldots, n\}, \mathcal{I})$ of bidders and feasible allocations forms a matroid (see Section 2.1.1) . Recall that the *rank* of a matroid is the size of its bases, and the *packing number* of a matroid is its maximum number of disjoint bases.

Theorem 5.8 (Supply-Limiting Mechanism for I.I.D. Matroids)**.** *For every matroid environment with $n \geq 2$ i.i.d. regular bidders, rank r and packing number κ:*

1. *If $\kappa \geq 2$ then the VCG mechanism gives a prior-independent $\frac{1}{2}$-approximation to optimal expected revenue.*

2. *If* $\kappa = 1$ *then the supply-limiting mechanism* $\mathrm{VCG}^{\leq \lfloor r/2 \rfloor}$ *gives a prior-independent* $\frac{1}{4}$*-approximation to optimal expected revenue.*

The proof is by instantiating the general reduction, where the restriction and augmentation consist roughly of removing and adding back a suitable basis of bidders, and the following Bulow-Klemperer-style theorem for i.i.d. matroid environments from [26].

Theorem 5.9 (B-K for I.I.D. Matroid Environments)**.** *[26] For every matroid environment with i.i.d. regular bidders, the expected revenue of VCG with an additional matroid basis of bidders is at least as high as the optimal expected revenue.*

Now we instantiate our general reduction to use Theorem 5.9 to prove Theorem 5.8.

Proof. (of Theorem 5.8) First note that if the matroid's packing number κ is 1 then its rank r is at least 2. We reinterpret the supply-limiting mechanism as follows:

1. **Restriction**: If $\kappa = 1$, first intersect the original matroid with a $\lfloor \frac{r}{2} \rfloor$-uniform matroid to get a new matroid environment with packing number $\kappa' \geq 2$.
   Now that the packing number is at least 2 we can remove a basis of bidders from the environment, where the basis is chosen randomly from a set of disjoint bases.

2. **Augmentation**: Add back a basis of bidders (without changing back the supply limit).

3. **VCG**: Run the VCG mechanism.

To analyze this mechanism we first upper-bound the loss due to restriction. If $\kappa = 1$, the first step of the restriction incurs a loss of factor $r/\lfloor \frac{r}{2} \rfloor$, and the second step incurs a loss of factor $n/(n - \lfloor \frac{r}{2} \rfloor)$. Since by assumption n is even, the total worst-case loss is 4. If $\kappa \geq 2$, since at most half of the bidders are removed, the loss factor is at most 2.

The expansion step is justified by applying Theorem 5.9 to the restricted environment. □

5.5 Supply-Limiting Mechanism for Parallel Multi-Unit Auctions

It seems restrictive to assume that in a k-unit auction, bidders' valuation distributions are identical. One could consider a setting where every bidder has a publicly-observable attribute j, say age bracket, which determines her (regular) distribution F_j. In other words, the values are i.i.d. for bidders with the same attribute, and are independent but not necessarily identically distributed for bidders with different attributes. We aim to derive supply-limiting mechanism for such an attribute-based setting. It turns out that it is more convenient to work with the slightly more general class of parallel multi-unit auction environments.

5.5.1 Parallel Multi-Unit Auctions

There is a set J of non-identical items. A single-dimensional *parallel multi-unit* environment consists for each item $j \in J$ a k_j-unit auction over n_j bidders with $n_j \geq k_j$. Bidders for item j have values drawn i.i.d. from a regular distribution F_j. Let $n = \sum_j n_j$ be the total number of bidders, and $m = \sum_j k_j$ be the total number of units of items available. These auctions are related by a global supply limit ℓ with $l \leq n$ In other words, a feasible allocation is a set of bidders containing at most ℓ bidders overall (a global supply limit) and at most k_j bidders for item j (a local supply limit). We say that a parallel multi-unit auction is non-singular if $n_j \geq 2$ for all j.

Note that the concept of item in parallel multi-unit auctions corresponds to the concept of attribute in the above-mentioned attribute-based setting. The difference between the two settings is that there is no supply limit for bidders of an attribute in the attribute-based setting.

Now consider the following supply-limiting mechanism. For every item j, impose a local supply limit of $\min\{k_j, \lfloor \frac{n_j}{2} \rfloor\}$ on the maximum number of winners of j, and run VCG. The following states the guarantee for the case that n_j's are even; the case of odd n_j can be handled at a small loss.

Theorem 5.10 (Supply-Limiting for Parallel Multi-Unit Auctions)**.** *For every non-singular parallel multi-unit auction with even n_j's, the supply-limiting VCG mechanism gives a prior-independent $\frac{1}{4}$-approximation to the optimal expected revenue.*

To prove this theorem, we need to first identify the right version of Bulow-Klemperer-style statement to work with.

5.5.2 A Bulow-Klemperer-Style Theorem

A parallel multi-unit auction is a particular case of a matroid environment, with the underlying matroid being the intersection of a partition matroid with an ℓ-uniform matroid. As such, the Bulow-Klemperer-style theorem for non-i.i.d. matroid environments by Hartline and Roughgarden [45] applies. However, this theorem requires augmenting the environment with an additional "duplicate" bidder for every original bidder, and adding the strong constraint that at most one of each pair of original and duplicate can win at a time. The following theorem shows that it is sufficient to augment the environment with only k_j additional bidders per item j, without imposing strong constraint on feasibility. We defer the proof of this theorem to Section 5.5.3.

Theorem 5.11 (Bulow-Klemperer-Style Theorem for Parallel Multi-Unit Auctions)**.** *For every parallel multi-unit auction with k_j units for item j, the expected revenue of VCG with k_j additional bidders per item j is a $\frac{1}{2}$-approximation to the optimal expected revenue of the original environment.*

Now we use this theorem to prove Theorem 5.10.

Proof. (of Theorem 5.10) We reinterpret the supply-limiting mechanism as follows.

1. **Restriction**: For every item j, remove $\min\{k_j, \frac{n_j}{2}\}$ bidders for item j from the environment.

   - This results in a parallel multi-unit auction with $\min\{k_j, \frac{n_j}{2}\}$ units of item j for $\max\{\frac{n_j}{2}, n_j - k_j\}$ bidders.

2. **Augmentation**: For every item j, add back $\min\{k_j, \frac{n_j}{2}\}$ bidders with item j, but without changing the number of units of item j.

3. **VCG**: Run the VCG mechanism.

In the restriction step, we remove at most half of the bidders. So by Lemma 2.15, the optimal expected revenue of the restricted environment is at least half of that of the original environment. Applying Theorem 5.11 to the restricted environment, the expected revenue of the VCG mechanism is at least half of the optimal expected revenue of the restricted environment. □

We remark that this supply-limiting mechanism is considerably simpler than Myerson's optimal mechanism for this setting, which requires computing different virtual value functions for different distributions.

5.5.3 Proof of the Bulow-Klemperer-Style Theorem for Parallel Multi-Unit Auctions

To prove Theorem 5.11, we need to show that the expected revenue of VCG in the parallel multi-unit auction after augmentation is a $\frac{1}{2}$-approximation to the optimal expected revenue in the original environment. As shown by Hartline and Roughgarden (Lemma 4.5 of [45]), a $\frac{1}{2}$-approximation would follow if we prove the following *commensuration conditions*:

- (C1) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in \mathrm{VCG}(\mathbf{v}) \setminus \mathrm{OPT}(\mathbf{v})} \phi_i] \geq 0$, and
- (C2) $\mathbb{E}_{\mathbf{v}}[\sum_{i \in \mathrm{VCG}(\mathbf{v}) \setminus \mathrm{OPT}(\mathbf{v})} p_i(\mathbf{v})] \geq \mathbb{E}_{\mathbf{v}}[\sum_{i \in \mathrm{OPT}(\mathbf{v}) \setminus \mathrm{VCG}(\mathbf{v})} \phi_i]$

where ϕ_i is the virtual value of bidder i, $\mathbf{v}$ denotes a valuation profile of both original and augmenting bidders, and $\mathrm{OPT}(\mathbf{v}), \mathrm{VCG}(\mathbf{v})$ denote the winning bidders chosen by the optimal mechanism in the original environment and VCG in the augmented environment, respectively.

The proof of (C2) in Hartline and Roughgarden [45] can be directly applied to our setting. However, proving (C1) in our setting turns out to be more technically

---

Algorithm 5.2 An Auxiliary Allocation Procedure

Maximize welfare over all bidders and all units, with the additional constraint that all bidders in O_j must be also chosen.

---

challenging. In particular, the random sets VCG($\mathbf{v}$) and OPT($\mathbf{v}$) are dependent in more complicated ways. To handle such dependency issues, we introduce an auxiliary allocation procedure as an intermediary for the analysis, and make careful use of the FKG inequality along the way (see [5]). The proof relies on the fact that in our setting, the VCG mechanism (as well as the auxiliary procedure) is a simple greedy process, which takes the top bidders up to the local supply limit for each item, and then sorts this set and take the best ones up to the global supply limit.

In the remainder of this section we prove (C1). This is sufficient to complete the proof of the theorem.

Fix an item j. We let B_j contain the n_j original bidders for item j as well as the k_j augmented bidders for item j. We aim to prove a stronger version of (C1) where i ranges only over B_j. We condition the rest of the analysis on fixed values of all bidders (including augmenting bidders) for all items except j, and fixed values of the original bidders for item j. Now only the values of the augmenting bidders for item j, which we denote by $\mathbf{v}'$, are still random.

As the optimal mechanism does not rely on the augmenting bidders, the winning set of OPT is now fully determined, which we denote by OPT as well. Let O_j denote those bidders from B_j that win in OPT. To analyze the VCG mechanism, we consider the auxiliary procedure in Algorithm 5.2.

Let $T(\mathbf{v}')$ be the bidders for item j that win in this procedure. (So $O_j \subseteq T(\mathbf{v}') \subseteq B_j$.) Let $\text{VCG}_j(\mathbf{v}')$ be the bidders for item j that win in VCG. The following claim relates the auxiliary procedure to VCG.

Claim 5.12. For all values $\mathbf{v}'$ of augmenting bidders for item j, $\text{VCG}_j(\mathbf{v}') \backslash T(\mathbf{v}')$ always have nonnegative virtual values.

Proof. Compare the auxiliary procedure to VCG. In VCG, we have the additional freedom of replacing bidders in O_j by others. By a property of the greedy process of VCG, each replacement will be either from B_j with even higher values and therefore higher virtual values, or from bidders for other items. As all bidders in O_j already had nonnegative virtual values, the replacements that come from B_j (i.e., bidders in $\mathrm{VCG}_j(\mathbf{v}')\backslash T(\mathbf{v}')$) will also have nonnegative virtual values. □

The following claim relates the auxiliary procedure to OPT.

Claim 5.13. $T(\mathbf{v}')\backslash O_j$ have nonnegative total virtual value in expectation over all values $\mathbf{v}'$ of augmented bidders for item j.

Proof. Let $\psi_i(\mathbf{v}')$ for $i = 1, \ldots, k_j$ be the i-th highest virtual value (or value) of a bidder in $B_j\backslash O_j$. Let 0-1 variable $1_i(\mathbf{v}')$ indicate whether the i-th highest bidder in $B_j\backslash O_j$ wins in the auxiliary procedure, and let $p_i = \mathbb{E}_{\mathbf{v}'}[1_i(\mathbf{v}')]$ be the corresponding probability. It follows that the total virtual value of bidders in $T(\mathbf{v}')\backslash O_j$ is $\sum_{i=1}^{k_j} \mathbb{E}_{\mathbf{v}'}[\psi_i(\mathbf{v}') \cdot 1_i(\mathbf{v}')]$.

By the greedy nature of the auxiliary procedure, one can verify that if the i-th highest bidder in $B_j\backslash O_j$ wins in the auxiliary procedure, and some bidders in $B_j\backslash O_j$ increase their values, then the i-th highest bidder in $B_j\backslash O_j$ still wins. Formally, for two valuation profiles $\mathbf{v}'$ and $\mathbf{v}''$ of the augmenting bidders for item j with $v'_i \leq v''_i$ for all i, we have $1_i(\mathbf{v}') \leq 1_i(\mathbf{v}'')$. This positive correlation allows us to apply the FKG inequality [5], and we have:

$$\begin{aligned}
\sum_{i=1}^{k_j} \mathbb{E}_{\mathbf{v}'}[\psi_i(\mathbf{v}') \cdot 1_i(\mathbf{v}')] &\geq \sum_{i=1}^{k_j} \mathbb{E}_{\mathbf{v}'}[\psi_i(\mathbf{v}')] \cdot \mathbb{E}_{\mathbf{v}'}[1_i(\mathbf{v})] \\
&= \sum_{i=1}^{k_j} \mathbb{E}_{\mathbf{v}'}[\psi_i(\mathbf{v}')] \cdot p_i \\
&= \sum_{i=1}^{k_j} \left(\mathbb{E}_{\mathbf{v}'}[\sum_{i'=1}^{i} \psi_{i'}(\mathbf{v}')] \cdot (p_i - p_{i+1}) \right)
\end{aligned}$$

where we let $p_{k+1} = 0$.

It is easy to see that p_i is decreasing in i. Therefore it suffices to prove that $\sum_{i'=1}^{i} \mathbb{E}_{\mathbf{v}'}[\psi_{i'}(\mathbf{v}')] \geq 0$ for all i. Fixing i, this is the total virtual value from the top i bidders from $B_j\backslash O_j$. The expected total virtual value of the first i augmented bidders (sorted by identity) for item j exactly equals 0. It follows that $\psi_1, \ldots, \psi_i$ as the top i virtual values can only have a nonnegative expected total sum. $\square$

Now by Claims 5.12 and 5.13, $\mathbb{E}_{\mathbf{v}'}[\sum_{i \in \mathrm{VCG}_j(\mathbf{v}') \setminus O_j} \phi_i] \geq 0$. Summing over all $\mathbf{v}'$, all values of bidders for other items, and all items j, we have $\mathbb{E}_{\mathbf{v}}[\sum_{i \in \mathrm{VCG}(\mathbf{v}) \setminus \mathrm{OPT}(\mathbf{v})} \phi_i] \geq 0$, which verifies condition (C1).

5.6 I.I.D. Matching Environments: Overview of Results

5.6.1 Supply-Limiting Mechanisms

In this section we present our main result of this chapter— a supply-limiting mechanism for i.i.d. matching environments. More precisely, we present three alternative supply-limiting mechanisms, all VCG-based, with different approximation factors depending on the number of bidders n and number of items m of the i.i.d. matching environment. The relation between the number of bidders n and total number of items m in the environment at hand determines which supply-limiting mechanism is most suitable.

We denote the revenue from the optimal mechanism for n bidders by $\mathrm{OPT}(n)$, and the revenue from the supply-limiting VCG mechanism for n bidders by $\mathrm{VCG}^{\leq \ell}(n)$, sometimes omitting ℓ from the notation when l is not binding, i.e., $\ell \geq \min\{n, m\}$. Note that $\mathrm{OPT}(n)$ and $\mathrm{VCG}^{\leq \ell}(n)$ are random variables over the sample space of bidder valuation profiles $\mathbf{v}$. All expectations below are over $\mathbf{v}$.

Again we only state theorems for n that is a multiple of 2 (or 3 if appropriate) to keep the statements clean. The case of general n can be handled at a small loss.

Theorem 5.14. *For every matching environment with $n \geq 2$ i.i.d. regular bidders and $m \leq n/2$ items, where n is even, $\mathbb{E}[\mathrm{VCG}(n)] \geq (1 - \frac{m}{n}) \cdot \mathbb{E}[\mathrm{OPT}(n)]$. Here $1 - \frac{m}{n} \geq \frac{1}{2}$.*

Theorem 5.15. *For every matching environment with $n \geq 2$ i.i.d. regular bidders and $m \geq n/2$ items, where n is even, $\mathbb{E}[\mathrm{VCG}^{\leq n/2}(n)] \geq \frac{n}{4m}\mathbb{E}[\mathrm{OPT}(n)]$.*

Theorem 5.16. *For every matching environment with $n \geq 3$ i.i.d. regular bidders and m items, where n is a multiple of 3,* $\mathbb{E}[\mathrm{VCG}^{\leq n/3}(n)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$.

Intuitively, achieving good approximation guarantees becomes more difficult as the number of items grows relative to the number of bidders, since the natural competition among the bidders in the environment is dispersed across different items. Accordingly, when number of items is less than half of the number of bidders, we show that simply applying VCG achieves a $\frac{1}{2}$-approximation to the optimal expected revenue (Theorem 5.14). When the number of items is more than half of the number of bidders but still proportional to it, applying VCG while artificially limiting the supply to half of the number of bidders achieves a $\frac{n}{4m}$-approximation, in particular a $\frac{1}{4}$-approximation when $m = n$ (Theorem 5.15). Finally, when the number of items is possibly much larger than the number of bidders, limiting the supply still achieves a constant-factor approximation but with a larger constant. We find that setting the supply limit to a third of the number of bidders guarantees a $\frac{1}{27}$-approximation (Theorem 5.16). We believe this approximation factor can be further improved, and leave this as an open problem.

We reduce proving the above theorems to proving appropriate Bulow-Klemperer-style theorems. In Section 5.6.2 we state these theorems and show how the main theorems are implied by them via the general reduction. The proofs of the Bulow-Klemperer-style theorems themselves are the main technical contribution, and appear in Sections 5.7, 5.8, and 5.9, respectively.

Finally, in Section 5.10, we extend the main results to a multi-unit matching environment, where every item can have multiple units.

5.6.2 Reduction to Bulow-Klemperer-Style Theorems

Our proof approach is based on the general reduction of proving prior-independent approximations to proving Bulow-Klemperer-style theorems. In particular, we reduce proving Theorems 5.14, 5.15, and 5.16 to the following corresponding Bulow-Klemperer-style theorems respectively. We defer the proof of these Bulow-Klemperer-style theorems to later sections.

Theorem 5.17. *For every matching environment with n i.i.d. regular bidders and m items,* $\mathbb{E}[\text{VCG}(n+m)] \geq \mathbb{E}[\text{OPT}(n)]$.

Now we show that this theorem implies Theorem 5.14.

Proof. (of Theorem 5.14) We need to show $\mathbb{E}[\text{VCG}(n)] \geq (1 - \frac{m}{n}) \cdot \mathbb{E}[\text{OPT}(n)]$ when $m \leq n/2$. We reinterpret the VCG mechanism as follows.

1. **Restriction**: Remove m bidders from the environment.
   - This results in a matching environment with $n - m$ bidders and m items.
2. **Augmentation**: Add back m bidders.
   - This results in a matching environment with n bidders and m items.
3. **VCG**: Run the VCG mechanism.

By monotonicity and Lemma 2.15, restricting the environment does not hurt the optimal expected revenue too much, i.e., $\mathbb{E}[\text{OPT}(n-m)] \geq (1 - \frac{m}{n}) \cdot \mathbb{E}[\text{OPT}(n)]$. Applying Theorem 5.17 to the restricted environment with $n - m$ bidders and m items, gives $\mathbb{E}[\text{VCG}(n)] \geq \mathbb{E}[\text{OPT}(n-m)]$ completes the proof. □

Theorem 5.18. *For every matching environment with n i.i.d. regular bidders and $m \geq n$ items,* $\mathbb{E}[\text{VCG}^{\leq n}(2n)] \geq \frac{n}{m}\mathbb{E}[\text{OPT}(n)]$.

Now we show that this theorem implies Theorem 5.15.

Proof. (of Theorem 5.15) We need to show $\mathbb{E}[\text{VCG}^{\leq n/2}(n)] \geq \frac{n}{4m}\mathbb{E}[\text{OPT}(n)]$ when $m \geq n/2$. We reinterpret the supply-limiting mechanism as follows.

1. **Restriction**: Remove $n/2$ bidders from the environment.
   - This results in a matching environment with $n/2$ bidders and m items.
2. **Augmentation**: Add back $n/2$ bidders without changing supply. Now at most $\frac{n}{2}$ bidders can win.

- This results in a matching environment with n bidders, m items, and a supply limit of $\frac{n}{2}$.

3. **VCG**: Run the VCG mechanism.

As above, the proof is by the inequality chain $\mathbb{E}[\mathrm{VCG}^{\leq n/2}(n)] \geq \frac{n}{2m}\mathbb{E}[\mathrm{OPT}(n/2)] \geq \frac{n}{4m}\mathbb{E}[\mathrm{OPT}(n)]$, where the second inequality is by Lemma 2.15, and the first inequality is by applying Theorem 5.18 to the restricted environment. □

Theorem 5.19. *For every matching environment with n i.i.d. regular bidders and m items,* $\mathbb{E}[\mathrm{VCG}^{\leq n}(3n)] \geq \frac{1}{9}\mathbb{E}[\mathrm{OPT}(n)]$.

Now we show that this theorem implies Theorem 5.16.

Proof. (of Theorem 5.16) We need to show $\mathbb{E}[\mathrm{VCG}^{\leq n/3}(n)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$. We reinterpret the supply-limiting mechanisms as follows.

1. **Restriction**: Remove $\frac{2}{3}n$ bidders from the environment.

   - This results in a matching environment with $\frac{1}{3}n$ bidders and m items.

2. **Augmentation**: Add back $\frac{2}{3}n$ bidders. Now at most $\frac{n}{3}$ bidders can win. (This constraint can be not binding if m is small.)

   - This results in a matching environment with n biders, m items, and a supply limit of $\frac{n}{3}$.

3. **VCG**: Run the VCG mechanism.

As above, the proof is by the inequality chain $\mathbb{E}[\mathrm{VCG}^{\leq n/3}(n)] \geq \frac{1}{9}\mathbb{E}[\mathrm{OPT}(n/3)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$, where the second inequality is by Lemma 2.15, and the first inequality is by applying Theorem 5.19 to the restricted environment. □

The first of the above Bulow-Klemperer-style theorems states that for matching environments with m items, the expected revenue of VCG with m additional bidders is at least as high as the optimal expected revenue. This generalizes the original

Bulow-Klemperer theorem to the more complex multi-dimensional matching setting. If $m \gg n$ however, the required resource augmentation — adding m bidders when originally there are only n — is substantial, which will cause our reduction to give weak bounds.

Our second and third Bulow-Klemperer-style theorems address this issue by adding only $O(n)$ bidders, with an additional supply limit of n. Theorem 5.18 provides a good approximation factor when m is a small multiple of n. Theorem 5.19 guarantees a $\frac{1}{9}$-approximation for any values of n, m.

5.7 BK-Style Theorem for Matching with m More Bidders

In this section we prove Theorem 5.17, i.e., $\mathbb{E}[\text{VCG}(n+m)] \geq \mathbb{E}[\text{OPT}(n)]$. In Section 5.7.1 we identify an upper bound on the optimal expected revenue in the original environment, and a lower bound on the revenue of the VCG mechanism in the augmented environment. The advantage of this step is that these bounds are relatively simple to analyze and are already similar in form, though not identical. In Section 5.7.2 we carefully relate the two bounds, thus establishing Theorem 5.17.

5.7.1 Basic Upper and Lower Bounds

Let $\text{Vic}_j(n+1)$ be the revenue from selling item j to $n+1$ i.i.d. bidders with value-distribution F_j using the Vickrey (second-price) auction. We use the concept of representative environment to show that the optimal expected revenue from selling all items to n bidders in an i.i.d. matching environment is upper-bounded by the expected revenue from selling each item to a separate group of $n+1$ single-dimensional bidders.

Claim 5.20 (Upper Bound on Optimal Expected Revenue). For every matching environment with n i.i.d. regular bidders, $\mathbb{E}[\text{OPT}(n)] \leq \sum_j \mathbb{E}[\text{Vic}_j(n+1)]$.

Proof. We leverage the notion of representative environments of Chawla et al. [19] (see Section 5.2) to prove the claim.

Given the matching environment, consider the corresponding complete bipartite graph with bidders on one side and items on the other, and the bidders' values for items drawn from distributions $\{F_j\}_j$ as edge weights; recall that feasible allocations correspond to matchings. By Lemma 5.3, the optimal expected revenue in the matching environment is upper-bounded by the optimal expected revenue in its single-dimensional counterpart, the corresponding representative environment.

We now relax the feasibility constraints, by which we only increase the optimal expected revenue. We define a new environment in which feasible allocations are all subsets of edges such that at most one edge is incident to an item-node (but unlike a matching, multiple edges can be incident to a bidder-node). Observe that the new environment is equivalent in terms of revenue to a collection of m single-item auctions, where in the j-th auction item j is auctioned to n single-dimensional bidders whose values are drawn i.i.d. from the regular distribution F_j. By the original Bulow-Klemperer theorem (Theorem 2.14), the optimal expected revenue from the j-th auction is upper-bounded by $\mathbb{E}[\text{Vic}_j(n+1)]$. Summing up over all items completes the proof. □

The revenue from the VCG mechanism is the sum of VCG payments for allocated items. For an allocated item j, we lower-bound the VCG payment for the winner of item j. (In the current case of $m \leq n$, every item is allocated. But this is not true if $m > n$ which is needed in next sections.)

Claim 5.21 (Lower Bound on VCG Revenue). For every matching environment, the VCG payment for an allocated item j is at least the value of any unallocated bidder for j.

Proof. If bidder i wins item j, then the VCG payment for item j is equal to the externality that i imposes on the rest of the bidders by winning j. Since i prevents any unallocated bidder from getting item j, the payment is at least the unallocated bidder's value for j. □

Algorithm 5.3 Selling Item j by Deferred Allocation

Given a matching environment with current supply limit l: ($l = \min\{m, n\}$ if no explicit supply limit was set)

1. Find a welfare-maximizing feasible allocation (a maximum weighted matching) of all items other than item j with a supply limit of $l - 1$ to a subset of the bidders.
   Let U be the set of $n + 1$ bidders who are not allocated in this allocation.
2. Run the Vickrey auction to sell item j to bidder set U.

In our matching context, the upper and lower bounds above turn out to share a similar form. On one hand, by definition of the Vickrey auction, the upper bound $\mathbb{E}[\text{Vic}_j(n+1)]$ on the expected revenue from separately auctioning item j is equal to the second-highest value for item j among $n+1$ bidders with values for item j drawn independently from F_j. On the other hand, the lower bound on the VCG payment for item j in the augmented environment is equal to the highest value for item j among n unallocated bidders with values drawn independently from F_j. We are using here the fact that since the augmented environment includes m more bidders, all items are allocated and exactly n out of $n + m$ bidders are unallocated.

From this it may appear as if we have already shown that the lower bound exceeds the upper bound. However, a dependency issue arises — conditioned on the event that a bidder in the augmented environment is unallocated by VCG, her value for item j is no longer a random sample from F_j. We address this issue in the next subsection by introducing a *deferred allocation* selling procedure.

5.7.2 Relating the Upper and Lower Bounds via Deferred Allocation

Algorithm 5.3 describes a deferred allocation procedure for selling item j.

In Figure 5.7.1, we depict an example of deferred allocation along with Vickrey auction and VCG.

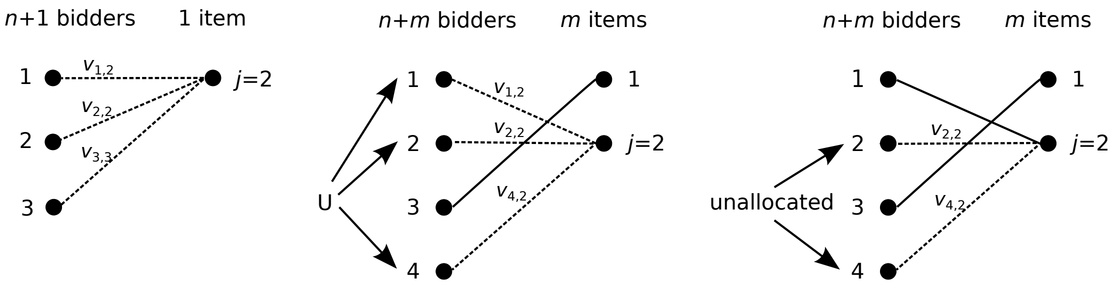

(a) Vickrey for Item j (b) Deferred Allocation for Item j (c) VCG

Figure 5.7.1: Relating Bounds via Deferred Allocation with $n = m = j = 2$

We now show how deferred allocation helps us get around this dependency issue. Consider the revenue from selling item j to bidder set U by the deferred allocation procedure described in Algorithm 5.3. We use this revenue to relate the upper and lower bounds found in the previous section.

Claim 5.22 (Relating to Upper Bound)*.* The revenue from selling item j by the deferred allocation procedure for item j is equal in expectation to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$.

Proof. Observe that the revenue from selling item j to bidder set U by the Vickrey auction is the second-highest value of a bidder in U for j. Since we exclude item j in step (1) of the deferred allocation procedure and allocate it only in step (2), the allocation in step (1) does not depend on the bidders' values for j. Therefore, the values of the unallocated bidders in U for item j are still independent random samples from F_j. The expected second-highest among $n+1$ values drawn independently from F_j is equal to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$. □

To relate this to the lower bound in Claim 5.21, we need the following stability property.

Claim 5.23 (Stability)*.* For every value profile of the augmented matching environment, the set of bidders left unallocated by VCG is U less at most one bidder.

Proof. Given the augmented matching environment, compare VCG and the deferred allocation procedure, both calculate maximum bipartite weighted matchings between bidders and items. The difference is that the sizes of the two matchings differ by at most one, and the item set differs by at most one node, item j. The symmetric difference of the two matchings is either an empty set, or a single alternating path[2]. It follows that the bidder sets matched in the two matchings differ by at most one node too, and our claim follows. □

Using this claim we can lower-bound the VCG payment for item j in the augmented environment.

Claim 5.24 (Relating to Lower Bound). For every value profile of the augmented matching environment, the VCG payment for item j is at least the revenue from selling item j by deferred allocation.

Proof. The revenue from selling item j by deferred allocation is the second-highest value of a bidder in U for j. Let i_1, i_2 be the two bidders in U who value item j the most. By definition, these bidders are left unallocated by the deferred allocation procedure, and by the previous claim, one of them (say i_1) is also unallocated by the VCG mechanism. Recall that an unallocated bidder's value for item j gives a lower bound on the VCG payment for j (Claim 5.21). So the VCG payment for j is at least $v_{i_1,j}$, which in turn is at least the second-highest value of a bidder in U for item j. □

Putting everything together, we can now complete the proof of the Bulow-Klemperer-style theorem.

Proof. (of Theorem 5.17) We need to show that $\mathbb{E}[\mathrm{VCG}(n+m)] \geq \mathbb{E}[\mathrm{OPT}(n)]$. By Claim 5.24, the VCG payment for item j in the augmented environment is at least the revenue from selling item j by deferred allocation, which by Claim 5.22 is equal in expectation to $\mathbb{E}[\mathrm{Vic}_j(n+1)]$. Summing up over all items, the total expected VCG revenue in the augmented environment is at least $\sum_j \mathbb{E}[\mathrm{Vic}_j(n+1)]$, and by Claim 5.20 this upper-bounds $\mathbb{E}[\mathrm{OPT}(n)]$ □

[2] We assume there is a unique maximum-weighted matching. This holds with probability 1 as all distributions have smooth density functions.

5.8 Weak BK-Style Theorem for Matching with n More Bidders

In this section we prove Theorem 5.18, i.e., $\mathbb{E}[\mathrm{VCG}^{\leq n}(2n)] \geq \frac{n}{m}\mathbb{E}[\mathrm{OPT}(n)]$. The proof is similar to that of Theorem 5.17. Here we highlight the necessary changes.

The upper bound on the optimal expected revenue remains $\sum_j \mathbb{E}[\mathrm{Vic}_j(n+1)]$ (Claim 5.20). This is not a very strong upper bound. A stronger upper bound will be used in Section 5.9 to achieve stronger bounds. As for the lower bound, it is no longer the case that in the augmented environment all items are allocated, and so we make use of a generalization of Claim 5.21 — the VCG payment for item j is lower-bounded not only by the value of any unallocated bidder for j itself, but also by the value of any unallocated bidder for any unallocated item. We call the highest of the latter among all unallocated bidders and items the *global* lower bound on VCG payments, and denote it by G. Note that since VCG is now applied with a supply limit of n, exactly n out of the $2n$ bidders in the augmented environment are unallocated.

We use the same deferred allocation selling procedure in Algorithm 5.3, and Claims 5.22, 5.23, and 5.24 still hold. For Theorem 5.17, every item is allocated, and these claims allow us to do a per-item charging argument, i.e., to charge the upper bound for each item j to the lower bound for item j, for every item j. However, for Theorem 5.18, we need an additional charging argument, since only n out of m items are allocated by VCG.

1. If item j is allocated by VCG, then as above the VCG payment for it is at least the revenue from selling j by deferred allocation.

2. If item j is not allocated by VCG, then the VCG payment for any allocated item $j' \neq j$ is at least the global lower bound G, and so is at least the revenue from selling j by deferred allocation (cf. Claim 5.24).

By the above, we can charge our upper bound $\mathbb{E}[\mathrm{Vic}_j(n+1)]$, which equals to the total expected revenue from selling all m items by deferred allocation, to the VCG payments for the n allocated items, thus obtaining an approximation factor of $\frac{n}{m}$.

5.9 BK-Style Theorem for Matching with $O(n)$ More Bidders

In this section we prove Theorem 5.19, i.e., $\mathbb{E}[\mathrm{VCG}^{\leq n}(3n)] \geq \frac{1}{9}\mathbb{E}[\mathrm{OPT}(n)]$.

5.9.1 Notations and Definitions

We first present several additional notations and definitions needed for the proof.

Environments Let Env be a matching environment with n i.i.d. regular bidders and m items, and let Env' be the augmented matching environment with $2n$ additional bidders and supply limit n. Let $\mathrm{Env}^{\mathrm{rep}}$ be the representative environment corresponding to Env. Environment $\widehat{\mathrm{Env}}$ is obtained from $\mathrm{Env}^{\mathrm{rep}}$ by adding one additional bidder per item, and replacing the unit-demand constraint by the more relaxed constraint of limiting the supply to $2n$ items that can be allocated in total.[3] So $\widehat{\mathrm{Env}}$ is a parallel multi-unit auction where there are $n+1$ bidders for one copy of each item and global supply limit $2n$.

Valuation profiles $V, \mathbf{v}$ Let $V = \{V^1, \ldots, V^m\}$ be a collection of m sets, each containing $n+1$ random values. The values in set V^j are i.i.d. samples from F_j. The collection V corresponds to a valuation profile in environment $\widehat{\mathrm{Env}}$ up to naming of the bidders, where V^j contains the values of the bidders interested in item j. Let $\mathbf{v}$ be a vector of $3nm$ random values. Values $v_{1,j}, \ldots, v_{3n,j}$ are i.i.d. samples from F_j. The vector $\mathbf{v}$ corresponds to a valuation profile in environment Env', where $v_{i,j}$ is the value of bidder i for item j.

Random variables over V Over the sample space of V we define the following. For every j, random variables H_j, S_j are the highest and second-highest values for item j. Random variable N is the $(2n+1)$-highest among $\{H_1, \ldots, H_m\}$ if $m > 2n$

[3] Note there is deliberate slackness in this relaxation — to obtain a matroid environment it would have been sufficient to replace the unit-demand constraint by a supply limit of n; by further relaxing the supply limit to $2n$ we aid later analysis.

and 0 otherwise. Random set A contains every item j such that H_j is within the $\min\{2n, m\}$ highest among $\{H_1, \ldots, H_m\}$. Let $a_1, \ldots, a_{|A|}$ denote the items in A ordered by their H_j value from high to low. Observe that if VCG runs on $\widehat{\text{Env}}$ with valuation profile V, the set of $\min\{2n, m\}$ allocated items is equal to A.

Random variables over v Over the sample space of $\mathbf{v}$ we define the following. Consider running $\text{VCG}^{\leq n}$ on Env' with valuation profile $\mathbf{v}$. Random set B contains every item j such that j is allocated by $\text{VCG}^{\leq n}$. Random variable G (for *global*) is the highest value of an unallocated bidder for an unallocated item, and for every j, random variable L_j (for *local*) is the highest value of an unallocated bidder for item j.

Bounds

Claim 5.25 (Upper Bound on Optimal Expected Revenue). $\mathbb{E}[\text{OPT}(\text{Env})] \leq 2\mathbb{E}[\text{VCG}(\widehat{\text{Env}})]$, i.e., the expected revenue of the optimal mechanism for Env is upper bounded by twice the expected revenue of the VCG mechanism for $\widehat{\text{Env}}$.

Proof. We know that $\mathbb{E}[\text{OPT}(\text{Env})] \leq \mathbb{E}[\text{OPT}(\text{Env}^{\text{rep}})]$ (Lemma 5.3). Relaxing the unit-demand constraint in Env^{rep} while maintaining a supply limit of $2n$ only increases the optimal revenue. The result is a matroid environment (a parallel 1-unit environment with n bidders per item and global supply limit $2n$, to be precise), to which we add one bidder per item and apply the Bulow-Klemperer-style result in Theorem 5.11,[4] stating that the expected revenue of VCG on the resulting environment $\widehat{\text{Env}}$ is a $\frac{1}{2}$-approximation to the optimal expected revenue. This completes the proof. □

Claim 5.26 (Global and Local Upper Bounds). Given a valuation profile V for environment $\widehat{\text{Env}}$, the VCG payment for every allocated item $j \in A$ is the maximum among N (global upper bound) and S_j (local upper bound).

[4]We can't use the classic Bulow-Klemperer theorem here since representative bidders who are interested in different items are not i.i.d. Alternatively with some modification we could have used the Bulow-Klemperer theorem of Hartline and Roughgarden for non-i.i.d. bidders [45].

Proof. The VCG revenue from every allocated item j is upper bounded by the highest value for any unallocated item, and by the second price for j among the $n+1$ bidders interested in j. Due to the supply limit of $2n$ in $\widehat{\text{Env}}$ and by definitions of N and S_j, we can write this as $\max\{N, S_j\}$. □

Claim 5.27 (Global and Local Lower Bounds). Given a valuation profile $\mathbf{v}$ for environment Env', the VCG payment of the supply-limiting $\text{VCG}^{\leq n}$ mechanism for every allocated item $j \in B$ is at least the maximum among G (global lower bound) and L_j (local lower bound).

Proof. The claim is based on VCG payments reflecting externalities. Say bidder i wins item j, then if i were absent from the auction, a currently unallocated bidder could have won a currently unallocated item, without violating the supply limit of n and without interfering with the rest of the current allocation. So by definition of G, it gives a lower bound on bidder i's payment for item j. Similarly, if bidder i were absent from the auction, then any currently unallocated bidder could have won item j, and L_j is also a lower bound on i's payment for j. □

5.9.2 Relating Upper- and Lower- Bounds

Next we state a main lemma.

Lemma 5.28 (Main). $\mathbb{E}[\text{VCG}^{\leq n}(\text{Env}')] \geq \frac{2}{9}\mathbb{E}[\text{VCG}(\widehat{\text{Env}})]$.

We then use this lemma to complete the proof of the Bulow-Klemperer-style theorem for the general n, m case.

The first step in relating the bounds is to fix the valuation profile V for environment $\widehat{\text{Env}}$. This completely determines the outcome of the VCG mechanism over $\widehat{\text{Env}}$, and fixes the set of allocated items $A = \{a_1, \ldots, a_{|A|}\}$, the global upper bound N, and the local upper bound S_j for every item j.

Now consider a valuation profile for environment Env' that's compatible with V. For every item j, $n+1$ out of the $3n$ values for j are fixed, but the rest of the values as well as the attribution of values to bidders remain random. We denote such a valuation profile by $\mathbf{v}(V)$. From now on, our probabilistic arguments are all over the remaining randomness in $\mathbf{v}(V)$.

Relating the Global Bounds

Claim 5.29. $\Pr[G \geq N] \geq \frac{2}{3}$.

Corollary 5.30. *Since G is non-negative, $\mathbb{E}[G] \geq \frac{2}{3}N$.*

Proof. (of Claim 5.29)

If $m \leq 2n$, then by definition $N = 0$ and the claim holds trivially. Assume from now on $m > 2n$.

Recall that A contains the $2n$ most-valued items according to the fixed valuation profile V. By definition of N, it is upper-bounded by the highest value in V for any item in A. Clearly it is also upper-bounded by the highest value in $\mathbf{v}(V)$ for any item in A. It remains to relate G to this bound.

In environment Env' with valuation profile $\mathbf{v}(V)$, consider the random subset of bidders that contains, for every $j \in A$, the bidder with the highest value for j among all $3n$ bidders. Denote this random subset of bidders by A'. Notice that for every j, the bidder with the highest value for j is distributed uniformly at random among the $3n$ bidders. Therefore, A' corresponds to the random subset of bins chosen by throwing $2n$ balls into $3n$ bins uniformly at random. The following claim formalizes the intuition that the balls are likely to occupy many of the bins.

Claim 5.31. $\Pr[|A'| \leq n] \leq \frac{1}{3}$.

Proof. We upper bound the probability that $2n$ balls thrown uniformly at random into $3n$ bins all land in a subset of at most n bins. The following is a loose upper bound: $\binom{3n}{n} n^{2n}/(3n)^{2n}$, where $\binom{3n}{n}$ is the number of subsets of n out of $3n$ bins, n^{2n} is the number of possibilities to arrange $2n$ balls in n bins, and $(3n)^{2n}$ is the number of possibilities to arrange $2n$ balls in $3n$ bins (this upper bound is not tight due to over-counting in the numerator). Simplifying we get the expression $\binom{3n}{n}/3^{2n}$, which is at most $\frac{1}{3}$ for every n. □

We remark that as $n \to \infty$, the cardinality of A' becomes concentrated around its expectation, and $\Pr[|A'| \leq n] \to 0$.

In the likely event that $|A'| > n$, there is at least one bidder $i \in A'$ that is unallocated by the supply-limiting mechanism $\text{VCG}^{\leq n}$ due to its supply limit. Let j be an item in A such that i has the highest value for j among all bidders. Then j is necessarily unallocated by $\text{VCG}^{\leq n}$, as otherwise the VCG allocation is not welfare-maximizing. We conclude that if $|A'| > n$ then $v_{i,j}$ is the value of an unallocated bidder for an unallocated item and so $G \geq v_{i,j}$. Since $v_{i,j} \geq N$, the probability that $G \geq N$ is at least $\frac{2}{3}$, as required. □

Relating the Local Upper Bound to the Lower Bounds

For every item $j \in A$ allocated by VCG in $\widehat{\text{Env}}$, we can relate between the local upper bound S_j and the lower bounds G and L_j provided that item j has the following property.

Definition 5.32 (Profitable Item)**.** Given a valuation profile $\mathbf{v}(V)$ for environment Env', item j is *profitable* if when j is removed from Env' and $\text{VCG}^{\leq n-1}$ is applied, the two bidders with the highest and second-highest values for j are unallocated.

When item j is *profitable* we can show the following relation among the relevant bounds, depending on whether j is allocated by $\text{VCG}^{\leq n}$ on Env' and so $j \in B$, or not.

Claim 5.33. For every *profitable* item j, if $j \in B$ then $L_j \geq S_j$, otherwise $G \geq S_j$.

Furthermore we argue that an item is *profitable* with high probability.

Claim 5.34. For every item j, $\Pr[j \text{ is good}] \geq \frac{4}{9}$.

We now begin to prove Claim 5.33, which is based on the claim that one of the two bidders in Env' with highest and second-highest values for a *profitable* item j remains unallocated by $\text{VCG}^{\leq n}$. This follows from the definition of a *profitable* item together with a stability property of VCG allocations in a matching environment, by which the allocations of $\text{VCG}^{\leq n}$ and $\text{VCG}^{\leq n-1}$ without item j are almost the same.

Proof. (of Claim 5.33) Let U be the set of bidders who are left unallocated by the deferred allocation procedure of removing item j from Env′ and running VCG$^{\leq n-1}$. Since j is *profitable*, U includes the two bidders with highest and second-highest values H_j, S_j for j. The set of bidders who remain unallocated by VCG$^{\leq n}$ is either exactly U or U with one bidder removed, by the same alternating path argument as in Claim 5.23. We conclude that there is an unallocated bidder whose value for j is at least S_j. By definition of G and L_j as the highest value of an unallocated bidder for an unallocated item and for item j, respectively, depending on whether item j is allocated either $G \geq S_j$ or $L_j \geq S_j$, as required. □

To show that item j is *profitable* with constant probability for every j, we use a deferred allocation argument and utilize the ratio between the number of bidders in Env′ and the supply limit of VCG$^{\leq n}$.

Proof. (of Claim 5.34) Running VCG$^{\leq n-1}$ on environment Env′ after removing item j leaves at least $2n+1$ out of the $3n$ bidders unallocated, so $|U| \geq 2n+1$. Since item j does not take part in this deferred allocation procedure, its values are distributed uniformly among all bidders, and the probability that a certain value is attributed to a bidder in U is $|U|/3n \geq 2/3$. The probability that the two bidders with highest and second-highest values for item j are both in U is therefore at least $\frac{4}{9}$. □

5.9.3 Putting It All Together

Using the relations we established among the various bounds, we are now ready to prove our main lemma and theorem. We assume for simplicity that $m > n$ (otherwise the proof reduces to that of Theorem 5.17). Our proof is based on a subtle charging argument – we need to charge the payments for items $a_1, \ldots, a_{|A|}$ in A, which were allocated by VCG on $\widehat{\text{Env}}$, against the payments for the n items in B, which were allocated by VCG$^{\leq n}$ on Env′. We use a somewhat different argument for items in $A \cap B$ and in $A \setminus B$.

Proof. (of Lemma 5.28 (Main Lemma)) For every $k \in [2n]$, we define an auxiliary random variable X_k as follows.

$$\begin{cases} \text{if } k \leq |A| \text{ and } a_k \in B & \text{then} \quad X_k = \max\{G, L_{a_k}\} \\ \text{if } k \leq |A| \text{ and } a_k \notin B & \text{then} \quad X_k = G \\ \text{otherwise} & \qquad\quad\;\; X_k = 0 \end{cases}$$

We also define a 2-to-1 mapping β from the set $\{1, \ldots, 2n\}$ to the set of n items B, such that for every $k \leq |A|$, if $a_k \in B$ then $\beta(k) = a_k$. We now show that X_k serves as an intermediary between the upper and lower bounds.

For the lower bounds, observe that $X_k \leq \max\{G, L_{\beta(k)}\}$, since if $k \leq |A|$ and $a_k \in B$ then $X_k = \max\{G, L_{a_k}\} = \max\{G, L_{\beta(k)}\}$, and the other cases clearly hold as well. Summing up over k we get

$$\sum_{k=1}^{2n} X_k \leq \sum_{k=1}^{2n} \max\{G, L_{\beta(k)}\} \leq 2 \sum_{j \in B} \max\{G, L_j\} \leq 2\,\mathrm{VCG}^{\leq n}(\mathrm{Env}') \tag{5.9.1}$$

where the second inequality uses the property that β is a 2-to-1 mapping.

For the upper bounds, let $k \leq |A|$. Since $X_k \geq G$ we know that $\mathbb{E}[X_k] \geq \mathbb{E}[G] \geq \frac{2}{3}N$ (Lemma 5.29). Now combining the probability that a_k is a *profitable* item (Claim 5.34), with the fact that if a_k is *profitable* then $X_k \geq S_{a_k}$ (Lemma 5.33), we also have that $\mathbb{E}[X_k] \geq \frac{4}{9} S_{a_k}$. Summing up over $k \leq |A|$ we get

$$\sum_{k=1}^{|A|} \mathbb{E}[X_k] \geq \frac{4}{9} \sum_{k=1}^{|A|} \max\{N, S_{a_k}\} = \frac{4}{9} \mathbb{E}[\mathrm{VCG}(\widehat{\mathrm{Env}})]. \tag{5.9.2}$$

Combining Inequalities 5.9.1 and 5.9.2 completes the proof that $\mathbb{E}[\mathrm{VCG}^{\leq n}(\mathrm{Env}')] \geq \frac{2}{9}\mathbb{E}[\mathrm{VCG}(\widehat{\mathrm{Env}})]$. □

Proof. (of Theorem 5.19) The theorem follows directly from the main lemma and from the bounds we established (Lemma 5.25). □

5.10 Extension to Multi-Unit Matching Environments

A *multi-unit matching* environment is a multi-dimensional matching environment with k_j units of item j and a total of $m = \sum_j k_j$ units. We can also impose an additional global supply limit $\ell \leq m$ on the total number of allocated units.

Two out of the three supply-limiting mechanisms in Section 5.6 for i.i.d. matching environments apply directly to i.i.d. multi-unit matching as well. In fact, Theorems 5.14 and 5.15 hold without change for multiple units. Recall that Theorem 5.16 gives a constant approximation guarantee in the challenging case where the number of items m is much larger than the number of bidders n. In order to generalize this theorem to multi-unit matching, we introduce a slightly more general supply-limiting mechanism. Let $\mathrm{VCG}^{\leq\ell,\leq\ell_j}$ be the VCG mechanism with a global supply limit ℓ on the total number of allocated units, and local supply limits $\{\ell_j\}$ on the number of allocated units of every item j. We then have the following multi-unit version of Theorem 5.16.

Theorem 5.35 ($\frac{1}{27}$-Approximation for Multiple Units)**.** *For every multi-unit matching environment with $n \geq 3$ i.i.d. regular bidders, m total units and k_j units per item j, where n is a multiple of 3,* $\mathbb{E}[\mathrm{VCG}^{\leq n/3,\leq\lceil k_j/2\rceil}(n)] \geq \frac{1}{27}\mathbb{E}[\mathrm{OPT}(n)]$.

We prove this theorem via our general reduction, using the following multi-unit version of the Bulow-Klemperer-style result in Theorem 5.19.

Theorem 5.36 ($\frac{1}{9}$-Approximate B-K for Multiple Units)**.** *For every multi-unit matching environment with n i.i.d. regular bidders, m total units and k_j units per item j,* $\mathbb{E}[\mathrm{VCG}^{\leq n,\leq\lceil k_j/2\rceil}(3n)] \geq \frac{1}{9}\cdot\mathbb{E}[\mathrm{OPT}(n)]$.

In this section we prove Theorem 5.36, the multi-unit version of Theorem 5.19. In particular we show that:

$$\mathbb{E}[\mathrm{VCG}^{\leq n,\leq\lceil k_j/2\rceil}(3n)] \geq \frac{1}{9}\cdot\mathbb{E}[\mathrm{OPT}(n)].$$

The proof is similar to that of Theorem 5.19. One main novel component is an application of the Bulow-Klemperer-Style theorem for parallel multi-unit auctions (Theorem 5.11). In what follows we highlight the remaining differences from the proof of Theorem 5.19.

For simplicity, we assume throughout that $m \geq 2n$ (otherwise, one can always use the Bulow-Klemperer-style results in Theorems 5.17 and 5.18 instead of Theorem 5.36). This assumption simplifies the analysis since it guarantees that despite the local supply limits, there are enough units such that $\mathrm{VCG}^{\leq n, \leq \lceil k_j/2 \rceil}$ allocates to exactly n bidders. We also assume for simplicity that $k_j < n$ for every j (the case where $k_j = n$ is only simpler).

5.10.1 Extending the Weak Bound

The upper bound on the optimal expected revenue remains $\sum_j \mathbb{E}[\mathrm{Vic}_j(n+k_j)]$ (Lemma 5.20). For the lower bound we need the stronger claim in Claim 5.27, since it is no longer the case that in the augmented environment all units can be allocated. The claim is that the VCG payment for item j is lower-bounded by the value of any unallocated bidder for either item j itself (the local lower bound, denoted by L_j), or for any unallocated item (the global lower bound, denoted by G). Note that due to the supply limit, exactly n out of $2n$ bidders are allocated.

Step (1) in the deferred allocation procedure (Algorithm 5.3) is now: Find a welfare-maximizing allocation (i.e., a maximum matching) without item j and under the global supply limit $(n - k_j)$. Let U be the set of $n + k_j$ unallocated bidders. Claim 5.22 holds. We also need an observation that the set of bidders left unallocated by VCG with supply limit n in the augmented environment is U with at most k_j bidders removed. Claim 5.24 holds for items j that are allocated by $\mathrm{VCG}^{\leq n}$ in the augmented environment.

Putting everything together we use the following charging argument.

- If the VCG mechanism with supply limit n allocates item j, then the payment for j is at least L_j and so in expectation is at least $\frac{1}{k_j}\mathbb{E}[\mathrm{Vic}_j(n + k_j)]$.

- If the VCG mechanism with supply limit n allocates item j' but not item j, then the payment for j' is at least G and so in expectation is at least $\frac{1}{k_j}\mathbb{E}[\text{Vic}_j(n+k_j)]$.

Thus we can charge all m total units against the n allocated units, obtaining an approximation factor of $\frac{n}{m}$.

5.10.2 Upper and Lower Bounds

We use the same notation as in the proof of Theorem 5.19, but in some cases to denote somewhat different objects. We now state the differences.

Environments Let Env be a multi-unit matching environment with n i.i.d. regular bidders, $m \geq 2n$ total units, and $k_j \leq n$ units per item j. Let $\text{Env}', \text{Env}^{\text{rep}}$ be the augmented and representative environments, respectively. Environment $\widehat{\text{Env}}$ is obtained from Env^{rep} by adding n additional bidders per item, and replacing the unit-demand constraint by a supply limit of $2n$. So $\widehat{\text{Env}}$ is a k_j-unit, m-item environment with $2n$ bidders per item and supply limit $2n$.

Valuation profiles $V, \mathbf{v}$ Let V be a collection of sets V_j, each containing $2n$ i.i.d. samples from F_j. There is no change in $\mathbf{v}$.

Random variables over V Let S_j be the $(k_j + 1)$-th-highest value for item j. If $m > 2n$, the random variable N is the $(2n + 1)$-st-highest among a set containing the k_j highest values for every item j. Otherwise, $N = 0$. Let A be a multiset containing the $2n$ items (with repetitions) with the highest values among a set containing the k_j highest values for every item j. Let $a_1, \ldots, a_{2n}$ denote the items in A ordered by their value from high to low. Observe that if VCG runs on $\widehat{\text{Env}}$ with valuation profile V, the set of $2n$ allocated items is equal to A.

Random variables over $\mathbf{v}$ Consider running $\text{VCG}^{\leq n, \leq \lceil k_j/2 \rceil}$ on Env' with valuation profile $\mathbf{v}$. Let B be a multiset containing a copy of j for every allocated unit of j. Random variable G is the highest value of an unallocated bidder for an unallocated unit, and for every j, random variable L_j is the highest value of an unallocated bidder for item j.

Bounds

Bounds do not change. Note that the proof of the upper bound on the optimal expected revenue critically uses the single-dimensional Bulow-Klemperer-style result in Theorem 5.11 to justify the augmentation of the environment by adding $n \geq k_j$ bidders per item, with no restrictions of the form "at most one of every pair of duplicates can be allocated".

5.10.3 Relating the Upper and Lower Bounds

Relating the Global Bounds

Claim 5.37. $\Pr[G \geq N] \geq \frac{2}{3}$.

Corollary 5.38. *Since G is non-negative,* $\mathbb{E}[G] \geq \frac{2}{3}N$.

Proof. (of Claim 5.37)

In environment Env' with random valuation profile $\mathbf{v}(V)$, denote by A' the random subset of bidders that contains the highest bidders for items in A. Subset A' corresponds to the random subset of bins chosen by throwing $2n$ balls into $3n$ bins uniformly at random, under the restriction that some balls cannot fall in the same bins, because they correspond to values of different bidders for the same item. The probability that $|A'| \leq n$ is thus at most the probability calculated in Claim 5.31, i.e., $\Pr[|A'| \leq n] \leq \frac{1}{3}$. So with high probability, there exists a bidder $i \in A'$ that is unallocated by the supply-limiting mechanism $\text{VCG}^{\leq n, \leq \ell_j}$ due to the global supply limit.

Let j be one of the items in A for which bidder i is in A', i.e., i has one of the highest values for j. At least one unit of j is unallocated by $\text{VCG}^{\leq n, \leq \lceil k_j/2 \rceil}$. We conclude that if $|A'| > n$ then $v_{i,j}$ is the value of an unallocated bidder for an unallocated unit and so $G \geq v_{i,j}$. Since $v_{i,j} \geq N$, the probability that $G \geq N$ is at least $\frac{2}{3}$, as required. □

Relating the Local Upper Bound to the Lower Bounds

We redefine a *profitable* item as follows.

Definition 5.39 (Profitable Item)**.** Given a valuation profile $\mathbf{v}(V)$ for environment Env$'$, item j is *profitable* if when all units of j are removed from Env$'$ and VCG with global supply constraint of $n - \lceil k_j/2 \rceil$ and local supply constraints $\ell'_j \neq j$ is applied, then $\lceil k_j/2 \rceil + 1$ bidders whose values for j are among the $k_j + 1$ overall highest values for j remain unallocated.

When item j is profitable we can show the following.

Claim 5.40*.* For every profitable item j, if $j \in B$ then $L_j \geq S_j$, otherwise $G \geq S_j$.

Furthermore we argue that an item is profitable with high probability.

Claim 5.41*.* For every item j, Pr[j is good] $\geq \frac{4}{9}$.

Proof. (of Claim 5.40)

Let U be the set of bidders who are left unallocated by the deferred allocation procedure of removing item j from Env$'$ and running VCG with global supply constraint of $n - \lceil k_j/2 \rceil$ and local supply constraints $\ell'_j \neq j$. Since j is profitable, U includes $\lceil k_j/2 \rceil + 1$ bidders whose values for j are among the $k_j + 1$ overall highest values for j. The set of bidders who remain unallocated by VCG$^{\leq n, \leq \lceil k_j/2 \rceil}$ is U with at most $\lceil k_j/2 \rceil$ bidders removed. We conclude that there is an unallocated bidder whose value for j is at least S_j. By definition of G and L_j as the highest value of an unallocated bidder for an unallocated unit and for item j, respectively, depending on whether item j is allocated either $G \geq S_j$ or $L_j \geq S_j$, as required. □

We show that item j is profitable with constant probability for every j by a deferred allocation argument.

Proof. (of Claim 5.41) Running VCG with global supply constraint of $n - \lceil k_j/2 \rceil$ and local supply constraints $\ell'_j \neq j$ on environment Env$'$ after removing item j leaves at least $2n + \lceil k_j/2 \rceil$ out of the $3n$ bidders unallocated, i.e., $|U| \geq 2n + \lceil k_j/2 \rceil$. Since item j does not take part in this deferred allocation procedure, its $k_j + 1$ highest values are distributed uniformly among all bidders. The probability that at least $\lceil k_j/2 \rceil + 1$ of these values are in U is at least $\frac{4}{9}$. □

Proof of Main Lemma and Main Theorem

The statement of the main lemma and its proof do not change. Note that to define the 2-to-1 mapping from $\{1, \ldots, 2n\}$ to B used in the proof we apply the assumption that $m \geq 2n$ and so $|B| = n$. Proof of Theorem 5.36) follows directly in the same way.

Chapter 6

Sequential Posted-Price Mechanisms

In this chapter, we propose mechanisms based on the sequential posted-price mechanisms, prove that they give approximately-optimal revenue guarantee based on the "correlation gap," and also show how to adapt such mechanisms to give prior-independent approximations. This chapter is mainly based on [76].

6.1 Introduction

6.1.1 Settings and Mechanisms

We consider the single-dimensional downward-closed environments as defined in Section 2.1. For such environments, mechanisms such as Myerson's mechanism [61] or the VCG mechanism [73, 21, 38] have optimal revenue or welfare guarantees, but often suffer from having complicated formats or severe computational overhead. For example, even in single-item auctions, the need for the agents to commit to the auction process itself can be a significant burden [7, 48], and in combinatorial auctions, determining the allocation and payments of the VCG mechanism is a computationally hard problem [63].

On the other hand, consider Sequential Posted-price Mechanisms (SPMs in short), in which the seller makes take-it-or-leave-it price offers to agents one by one. Such mechanisms are easy to run for the sellers, leave little room for agents' strategic behavior, and keep the information elicitation from the agents at a minimum level. Therefore, not surprisingly, such mechanisms are very often favored in practice [48].

In this chapter, we want to achieve two goals.

1. First, as simple SPMs are in general not optimal, we want to quantify how much we lose by using the simple SPMs instead of the optimal mechanism.

2. Second, we want to understand how much information about the distribution is needed for us to set prices in an SPM to achieve approximately-optimal revenue.

We achieve these two goals by instantiating the three-step approach in Section 1.7.1 toward achieving prior-independence. Essentially,

1. We identify a class of mechanisms, with certain parameters.

   Here we consider the class of SPMs, with the parameters being prices set for bidders.

2. Given distribution information, prove that for appropriate choice of prices SPMs, we can achieve approximately-optimal revenue.

   This helps us achieve our first goal, bounding the relative loss of using simple SPMs compared to the optimal.

3. Without distribution information, find a way to set such prices, at a bounded loss in revenue.

   This helps us achieve our second goal, quantifying the information needed about the distributions.

6.1.2 Main Results

SPM vs. OPT In a recent work of Chawla et al. [19], it was shown for several contexts that the performance of a SPM (which we call greedy-SPM) approximates that of the optimal mechanism by a constant factor, where the factor is $\frac{1}{2}$ for matroid environments (which generalize k-unit auctions, certain matching markets etc. see Section 2.1.1), and $1-\frac{1}{e}$ for k-unit auctions. This is surprising, as SPMs can only offer prices to agents in a very restricted way, while the optimal mechanism can choose a price for each agent based on full information about all other agents. What is the underlying reason for the good performance of SPMs?

We give a theoretical explanation for this curious fact, based on a connection to the notion of correlation gap which we elaborate on in Section 6.1.3. Exploiting this connection, we give a tight analysis of the greedy-based SPM of Chawla et al. for several environments. In particular, we show that it gives an $(1 - \frac{1}{e})$-approximation for matroid environments (an improvement over the previous $\frac{1}{2}$-approximation), gives asymptotically a $(1 - \frac{1}{\sqrt{2\pi k}})$-approximation for the important sub-case of k-unit auctions (an improvement over the previous $(1 - \frac{1}{e})$-approximation), and gives a $\frac{1}{p+1}$-approximation for environments with p-independent set system constraints, which generalizes the result on intersection of p matroids in [19].

Prior-Independence We put SPMs into the prior-independent analysis framework to study how much information about the distribution is needed to set good prices in an SPM. For k-unit auctions with i.i.d. regular bidders, we prove that when k is large compared to n, a single sample from the distribution as the price in SPM gives a constant factor approximation. In the general case, a single sample fails to work, but $O(n)$ samples are sufficient for a constant factor approximation.

6.1.3 Technique: Analysis via Correlation Gap

Reducing Mechanism Design to Correlation Gap The notion of correlation gap was first formalized in Agrawal et al. [2]. Let $f(S)$ be a function that maps a subset S of a finite ground set N to a nonnegative real number. For $\mathcal{D}$ a distribution over 2^N with marginal probabilities $q_i = Pr_{S\sim\mathcal{D}}[i \in S]$, let $\mathcal{I}_\mathcal{D}$ be the independent distribution where each $i \in N$ is included in the set with the same marginal probability q_i, but independently. The correlation gap of f is defined as the supremum of $\frac{E_{S\sim\mathcal{D}}[f(S)]}{E_{S\sim\mathcal{I}(\mathcal{D})}[f(S)]}$ over all distributions $\mathcal{D}$, which in some sense bounds our "loss" in expected value of the function by ignoring correlation.

Loosely speaking, the approximation ratio of SPMs w.r.t. the optimal mechanism is related to correlation gap in the following way. The performance of a mechanism can often be related to the expectation of certain function f over a random set of agents. For an optimal mechanism, this random set corresponds to the set of winners, while for an SPM, this random set corresponds to the demand set, which is the set of agents whose values beat the prices set for them in the SPM. Notice that the winner set is highly-dependent, while the demand set is independent. By setting prices for agents in an SPM carefully such that these two random sets have the same marginal probabilities, we can apply the correlation gap of f to get a bound on the approximation ratio of the SPM w.r.t. the optimal mechanism.

Reduction for k-Unit Auctions To illuminate the idea, suppose we sell k items to a set of n agents $N = \{1, \ldots, n\}$ with valuations drawn i.i.d. from a regular distribution F, and our goal is to maximize expected revenue. Define set function f as $f(S) = \min(|S|, k)$ for $S \subseteq N$. Let q be the probability that Myerson's optimal mechanism sells to a particular agent (which is the same for every agent by symmetry). It can be shown that if the distribution is regular, the optimal way to sell to an agent with success probability q in an incentive compatible manner is to offer the deterministic price $p = F^{-1}(1-q)$. Therefore if we pretend that an agent pays p whenever she wins in the optimal mechanism, the total calculated revenue is only higher. In other words, the revenue of Myerson's mechanism is upper-bounded by $E_W[f(W)] \cdot p$, where W is the set of winners. On the other hand, let an SPM make take-it-or-leave-it offers at price p to every agent sequentially. Define demand set D as the set of agents whose values are at least p. Since at most k agents can be served, by definition of f, the revenue of the SPM is equal to $E_D[f(D)] \cdot p$. Note that W and D have the same marginal probability q for every bidder i, and D follows an independent distribution. Therefore if we can show that the correlation gap of f is at most β, then $E_D[f(D)] \geq \frac{1}{\beta} \cdot E_W[f(W)]$, and it follows that the revenue of SPM is a $\frac{1}{\beta}$-approximation to that of Myerson's mechanism.

Submodularity The set function f that arises in our context is the weighted rank function of the set system that encodes the feasibility constraints of the environment. For settings where constraints are modeled by matroids, the weighted rank functions are well-known to be monotone and submodular (see e.g., [68]). This fact enables us to invoke a result from [74, 2], which says that the correlation gap of a monotone and submodular function is at most $\frac{e}{e-1}$. It follows that for matriod environments, SPMs can approximate the optimal mechanism by a factor of $1 - \frac{1}{e}$. This result would be otherwise difficult to achieve without making use of our explicit connection to correlation gap and submodularity.

Recognizing submodularity is also helpful in other ways. In the analysis for k-unit auctions, we exploit the cross-convexity of the multi-linear extension of submodular functions to get a tight bound on the correlation gap of the corresponding weighted rank function.

Benefits of the Reduction The reduction to correlation gap gives us a structured way of analyzing greedy-SPM. It abstracts away all the mechanism design aspects of the problem, such that we can focus on the purely mathematical question of quantifying correlation gaps of weighted rank functions.

6.1.4 Related Work

Sequential posted-price mechanisms have also been a recent focus of study due to their simplicity and various appealing properties. Blumrosen and Holenstein [13] first compared SPMs to Myerson's mechanism for single-item auctions by an asymptotic analysis. Chawla et al. [19] studied SPMs in various auction contexts, proving that SPMs perform very well compared to Myerson's mechanism, which motivated our work. They also used SPMs as a building block to construct approximately-optimal mechanisms in multi-dimensional environments. Independent of our work, Chakraborty et al. [18] proved almost the same approximation guarantee for k-unit auctions. They also studied SPMs that adaptively choose prices and the ordering of agents. Babaioff et al. [9] studied adaptive SPMs in settings where agents' valuations are drawn i.i.d. from an unknown distribution. In other aspects, Sundararajan and Yan [70] studied the performance of SPMs when the sellers are risk-averse, and aim to maximize expected utility.

There is a vast literature on the study of submodular functions (see references in [74]). The correlation gap of monotone submodular functions was first bounded in [17], and it is also tightly related to the submodular welfare maximization problem [75]. In the context of auctions, Dughmi et al. [26] showed that in matroids environments, the revenue of Myerson's mechanism is submodular in the agent set.

6.2 Preliminaries

Correlation Gap and Submodularity Given a set function $f : 2^N \to [0, \infty)$ over a finite set N, let $\mathcal{D}$ be a distribution over 2^N with marginal probabilities $\mathbf{q} = (q_i)_{i \in N}$. Let $S \sim \mathcal{I}(\mathcal{D})$ denote that each $i \in N$ is included in S with probability q_i independently. Then the correlation gap [2] of f is $\sup_{\mathcal{D}} \frac{E_{S \sim \mathcal{D}}[f(S)]}{E_{S \sim \mathcal{I}(\mathcal{D})}[f(S)]}$. (let $\frac{0}{0} = 1$)

A set function $f : 2^N \to [0, \infty)$ is monotone if $f(S) \leq f(T)$ whenever $S \subseteq T$, and is submodular if $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ for all S, T.

Theorem 6.1. *[17, 2] The correlation gap of a monotone submodular function is at most* $\frac{e}{e-1}$.

6.3 Posted-Price vs. Optimal: A Reduction to Correlation Gap

We will focus on comparing SPMs to the optimal mechanism in the context of revenue maximization. Almost identical claims can be made for welfare and certain other objectives, which we discuss in Section 6.3.3.

6.3.1 A Single Bidder Optimization Problem

Before we embark on studying mechanisms that involve multiple bidders, it is crucial to first understand the following optimization problem that involves only one bidder.

Problem 6.2. Given an agent with valuation distribution F, and a target selling probability $0 < q < 1$, what price distribution $\mathcal{D}$ maximizes our expected revenue, i.e., $E_{p \sim \mathcal{D}}[p \cdot (1 - F(p))]$, subject to the constraint that the selling probability is exactly q, i.e., $E_{p \sim \mathcal{D}}[1 - F(p)] = q$?

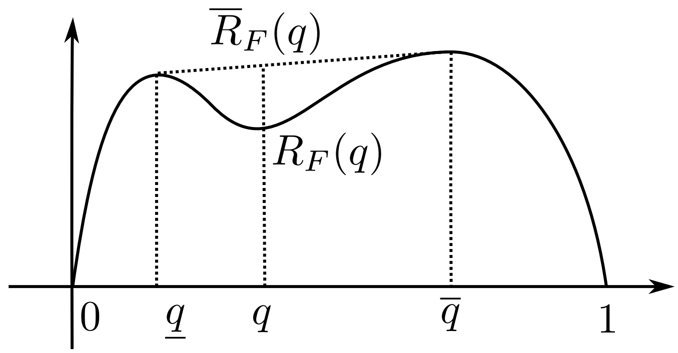

Figure 6.3.1: Revenue Curve and "Ironed" Revenue Curve

To study this problem, first suppose that we can only offer a deterministic price. Then for any selling probability q, our only choice is to offer the deterministic price $F^{-1}(1-q)$, and the revenue we get as a function of q is $R_F(q) = q \cdot F^{-1}(1-q)$.

Now suppose instead we are allowed to offer a random price, then we can do possibly better. To be specific, we can randomize between two prices $\underline{p}$ and $\bar{p}$ with selling probabilities $\underline{q} = 1 - F(\underline{p})$ and $\bar{q} = 1 - F(\bar{p})$ satisfying $\underline{q} \leq q \leq \overline{q}$, and in particular we draw $\underline{p}$ with probability $\frac{\overline{q}-q}{\overline{q}-\underline{q}}$ and draw $\overline{p}$ with probability $\frac{q-\underline{q}}{\overline{q}-\underline{q}}$ such that the selling probability is exactly equal to q. Then our revenue is equal to $\frac{\overline{q}-q}{\overline{q}-\underline{q}} \cdot R_F(\underline{q}) + \frac{q-\underline{q}}{\overline{q}-\underline{q}} \cdot R_F(\overline{q})$, which is possibly better than $R_F(q)$ (see Figure 7.2.2). Let $\overline{R}_F(q)$ be the maximum revenue one can get by randomizing between two prices this way. One can show that $\overline{R}_F$ equals to the concave closure of R_F, i.e., the minimum concave function that upper-bounds R_F. Moreover, the optimal distribution is in fact just the two-price distribution that gives $\overline{R}_F(q)$.

In the well-known special case that F is regular, i.e., $R_F(q)$ is concave in q, the two-price distribution degenerates to a single deterministic price $F^{-1}(1-q)$, and $\overline{R}_F(q) = R_F(q)$ in this case.

For the purpose of the rest of the paper, the following lemma summarizes this discussion.

Lemma 6.3 (Ironing Lemma)**.** *[61, 16] For all valuation distribution F and probability q, the price distribution $\mathcal{D}$ that maximizes $E_{p\sim\mathcal{D}}[p \cdot (1 - F(p))]$ subject to the constraint that $E_{p\sim\mathcal{D}}[1 - F(p)] = q$ is a two-price distribution, where this distribution* *as well as the revenue $\overline{R}_F(q)$ it gives us can be determined from F. Moreover, $\overline{R}_F(q)$ is a concave function.*

For notational convenience, we will use $\overline{R}_i$ to denote the $\overline{R}_F$ function for agent i.

6.3.2 Reduction Theorem: the Revenue Case

For each bidder i and $q_i \in [0, 1]$, we define effective price as $\hat{p}_i = \overline{R}_i(q_i)/q_i$. The following defines the greedy-SPM of Chawla et al. [19] (with modifications that are important for irregular distributions).

Definition 6.4. The greedy-SPM does the following:

1. For each agent i, calculate q_i, the winning probability of agent i in Myerson's mechanism. Remove agent i if $q_i = 0$.
2. For each agent i, draw a random price p_i from the optimal price distribution w.r.t. distribution F_i and selling probability q_i according to the discussion in Section 6.3.1.
3. Let $A = \emptyset$. For all agents i in decreasing order of effective prices $\hat{p}_i$, if serving agent i is feasible, i.e., $A + i \in \mathcal{I}$, offer price p_i to agent i, and add i into A if agent i accepts.

Theorem 6.5 (Reduction Theorem for Matroids)**.** *For matroid environments, if the correlation gap of the weighted rank function is at most β for all non-negative weights, then the expected revenue of greedy-SPM is a $\frac{1}{\beta}$-approximation to that of Myerson's optimal mechanism.*

Proof. In the following two claims, we relate the expected revenue of both Myerson's mechanism and greedy-SPM to the weighted rank function with effective prices $\hat{p}_i$ as weights, which we denote as $\hat{p}^*(\cdot)$.

Claim 6.6. Let W be the (random) set of winning agents in Myerson's mechanism. The expected revenue of Myerson's mechanism is upper-bounded by $E_W[\hat{p}^*(W)]$.

Proof. Let $q_i = Pr_W[i \in W]$ be the probability that agent i wins in Myerson's mechanism. By Lemma 6.3, the optimal way to sell to agent i with probability q_i gives expected revenue $\overline{R}_i(q_i)$. By linearity of expectation, the expected revenue of Myerson's mechanism is upper-bounded by $\sum_{i \in N} \overline{R}_i(q_i)$. To relate this to the effective prices, suppose in Myerson's mechanism, we get effective payment $\hat{p}_i$ whenever agent i wins. Then the total effective revenue is $E_W[\sum_{i \in W} \hat{p}_i]$. Also, each agent i wins with probability q_i in Myerson's mechanism, contributing $q_i \hat{p}_i = \overline{R}_i(q_i)$ to total effective revenue, and hence $\sum_{i \in N} \overline{R}_i(q_i)$ equals effective revenue $E_W[\sum_{i \in W} \hat{p}_i]$. Furthermore, since W is a feasible set, we can rewrite $E_W[\sum_{i \in W} \hat{p}_i]$ as $E_W[\hat{p}^*(W)]$, and our claim follows. □

Claim 6.7. Let demand set D be the (random) set of agents whose values beat the prices set for them. The expected revenue of greedy-SPM equals to $E_D[\hat{p}^*(D)]$.

Proof. Because valuation distributions of the agents are independent, each agent i is in the demand set D with probability q_i independently. Observe that ignoring agents not in the demand set, who do not win anyway, greedy-SPM effectively runs the greedy algorithm on the demand set D w.r.t. weights $\hat{p}_i$ subject to feasibility constraints. The expected effective revenue of greedy-SPM is hence equal to $E_D[\sum_{i \in greedy(D)} \hat{p}_i]$, which is equal to $E_D[\hat{p}^*(D)]$ by the optimality of the greedy algorithm for matroid. Note that whenever the random price p_i is offered to an agent, we get expected revenue $\overline{R}_i(q_i)$, while the expected effective revenue is $q_i \hat{p}_i$, also equal to $\overline{R}_i(q_i)$. Therefore the expected revenue of greedy-SPM equals to the expected effective revenue, which is $E_D[\hat{p}^*(D)]$. □

By our assumption that the correlation gap of the weighted rank function is at most β, we have $E_D[\hat{p}^*(D)] \geq \frac{1}{\beta} \cdot E_W[\hat{p}^*(W)]$, and our theorem follows by chaining this inequality with the above two claims.

□

For settings beyond matroids, we need the following technical condition for the reduction to work, which is a stronger condition than merely a bound on correlation gap.

Definition 6.8. We say that the greedy algorithm verifies a correlation gap of β for the weighted rank function of a set system, if for all nonnegative weights $(w_i)_{i\in N}$, and distribution $\mathcal{D}$ over 2^N, we have $E_{S\sim\mathcal{I}(\mathcal{D})}[\sum_{i\in greedy(S)} w_i] \geq \frac{1}{\beta}\cdot E_{S\sim\mathcal{D}}[w^*(S)]$.

Theorem 6.9 (Reduction Theorem in General)**.** *For any downward-closed environment, if the greedy algorithm verifies a correlation gap of β for the weighted rank function for arbitrary non-negative weights, then the expected revenue of greedy-SPM is a $\frac{1}{\beta}$-approximation to that of Myerson's optimal mechanism.*

Proof. Similarly, we upper-bound the revenue of Myerson by $E_W[\hat{p}^*(W)]$, and express the revenue of greedy-SPM as $E_D[\sum_{i\in greedy(D)}\hat{p}_i]$. The theorem follows by applying the assumption that greedy verifies a correlation gap of β. □

Remark 6.10. One crucial property about the greedy algorithm is that although we are running greedy on the all agents, but for *no matter what* demand set it turns out to be, greedy is also optimizing or approximately optimizing for this demand set. Most other approximation algorithms do not have this property.

6.3.3 Extension to Welfare and Other Objectives

We specify an objective by defining functions of the form $g_i(v,p)$ for agents. If agent i has true value v and is offered a price p with $v \geq p$, then agent i wins, and we gain objective value $g_i(v,p)$. Our goal is then to maximize the total objective value we collect from the agents. For maximizing welfare, revenue, and surplus, we set $g_i(v,p) = v$, $g_i(v,p) = p$, and $g_i(v,p) = v - p$, respectively. One can also define other objectives this way.

To adapt the definition of greedy-SPM and our reduction theorems, we need the following changes. We define $G_i(q)$ as the maximum expected objective value the seller can get by offering a deterministic price such that the agent wins with probability q. We can then derive an Ironing Lemma similarly, and also define $\overline{G}_i(q)$ similarly from $G_i(q)$. Then we use effective gain defined as $\overline{G}_i(q)/q$ to replace effective prices as weights, and the rest of the proof goes the same way.

6.4 Revenue and Welfare Guarantees of Greedy-SPM

Based on the reduction theorem, we give a tight analysis of greedy-SPM, and prove the guarantees in Theorem 6.11. By the reduction theorem, it suffices to study the correlation gaps of the weighted rank functions, and the greedy algorithm, which we do separately in the following subsections.

Theorem 6.11. *The expected revenue of greedy-SPM is a β-approximation to that of Myerson's optimal mechanism, and the expected welfare of (the welfare version of) greedy-SPM is a β-approximation to that of the VCG mechanism, where:*

- $\beta = 1 - \frac{1}{e}$ *for matroid environments*
  (an improvement over $\frac{1}{2}$)
- $\beta = 1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$ *for k-unit auctions*
  (an improvement over $1 - \frac{1}{e}$)
- $\beta = \frac{1}{p+1}$ *for p-independent environments*
  (a generalization from intersection of p matroids)

Remark 6.12. For matroid environments, as noticed in [19], if we run the VCG mechanism, and set reserves to be the same as the prices used in greedy-SPM, the revenue we get is as good as that of greedy-SPM, for any particular valuation profile. It follows that the VCG mechanism with such reserve prices has the same approximation guarantee for revenue.

6.4.1 Matroid Environments

By the reduction theorem, to establish an $1 - \frac{1}{e}$-approximation of greedy-SPM in matroid environments (see Section 2.1.1), it suffices to prove the following lemma.

Lemma 6.13. *The correlation gap of the weighted rank function of a matroid is at most $\frac{e}{e-1}$.*

Proof. This lemma follows from the fact that the weighted rank function of a matroid is monotone and submodular [68], and that the correlation gap of a monotone submodular function is at most $\frac{e}{e-1}$ [17, 2]. □

6.4.2 k-Unit Auctions

k-Unit auctions form an important sub-class of a matroid environments. The feasibility constraints of a k-unit auction is modeled by a k-uniform matroid. In the following, we precisely quantify the correlation gap of the weighted rank function of k-uniform matroids.

For a k-uniform matroid over n elements, the (unweighted) rank function is $f_n^k(S) = \min(|S|, k)$ for $S \subseteq N = \{1, \ldots, n\}$. We drop superscript and subscript when the context is clear. It is easy to verify that f is monotone and submodular. Define the multi-linear extension $Ef(\mathbf{q})$ for $\mathbf{q} \in [0,1]^n$ (in the sense of [17]) as the expectation of $f(S)$ where each $i \in N$ is included in S with probability q_i independently. As was shown in [17], or can be easily verified using definitions, if f is submodular, then Ef satisfies cross-convexity, in the sense that $\frac{\partial^2 Ef(\mathbf{q})}{\partial q_i \partial q_j} \leq 0$ for all $\mathbf{q} \in (0,1)^n$ and $i \neq j$.

For all n and $0 \leq k \leq n$, define $\Phi(n,k)$ as the minimum of $Ef_n^k(\mathbf{q})$ over all marginal probability vector $\mathbf{q}$ such that $\sum_{i \in N} q_i = k$. In the following lemma, we identify the probability vector $\mathbf{q}$ that minimizes $Ef(\mathbf{q})$ subject to this constraint, and show several useful properties about $\Phi(n,k)$. This lemma is interesting in itself, and in fact can be used to improve the analysis of an SPM in [70].

Lemma 6.14. *The following holds for $\Phi(n,k)$:*

(a) $\Phi(n,k) = Ef_n^k(\mathbf{q})$ *where $q_i = k/n$ for all $i \in \{1, \ldots, n\}$. In other words, $\Phi(n,k)$ is the expected value of* $\min(X,k)$*, where X is a binomial random variable with parameters n and k/n.*

(b) $\Phi(n,k)$ *monotonically increases with k, and monotonically decreases with n.*

(c) $\lim_{n\to\infty} \Phi(n,k) = k - \frac{k^{k+1}}{e^k k!} \approx k - \frac{k}{\sqrt{2\pi k}}$.

Proof. To prove (a), first for an arbitrary marginal probability vector $\mathbf{q} \in [0,1]^n$, consider vector $\overline{\mathbf{q}}$ that is the same as $\mathbf{q}$ except that the i-th and j-th components are averaged for some $i \neq j$, i.e., $\overline{q}_i = \overline{q}_j = (q_i + q_j)/2$. We show that $Ef(\overline{\mathbf{q}}) \leq Ef(\mathbf{q})$. Let $\mathbf{q}'$ be the same as $\mathbf{q}$ except with the i-th and j-th components switched, i.e., $q'_i = q_j$ and $q'_j = q_i$. By symmetry of f, $Ef(\mathbf{q}) = Ef(\mathbf{q}')$, and $\overline{\mathbf{q}}$ is the middle-point of $\mathbf{q}$ and $\mathbf{q}'$. By the cross-convexity of Ef, the value of Ef is convex in the line segment connecting $\mathbf{q}$ and $\mathbf{q}'$. Therefore $Ef(\overline{\mathbf{q}})$ is at most the average of $Ef(\mathbf{q})$ and $Ef(\mathbf{q}')$, or simply $Ef(\mathbf{q})$. Now starting with an arbitrary $\mathbf{q}$, by repeatedly averaging the maximum and minimum components of $\mathbf{q}$ this way, the value of $Ef(\mathbf{q})$ keeps decreasing, while all q_i's converge to k/n. By the continuity of $Ef(\mathbf{q})$ in $\mathbf{q}$, the value of $Ef(\mathbf{q})$ converges to the value of Ef at $q_i = k/n$ for all i. Therefore $Ef(\mathbf{q})$ is minimized at $q_i = k/n$ for all i. □

To show (b), it is obvious that $\Phi(n,k)$ is monotonically increasing in k, because $f_n^k(S)$ is increasing in k. It suffices to show that $\Phi(n,k)$ is monotonically decreasing in n. Recall that $\Phi(n,k)$ was defined to be the optimal value of a minimization problem. To relate $\Phi(n,k)$ to $\Phi(n+1,k)$, we cast the optimal solution for the optimization problem that underlies $\Phi(n,k)$, which is an n-dimensional independent distribution, to $(n+1)$-dimensional space, such that it gives a candidate solution to the minimization problem underlying $\Phi(n+1,k)$. To be specific, we observe that $\Phi(n,k)$ is equal to $Ef_{n+1}^k(\mathbf{q})$, where $\mathbf{q}$ is an $(n+1)$-dimensional vector with $q_i = k/n$ for $i = 1,\ldots,n$, and $q_{n+1} = 0$. By definition of $\Phi(n+1,k)$, $\Phi(n+1,k) \geq Ef_{n+1}^k(\mathbf{q}) = \Phi(n,k)$.

We derive the asymptotics for $\Phi(n, k)$ as follows, where the last step is by Stirling's approximation of factorials.

$$
\begin{aligned}
& \lim_{n\to\infty} \Phi(n,k) \\
= \; & \lim_{n\to\infty} \sum_{t=0}^{n} \binom{n}{t} \cdot \left(\frac{k}{n}\right)^t \cdot \left(\frac{n-k}{n}\right)^{n-t} \cdot \min(t,k) \\
= \; & \lim_{n\to\infty} \sum_{t=0}^{k-1} \binom{n}{t} \cdot \left(\frac{k}{n}\right)^t \cdot \left(\frac{n-k}{n}\right)^{n-t} \cdot t \\
& + k \cdot \left(1 - \sum_{t=0}^{k-1} \binom{n}{t} \cdot \left(\frac{k}{n}\right)^t \cdot \left(\frac{n-k}{n}\right)^{n-t}\right) \\
= \; & \sum_{t=0}^{k-1} \frac{k^t}{t!} \cdot \frac{1}{e^k} \cdot t + k \cdot \left(1 - \sum_{t=0}^{k-1} \frac{k^t}{t!} \cdot \frac{1}{e^k}\right) \\
= \; & k \cdot \left(1 - \frac{k^k}{e^k k!}\right) \approx k \cdot \left(1 - \frac{1}{\sqrt{2\pi k}}\right).
\end{aligned}
$$

Based on Lemma 6.14, we can first quantify the correlation gap of the unweighted rank function, and then extend it to the weighted case.

Lemma 6.15. *For $n, k \geq 1$, the correlation gap of the function $f(S) = \min(|S|, k)$ for $S \subseteq N = \{1, \ldots, n\}$ is exactly $\frac{k}{\Phi(k,n)}$.*

Proof. For any probability vector $\mathbf{q}$, let $\mathcal{O}_\mathbf{q}$ be the distribution over 2^N with marginal probabilities $\mathbf{q}$ that maximizes $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)]$. We first show that $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)]$ equals $\sum_i q_i$ if $\sum_i q_i \leq k$, and equals k otherwise. (1) Suppose $\sum_i q_i \leq k$. First note that $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)] \leq E_{S\sim\mathcal{O}_\mathbf{q}}[|S|] = \sum_i q_i$. Moreover, $\mathbf{q}$ can be seen as a point inside the integral polytope with (characteristic vectors of) feasible sets (sets of size at most k) as vertices. Then by standard polyhedral combinatorics [68], one can decompose this point as a convex combination of the vertices, which corresponds to a distribution over feasible sets with marginal probabilities $\mathbf{q}$. This distribution gives expected f value $\sum_i q_i$. (2) If $\sum_i q_i > k$, then by the monotonicity of $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)]$ in $\mathbf{q}$, $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)]$ is at least k. However it is also upper-bounded by k as f is upper-bounded by k. Therefore $E_{S\sim\mathcal{O}_\mathbf{q}}[f(S)] = k$ in this case. □

Suppose that $\mathbf{q}$ maximizes the "gap ratio" $\frac{E_{S\sim\mathcal{O}_{\mathbf{q}}}[f(S)]}{E_{S\sim\mathbf{q}}[f(S)]}$. We first show that $r = \sum_i q_i \le k$. If this is not the case, then by lowering the q_i's such that $\sum_i q_i = k$, $E_{S\sim\mathbf{q}}[f(S)]$ strictly decreases, while $E_{S\sim\mathcal{O}_{\mathbf{q}}}[f(S)]$ is still k. This gives a strictly higher gap ratio, contrary to that assumption that $\mathbf{q}$ maximizes the gap ratio.

Next we show that $r = k$. For $r \le k$, we can explicitly express the reciprocal of the gap ratio as:

$$\begin{aligned} & \frac{1}{r} \cdot \sum_{t=0}^{n} \binom{n}{t} \cdot \left(\frac{r}{n}\right)^t \cdot \left(\frac{n-r}{n}\right)^{n-t} \cdot \min(t,k) \\ = & \sum_{t=1}^{n} \binom{n-1}{t-1} \left(\frac{r}{n}\right)^{t-1} \left(\frac{n-r}{n}\right)^{n-t} \cdot \frac{\min(t,k)}{t} \end{aligned}$$

This is equal to the expectation of $\frac{\min(X+1,k)}{X+1}$ where X is the binomial random variable with parameters $n-1$ and r/n. It is also equal to $\int_0^\infty Pr[\frac{\min(X+1,k)}{X+1} \ge x]dx$. Note that for $x > 1$, $Pr[\frac{\min(X+1,k)}{X+1} \ge x] = 0$, and otherwise $Pr[\frac{\min(X+1,k)}{X+1} \ge x] = Pr[X+1 \le k/x]$, where $Pr[X+1 \le k/x]$ strictly decreases as r increases. Therefore the gap ratio is maximized at $r = k$.

Lemma 6.16. *For $n, k \ge 1$, the correlation gap of the weighted rank function of a k-uniform matroid of size n is at most $\frac{k}{\Phi(k,n)}$.*

Proof. Again let $f(S) = \min(|S|, k)$ for $S \subseteq N = \{1, \ldots, n\}$. Assume w.l.o.g. that $w_1 \ge w_2 \ge \cdots \ge w_n$, and let $w_{n+1} = 0$ for convenience. The weighted rank function $w^*(S)$ can be written as $\sum_{i\in N}(w_i - w_{i+1}) \cdot f(S \cap \{1, \ldots, i\})$, a conic combination of unweighted rank functions. The correlation gap of w^* is therefore witnessed by the correlation gap of $f(S \cap \{1, \ldots, i\})$ for some i, and hence it equals $\sup_{1\le i\le n} k/\Phi(i,k)$. By Lemma 6.14(b), $\Phi(i,k)$ is decreasing in i, and hence the correlation gap of w^* is $\frac{k}{\Phi(k,n)}$. □

Remark 6.17. We cannot generalize Lemma 6.15 or 6.16 to work for arbitrary matroids with rank k. For any k, consider the partition matroids with k parts, each of size n, where a feasible set can only have at most one element from each part. The rank of such a matroid is k, while the correlation gap is the same as that of a 1-uniform matroid over n elements, which approaches $e/(e-1)$ as n increases.

6.4.3 p-Independent Environments

There are interesting auction constraints that cannot be modeled by matroids, but can be modeled by p-independent set systems. In a set system $(N, \mathcal{I})$, a base of a subset $S \subseteq N$ is a maximal feasible subset of S. A set system $(N, \mathcal{I})$ is a p-independent system if for any non-empty subset S of N:

$$\frac{\text{maximum size of a base of } S}{\text{minimum size of a base of } S} \leq p.$$

For example, a matroid is 1-independent, and vice versa. The edge sets of (non-bipartite) matchings of a graph form a 2-independent system (but in general cannot be cast as the intersection of a constant number of matroids). The intersection of p matroids is p-independent. The feasible sets of agents in single-minded combinatorial auctions with bounded bundle size p form a p-independent system (see Section 2.1 for definition).

It is well known that the greedy algorithm gives a p-approximation for p-independent systems [51]. For our purpose, it suffices to prove the following lemma, by combining arguments of [17, 19].

Lemma 6.18. *The greedy algorithm verifies a correlation gap of $p + 1$ for p-independent system constraints.*

Proof. Fix marginal probabilities $\mathbf{q}$. In the dependent case, if S is drawn from a distribution $\mathcal{D}$ with marginal probabilities $\mathbf{q}$, let $\tilde{q}_i$ be the probability that i is in the optimal feasible subset of S (with arbitrary fixed tie-breaking). we can rewrite $E_{S\sim\mathcal{D}}[w^*(S)]$ as $\sum_{i\in N} \tilde{q}_i w_i$ by linearity of expectation.

Now consider the independent case, where each i is in S with probability q_i independently, which we denote by $S \sim \mathbf{q}$. Let $A = g(S)$ be the agents allocated by running the greedy algorithm on S. The expected performance of greedy is $E_{S\sim\mathcal{D}}[\sum_{i\in A} w_i]$. An equivalent way of looking at running the greedy algorithm on the random set S is the following: □

1. $A = \emptyset$

2. visit all agents $i \in N$ in decreasing order of weights:

   (a) if $A + i \in \mathcal{I}$, we check if i is in S, and add i into A if yes.

   (b) if $A + i \notin \mathcal{I}$, we ignore i.

3. output A

Proof. Let the random set U be the set of agents that are ignored by greedy. Consider the quantity $Q = E_{S\sim\mathbf{q}}[\sum_{i\in A} w_i + \sum_{i\in U} \tilde{q}_i w_i]$. For every agent i, if she is checked by greedy, she contributes $q_i w_i$ to Q. In particular, with probability q_i, i is in S, and we get weight w_i. On the other hand, if she is ignored, she contributes $\tilde{q}_i w_i$ to Q. Therefore,

$$Q = E_{S\sim\mathbf{q}}[\sum_{i\in A} w_i + \sum_{i\in U} \tilde{q}_i w_i] \geq \sum_{i\in N} \tilde{q}_i w_i = E_{S\sim\mathcal{D}}[w^*(S)].$$

Next we show that $w(A) \geq \frac{1}{p}\sum_{i\in U} \tilde{q}_i w_i$, and our theorem would follow as:

$$\begin{aligned} E_{S\sim\mathbf{q}}[\sum_{i\in greedy(S)} w_i] &= E_{S\sim\mathbf{q}}[\sum_{i\in A} w_i] \\ &\geq \frac{1}{p+1} E_{S\sim\mathcal{D}}[w^*(S)]. \end{aligned}$$

Let A contain $i_1, i_2, \ldots, i_l$ in the order of inclusion into A by greedy. Partition U into B_j's for $j = 1, \ldots, l$, where B_j is the set of agents ignored by greedy after $i_1, \ldots, i_j$ have been added into A. Therefore $w_i \leq w_{i_j}$ for $i \in B_j$. Consider the set $\{i_1, \ldots, i_j\} \cup B_1 \cup \cdots \cup B_j$. At any time step, greedy's solution set is always a maximal feasible subset of the agents visited so far. Therefore $\{i_1, \ldots, i_j\}$ is a base of $\{i_1, \ldots, i_j\} \cup B_1 \cup \cdots \cup B_j$. By the definition of p-independence, the maximal base of $\{i_1, \ldots, i_j\} \cup B_1 \cup \cdots \cup B_j$ has size at most $p \cdot j$, and it follows that $\sum_{i\in B_1\cup\cdots\cup B_j} \tilde{q}_i \leq p \cdot j$.

Now our claim $\sum_{i\in A} w_i \geq \frac{1}{p}\sum_{i\in U} \tilde{q}_i w_i$ follows from the following inequalities: (let $w_{i_{l+1}} = 0$)

$$\begin{aligned}
\sum_{i\in U} \tilde{q}_i w_i &= \sum_{1\le j\le l} \sum_{i\in B_j} \tilde{q}_i w_i \\
&\le \sum_{1\le j\le l} \sum_{i\in B_j} \tilde{q}_i w_{i_j} \\
&= \sum_{1\le j\le l} \sum_{i\in B_1\cup\cdots\cup B_j} \tilde{q}_i (w_{i_j} - w_{i_{j+1}}) \\
&\le \sum_{1\le j\le l} p\cdot j\cdot (w_{i_j} - w_{i_{j+1}}) \\
&= p\cdot \sum_{1\le j\le l} w_{i_j} = p\cdot \sum_{i\in A} w_i.
\end{aligned}$$

□

This ratio of $p+1$ is tight, up to lower order terms.

Proposition 6.19. *For any sufficiently large positive integer p, there is a p-independent set system with correlation gap at least $\frac{p}{\log p}$.*

Proof. To define the set system $(N, \mathcal{I})$, let Y be the set of all strings $a_1 a_2 \ldots a_n$ of length n over the alphabet $\{1, \ldots, n\}$. For every $i \in \{1, 2, \ldots, n\}$ and $b \in \{1, \ldots, n\}$, we denote by $[a_i = b]$ the "miniset" that contains all strings from Y with the i-th letter a_i being b. Then N is the set of all such minisets. To define the feasible subsets $\mathcal{I}$, a subset S of minsets from N is feasible if and only if no two minisets in S intersect. Note that two different minisets $[a_i = b]$ and $[a_{i'} = b']$ intersect if and only if $i \neq i'$. It is easy to verify that this set system is n-independent. Finally, we assign unit weights to every miniset.

We choose a random subset S of N in two ways. In the dependent case, an index i from $\{1, \ldots, n\}$ is chosen at random, and S contains the miniset $[a_i = b]$ for all $b \in \{1, \ldots, n\}$. Clearly all such S's are feasible, and the rank function has expected value n.

In the independent case, for all i, b, we include every miniset $[a_i = b]$ in S with probability $1/n$ independently. For all i, let X_i be the number of minisets in S that have the form $[a_i = b]$ for some b. Then the rank function is equal to $\max_i X_i$. To give a rough estimate of $E[\max_i X_i]$, note that for all i,

$$\begin{aligned}
& Pr[X_i \geq \frac{1}{2}\log n] \\
= & \sum_{k=\frac{1}{2}\log n}^{n} \binom{n}{k} \left(\frac{1}{n}\right)^k \left(1 - \frac{1}{n}\right)^{n-k} \\
\leq & \sum_{k=\frac{1}{2}\log n}^{n} \left(\frac{n \cdot e}{k}\right)^k \frac{1}{n^k} \leq n \cdot \left(\frac{e}{\frac{1}{2}\log n}\right)^{\frac{1}{2}\log n} \\
= & \frac{n}{2^{\Omega(\log n \cdot \log\log n)}}.
\end{aligned}$$

Therefore for sufficiently large n, $Pr[\max_i X_i \geq \frac{1}{2}\log n] \leq 1 - \left(1 - \frac{n}{2^{\Omega(\log n \cdot \log\log n)}}\right)^n \leq \frac{1}{n}$, and hence $E[\max_i X_i] \leq Pr[\max_i X_i \geq \frac{1}{2}\log n] \cdot n + \frac{1}{2}\log n \leq \log n$. It follows that the correlation gap is at least $\frac{n}{\log n}$ for sufficiently large n. □

6.5 Prior-Independence

In previous sections, it was crucial that we have knowledge about the prior distributions, so that we can calculate the prices to use in SPMs. In this section, we study how this assumption can be removed to achieve prior-independent approximation guarantees.

6.5.1 A Convex Program

In greedy-SPM, we need to compute the winning probabilities of the agents in Myerson's mechanism, which is potentially computationally hard. This was addressed in Chawla et al. by a sampling-based approach, which estimates the winning probabilities by repeatedly running Myerson's mechanism for sufficiently many times.

We note that the winning probabilities give a feasible solution to the following convex program, whose optimal value gives an upper bound on the revenue of Myerson's mechanism.

$$\begin{aligned} &\text{maximize} \sum_{i \in N} \overline{R}_i(q_i) \\ &\text{subject to} \\ &\sum_{i \in S} q_i \le rank(S) \quad \text{for all } S \\ &q_i \ge 0 \quad \text{for all } i \end{aligned}$$

Theorem 6.20. *For matroid environments, let q_i's be an optimal solution to the above convex program, then the greedy-SPM using q_i's in step 1 (see Definition 6.4) gives a $\frac{1}{\beta}$-approximation to optimal expected revenue, where β is the correlation gap of the weighted rank function of the matroid.*

Proof. The proof is by literally applying previous analysis, along with two observations. First, we did not really work with the optimal expected revenue as a benchmark, but used the upper-bound $\sum_{i \in N} \overline{R}_i(q_i)$ in the analysis. Second, given feasible q_i's, the optimal dependent distribution with marginals q_i's is a distribution defined over sets that are feasible in the matroid. □

Essentially, identifying q_i's that give us approximately-optimal SPMs is equivalent to solving the above convex program approximately.

I.I.D. k-Unit Auctions

The convex program is particularly simple for a k-unit auction over n bidders whose valuations are drawn i.i.d. from a distribution F. Let $\overline{R}$ be the ironed revenue curve of F. Then the convex program is:

$$\begin{aligned} &\text{maximize} \quad n \cdot \overline{R}(q) \\ &\text{subject to} \quad 0 \le nq \le k \end{aligned}$$

Note that this convex program is simply maximizing a concave function over a bounded interval, which can be done using a binary search.

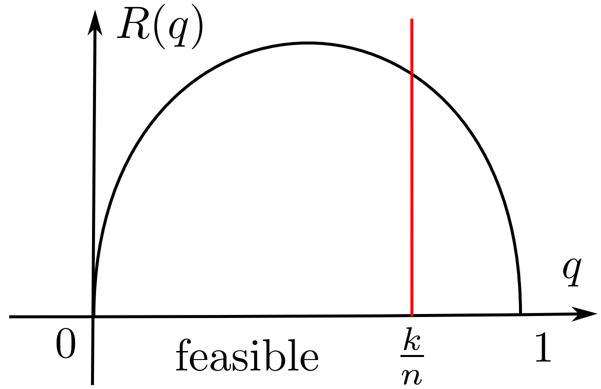

Figure 6.5.1: Convex Program for Bounding Optimal Revenue

6.5.2 Prior-Independence via Sampling

We focus on k-unit auctions over n bidders whose valuations are drawn i.i.d. from a distribution F, where F is assumed to be regular. Regularity was defined in Section 2.3, and is important for prior-independence (see Example 3.5). To obtain a prior-independent approximation guarantee, a natural idea is again to take sample bids to learn about the distribution, which is the approach we used in Chapter 4. In this section, we quantify how many samples are needed, depending on whether k is large or small compared to n.

Remark 6.21. One might argue that if we run SPM in a practical setting like Buy-It-Now in eBay, we do not get to take bid samples, as we only get binary feedbacks based on whether the bidders accept the price offers or not. Here we emphasize that we should not take our results too literally. For example, an eBay seller who runs a Buy-It-Now auction definitely does not simply set a price without any idea about the bidders' potential values. Instead, he or she will look for sources to get an estimate for bidders' valuation information. It is not possible to precisely capture how the seller obtains such information. But we can use bid sampling to roughly capture this process as part of our mechanism, which allows us to quantify the amount of information we need from the distribution to set good prices for an SPM.

6.5.3 k-Unit Auctions with Large k

First we look at the large k case, where we assume that the ratio of $\rho = \frac{k}{n}$ is $\Omega(1)$. Consider the following mechanism:

Definition 6.22 (Single-Sample SPM)**.** Given a k-unit auctions with $n \geq 2$ i.i.d. regular bidders, the Single-Sample SPM mechanism does the following:

(1) Ask the first bidder's bid, call it p.

(2) Run SPM with price p over the other bidders.

Theorem 6.23. *For k-unit auctions with n i.i.d. regular bidders, the Single-Sample SPM mechanism gives a prior-independent $\frac{\Phi(k,n-1)}{2n-2}$-approximation if $k < n$, and a prior-independent $\frac{1}{2}(1 - \frac{1}{n})$-approximation if $k = n$. In particular, for $k = \Omega(n)$, the approximation ratio is $\Omega(1)$.*

Proof. For the case of $k = n$, the auction environment is the same as a digital goods auction, where Single-Sample SPM is the same as the Single Sample mechanism in Chapter 4 for digital goods auctions, where we lose a factor of $\frac{1}{2}$ due to the use of a random price instead of the optimal price, and a factor of $1 - \frac{1}{n}$ due to the loss of the sample bidder. We let $k < n$ in the rest of the proof.

Let $q = 1 - F(p)$, which is distributed according to the uniform distribution over $[0, 1]$. Then essentially we are using this random q as a candidate solution to the following convex program:

- maximize $(n-1)R(q)$ subject to $0 \leq q \leq \frac{k}{n-1}$

This q may not be feasible. But when it is feasible, the objective value we get for the convex program is $(n-1)R(q)$. Therefore the expected objective we get from a feasible q is $E_q[(n-1)R(q) \cdot 1_{q \leq \frac{k}{n-1}}]$, which corresponds to $n-1$ times the area of the shaded region in Figure 6.5.2. Let q^* be the optimal solution to the convex program, which corresponds to the height of the shaded region. Then by concavity of the $R(\cdot)$ function, this shade area is at least half of the base $\frac{k}{n-1}$ times the height $R(q^*)$. It follows that a random q gives at least $\frac{k}{2(n-1)}$ fraction of optimal objective value. Our

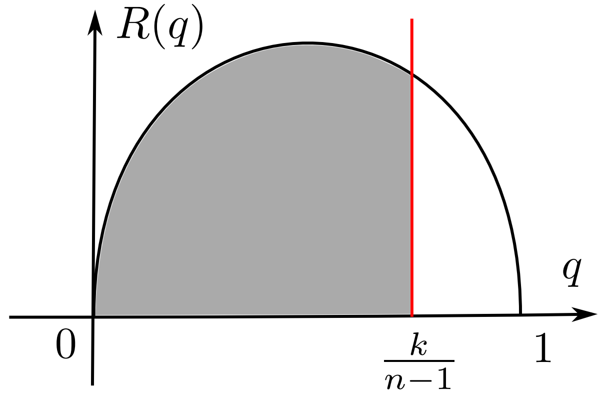

Figure 6.5.2: Random q As A Solution to the Convex Program

theorem follows by applying the fact that the expected revenue of SPM at price p is at least $\frac{\Phi(k,n-1)}{k}$ fraction of the optimal objective value of the program, which in turn upper-bounds the optimal expected revenue. □

In other words, when k is large compared to n, a single sample from the distribution is sufficient enough information to achieve approximately optimal revenue.

6.5.4 k-Unit Auctions with Small k

When k is small, the following example shows that Single-Sample SPM does not give a prior-independent constant factor approximation.

Example 6.24. Consider a single-item auction over n bidders, where the distribution F is the equal-revenue distribution with parameter n. I.e., $Pr_{v\sim F}[v \geq x] = \frac{1}{x}$ for $x \in [1, n)$ and $Pr_{v\sim F}[v = n] = \frac{1}{n}$. The optimal mechanism is the Vickrey auction with reserve n. Its expected revenue is at least $(1 - \frac{1}{e}) \cdot n$, because with probability at least $1 - \frac{1}{e}$, some bidder has value n. On the other hand, consider Single-Sample SPM. Let $q = 1 - F(p)$, where p is the sample bid. With probability $\frac{1}{\sqrt{n}}$, q is at most $\frac{1}{\sqrt{n}}$, which corresponds to a price that is at most n. With probability $1 - \frac{1}{\sqrt{n}}$, q is larger than $\frac{1}{\sqrt{n}}$, which corresponds to a price of at most $\sqrt{n}$. It follows that the total expected revenue is at most $\frac{1}{\sqrt{n}} \cdot n + (1 - \frac{1}{\sqrt{n}}) \cdot \sqrt{n}$, which is a vanishingly small fraction of the optimal expected revenue.

However, if we take more samples to learn about the distribution, we can achieve a prior-independent $\Omega(1)$-approximation.

Definition 6.25 (Many-Sample SPM for $k = o(n)$)**.** For k-unit auctions with n i.i.d. regular bidders where $k = o(n)$, the Many-Sample SPM mechanism does the following:

- Take the first half of the bidders' bids $v_1, \dots, v_{\lfloor n/2 \rfloor}$ as samples.
- Run SPM with the k-th highest price $v_{(k)}$ among $v_1, \dots, v_{\lfloor n/2 \rfloor}$ over the rest of the bidders.

Theorem 6.26. *When $k = o(n)$, Many-Sample SPM gives a prior-independent $\Omega(1)$-approximation.*

Proof. Let n be even. Otherwise we ignore the last bidder in the analysis of revenue of SPM.

Consider a k-unit auction over a set of $n/2$ bidders. The optimal expected revenue of this auction is at least a constant fraction of that of the original environment. We claim that SPM with a price p that corresponds to sale probability in the range of $[\frac{k}{n}, \frac{2k}{n}]$ gives a constant factor approximation to the optimal revenue for this setting. First, by Theorem 6.1 of [26], the VCG mechanism gives a prior-independent $(1-\frac{2k}{n})$-approximation, where $1 - \frac{2k}{n} = 1 - o(1)$ as $k = o(n)$. Then notice that in the VCG mechanism, every bidder wins with probability $\frac{2k}{n}$. It follows that $\frac{2k}{n}$ is a constant factor approximately optimal solution to the upper-bounding convex program for this k-unit auction. By concavity of the objective of the convex program, any q in the range of $[\frac{k}{n}, \frac{2k}{n}]$ is also a constant factor approximately optimal solution, and our claim follows.

In the many-sample SPM mechanism, with positive constant probability, the k-th sample $v_{(k)}$ has a sale probability q that is in the range of $[\frac{k}{n}, \frac{2k}{n}]$, and it is used in the non-sample set of $n/2$ bidders. When this is the case, by the above claim, we achieve a constant factor approximation to optimal expected revenue of the original environment. $\square$

Part III

Prior-Free Mechanisms

Chapter 7

A Theory of Optimal Envy-Free Pricing

7.1 Overview

Motivation In previous chapters, we developed several techniques for designing mechanisms with prior-independent approximation guarantees[1]. However, the prior-independence approach also has its limitations.

First, the guarantee only holds in expectation. In particular, it is possible that for a valuation profile that is intuitively easy to extract good revenue from, the mechanism does poorly. It would be ideal if for every valuation profile, we can also achieve some kind of point-wise guarantee, making the guarantee more robust.

Second, we used the regularity assumption on valuation distributions. Although this is a standard assumption in auction theory, there do exist distributions that are not regular, such as the bimodal distributions. We should try to understand whether and to what extent can this assumption be relaxed.

An Approach Based on An Envy-Free Benchmark In this and the following chapters, we deal with these two limitations by designing mechanisms that

[1]In this chapter, approximation ratios will be at least 1 instead of at most 1, to be consistent with the original paper of Hartline and Yan [46].

approximate what we call the envy-free optimal revenue benchmark EFO($\mathbf{v}$), for every input profile $\mathbf{v}$. Such an approximation guarantee holds for every input profile, and is hence more robust. Moreover, we prove that a mechanism that approximates the EFO benchmark is also automatically a prior-independent constant factor approximation w.r.t. i.i.d. tail-regular distributions, for matroid environments. Here tail-regularity is a much weaker assumption than regularity. It follows that a mechanism that approximates the EFO benchmark achieves simultaneously average-case and worst-case guarantees.

Envy-Freeness The key in our approach lies in the notion of envy-freeness. The main challenge in designing revenue-maximizing mechanisms comes from the fact that truthfulness, or incentive compatibility, binds across different input valuation profiles, and we have to trade-off the performance of a mechanism over different inputs. On the other hand, the notion of envy-freeness is defined per valuation profile. In particular, an outcome is envy-free if no agent prefers the treatment of another. It follows that optimal revenue from an envy-free outcome is well-defined for every valuation profile, which is in fact how the EFO benchmark is defined. The EFO benchmark also seems to be a natural one. For example for the simple setting of digital goods auction, the benchmark corresponds to the benchmark of best single-price revenue, which has been extensively studied (see e.g., Goldberg et al. [36]).

Interestingly, envy-freeness and incentive compatibility are similar, both structurally, and revenue-wise. First of all, we give a characterization of envy-free outcomes and their optima which are structurally equivalent to the characterization of Myerson [61] of Bayesian optimal mechanisms applied to the empirical distribution given by the actual profile of agent values. In particular, the envy-free optimal outcome is a virtual-surplus maximizer. Second, for a given virtual surplus maximizer, we prove that the maximum envy-free revenue and incentive compatible revenue are closely related, roughly by a constant factor of 2.

There are several implications of the connections between envy-freeness and incentive compatibility. First, their revenue relationship implies that mechanisms that approximate the envy-free benchmark automatically give prior-independent approximation guarantees. Second, the structural similarity is very helpful in designing mechanisms that actually approximate the envy-free benchmark.

Prior-Free Mechanisms We exploit the connection between envy-freeness and incentive compatibility to design mechanisms that approximate the envy-free benchmark. For matroid environment, we use a reduction-based approach, which reduces the problem to position auctions and then to multi-unit auctions. For general downward-closed settings, we prove that the random sampling empirical Myerson mechanism gives a constant factor approximation.

7.1.1 Outline

This chapter studies the notion of envy-freeness, and its connection to incentive compatibility. In Section 7.2 we characterize envy-free outcomes and their optima. In Section 7.3 we compare incentive-compatible and envy-free revenues. In Section 7.4 we describe the EFO benchmark and study its implication to prior-independence.

Chapter 8 establishes mechanisms that approximate the EFO benchmark. In Section 8.1 we describe a reduction-based approach where we show that, matroid environments reduce to position auctions, which reduce to multi-unit auctions, which reduce (with a loss of a factor of two) to digital goods auctions. In Section 8.2 we describe a random sampling mechanism and analyze this mechanism in downward-closed environments.

7.1.2 Related Work

Our work relies on the theory of optimal auctions as defined by Myerson [61] and refined by Bulow and Roberts [16]. In particular, Myerson showed that Bayesian optimal mechanisms are virtual surplus optimizers and Bulow and Roberts show that the virtual value of an agent in this virtual surplus maximization can be viewed as the marginal revenue as given by an agent's value distribution. This connection between envy-freeness and incentive compatibility is implicit in Jackson and Kremer [50] where it is shown that they are equivalent in the limit, in contrast we show that they are structurally equivalent generally for finite cases as well.

The random sampling mechanism for digital goods was first studied by Goldberg et al. [35]. The asymptotic performance of the mechanism was given by Segal [69] and Baliga and Vohra [11] and the convergence rate was studied by Balcan et al. [10]. In contrast, Goldberg et al. [37] consider the non-asymptotic behavior of the random sampling mechanism and show that its performance is a (large) constant factor from a prior-free benchmark that in retrospect coincides with ours. Alaei et al. [4] give a nearly tight analysis that shows that the random sampling mechanism is a 4.68-approximation (the lower-bound is 4).

There have been numerous attempts to design good prior-free mechanisms for digital goods outside the random sampling paradigm. Two notable approaches include an approximate reduction to the "decision problem"[2] by Goldberg et al. [37] and an approach based on statistical estimates that are non-manipulable with high probability by Goldberg and Hartline [34]. Hartline and McGrew [43] extend the former approach and obtain an approximation factor of 3.25. Finally, Ichiba and Iwama [49] show that a convex combination of these approaches gives an approximation factor of 3.12. The 6.24-approximation we obtain is the instantiation of our reduction with the 3.12 approximation of Ichiba and Iwama [49].

The digital goods mechanisms described above were analyzed in comparison to a natural single-priced benchmark. Hartline and Roughgarden [44] suggest that approximation of a prior-free benchmark should imply approximation of the Bayesian optimal mechanism for any i.i.d. distribution. Benchmarks for which such an implication holds are well grounded in the classical theory of Bayesian optimal auctions. They propose the performance of the best, in hindsight, Bayesian optimal mechanism as a benchmark. For multi-unit auctions, they characterize this benchmark as two-priced lotteries. With this benchmark, Devanur and Hartline [22] extended the analysis of Alaei et al. [4] to multi-unit auctions. In contrast, our benchmark, the envy-free optimal revenue, can be viewed as a relaxation of the Hartline-Roughgarden benchmark

[2]Given a target profit, the decision problem is to construct a mechanism that obtains that target profit when it is attainable.

that is structurally simpler and analytically tractable in general downward-closed environments.

Subsequent to our study, Ha and Hartline [40] generalized the statistical-estimation-based approach of Goldberg and Hartline [34] to design a 32-approximation of the envy-free benchmark in downward-closed environments. A generalization of the reduction-to-the-decision-problem approach of Goldberg et al. [37] yields a 19-approximation [39].

7.2 A Theory of Optimal Envy-Free Outcomes

In this section we derive a theory of optimal envy-free outcomes. It is most natural to study envy-freeness in the context of position auctions.

7.2.1 Position Auctions

Position auctions are a generalization of multi-unit auctions that has recently been under intense scrutiny due to its application to auctions for selling advertisements on Internet search engines [72, 27].

Definition 7.1 (Position Auctions)**.** There are n agents and n positions. We assume that bidders are ordered in non-increasing order by their values. I.e., $v_1 \geq \cdots \geq v_n$. The positions have non-increasing weights $w_1 \geq \cdots \geq w_n \geq 0$, where we let $\mathbf{w} = (w_1, \ldots, w_n)$. If an agent i is assigned position j she is served with probability w_j and her value for this assignment is $v_i w_j$.

A allocation is a (possibly random) assignment of agents to positions. We let $\mathbf{x} = (x_1, \ldots, x_n)$ denote the expected allocation, where x_i is the probability of service of the i-th highest bidder. We will only talk about expected allocation that can be realized, and will not distinguish between allocation and expected allocation in the rest of the chapter.

7.2.2 Envy-Free Outcomes

Definition 7.2 (Envy-Freeness)**.** An allocation $\mathbf{x}$ with payments $\mathbf{p}$ is *envy-free* for valuation profile $\mathbf{v}$ if no agent prefers the outcome of another agent to her own. Formally,

$$\forall i, j \in \{1, \ldots, n\},\ v_i x_i - p_i \geq v_i x_j - p_j.$$

We also require individual rationality, i.e., $\forall i \in \{1, \ldots, n\},\ v_i x_i - p_i \geq 0$.

One could try to "absorb" individual rationality into the definition of envy-freeness, by introducing a dummy agent $n+1$ with $v_{n+1} = x_{n+1} = p_{n+1} = 0$. However, to avoid the possible confusion from having two types of agent, we will not take this approach. But for notational convenience, we will use $v_{n+1} = x_{n+1} = p_{n+1} = 0$. (But i still ranges from 1 to n.)

We first characterize envy-free outcomes in terms of the allocation. For a given allocation $\mathbf{x}$ there may be several pricings $\mathbf{p}$ for which the allocation is envy-free. Since our objective is revenue maximization we will characterize the $\mathbf{p}$ corresponding to $\mathbf{x}$ that gives the highest total revenue. The proof of this characterization as it is nearly identical to that of the analogous (and standard) characterization of incentive compatible mechanisms; we include it for completeness.

Definition 7.3. An allocation is *swap monotone* if the allocation probabilities have the same order as the valuations of the agents, i.e., $x_i \geq x_{i+1}$ for every agent $i \in \{1, \ldots, n\}$. (Recall agents are ordered with $v_i \geq v_{i+1}$.)

Lemma 7.4. *In position auctions, an allocation* $\mathbf{x}$ *admits payments* $\mathbf{p}$ *such that* $(\mathbf{x}, \mathbf{p})$ *is envy-free if and only if* $\mathbf{x}$ *is swap monotone. If* $\mathbf{x}$ *is swap monotone, then the maximum payments for which* $(\mathbf{x}, \mathbf{p})$ *is envy-free and individual rational satisfy, for every agent* i*,*

$$\begin{aligned} p_i &= \sum\nolimits_{j=i}^{n} (v_j - v_{j+1}) \cdot (x_i - x_{j+1}) \\ &= \sum\nolimits_{j=i}^{n} v_j \cdot (x_j - x_{j+1}). \end{aligned}$$

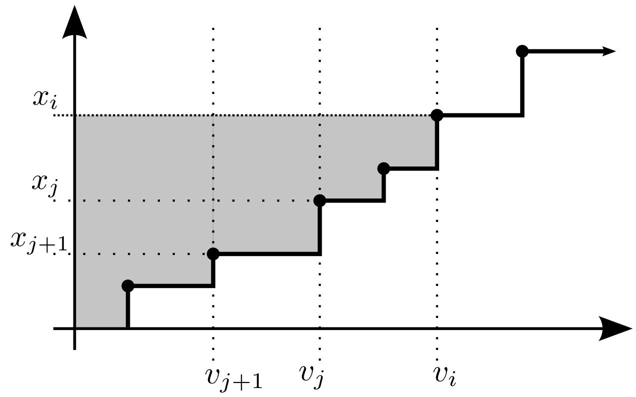

The solid curve depicts a swap monotone allocation as a function of the values (points). The shaded area corresponds to agent i's payment p_i from Lemma 7.4.

Figure 7.2.1: Reading Off Payment from Allocation Curve

See Figure 7.2.1 for a graphical exposition of the lemma.

Proof. In one direction, suppose $\mathbf{x}$ admits $\mathbf{p}$ such that $(\mathbf{x}, \mathbf{p})$ is envy-free, and we prove that $\mathbf{x}$ is swap monotone. By definition, $v_i x_i - p_i \geq v_i x_j - p_j$ and $v_j x_j - p_j \geq v_j x_i - p_i$. By summing these two inequalities and rearranging, $(x_i - x_j) \cdot (v_i - v_j) \geq 0$, and hence $\mathbf{x}$ is swap monotone.

In the other direction, suppose $\mathbf{x}$ is swap monotone. Let $\mathbf{p}$ be given as in the lemma. We verify that $(\mathbf{x}, \mathbf{p})$ is envy-free. There are two cases: if $i \leq j$, we have:

$$p_i - p_j = \sum_{k=i}^{j-1} v_k \cdot (x_k - x_{k+1}) \leq v_i \cdot \sum_{k=i}^{j-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j),$$

and if $i \geq j$, we have:

$$p_i - p_j = -\sum_{k=j}^{i-1} v_k \cdot (x_k - x_{k+1}) \leq -v_i \cdot \sum_{k=j}^{i-1} (x_k - x_{k+1}) = v_i \cdot (x_i - x_j).$$

Rearranging the results of these calculations we have the definition of envy-freeness.

Any envy-free prices satisfy that: $p_i \leq v_i(x_i - x_{i+1}) + p_{i+1}$ as agent i does not envy agent $i+1$ for $i \in \{1, \ldots, n-1\}$, and $p_n \leq v_n x_n$ by individually rationality. The payment identity above happens to make these inequalities equal. It follows that it gives the maximum payments.. □

Remark 7.5. A similar characterization readily exists for minimum envy-free revenue, which might be also of independent interest.

Importantly, the above characterization leaves us free to speak of the (maximum) envy-free revenue of any swap monotone allocation $\mathbf{x}$ on values $\mathbf{v}$, which we denote by $\mathrm{EF}^{\mathbf{x}}(\mathbf{v})$. For any $\mathbf{v}$ and any symmetric environment we will now solve for the envy-free optimal revenue, denoted by $\mathrm{EFO}(\mathbf{v})$.

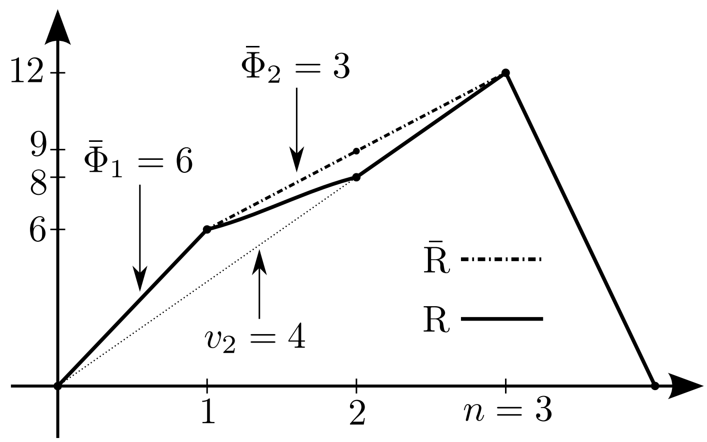

R and $\bar{\mathrm{R}}$ are the revenue curve and ironed revenue curve of the valuation profile $(6, 4, 4)$. The ironed virtual value of the high-value agent is 6, and the ironed virtual value of the two low-value agents are both $(12 - 6)/2 = 3$. E.g., the optimal EF revenue in the $k = 2$ unit auction is $\bar{\mathrm{R}}(2) = 9$.

Figure 7.2.2: An Example of Ironing

We will characterize the envy-free optimal revenue in terms of properties of the valuation profile $\mathbf{v}$. Given a valuation profile $\mathbf{v}$ we denote the *revenue curve* by $\mathrm{R}^{\mathbf{v}}(i) = i \cdot v_i$ for $i = \{1, \ldots, n\}$ (recall v_i's are indexed in decreasing order). For convenience we also let $\mathrm{R}^{\mathbf{v}}(0) = \mathrm{R}^{\mathbf{v}}(n+1) = 0$. The *ironed revenue curve*, denoted $\bar{\mathrm{R}}^{\mathbf{v}}(i)$, is the minimum concave function that upper-bounds $\mathrm{R}^{\mathbf{v}}$. Likewise, define the *virtual valuation function* $\Phi^{\mathbf{v}}(v) = \mathrm{R}^{\mathbf{v}}(i) - \mathrm{R}^{\mathbf{v}}(i-1)$ and the *ironed virtual valuation function* $\bar{\Phi}^{\mathbf{v}}(v) = \bar{\mathrm{R}}^{\mathbf{v}}(i) - \bar{\mathrm{R}}^{\mathbf{v}}(i-1)$, where $i \in \{1, \ldots, n+1\}$ is such that $v \in [v_i, v_{i-1})$. (We set $v_0 = \infty$ for notational convenience.) See Figure 7.2.2.

$\mathrm{R}^{\mathbf{v}}(i)$ is the best envy-free revenue one can get from serving exactly i agents at the same price deterministically. Consider a 2-unit auction example with one high-value agent with value 6 and two low-value agents with value 4. It is envy-free to serve one high-value agent and one low-value agent at price 4, achieving revenue $\mathrm{R}(2) = 8$. Interestingly, this is not optimal. The following allocation and payments are also envy-free: serve the high-value agent with probability 1 at price 5, and serve a low-value agent chosen at random at price 4. Both units are always sold and the total revenue is $\bar{\mathrm{R}}(2) = 9$. In what follows we will derive that this revenue is optimal among all envy-free outcomes.

Lemma 7.6. *The (maximum) envy-free revenue of a swap monotone allocation* $\mathbf{x}$ *satisfies:*

$$\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) = \sum\nolimits_{i=1}^{n} \mathrm{R}^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}) = \sum\nolimits_{i=1}^{n} \Phi^{\mathbf{v}}(v_i) \cdot x_i.$$

In a sense, the lemma states that two ways of accounting for revenue are equivalent. In one way, for each agent i, with probability $x_i - x_{i+1}$, we allow the the top i agents to win, extracting an envy-free revenue of $\mathrm{R}^{\mathbf{v}}(i)$. In another way, for each agent i, the virtual value of the i-th agent $\Phi^{\mathbf{v}}(v_i) = \mathrm{R}^{\mathbf{v}}(i) - \mathrm{R}^{\mathbf{v}}(i-1)$ can be seen as the marginal revenue contribution from allocating to agent i (in addition to agents $1, \ldots, i-1$).

Proof. The formal proof is by the following equalities:

$$\begin{aligned}
\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum\nolimits_{i=1}^{n} p_i = \sum\nolimits_{i=1}^{n} \sum\nolimits_{j=i}^{n} v_j \cdot (x_j - x_{j+1}) \\
&= \sum\nolimits_{i=1}^{n} i v_i \cdot (x_i - x_{i+1}) = \sum\nolimits_{i=1}^{n} \mathrm{R}(i) \cdot (x_i - x_{i+1}) \\
&= \sum\nolimits_{i=1}^{n} (\mathrm{R}(i) - \mathrm{R}(i-1)) \cdot x_i = \sum\nolimits_{i=1}^{n} \Phi^{\mathbf{v}}(v_i) \cdot x_i. \square
\end{aligned}$$

An implication of the characterization of the envy-free revenue of a pricing as its *virtual surplus*, i.e., $\sum_i \Phi(v_i) x_i$, suggests that to maximize revenue, the allocation should maximize virtual surplus subject to swap monotonicity (and feasibility). For monotone virtual valuation functions, the maximization of virtual surplus results in a swap monotone allocation. In general, the allocation that maximizes *ironed* virtual surplus is both swap monotone and revenue optimal among all swap monotone allocations.

Lemma 7.7. *In a position auction, the allocation that maximizes ironed virtual surplus with ties broken randomly is swap monotone.*

Proof. If $\bar{\Phi}(v_i) > \bar{\Phi}(v_j)$ then $x_i \geq x_j$; otherwise, swapping x_i for x_j would have higher ironed virtual surplus. If $\bar{\Phi}(v_i) = \bar{\Phi}(v_j)$, then $x_i = x_j$ because of random tie breaking. □

Theorem 7.8. *For any valuation profile* $\mathbf{v}$*, the allocation* $\mathbf{x}$ *that maximizes ironed virtual surplus w.r.t.* $\bar{\Phi}^{\mathbf{v}}$ *maximizes envy-free revenue among all swap-monotone allocations. I.e.,* $\mathrm{EFO}(\mathbf{v}) = \mathrm{EF}^{\mathbf{x}}(\mathbf{v})$.

This theorem is proved by a useful lemma that relates revenue to ironed virtual surplus.

Lemma 7.9. *For any swap-monotone allocation* $\mathbf{x}$ *on valuation profile* $\mathbf{v}$*,*

$$\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) \leq \sum\nolimits_{i=1}^{n} \bar{\Phi}^{\mathbf{v}}(v_i) \cdot x_i = \sum\nolimits_{i=1}^{n} \bar{\mathrm{R}}^{\mathbf{v}}(i) \cdot (x_i - x_{i+1}),$$

with equality holding if and only if $x_i = x_{i+1}$ *whenever* $\bar{\mathrm{R}}^{\mathbf{v}}(i) > \mathrm{R}^{\mathbf{v}}(i)$.

Proof. To show the inequality, we have:

$$\begin{aligned}
\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) &= \sum\nolimits_{i=1}^{n} \mathrm{R}(i) \cdot (x_i - x_{i+1}) \\
&= \sum\nolimits_{i=1}^{n} \bar{\mathrm{R}}(i) \cdot (x_i - x_{i+1}) \\
&\quad - \sum\nolimits_{i=1}^{n} (\bar{\mathrm{R}}(i) - \mathrm{R}(i)) \cdot (x_i - x_{i+1}) \\
&\leq \sum\nolimits_{i=1}^{n} \bar{\mathrm{R}}(i) \cdot (x_i - x_{i+1}),
\end{aligned}$$

where we use the fact that $\bar{\mathrm{R}}(i) \geq \mathrm{R}(i)$ and $x_i \geq x_{i+1}$. Clearly the equality holds if and only if $x_i = x_{i+1}$ whenever $\bar{\mathrm{R}}(i) > \mathrm{R}(i)$. □

Proof of Theorem 7.8. Consider $\mathbf{x}$ that optimizes ironed virtual surplus with random tie breaking and also consider any other swap monotone $\mathbf{x}'$. Note that whenever $\bar{\mathrm{R}}(i) > \mathrm{R}(i)$, we have $\bar{\Phi}^{\mathbf{v}}(v_i) = \bar{\Phi}^{\mathbf{v}}(v_{i+1})$ for which random tie-breaking implies $x_i = x_{i+1}$. Therefore $\mathbf{x}$ satisfies Lemma 7.9 with equality and it is optimal for the summation of the equality, whereas $\mathbf{x}'$ satisfies it with inequality and may not be optimal for the summation. Thus $\mathrm{EF}^{\mathbf{x}}(\mathbf{v}) \geq \mathrm{EF}^{\mathbf{x}'}(\mathbf{v})$ and $\mathbf{x}$ is revenue optimal. □

As an example of this theorem, consider the position auction environment with weights $w_1 \geq w_2 \geq \ldots \geq w_n$. An ironed virtual surplus maximizer assigns agents with higher ironed virtual values to positions with larger weights, breaking ties randomly, ignoring agents with negative ironed virtual values. The ironed virtual surplus, and thus revenue, is $\sum_{\{i\,:\,\bar{\Phi}(v_i)\geq 0\}} \bar{\Phi}(v_i) \cdot w_i$, which can be read off the revenue curve, e.g., Figure 7.2.2.

Importantly, ironed virtual surplus maximization for position auctions is ordinal, i.e., only the order of the ironed virtual values matters. The optimal envy-free outcome can then rephrased as follows: First, tentatively assign the agents to slots in order of their values. Second, randomly permute the order of each group of agents with equal ironed virtual surplus. In Section 8.1 we will discuss consequences for environments for which surplus maximization is ordinal.

7.3 Incentive Compatibility vs. Envy-Freeness

7.3.1 Permutation Environments

In this section, we aim to connect envy-freeness to incentive compatibility in downward-closed environments. However, envy-freeness only makes sense for environment like position auctions where agents' roles are symmetric, while downward-closed environments are in general not symmetric. To study both envy-freeness and incentive compatibility on a common ground, we introduce the class of permutation environments, which is a symmetric version of downward-closed environments. Another motivation for permutation environments is that guarantee in permutation environments extends to i.i.d. environments, which we will elaborate on in Section 7.4.

Definition 7.10 (Permutation Environments)**.** Given a downward-closed environment with agent set $N = \{1, \ldots, n\}$ and feasible sets $\mathcal{I}$, in the corresponding permutation environment, a permutation π of agents is drawn uniformly at random, and the feasible sets are given by $\mathcal{X}_\pi = \{\{\pi(i) \ : \ i \in S\} \ : \ S \in \mathcal{I}\}$.

In particular, given a truthful mechanism for an original environment, in the permutation environment, we first randomly permute the roles of agents, and then execute the mechanism.

Next we discuss what envy-freeness and incentive compatibility mean in permutation environments.

Envy-freeness An allocation specifies for each permutation, which agents get served, while a payment profile specifies for each permutation, how much each agent pays. However for the purpose of envy-freeness, we only care about allocation and payment in expectation over the random permutations. We let $\mathbf{x} = (x_1, \ldots, x_n)$ denote the expected allocation, where x_i denotes the probability of service of the i-th highest bidder, and let $\mathbf{p} = (p_1, \ldots, p_n)$ denote the expected payments, with expectation over random permutations. An outcome that consists of an allocation and a payment profile is envy-free if:

$$\forall i, j \in \{1, \ldots, n\}, \ v_i x_i - p_i \geq v_i x_j - p_j,$$

and is individual rational if:

$$\forall i, j \in \{1, \ldots, n\}, \ v_i x_i - p_i \geq 0.$$

The characterization of envy-free outcomes and their optima extends straightforwardly to permutation environments.

Incentive Compatibility The definition of a truthful mechanism is the same as before. We repeat here for the sake of comparison.

A mechanism is given by an allocation rule $\mathbf{x}(\mathbf{v})$ and a payment rule $\mathbf{p}(\mathbf{v})$. Note that here $\mathbf{v}$ is not ordered. A mechanism is *incentive compatible* or truthful if no agent prefers the outcome when misreporting her value to the outcome when reporting the truth. Formally,

$$\forall i, z, \mathbf{v}, \quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i x_i(z, \mathbf{v}_{-i}) - p_i(z, \mathbf{v}_{-i}),$$

where $(z, \mathbf{v}_{-i})$ is obtained from $\mathbf{v}$ with v_i replaced by z.

A mechanism is individual rational if

$$\forall i, z, \mathbf{v}, \quad v_i x_i(\mathbf{v}) - p_i(\mathbf{v}) \geq 0.$$

We recall the well-known characterization of IC mechanisms and their revenue as follows.

Definition 7.11 (Value Monotonicity)**.** An allocation rule is *value monotone* if the probability that an agent is served is monotone non-decreasing in her value, i.e., $x_i(z, \mathbf{v}_{-i})$ is non-decreasing in z for all agents i.

Theorem 7.12. *(Myerson [61]) An allocation rule* $\mathbf{x}(\cdot)$ *admits a payment rule* $\mathbf{p}(\cdot)$ *such that* $(\mathbf{x}, \mathbf{p})$ *is incentive compatible if and only if* $\mathbf{x}(\cdot)$ *is value monotone, and the non-negative and individual rational payment rule is uniquely determined by:*

$$p_i(\mathbf{v}) = v_i x_i(\mathbf{v}) - \int_0^{v_i} x_i(z, \mathbf{v}_{-i}) dz.$$

An Expected View of Truthful Mechanisms Note that the envy-freeness constraint is imposed loosely in expectation over permutations, while incentive compatibility is imposed for every random permutation that assigns agents to roles in the set system. For the purpose of revenue study in downward-closed permutation environments, it is more convenient to look at the expected behavior of a truthful mechanism in expectation over random permutations. In particular, for a valuation profile $\mathbf{v}$, we will redefine $\mathbf{x}(\mathbf{v}) = (x_1(\mathbf{v}), \ldots, x_n(\mathbf{v}))$, where $x_i(\mathbf{v})$ is the probability that the i-th highest agent gets allocated in permutation environment, and $\mathbf{p}(\mathbf{v}) = (p_1(\mathbf{v}), \ldots, p_n(\mathbf{v}))$, where $p_i(\mathbf{v})$ is the expected payment that the i-th highest agent pays over random permutations.

The same revenue formula in Theorem 7.12 applies[3]. Because the payments are uniquely determined by the allocation rule, for any $\mathbf{x}(\cdot)$, we let $\mathrm{IC}^{\mathbf{x}}(\mathbf{v})$ denote the IC revenue from running $\mathbf{x}(\cdot)$ over $\mathbf{v}$, where $\mathbf{v}$ is w.l.o.g. assumed to be ordered.

For brevity, we will simply call such $\mathbf{x}(\cdot), \mathbf{p}(\cdot)$ as allocation rule and payment rule, with the understanding that they only describe the expected behavior of an underlying truthful mechanism.

7.3.2 Revenue Comparison

We now compare envy-free revenue to incentive-compatible revenue for ironed virtual surplus optimizers in permutation environments. We show that these quantities are often within a factor of two of each other.

In the following we use $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ and $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ to denote the IC and EF revenue from agent i by applying the ironed virtual surplus maximizer for $\bar{\Phi}$, respectively.

First we lower-bound $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ by half of $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$, under a technical condition on $\bar{\Phi}$ and $\mathbf{v}$.

Lemma 7.13. *For every downward-closed permutation environment, every valuation profile* $\mathbf{v}$*, and* $\bar{\Phi}$*, the ironed virtual valuation function corresponding to some* $\mathbf{v}'$ *obtained from* $\mathbf{v}$ *by setting a subset of agents' values to be 0, we have that* $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v}) \geq \frac{1}{2}\,\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ *for all* i*.*

Proof. Let $\mathbf{x}(\cdot)$ denote the allocation rule of the ironed virtual surplus optimizer for $\bar{\Phi}$. By the assumption of the lemma, for all j, $\bar{\Phi}(z)$ is constant for all $z \in [v_{j+1}, v_j)$, and hence the IC allocation rule in fact maps each $z \in [v_{j+1}, v_j)$ to $x_i(v_{j+1}, \mathbf{v}_{-i})$.

By Lemma 7.12, $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$ which, referring to Figure 7.3.1, equals the area above the IC curve and below the horizontal dotted line. On the other hand, $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v})$ is equal to $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$, which similarly corresponds to the area above the EF curve and below the horizontal dotted line. It suffices to prove that: $x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}) \geq \frac{1}{2} \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$. This is equivalent to $x_i(\mathbf{v}) + x_{j+1}(\mathbf{v}) \geq 2x_i(v_{j+1}, \mathbf{v}_{-i})$.

[3]Here we assume that $x_i(z, \mathbf{v}_{-i})$ refers to the allocation probability of the agent with value z, not the bidder with the i-th highest value among profile $(z, \mathbf{v}_{-i})$.

Here the right-hand side can be interpreted as the total winning probability of agents i and $j+1$ after agent i drops her value. So the last inequality says that the total winning probability of agent i and $j+1$ can only decrease if agent i lowers her bid to v_{j+1}. To prove this, we fix the permutation that maps agents to roles of the set system, and show that the number of winning agents from i and $j+1$ can only be lower after agent i decreases her value. There are two cases to verify: (1) Agent i wins after the decrease. Then before the decrease, agent i had higher value, and the optimal feasible set would be the same. (2) Agent $j+1$ wins and agent i loses after the decrease. Then before the decrease, at least one of agents i and $j+1$ would win. □

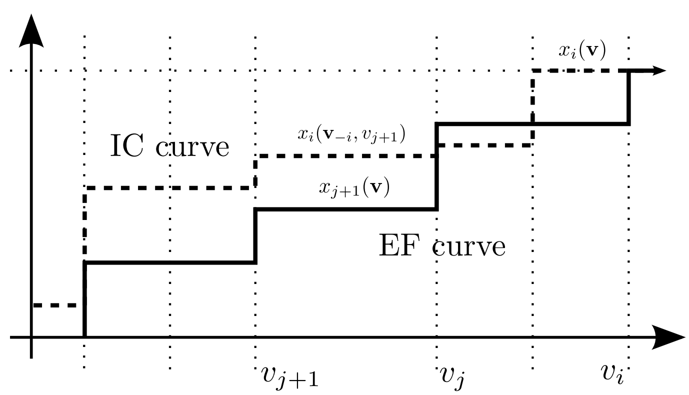

Depiction of EF allocation and IC allocation rule from which the payments for agent i are computed. The EF allocation curve maps each value in $[v_{j+1}, v_j)$ to $x_{j+1}(\mathbf{v})$, and the IC allocation curve maps each z to $x_i(z, \mathbf{v}_{-i})$.

Figure 7.3.1: Comparison of EF Allocation Curve and IC Allocation Curve

In matroid permutation environments, envy-free revenue upper-bounds incentive-compatible revenue.

Lemma 7.14. *For every matroid permutation environment, valuation profile* $\mathbf{v}$*, and ironed virtual valuation function* $\bar{\Phi}$*, for all agent* i*,* $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v}) \geq \mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$.

Proof. Recall that $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v}) = \sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_{j+1}(\mathbf{v}))$ and $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v}) = \int_0^{v_i}(x_i(\mathbf{v}) - x_i(z, \mathbf{v}_{-i}))dz$. By the monotonicity of $x_i(z, \mathbf{v}_{-i})$ in z, $\mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ is upper-bounded by $\sum_{j=i}^{n}(v_j - v_{j+1}) \cdot (x_i(\mathbf{v}) - x_i(v_{j+1}, \mathbf{v}_{-i}))$. It suffices to prove that

$x_{j+1}(\mathbf{v}) \le x_i(v_{j+1}, \mathbf{v}_{-i})$. Here the right hand side can be also seen as the winning probability of agent $j+1$ after agent i drops value. To prove this inequality, ironed virtual surplus maximizers are greedy algorithms in matroid permutation settings, and if agent i decreases her bid to v_{j+1}, agent $j+1$ is less likely to be blocked by i who was earlier in the greedy order, and is hence more likely to be allocated. □

There are downward-closed permutation environments where the envy-free optimal revenue does not upper-bound the incentive-compatible revenue of all virtual surplus maximizers.

Lemma 7.15. *There exists a downward-closed permutation environment and valuation profile* $\mathbf{v}$*, such that if* $\bar{\Phi} = \bar{\Phi}^{\mathbf{v}}$ *is the ironed virtual valuation function of* $\mathbf{v}$*, then* $\mathrm{IC}^{\bar{\Phi}}(\mathbf{v}) > \mathrm{EF}^{\bar{\Phi}}(\mathbf{v})$.

Proof. Let there be $n+1$ agents. The "1 vs n" set system has two maximum feasible sets, one is a singleton set and the other one has size n. These two sets are disjoint. We define the valuation profile by specifying the virtual valuations. There are n "small" agents with virtual values $v+\epsilon, v+2\epsilon, \ldots, v+n\epsilon$ respectively, and one "big" agent with virtual value $nv + \frac{n(n+1)}{2}\epsilon - \epsilon^2$ for some small positive ϵ. The choice of the ϵ terms is such that for the sum of the virtual valuations of the first n agents to beat the big agent, no small agent can lower her virtual value to some other agent's virtual value. We will ignore ϵ terms from now on. Correspondingly, one can calculate the revenue curve, and then derive the valuations of the agents: the valuation of the big agent is nv, and the small agents have values $\frac{n+1}{2}v$, $\frac{n+2}{3}v$,..., $\frac{2n}{n+1}v$, ignoring ϵ terms. The allocation rule is the ironed virtual surplus optimizer w.r.t. this valuation profile. Note that a reserve of $\frac{2n}{n+1}v$ is set because any value lower than this corresponds to a negative ironed virtual value.

Observe that every agent wins if and only if she is assigned to the size n set, which happens with probability $n/(n+1)$. Therefore the EF revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot (n+1) = \frac{2n^2}{n+1}v$. To calculate the IC revenue, with probability $n/(n+1)$, the big agent is assigned to the size n set, and every of the n winning agents pays the reserve $\frac{2n}{n+1}v$. Also with probability $1/(n+1)$, the big agent is assigned to the singleton set, and every agent has to pay her own value, which sums up to $\Theta(nv\log(n))$. Therefore the IC revenue is $\frac{2n}{n+1}v \cdot \frac{n}{n+1} \cdot n + \frac{1}{n+1}\Theta(nv\log(n))$, which is larger than EF revenue for sufficiently large n. □

7.4 The Envy-Free Optimal Revenue Benchmark

As discussed previously, no incentive-compatible mechanism obtains an optimal profit point-wise on all possible valuation profiles. Therefore, to obtain point-wise guarantees, the literature on prior-free mechanism design looks for the incentive compatible mechanism that minimizes, over valuation profiles, its worst-case ratio to a given performance benchmark. It is important to identify a good benchmark for such an analysis to be meaningful.

If the designer had a prior distribution over the agent valuations then she could design the mechanism that maximizes revenue in expectation over this distribution. This is the approach of Bayesian optimal mechanism design as characterized by Myerson [61] and refined by Bulow and Roberts [16]. Given a distribution F, virtual values and revenue curves can be derived. The optimal mechanism is the one that maximizes ironed virtual surplus.

Theorem 1. *(Myerson [61]) When values are i.i.d. from distribution F the optimal mechanism,* ICO^F*, is the ironed virtual surplus optimizer for $\bar{\Phi}$ corresponding to F.*

If the agent values are indeed drawn from a prior distribution, but the designer is unaware of the distribution, then a reasonable objective might be to design a mechanism that is a good approximation to the optimal mechanism for any unknown distribution that satisfies some natural assumption. This is the prior-independent objective.

One important criterion for a prior-free benchmark is that its approximation should imply prior-independent approximation: if a mechanism is a constant approximation to the benchmark, then for a relevant class of distributions, it should be a constant approximation to the Bayesian optimal mechanism under any distribution from the class.

For matroid permutation environments, Lemma 7.14 implies that for any values $\mathbf{v}$ the optimal envy-free revenue $\mathrm{EFO}(\mathbf{v})$ (which is at least the envy-free revenue of any ironed virtual surplus optimizer) is at least the incentive compatible revenue of any ironed virtual surplus optimizer. By Theorem 1, the Bayesian optimal mechanism is an ironed virtual surplus optimizer so $\mathrm{EFO}(\mathbf{v})$ upper-bounds its revenue. Consequently, a prior-free β-approximation to EFO is also a prior-independent β-approximation for all distributions.

Unfortunately, even for simple the digital goods auctions it is not possible to obtain a prior-free constant approximation to the EFO benchmark, as shown by Goldberg et al. [36]. This impossibility arises because it is not possible to approximate the highest value v_1. For essentially the same reason, it is not possible to design a prior-independent constant approximation for all distributions. We therefore restrict attention to the large family of distributions with tails that are not too irregular.

Definition 7.16 (Tail Regularity)**.** A distribution F is *n-tail regular* if in a single-item auction over n agents with values drawn i.i.d. from F, the expected revenue of the Vickrey auction is a 2-approximation to that of the optimal mechanism.

The definition of tail regularity is implied by Myerson's regularity assumption via the main theorem of Bulow and Klemperer [15]. The intuition for the definition is the following. For single-item auctions over n agents, most of the actions happen in the tail of the distribution, i.e, values v for which $F(v) \approx 1 - O(1)/n$; therefore, irregularity of the rest of the distribution does not have much consequence on revenue. Tail regularity, then, restates the Bulow-Klemperer consequence, as a constraint on the tail of the distribution and leaves the rest unconstrained.

We now define the benchmark for prior-free mechanism design. Approximation of this benchmark guarantees prior-independent approximation of all n-tail-regular distributions.

Definition 7.17. The *envy-free benchmark* is $\mathrm{EFO}^{(2)}(\mathbf{v}) = \mathrm{EFO}(\mathbf{v}^{(2)})$ where $\mathbf{v}^{(2)} = (v_2, v_2, v_3, \ldots, v_n)$.

Theorem 7.18. *For every n-agent matroid permutation environment, n-tail-regular distribution F, and β-approximation mechanism to* $\mathrm{EFO}^{(2)}$*, the expected revenue of* *the mechanism with valuations* $\mathbf{v}$ *drawn i.i.d. from* F *is a* 3β*-approximation to the optimal mechanism for* F*.*

Proof. By the reduction from matroid permutation environments to multi-unit auctions in Section 8.1, it is sufficient to prove the statement for k-unit auctions. Let $\bar{\Phi}$ be the ironed virtual surplus maximizer for F. We first upper-bound IC revenue from bidders 2 to n:

$$\begin{aligned}\mathrm{IC}^F_{2\ldots n}(\mathbf{v}) &\leq \mathrm{EF}^{\bar{\Phi}}_{2\ldots n}(\mathbf{v})\\ &\leq \mathrm{EFO}(v_2, v_3, \ldots, v_n, 0)\\ &\leq \mathrm{EFO}(\mathbf{v}^{(2)})\\ &= \mathrm{EFO}^{(2)}(\mathbf{v})\end{aligned}$$

Here the first inequality is by Lemma 7.14, and the last equality is by definition of $\mathrm{EFO}^{(2)}$.

To see the second inequality, both left hand side and right hand side correspond to the maximum envy-free revenue from $2\ldots n$ that correspond to some outcome with at most k items allocated and no envy among $2\ldots n$, except that in the right hand side, the outcome that maximizes this revenue is chosen.

To see the third inequality, note that the revenue curve of $\mathbf{v}^{(2)}$ dominates that of $(v_2, v_3, \ldots, v_n, 0)$, and hence by Lemma 7.6, for every allocation, the envy-free revenue for $\mathbf{v}^{(2)}$ can only be higher.

Next we upper-bound IC revenue from bidder 1, where the expectation is over i.i.d. draws from distribution F.

$$\begin{aligned}E[\mathrm{IC}^{\bar{\Phi}}_1(\mathbf{v})] &\leq E[\mathrm{IC}^{\bar{\Phi}}(\mathbf{v}) \text{ for single item auction}]\\ &\leq 2E[v_2]\\ &\leq 2E[\mathrm{EFO}^{(2)}(\mathbf{v})]\end{aligned}$$

Here the second inequality is by the tail regularity assumption. The third inequality is because $\mathrm{EFO}^{(2)}(\mathbf{v}) \geq v_2$.

To see the first inequality, consider the mechanism that first runs $\bar{\Phi}$, and then only allows the highest bidder (bidder 1) to win. The IC payment of bidder 1 is the maximum of the second highest bid and the threshold bid for bidder 1 to win in $\bar{\Phi}$. This is as much as the threshold (or revenue) from bidder 1 in $\bar{\Phi}$ as in the left hand side. On the other hand, the IC revenue of this mechanism is at most that of the optimal single-item auction, which equals to the right hand side.

Together, we have that $E[\mathrm{IC}^{\bar{\Phi}}(\mathbf{v})] \leq 3E[\mathrm{EFO}^{(2)}(\mathbf{v})]$.

□

It is useful to compare the EFO benchmark to ones proposed in the literature that are based on the VCG mechanism with the best (for the particular valuation profile $\mathbf{v}$) reserve price, e.g., [37, 45]. The VCG mechanism with a reserve price first rejects all agents whose values to not meet the reserve, then it serves the remaining agents to maximize welfare.

For matroid permutation environments, the EFO benchmark is a stronger benchmark than VCG-with-reserve. On one hand, the VCG-with-reserve benchmark can be expressed as an ironed virtual surplus optimizer, and so by Lemma 7.14, EFO is no smaller. On the other hand, the following lemma shows that EFO can be (almost) a logarithmic factor larger than VCG-with-reserve. Therefore, the EFO-based benchmark results in stronger approximation guarantees for matroid permutation environments.

Lemma 7.19. *There exists a distribution F and an n-agent matroid environment for which VCG with any reserve price is an $\Omega(\log n/\log\log n)$-approximation to the optimal mechanism for F.*

Proof. Fix some number m.

The matroid we use is a partition matroid. In general in a partition matroid, the ground set is partitioned into a number of disjoint sets, or sectors, where each sector is associated with a capacity number. A set is feasible if and only if its intersection with each sector does not exceed the capacity number of the sector.

Now we define the partition matroid we use. For each $k \in \{1, \ldots, m\}$, a type k sector contains m^{3k-1} elements or agents and has capacity one. There are m^{2m-2k}

disjoint type k sectors for each $k \in \{1, \ldots, m\}$. So total number of agents n is at most $m^{O(m)}$. Hence m is at least of order $\frac{\log n}{\log \log n}$.

Next we define the "sydney opera house distribution", named after the zig-zag shape of the revenue function. The distribution F is such that the value is distributed according to uniform distribution $[m^{2k+1} - \epsilon, m^{2k+1} + \epsilon]$ with probability $\frac{1}{m^{3k}} - \frac{1}{m^{3k+3}}$ for $k \in \{0, \ldots, m-1\}$, and with probability $\frac{1}{m^{3k}}$ for $k = m$. Here we take ϵ to be some sufficiently negligible positive amount, and we will often omit ϵ related terms. So for each k the revenue function R at $\frac{1}{m^{3k+3}}$ has left limit $R(\frac{1}{m^{3k+3}}-) = \frac{m^{2k+3}}{m^{3k+3}} = \frac{1}{m^k}$, and right limit $R(\frac{1}{m^{3k+3}}+) = \frac{m^{2k+1}}{m^{3k+3}} = \frac{1}{m^{k+2}}$. Hence the ironed virtual valuation between quantile $\frac{1}{m^{3k+3}}$ to quantile $\frac{1}{m^{3k}}$ is $\frac{\frac{1}{m^{k-1}} - \frac{1}{m^k}}{\frac{1}{m^{3k}} - \frac{1}{m^{3k+3}}} \approx m^{2k+1}$. Note that the ironed virtual valuation is equal to valuation, ignoring minor terms.

To calculate the revenue of Myerson's mechanism, for a type k sector, there are m^{3k-1} agents. With probability at least $1 - (1 - \frac{1}{m^{3k}})^{m^{3k-1}} \approx \frac{1}{m}$, the highest agent is in quantile range $(0, \frac{1}{m^{3k}})$, with ironed virtual valuation at least m^{2k+1}. So the expected ironed virtual valuation from a type k sector is at least m^{2k}. Multiplied by the number of type k sectors the total ironed virtual valuation, and hence expected revenue is at least $\sum_k m^{2k} \cdot m^{2m-2k} = m \cdot m^{2m}$.

To calculate the revenue of VCG with some reserve r, suppose w.l.o.g. $r \approx m^{2k'+1}$ for some k'. For a type k sector with m^{3k-1} agents, the dominant amount of revenue is obtained from the following two cases:

1. When there are at least two agents with value at least $m^{2k+1} - \epsilon$ (i.e. in quantile $\frac{1}{m^{3k}}$), the lower of which has value at most $m^{2k+1} + \epsilon$. This happens with probability roughly $\frac{1}{m^2}$, and gives revenue m^{2k+1}. Therefore the expected revenue we get from this case is $\frac{1}{m^2} \cdot m^{2k+1}$, which multiplied by the number of type k sector, is $O(m^{2m-1})$.

2. When $k = k'$, and there is at least one agent who beats the reserve m^{2k+1}. This happens with probability at most $\frac{1}{m}$. Therefore the expected revenue from this case is $m^{2k+1} \cdot \frac{1}{m}$, which multiplied by the number of type k sectors is $O(m^{2m})$.

Summing over all k, the total expected revenue of VCG with reserve r is at most $O(m^{2m}) + m \cdot O(m^{2m-1}) = O(m^{2m})$, which is less than that of Myerson's mechanism by a factor of $\Omega(m) = \Omega(\log n / \log \log n)$.

□

For matroid permutation environments we have given a Bayesian justification for the envy-free benchmark $\mathrm{EFO}^{(2)}$. A mechanism that approximates this prior-free benchmark simultaneously approximates the Bayesian optimal mechanisms for most i.i.d. distributions (i.e., those satisfying the relatively unrestrictive tail-regularity assumption). We do not have a formal justification for the envy-free benchmark in downward-closed permutation environments because of Lemma 7.15. Such a justification would follow from the following conjecture:

Conjecture 7.20. *For every downward-closed permutation environment and valuation profile* $\mathbf{v}$*, and ironed virtual valuation function* $\bar{\Phi}$*, for all agent* i*,* $\mathrm{EF}_i^{\bar{\Phi}}(\mathbf{v}) \geq \rho \cdot \mathrm{IC}_i^{\bar{\Phi}}(\mathbf{v})$ *for some positive constant* ρ*.*

Chapter 8

Prior-Free Mechanisms

In this chapter, we give two approaches for designing prior-free mechanisms that approximate the envy-free revenue benchmark.

Our first approach is via reduction. A digital goods auction (a.k.a., unlimited supply) is one where the mechanism has no inter-agent feasibility constraint. A multi-unit auction (a.k.a., limited supply) is one where the mechanism has a constraint on the number of agents that can be simultaneously served, e.g., multiple units of a single item. We give a reduction from multi-unit auctions to digital goods auctions that loses at most a factor of two in approximation factor. We then give a lossless reduction from position auctions and matroid environments to multi-unit auctions. To obtain these reductions we establish a structural relationship between these environments. Given the 3.12-approximation for digital goods given by Ichiba and Iwama [49], our reduction implies 6.24-approximations for multi-unit, matroid environment, and position auctions.

Our second approach is via a generalization of the random sampling mechanism of Goldberg et al. [35] and Baliga and Vohra [11]. The mechanism takes the following form. The agents are randomly partitioned into a market and a sample. The sample is then used for market analysis and its empirical distribution is calculated. The optimal mechanism for the empirical distribution of the sample is then run on the market. This prior-free mechanism is one of the most fundamental, and we extend the analysis techniques derived for it in digital goods auctions to multi-unit auctions and (more generally) downward-closed environments. The approximation factors we obtain are 12.5 for multi-unit (and by the reduction above, constrained-matching and position auction environments) and 189 for downward-closed environments.

8.1 Multi-unit, Position, and Matroid Permutation Environments

In this section we consider matroid permutation environments, position auctions, and multi-unit auctions. We show that these environments are closely related. In fact, the optimal mechanisms that are incentive compatible (or envy-free, resp.) across these environments give the same expected allocation, and a wide class of mechanisms give the same approximation factor. As an example, we will focus on approximating the envy-free benchmark $\text{EFO}^{(2)}$ (Definition 7.17) with a prior-free mechanism. Our solution will be via a two-step reduction: we reduce matroid permutation to position auctions, which we then reduce to multi-unit auctions.

Recall that in a multi-unit auctions it is feasible to serve any set of agents of cardinality at most some given k. In position auction environments there are weights $w_1 \geq w_2 \geq \cdots \geq w_n$ for positions and feasible outcomes are partial assignments of agents to positions. In matroid permutation environments there is a feasibility constraint given by independent sets of a matroid, but the roles of the agents are assigned by random permutation.

The property of these three environments that enables this reduction is that in each environment the greedy algorithm on ironed virtual values (with ties broken randomly) obtains the maximum ironed virtual surplus. The greedy algorithm works as follows: order the agents by ironed virtual value and serve each agent in this order if her ironed virtual value is positive and if doing so is feasible given the set of agents previously served. Notice that the only information needed to perform this maximization is the ordering on the agents' ironed virtual values (but not their magnitudes).

Definition 8.1. The *characteristic weights* $w_1 \geq w_2 \geq \cdots \geq w_n$ of a matroid environment are as follows: choose any valuation profile $\mathbf{v}$ with all distinct values, assign the agents to elements in the matroid via a random permutation, run the greedy algorithm w.r.t. $\mathbf{v}$, and define w_i to be the probability that i-th largest valued agent is served.

8.1.1 Reduction for Ironed Virtual Surplus Maximizers

We first connect ironed virtual surplus optimization in the three environments.

Proposition 8.2. *The ironed virtual surplus maximizing allocations have the same expected allocation and virtual surplus in the following environments:*

1. *a matroid permutation environment with characteristic weights* $\mathbf{w}$*,*
2. *a position auction with weights* $\mathbf{w}$*,*
3. *a convex combination of multi-unit auctions where k units are available with probability $w_k - w_{k+1}$ for $k \in \{1, \ldots, n\}$ and $w_{n+1} = 0$.*

Proof. Fix a tie-breaking rule, which induces an ordering on the agents. Consider the greedy algorithm on the agents with non-negative $\bar{\Phi}$ values according to this ordering. The j-th agent with non-negative $\bar{\Phi}$ value in this ordering (1) gets allocated with probability w_j in the matroid permutation setting by definition of characteristic weights, (2) gets assigned to position j in the position auction and hence gets allocated with probability w_j, and, (3) gets allocated in the k-unit auction for each $k \geq j$, and hence has probability $\sum_{k \geq j}(w_k - w_{k+1}) = w_j$ of being served in the convex combination setting. Taking expectation over all tie-breaking orders, agent i has the same probability of being served in the three settings. □

The following corollary is immediate.

Corollary 8.3. *For every valuation profile* $\mathbf{v}$ *and weights* $\mathbf{w}$*, the envy-free optimal revenue is the same in each of the environments of Proposition 8.2.*

We now illustrate how to use Proposition 8.2 to show that an incentive compatible prior-free approximation mechanisms for multi-unit auctions can be adapted to give the same approximation factor in position auctions and matroid permutation environments. Consider the following incentive compatible mechanism.

Definition 8.4. The *Random Sampling Empirical Myerson (*RSEM*)* mechanism does the following: (discussion of payments omitted)

1. randomly partition the population of agents $N = \{1, \ldots, n\}$ into two sets by flipping a fair coin for each agent,
2. designate the set containing the highest-valued agent as the market M and the other set as the sample S,
3. calculate the ironed virtual surplus function $\bar{\Phi}^S$ for the sample S, and,
4. serve a feasible subset of M to maximize ironed virtual surplus with respect to $\bar{\Phi}^S$ and reject all other agents.

Lemma 8.5. *For every downward-closed environment, RSEM is incentive compatible.*

Proof. We verify that RSEM is monotone. An agent in S loses unless she raises her bid to beat the highest-valued agent (in which case the roles of S and M are reversed). An agent in M wins when the virtual surplus maximizing set contains the agent. If she raises her bid, she (weakly) increases her virtual value thus increasing the virtual surplus of any set containing her, while the virtual surplus of other sets remain the same. Therefore, she continues to win. By Theorem 7.12 monotonicity implies that, with the appropriate payments, RSEM is incentive compatible. □

The following theorem is from Devanur et al. [24].

Theorem 8.6. *In multi-unit auctions, RSEM is a* 12.5*-approximation to the envy-free benchmark* $\mathrm{EFO}^{(2)}(\mathbf{v})$*.*

RSEM is incentive compatible because it is essentially an ironed virtual surplus optimizer on the set M, and furthermore, it is incentive compatible even if the permutation that assigns agents to the set system is fixed. As a final corollary of Proposition 8.2, we can view RSEM's revenue in the matroid permutation or position auction environment as the analogous convex combination of its revenue in multi-unit auction environments.

Corollary 8.7. *In matroid permutation environments and position auctions, RSEM is a prior-free* 12.5*-approximation to the envy-free revenue* $\mathrm{EFO}^{(2)}(\mathbf{v})$*.*

8.1.2 General Reduction

The following prior-free approximations are essentially the best known for digital goods and multi-unit auctions. Notably, the mechanism from Corollary 8.11 below, is not based on ironed virtual surplus maximization and therefore Proposition 8.2 cannot be applied to construct a matroid permutation or position auction mechanism from it.

Lemma 8.8. *[49] In digital goods auctions, there is a prior-free incentive compatible* 3.12*-approximation to* $\mathrm{EFO}^{(2)}(\mathbf{v})$*.*

We now give an approximate reduction from multi-unit auctions to digital good auctions. This construction and the proof that the resulting mechanisms incentive compatibility are standard. See, e.g., Myerson [61], Goldberg et al. [37], and Aggarwal and Hartline [1].

Definition 8.9 (Multi-Unit Reduction)**.** Given a k-agent digital goods auction mechanism, we construct the following mechanism for k-unit auctions:

1. Simulate the k-unit Vickrey auction.
2. Simulate the mechanism for k-agent digital goods auction on the k winners of the Vickrey auction.
3. Serve the agents who win in both stages and charge them the maximum of their simulation payments; reject all other agents.

Theorem 8.10. *Given a mechanism for digital goods auctions that β-approximation to the envy-free benchmark (resp. Bayesian optimal mechanism), the mechanism for multi-unit auctions from the reduction is a 2β-approximation to the envy-free benchmark (resp. Bayesian optimal mechanism).*

Proof. The digital goods auction is a β-approximation the envy-free benchmark on the top k agents. The envy-free benchmark on the top k agents is equal to the VCG-with-reserve benchmark for the full set of agents (both are equal to $\max_{i \leq k} \mathrm{R}(i)$). Theorem 10 of Devanur and Hartline [22] essentially states that the VCG-with-reserve benchmark is a 2-approximation to the envy-free benchmark. Therefore, the multi-unit auction from the reduction is a 2β-approximation to the envy-free benchmark. □

Corollary 8.11. *In multi-unit auctions, there is an incentive compatible prior-free* 6.24*-approximation to* $\mathrm{EFO}^{(2)}(\mathbf{v})$.

We now show how to construct, from a mechanism for multi-unit auction, a position auction mechanism and a matroid permutation mechanism that has the same expected allocation as a convex combination of a series of mechanisms for multi-unit auctions (as in Proposition 8.2). The challenge here is the distinct interfaces to the environment: in multi-unit auctions we are given a supply constraint k and we need to specify a set of at most k winners, whereas in position auctions, we are given weights and need to output a partial assignment of agents to positions.

Definition 8.12 (Position Auction Reduction)**.** Given a series of k-unit auction mechanisms for $k \in \{1, \ldots, n\}$, we construct the following mechanism for the position auction environment with weights $\mathbf{w}$:

1. Introduce n dummy agents and n dummy positions into the system, indexed by $\{n+1, \ldots, 2n\}$. Correspondingly, we pad weights $\mathbf{w}$ and valuation profile $\mathbf{v}$ with zeros such that they have dimension $2n$.

2. For each $k \in \{1, \ldots, n\}$, simulate the mechanism for k-unit auctions on valuation profile $\mathbf{v}$, and give the unallocated leftover units to the dummy agents arbitrarily for free. Let the resulting allocation of all $2n$ agents be $\mathbf{x}^{(k)}$.

3. Calculate the probability that each agent is served in the convex combination: $x_i = \sum_{k=1}^{n} x_i^{(k)}(w_k - w_{k+1})$, for $i \in \{1, \ldots, 2n\}$.

4. Solve for a set of permutation matrices $P_t \in \{0,1\}^{2n\times 2n}$ and nonnegative weights r_t with $\sum_t r_t = 1$ such that $\sum_t r_t \cdot P_t \cdot \mathbf{w} = \mathbf{x}$.

5. With probability r_t, assign agents to positions according to the permutation specified by P_t.

6. Discard dummy agents and dummy position assignments.

To justify step 4, one can verify that $\mathbf{w}$ majorizes $\mathbf{x}$ in the sense that $\sum_{i=1}^{k} w_i \geq \sum_{i=1}^{k} x_i$ for $k \in \{1, \ldots, 2n\}$, with equality holding for $k = 2n$. Therefore by a theorem of Rado [65], the desired permutation matrices and weights exist, and con be computed efficiently. The following consequences are immediate.

Lemma 8.13. *The resulting mechanism for position auction with weights* $\mathbf{w}$ *obtained from the above reduction has the same expected allocation as the convex combination of k-unit auctions with $(w_k - w_{k+1})$'s as probabilities.*

Lemma 8.14. *Given an incentive compatible multi-unit auctionmechanism , the mechanism from the position auction reduction is also incentive compatible.*

Definition 8.15 (Matroid Permutation Reduction)**.** Given a matroid permutation environment with characteristic weights $\mathbf{w}$, and a position auction mechanism for weights $\mathbf{w}$, we construct the following mechanism for the matroid permutation environment:

1. We run the mechanism for position auctions and for $i = 1, \ldots, n$, let j_i be the position assigned to agent i, or $j_i = \perp$ if i is not assigned a position.

2. Reject all agents i with $j_i = \perp$.

3. Run the greedy algorithm in the matroid permutation environment with agent i's value reset to j_i.

The following conclusions are immediate.

Lemma 8.16. *The resulting mechanism for the matroid permutation environment obtained from the above reduction has the same expected allocation as the mechanism for position auction.*

Lemma 8.17. *Given an incentive-compatible position auction mechanism , the mechanism from the matroid permutation reduction is weakly incentive compatible, where weak incentive compatibility is defined where agents' utilities are defined in expectation over random permutations, instead of per permutation.*

Theorem 8.18. *The factor β to which there is a prior-free incentive-compatible (weakly incentive-compatible for matroid) approximation of* $\mathrm{EFO}^{(2)}(\mathbf{v})$ *is the same for multi-unit, position, and matroid permutation environments.*

Corollary 8.19. *There is a prior-free incentive-compatible* 6.24*-approximation to* $\mathrm{EFO}^{(2)}(\mathbf{v})$ *in position auctions, and a weakly-incentive compatible one for matroid permutation environments.*

There are two weakness in the reductions implied by Theorem 8.18 in comparison to those implied by Proposition 8.2. Recall that for the latter, ironed virtual surplus maximizations are via the greedy algorithm, and so the reductions were algorithmically trivial. In contrast, Theorem 8.18 requires knowledge of the characteristic weights to run the construction, these weights may be hard to compute. In addition the mechanism that results from the matroid permutation reduction is only weakly incentive compatible.

8.2 Downward-closed Permutation Environments

In this section, we will show that a variant of RSEM (recall Definition 8.4) approximates the envy-free benchmark by a constant factor.

Definition 8.20 (RSEM'). The variant RSEM$'$ is identical to RSEM except Step 4:

4′. find the feasible subset W of N (the full set of agents) to maximize ironed virtual surplus with respect to $\bar{\Phi}^S$, serve agents in $M \cap W$ (the winners from the market M) only, and reject all other agents.

The proof we give that RSEM′ is a good approximation to the envy-free benchmark is based on the fact that with large probability the sample and market satisfy a natural balanced condition. This condition requires that for all prefixes of the agents sorted by value that a good fraction of these agents are in both of market and sample. The proof then has four main steps: show the probability of balance is high, show that balance implies that the loss in ironing in the "wrong" way can be bounded, show that balance implies that the IC revenue of RSEM′ (on the market M) is close to the optimal EF revenue for the sample S, and show the expected optimal EF revenue of the sample is close to the envy-free benchmark.

8.2.1 Balanced Partitioning

We now show that with high probability the partitioning of the agents into the market and sample satisfies a natural balanced property. Recall that, by definition of the mechanism, agent 1 is in M. This balanced property is a double-sided version of the balanced property introduced by Feige et al. [29].

Definition 8.21. A partitioning (S, M) of agents $N = \{1, \ldots, n\}$ is *balanced* if $1 \in M$ and $2 \in S$ and for any set of three or more of the highest valued agents both the market and sample contain at least a quarter of agents in the set. I.e., for $i \geq 3$, $|S \cap \{1, \ldots, i\}| \geq i/4$ and $|M \cap \{1, \ldots, i\}| \geq i/4$.

Lemma 8.22. *Conditioning on $1 \in M$, a random partitioning (S, M) of N is balanced with probability at least 0.339.*

Proof. Conditioning on $1 \in M$ and $2 \in S$, the probability that either part is imbalanced can be calculated to be at most 0.161 by a simple probability of ruin analysis which comes from Feige et al. [29] (details given below). By the union bound, both parts are balanced with probability at least 0.678. Agent 2 is in S with probability $1/2$ so the probability of balance conditioned on agent 1 in M is at least 0.339.

The following analysis from Feige et al. [29] shows that the probability that S is imbalanced is at most 0.161. Consider the random variable $Z_i = 4\,|S \cap \{1, \ldots, i\}| - i$; the balanced condition is equivalent to $Z_i \geq 0$ for all $i \geq 3$. By the conditioning $i = 2$ and $S \cap \{1, 2\} = \{2\}$ imply that $Z_2 = 2$. View Z_i as the positions of a random walk on the integers that starts from position two and takes three steps forward (at step i with $i \in S$) or one step back (at step i with $i \notin S$), each with probability one half. If this random walk ever arrives at position -1 the partitioning is imbalanced. This probability r of ever taking one step back in such a random walk can be calculated as the root of $r^4 - 2r + 1$ on interval $(0, 1)$ which is about 0.544. The probability of imbalance is then $r^3 \leq 0.161$ (i.e., if we ever take three steps back when starting from position two). By symmetry, the probability of imbalance in the market M is also at most 0.161. □

8.2.2 Sub-Optimal Ironing

We need to give a detailed analysis of what happens in terms of envy-free revenue when we optimize for the wrong virtual values. To do that we will define and consider the *effective revenue curve*, $\widetilde{\mathrm{R}}$, and *perceived revenue curve*, $\widehat{\mathrm{R}}$. Intuitively, $\widehat{\mathrm{R}}$ corresponds to the revenue we think we get when optimizing $\bar{\Phi}^S$ on $\mathbf{v}$, and $\widetilde{\mathrm{R}}$ corresponds to the revenue curve we actually end up with.

Definition 8.23 (Effective revenue curve $\widetilde{\mathrm{R}}$)**.** For values $\mathbf{v}$ and ironed virtual valuations $\bar{\Phi}^S$ for S: group agents with equal nonnegative $\bar{\Phi}^S$ values into consecutive classes $\{1, \ldots, n_1\}$, $\{n_1 + 1, \ldots, n_2\}$, …, $\{n_{t-1} + 1, \ldots, n_t\}$ and define the *effective revenue curve* $\widetilde{\mathrm{R}}$ from $\mathrm{R} = \mathrm{R}^{\mathbf{v}}$ by connecting the points $(0, 0)$, $(n_1, \mathrm{R}(n_1))$, …, $(n_t, \mathrm{R}(n_t))$ and then extending horizontally to $(n, \mathrm{R}(n_t))$, i.e., ironing the values in each class.

Figure 8.2.1a depicts an example of the effective revenue curve. The three rays from the origin, which correspond to values at which $\bar{\Phi}^S$ makes a piece-wise jump, divide the first orthant into four regions. For every region, every point $(i, \mathrm{R}(i))$ in the region (which corresponds to value v_i) has the same $\bar{\Phi}^S$ value. In each region these points get "ironed", and hence the line segment in $\widetilde{\mathrm{R}}$.

Lemma 8.24. $\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}) = \sum_{i=1}^{n} \widetilde{\mathrm{R}}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v}))$.

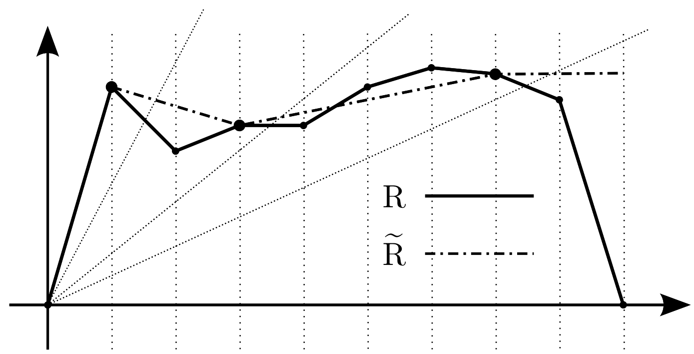

(a)

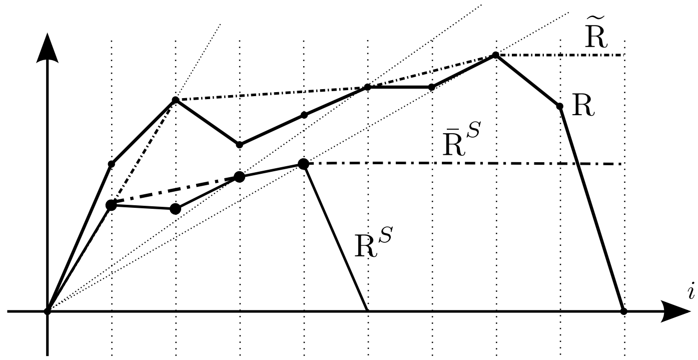

(b) Revenue dominance

Figure 8.2.1: Effective Revenue Curves and Revenue Dominance

Proof. We have the following equalities:

$$\begin{aligned}\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}) &= \sum\nolimits_{i=1}^{n} \mathrm{R}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v})) \\ &= \sum\nolimits_{i=1}^{n} \widetilde{\mathrm{R}}(i) \cdot (x_i^S(\mathbf{v}) - x_{i+1}^S(\mathbf{v}))\end{aligned}$$

Here the first equality is by Lemma 7.9. To justify the second equality, note that whenever $\widetilde{\mathrm{R}}(i) \neq \mathrm{R}(i)$, there are two cases: (1) i is in $\{n_{j-1}+1, \ldots, n_j - 1\}$ for some j, and so v_i and v_{i+1} have the same $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v})$; and (2) i is bigger than n_t, and so v_i and v_{i+1} both have negative $\bar{\Phi}^S$ value, and hence $x_i^S(\mathbf{v}) = x_{i+1}^S(\mathbf{v}) = 0$. □

For a set of agents S, let $\mathbf{v}_S$ denote $(\mathbf{v}_S, \mathbf{0}_{N-S})$, i.e., the valuation profile (of n agents) obtained from $\mathbf{v}$ by decreasing the values of agents outside S to 0. Note that $\mathbf{v} = \mathbf{v}_N$. Let R^S and $\bar{\mathrm{R}}^S$ be the revenue curve and ironed revenue curve of the valuation profile $\mathbf{v}_S$ respectively.

Lemma 8.25. *For all* $1 \leq i \leq n$, $\widetilde{\mathrm{R}}(i) \geq \bar{\mathrm{R}}^S(i)$.

Proof sketch. Figure 8.2.1 depicts the relationship between the revenue curves. Observe that revenue curve R dominates R^S in the sense that for every slope t, the intersection of the ray $y = tx$ with R is farther away from the origin than its intersection with R^S. Transforming R and R^S to the effective revenue curves using the same ironed virtual valuation function $\bar{\Phi}^S$ do not change such dominance relationship, and moreover, because $\bar{\mathrm{R}}^S$ is non-decreasing and concave, it follows that vertical dominance also holds, i.e., $\widetilde{\mathrm{R}}(i) \geq \bar{\mathrm{R}}^S(i)$ for all i. □

Definition 8.26 (Perceived revenue curve $\widehat{\mathrm{R}}$)**.** The *perceived revenue curve* for $\bar{\Phi}^S$ on $\mathbf{v}$ is given by $\widehat{\mathrm{R}}(i) = \sum_{j=1}^{i} \bar{\Phi}^S(v_i)$ for $i \in N$.

Let $\hat{\mathbf{v}}$ be the valuation profile corresponding to $\widehat{\mathrm{R}}$, i.e., $\hat{v}_i = \widehat{\mathrm{R}}(i)/i$, and let $\mathbf{x}^{\hat{\mathbf{v}}}$ be the ironed virtual surplus maximizer for $\bar{\Phi}^{\hat{\mathbf{v}}}$.

Lemma 8.27. $x_i^S(\mathbf{v}) = x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})$.

Proof. Compare running the ironed virtual surplus maximizer $\mathbf{x}^S$ for $\bar{\Phi}^S$ on $\mathbf{v}$ with running $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\bar{\Phi}^{\hat{\mathbf{v}}}$ on $\hat{\mathbf{v}}$, the ironed virtual valuation of agent i in either case is equal to $\bar{\Phi}^S(v_i)$. Therefore these two ironed virtual surplus optimizers will choose the same allocation, and the lemma follows. □

Lemma 8.28. *Given a balanced partitioning* (S, M)*, then* $\bar{\mathrm{R}}^S(i) \geq \frac{1}{4}\widehat{\mathrm{R}}(i) \geq \frac{1}{4}\bar{\mathrm{R}}^S(i)$ *for all* $1 \leq i \leq n$.

Proof. For each i, $\widehat{\mathrm{R}}(i)$ is the sum of the i largest ironed virtual values in N with respect to $\bar{\Phi}^S$ while $\bar{\mathrm{R}}^S(i)$ is the sum of the i largest with respect to S. Therefore $\widehat{\mathrm{R}} \geq \bar{\mathrm{R}}^S(i)$. Since (S, M) is double-side balanced, applying Lemma 8.30, we also have that for all i, $\bar{\mathrm{R}}^S(i) \geq \frac{1}{4}\widehat{\mathrm{R}}(i)$. □

Now we are ready prove the following key lemma:

Lemma 8.29. *For every downward-closed permutation environment, valuation profile* $\mathbf{v}$*, and balanced partitioning* (S, M)*,* $\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) \geq \frac{1}{4}\,\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_S) = \frac{1}{4}\,\mathrm{EFO}(\mathbf{v}_S)$.

Proof. Let $\mathbf{x}^S$ and $\mathbf{x}^{\hat{\mathbf{v}}}$ be short-hands for the ironed virtual surplus optimizers with ironed virtual valuation functions defined for $\mathbf{v}_S$ and $\hat{\mathbf{v}}$, respectively. The proof is by the following inequalities:

$$
\begin{aligned}
\mathrm{EF}^{\bar{\Phi}^S}(\mathbf{v}_N) &= \sum\nolimits_i \widetilde{\mathrm{R}}(i) \cdot (x_i^S(\mathbf{v}_N) - x_{i+1}^S(\mathbf{v}_N)) \\
&= \sum\nolimits_i \widetilde{\mathrm{R}}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\
&\geq \tfrac{1}{4} \cdot \sum\nolimits_i \widehat{\mathrm{R}}(i) \cdot (x_i^{\hat{\mathbf{v}}}(\hat{\mathbf{v}}) - x_{i+1}^{\hat{\mathbf{v}}}(\hat{\mathbf{v}})) \\
&\geq \tfrac{1}{4} \cdot \sum\nolimits_i \widehat{\mathrm{R}}(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)) \\
&\geq \tfrac{1}{4} \cdot \sum\nolimits_i \bar{\mathrm{R}}^S(i) \cdot (x_i^S(\mathbf{v}_S) - x_{i+1}^S(\mathbf{v}_S)).
\end{aligned}
$$

Here the first two equalities are guaranteed by our definitions of $\widetilde{\mathrm{R}}$ and $\widehat{\mathrm{R}}$. The first inequality is by Lemma 8.25 and Lemma 8.28, the second inequality is by the optimality of $\mathbf{x}^{\hat{\mathbf{v}}}$ for $\hat{\mathbf{v}}$, and the third inequality is by Lemma 8.28 again. □

8.2.3 Market Revenue vs. Sample Revenue

We now show that conditioned on a balanced partitioning of the agents into a market and sample, that the revenue of RSEM′ from the market is close to the envy-free optimal revenue from the sample. The revenue of RSEM′ is precisely $\mathrm{IC}_M^S(\mathbf{v}_N)$, i.e., the revenue we get from the agents in M when using the virtual value functions from S and optimizing virtual values over the full set of agents N. We wish to compare this revenue to the envy-free optimal revenue on the sample, $\mathrm{EFO}(\mathbf{v}_S)$.

Lemma 8.30. *Given a balanced partitioning* (S, M)*, for every non-increasing sequence* $a_1, \ldots, a_n$ *of nonnegative reals and all* $i \in N$*,* $\sum_{j \in M \cap \{1,\ldots,i\}} a_j \geq \frac{1}{4} \sum_{j \in \{1,\ldots,i\}} a_j$.

Corollary 8.31. *Given a balanced partitioning* (M, S)*, we have* $\mathrm{EF}_M^S(\mathbf{v}_N) \geq \frac{1}{4} \cdot \mathrm{EF}_N^S(\mathbf{v}_N)$.

Proof. Notice that the EF payments of bidders in $\mathrm{EF}_N^S(\mathbf{v}_N)$ forms a non-increasing sequence. The corollary follows from plugging these EF payments into Lemma 8.30. □

Lemma 8.32. *Given a balanced partitioning* (M, S)*,* $\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v}) \geq \frac{1}{32}\,\mathrm{EFO}(\mathbf{v}_S)$*.*

Proof. The proof is given by the following sequence of inequalities:

$$
\begin{aligned}
\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v}) &= \mathrm{IC}_M^S(\mathbf{v}_N) && \text{(Definition 8.20)} \\
&\geq \tfrac{1}{2}\,\mathrm{EF}_M^S(\mathbf{v}_N) && \text{(Lemma 7.13)} \\
&\geq \tfrac{1}{8}\,\mathrm{EF}_N^S(\mathbf{v}_N) && \text{(Corollary 8.31)} \\
&\geq \tfrac{1}{32}\,\mathrm{EF}_S^S(\mathbf{v}_S) && \text{(Lemma 8.29)} \\
&= \tfrac{1}{32}\,\mathrm{EFO}(\mathbf{v}_S).
\end{aligned}
$$
□

8.2.4 Expected Sample Revenue vs. The Envy-Free Benchmark

We now show that the expected envy-free revenue of the sample compares favorably with the envy-free benchmark; this is the last ingredient in the proof of Theorem 8.35. We will make this argument conditioned on a balanced partitioning; however, the result is true for any symmetric conditioning (including none at all).

Lemma 8.33. *For a partitioning* (S, M) *of* N*, we have that* $\mathrm{EFO}(\mathbf{v}_S) + \mathrm{EFO}(\mathbf{v}_M) \geq \mathrm{EFO}(\mathbf{v}_N)$*.*

Proof. $\mathrm{EFO}(\mathbf{v}_N) = \mathrm{EFO}(\mathbf{v}_{S\cup M})$ is the maximum revenue we can get from $S \cup M$ subject to the envy free constraints. Let agents in M contribute total revenue R to $\mathrm{EFO}(\mathbf{v}_N)$. By setting the agents in S to have zero valuations to obtain valuation profile $\mathbf{v}_S$, we basically removed envy-freeness constraints between agents in S and agents in M. With fewer envy-freeness constraints, the maximum envy-free revenue we can get from M, i.e., $\mathrm{EFO}(\mathbf{v}_M)$, can only be larger. Similarly, the total revenue that S contributes to $\mathrm{EFO}(\mathbf{v}_N)$ is at most $\mathrm{EFO}(\mathbf{v}_S)$, and our lemma follows. □

Lemma 8.34. $\mathbf{E}[\mathrm{EFO}(\mathbf{v}_S) \mid \mathcal{B}] \geq \frac{1}{2}\,\mathrm{EFO}^{(2)}(\mathbf{v})$ *where* $\mathcal{B}$ *is the event that* S *and* M *are balanced.*

Proof. The lemma follows from Lemma 8.33, symmetry, and the fact that agent 1 is always in M. □

8.2.5 Putting It All Together

Theorem 8.35. *For downward-closed permutation environments,* $\mathbf{E}\left[\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v})\right] \geq \frac{1}{189}\,\mathrm{EFO}^{(2)}(\mathbf{v})$ *for all* $\mathbf{v}$.

Proof. Let $\mathcal{B}$ denote the event that the market and sample are balanced. Lemma 8.22 states that the probability that the partition is balanced is at least:

$$\mathbf{Pr}[\mathcal{B}] \geq 0.339.$$

The expected IC revenue of RSEM' is at least its revenue conditioned on event $\mathcal{B}$. I.e.,

$$\mathbf{E}\left[\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v})\right] \geq \mathbf{Pr}[\mathcal{B}]\,\mathbf{E}\left[\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v}) \mid \mathcal{B}\right].$$

Lemma 8.32 states that for every $\mathbf{v}$ that the balance condition implies the IC revenue of RSEM' is at least a $\frac{1}{32}$ fraction of the EF optimal revenue on the sample. Taking expectations,

$$\mathbf{E}\left[\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v}) \mid \mathcal{B}\right] \geq \tfrac{1}{32}\mathbf{E}[\mathrm{EFO}(\mathbf{v}_S) \mid \mathcal{B}].$$

Lemma 8.34 states that the EF optimal revenue on the sample is at least half the envy-free benchmark, in expectation and conditioned on a balanced partitioning. I.e.,

$$\mathbf{E}[\mathrm{EFO}(\mathbf{v}_S) \mid \mathcal{B}] \geq \tfrac{1}{2}\,\mathrm{EFO}^{(2)}(\mathbf{v}).$$

Combining the above inequalities we conclude that the IC revenue of RSEM' is at least a 189-approximation to the envy-free benchmark. I.e.,

$$\mathbf{E}\left[\mathrm{IC}^{\mathrm{RSEM}'}(\mathbf{v})\right] \geq \tfrac{1}{189}\,\mathrm{EFO}^{(2)}(\mathbf{v}).$$

□

Part IV

Conclusion

Chapter 9

Conclusions and Future Directions

9.1 Summary

In this thesis, we have introduced the prior-independent analysis framework, and identified several prior-independent mechanisms. Our mechanisms are based on welfare-maximizing with reserves, supply-limiting, sequential posted-price, and Myerson's mechanism with respect to empirical distribution. Except the last one, our mechanisms enjoy the following features that are important in practice:

Simplicity These mechanisms are in general simple to run, which do not incur the burden of obtaining distribution information outside of the auction.

Naturalness These mechanisms are also based on natural ideas such as welfare-maximization, reserve pricing, supply-limiting, etc. In particular from the bidders' point of view, what the mechanisms do are natural and can be expected.

Robustness The approximation guarantee holds for a wide class of distributions.

Near-optimality The approximation ratios we prove are in most cases very good. In several cases, our approximation ratios approach 1 in the limit as either the number of bidders or the number of items goes to infinity.

Indeed, mechanisms based on similar strategies are commonly deployed in practice, and prior-independence gives us a formal way to justify that we are not losing much revenue compared to the optimal mechanism, by sticking to simple strategies for practical reasons.

9.2 Open Questions

Besides improving over the various constant approximation factors in this paper, there are several questions that we left open. We list a few important ones in the following:

1. For single-dimensional downward-closed environments, is there a prior-independent mechanism w.r.t. i.i.d. regular distributions?

2. For single-dimensional downward-closed permutation environments, does EF revenue (approximately) upper-bound IC revenue for virtual surplus maximizers? (See Conjecture 7.20.) We only know that this is true for matroid environments. If this can be proved, then approximating the EFO benchmark for downward-closed environments would imply a prior-independent approximation guarantee. It would also resolve the previous open problem.

3. To what extent does the supply-limiting mechanism work for matching problems with non-i.i.d. bidders?

9.3 Prior-Independence More Broadly

We only applied prior-independence to revenue-maximizing auction mechanisms. However, the concept is clearly a general one, and is not limited to mechanism design.

Example 9.1 (Online Matching)**.** As one example, in Mahdian and Yan [58], we prove that for the classic online matching problem of Karp et al. [53], the RANKING algorithm gives a prior-independent 0.696-approximation w.r.t. i.i.d. distributions (in the stochastic matching model of Feldman et al. [31]), which is an improvement over the $1 - \frac{1}{e} \approx 0.632$ approximation in worst case [53], and almost matches with the best known ratio of 0.702 based on prior-dependent approximation by Manshadi et al. [59].

Now our question is, for what other problems can prior-independence serve as a suitable framework? We are certain that there are many potential applications of prior-independence. However there does not seem to be a unified way of applying prior-independence, and how prior-independence can be useful for a specific area domain will necessarily depend on domain-specific knowledge and ideas. On the other hand, this means there are plenty of opportunities in this direction for future research.

In general, there is a vast gap between theory and practice of algorithm analysis. To obtain theoretical guarantees that better reflect practical performance of algorithms, we need to model inputs to the algorithms in more realistic ways. Prior-independence seems to be a good general approach in this direction, and we believe that fruitful results are awaiting discovery.

Bibliography

[1] G. Aggarwal and J. Hartline. Knapsack auctions. In *Proc. 17th ACM Symp. on Discrete Algorithms (SODA)*, pages 1083–1092, 2006.

[2] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye. Correlation robust stochastic optimization. In *Proc. 21st ACM Symp. on Discrete Algorithms (SODA)*, pages 1087–1096, 2010.

[3] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *CoRR*, abs/0911.2974, 2009.

[4] Saeed Alaei, Azarakhsh Malekian, and Aravind Srinivasan. On random sampling auctions for digital goods. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 187–196, 2009.

[5] Noga Alon and Joel H. Spencer. *The probabilistic method.* Wiley, New York, 1992.

[6] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 482–491, 2001.

[7] L. M. Ausubel and P. Milgrom. The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 1, pages 17–40. MIT Press, 2006.

[8] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proc. 18th ACM Symp. on Discrete Algorithms (SODA)*, pages 434–443, 2007.

[9] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Innovations in Computer Science - ICS 2010, Beijing, China*, pages 166–178, 2011.

[10] Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci*, 74(8):1245–1270, 2008.

[11] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3, 2003.

[12] S. Bikhchandani, S. de Vries, J. Schummer, and R. Vohra. Linear programming and Vickrey auctions. *IMA Volume in Mathematics and its Applications, Mathematics of the Internet: E-auction and Markets*, 127:75–116, 2001.

[13] L. Blumrosen and T. Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *Proc. 10th ACM Conf. on Electronic Commerce (EC)*, page 49, 2008.

[14] Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. In *Proc. 36th ACM Symp. on Theory of Computing (STOC)*, pages 39–48, 2005.

[15] J. Bulow and P. Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

[16] J. Bulow and J. Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[17] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set function subject to a matroid constraint (extended abstract). In *Proceedings of Integer Programming and Combinatorial Optimization, 12th*

International IPCO Conference, volume 4513 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2007.

[18] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan. Approximation schemes for sequential posted pricing in multi-unit auctions. In *Workshop on Internet and Network Economics (WINE)*, 2010.

[19] Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing (STOC)*, pages 311–320, 2010.

[20] Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. Private communication. 2011.

[21] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[22] N. Devanur and J. D. Hartline. Limited and online supply and the Bayesian foundations of prior-free mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 41–50, 2009.

[23] Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. Prior-independent multi-parameter mechanism design. In *Workshop on Internet and Network Economics (WINE)*, pages 122–133, 2011.

[24] Nikhil Devanur, Jason Hartline, and Qiqi Yan. Envy, truth, and profit. journal paper in submission, 2012.

[25] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*, pages 129–138, 2010.

[26] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 243–252, 2009.

[27] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, 2007.

[28] Christian Ewerhart. Optimal design and ρ-concavity. working paper, 2011.

[29] U. Feige, A. Flaxman, J. Hartline, and R. Kleinberg. On the Competitive Ratio of the Random Sampling Auction. In *Proc. 1st Workshop on Internet and Network Economics*, pages 878–886, 2005.

[30] Uriel Feige. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput*, 39(1):122–142, 2009.

[31] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating 1-1/e. *Proceedings of 50th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 0:117–126, 2009.

[32] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Clifford Stein. Online stochastic packing applied to display ad allocation. In *ESA (1)*, pages 182–194, 2010.

[33] Gagan Goel and Aranyak Mehta. Online budgeted matching in random input models with applications to adwords. In *Proc. 19th ACM Symp. on Discrete Algorithms (SODA)*, pages 982–991, 2008.

[34] A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *Proc. 14th ACM Symp. on Discrete Algorithms (SODA)*, pages 215–222, 2003.

[35] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms (SODA)*, pages 735–744, 2001.

[36] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[37] Andrew Goldberg, Jason Hartline, Anna Karlin, Mike Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

[38] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[39] B. Ha and J. Hartline. Biased random sampling and profit extraction. Manuscript, 2011.

[40] B. Ha and J. Hartline. Mechanism design via profit extraction, consensus, and cross-check. In *Proc. 23rd ACM Symp. on Discrete Algorithms (SODA)*, pages 887–895, 2012.

[41] J. Hartline. Approximation in mechanism design. Manuscript, 2011.

[42] J. Hartline and A. Karlin. Profit maximization in mechanism design. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.

[43] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proc. 7th ACM Conf. on Electronic Commerce (EC)*, pages 175–182, 2005.

[44] J. D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 75–84, 2008.

[45] J. D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce (EC)*, pages 225–234, 2009.

[46] Jason Hartline and Qiqi Yan. Envy, truth, and profit. In *Proc. 13th ACM Conf. on Electronic Commerce (EC)*, pages 243–252, 2011.

[47] Jason D. Hartline, Vahab S. Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *17th International World Wide Web Conference*, pages 189–198, 2008.

[48] C. Holahan. Auctions on eBay: A dying breed. Businessweek, 2008.

[49] T. Ichiba and K. Iwama. Averaging techniques for competitive auctions. In *Proceedings of the Seventh Workshop on Analytic Algorithmics and Combinatorics*, pages 74–81, 2010.

[50] M. Jackson and I. Kremer. Envy-freeness and implementation in large economies. *Review of Economic Design*, 11:185–198, 2007.

[51] T.A. Jenkyns. The efficiency of the greedy algorithm. In *Proceedings of the 7th South Eastern Conference on Combinatorics, Graph Theory and Computing*, pages 341–350. ACM, 1976.

[52] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proc. 42nd ACM Symp. on Theory of Computing (STOC)*, pages 587–596, 2011.

[53] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of Computing (STOC)*, pages 352–358, 1990.

[54] R. Kirkegaard. A short proof of the bulow-klemperer auctions vs. negotiations result. *Journal of Economic Theory*, 28(2):449–452, 2006.

[55] V. Krishna. *Auction Theory, Second Edition.* Academic Press, 2009.

[56] John O. Ledyard. Optimal combinatoric auctions with single-minded bidders. In *ACM Conference on Electronic Commerce*, pages 237–242, 2007.

[57] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce (EC)*, pages 96–102. ACM Press, 1999.

[58] Mohammad Mahdian and Qiqi Yan. Online bipartite matching with random arrivals: An approach based on strongly factor-revealing LPs. In *Proc. 42nd ACM Symp. on Theory of Computing (STOC)*, pages 597–606, 2011.

[59] Vahideh H. Manshadi, Shayan Oveis Gharan, and Amin Saberi. Online stochastic matching: Online actions based on offline statistics. In *Proc. 22nd ACM Symp. on Discrete Algorithms (SODA)*, pages 1285–1294, 2011.

[60] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, Cambridge, UK, 1995.

[61] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1): 58–73, 1981.

[62] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

[63] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC)*, pages 242–252. ACM Press, 2000.

[64] J. G. Oxley. *Matroid Theory.* Oxford, 1992.

[65] R. Rado. An equality. *J. London Math. Soc.*, 27:1–6, 1952.

[66] Tim Roughgarden, Inbal Talgam-Cohen Jason Hartline, and Qiqi Yan. Supply-limiting mechanisms. In *Proc. 14th ACM Conf. on Electronic Commerce (EC)*, 2012.

[67] Michael E. Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains.

[68] A. Schrijver. *Combinatorial Optimziation: Polyhedra and Efficiency.* Springer, 2003.

[69] I. Segal. Optimal pricing mechanisms with unknown demand. In *American Economic Review 93*, 2003.

[70] Mukund Sundararajan and Qiqi Yan. Robust mechanisms for risk-averse sellers. In *Proc. 12th ACM Conf. on Electronic Commerce (EC)*, pages 139–148, 2010.

[71] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 608–619, 2003.

[72] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[73] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37, 1961.

[74] J. Vondrák. *Submodularity in combinatorial optimization.* PhD thesis, Charles University, Prague, 2007.

[75] J. Vondrák. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proc. 39th ACM Symp. on Theory of Computing (STOC)*, pages 67–74, 2008.

[76] Qiqi Yan. Mechanism design via correlation gap. In *Proc. 22nd ACM Symp. on Discrete Algorithms (SODA)*, pages 710–719, 2011.