APPROXIMATIONS IN BAYESIAN MECHANISM DESIGN FOR MULTI-PARAMETER SETTINGS

By

David Malec

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Computer Sciences)

at the

UNIVERSITY OF WISCONSIN–MADISON

2014

Date of final oral examination: November 6, 2013

This dissertation is approved by the following members of the Final Oral Committee:
Shuchi Chawla (Advisor), Associate Professor, Computer Sciences
Dieter van Melkebeek, Professor, Computer Sciences
Jin-Yi Cai, Professor, Computer Sciences
Marzena Rostek, Associate Professor, Economics
Jason D. Hartline, Associate Professor, Electrical Engineering and Computer Science,
Northwestern University

Acknowledgments

Going from a first-year graduate student who had never heard of my eventual thesis topic to defending has been a long journey with many challenges – both expected and unexpected – along the way. There is no way I could have made it without help and support from a great many people; while there is no way I could list everyone who deserves it, I want to list a few of the people who helped me along the way.

I want to thank Shuchi Chawla, my advisor, for her guidance on my journey. Learning how to ask the right questions is a never-ending process, but I learned a great deal about how to cut to the heart of a problem in our many research meetings; I truly appreciate all of the time put into helping me find my way in research. I similarly want to thank Jason Hartline and Nikhil Devanur for their time as my mentors when visiting Northwestern University and Microsoft Research, respectively. They have each given me unique perspectives on how to find good questions, and helped me immensely along the way.

I also want to thank my committee; their support, encouragement, and efforts throughout this process has meant a great deal to me.

Many thanks to all of my coauthors. There is nothing quite like finally catching the essence of a problem after hours and hours of searching. I thank each of my coauthors for having the opportunity to share both the hunt for knowledge and wonder of discovery with them.

I also want to thank Jason for hosting me for an entire Fall at Northwestern University. I have many fond memories of my time there. Thanks to Jason, Nicole Immorlica, and Azarakhsh Malekian for the wonderful mechanism design seminar they ran as well as their collaboration on many interesting projects; I learned so much from the experience. I want to thank Tom, Greg, Manolis, Nima, Michele, Darrell, Michael, and all of the other students and visiting researchers as well for creating a fun, vibrant, and welcoming research environment.

Many thanks to all of my fellow theory students at Madison. All of you have offered friendship, advice, and help on countless occasions. In particular, thanks to Jeff and Matt for all of their advice over the years; William and Siddharth for always being ready grab a cup of coffee and remind me to keep things in perspective; Balu for our many collaborations over the years, and helping me remember to constantly challenge ideas and assumptions; and Tyson for the boundless enthusiasm he brought to any discussion or task. The many shared cups of coffee and tea, lunches on the terrace, and periodic expeditions to Maharani will all be fondly remembered.

I want to thank everyone involved in the training for the ACM International Collegiate Programming Contest at Madison. Both as a student and as a coach, my involvement with the contest preparations was a joy. Thanks to Dieter van Melkebeek for the time he puts in to keep it going; to Matt, Scott, and Zef for all the time they put into coaching; and to all the amazing students who have put time into training over the years. I learned a great deal from the experience, and hope I passed along at least some small fraction of that to others.

Last but definitely not least, thanks to all of my friends and to my family for supporting me and keeping me sane through this process. A special thanks to Dalibor, Tristan, Brandon, Vas, and Eric for putting up with me as a housemate over the years; living in CS East was a special part of my time in Madison. Also, thanks to Casey, Lauren, and Pavel for their friendship during undergrad and continuing to keep track of me, and help me keep track of myself, during grad school. If I tried to list all of the friends I made over my years in Madison and who made those years what they were, I know I'd miss a more than a few; know, however, that if we ever shot pool; got coffee at Indie; went to a show or play; enjoyed the nightlife of Madison on a Wednesday; or shared some turducken or other craziness at CS East, then your name belongs here and your friendship is truly appreciated.

Madison has been an incredible place to spend the past several years of my life. The city, University, and Computer Sciences department here are all amazing in their own way; in particular, I believe the sense of community in the department here is quite special, and I feel lucky to have had the chance to be a part of that community. I will always treasure my time in Madison and the friends I have made along the way; and I hope Madison always continues to be ever-changing but unique and wonderful place I have known.

Contents

Contents **iii**

List of Tables **v**

List of Figures **vi**

Abstract **vii**

1 Introduction **1**
- 1.1 Approximation in Bayesian multi-parameter mechanism design . . . 2
- 1.2 Contributions . . . 7
- 1.3 Related works . . . 11
- 1.4 Bibliographic notes . . . 15

2 Definitions and notations **17**
- 2.1 Bayesian settings for mechanism design . . . 17
- 2.2 Mechanism design desiderata . . . 22
- 2.3 Optimal single-parameter mechanism design . . . 23

3 Sequential pricings for single-parameter settings **26**
- 3.1 Preliminaries . . . 28
- 3.2 Order-oblivious sequential pricings . . . 32
- 3.3 Order-specifying sequential pricings . . . 42
- 3.4 Computing the near-optimal posted-price mechanisms . . . 50
- 3.5 Revenue maximization through VCG mechanisms . . . 52
- 3.6 Discussion . . . 53

4 Multi-service settings with unit-demand agents **55**
- 4.1 Reducing multi-service instances to single-service instances . . . 57
- 4.2 Approximation through order-oblivious sequential pricings . . . 60
- 4.3 Approximation through order-specifying sequential pricings . . . 66
- 4.4 Discussion . . . 68

5 The power of randomization in multi-service settings **70**

5.1 Warm up: Single-agent setting 73
5.2 Multi-agent settings 78
5.3 Common base value correlation 87
5.4 A gap example 97
5.5 Conclusions and open problems 100

6 Settings with budget-constrained agents **101**
6.1 Maximizing revenue 103
6.2 Maximizing welfare 121
6.3 Incentive compatibility for budgets 127

References **130**

List of Tables

1.1 A selection of approximation factors for single-dimensional settings through posted pricings . . . 16
1.2 A selection of approximation factors for multi-dimensional unit-demand settings. . . . 16
1.3 A selection of approximation factors between deterministic and randomized mechanisms for multi-dimensional unit-demand settings. . . . 16
1.4 A selection of approximation factors for single-service settings with budget-constrained agents. . . . 16

3.1 A summary of approximation factors for the BSMD achievable through OPMs. Here k is the rank of $\mathcal{S}$. . . . 54
3.2 A summary of approximation factors for the BSMD achievable through SPMs. . . . 54

4.1 A summary of approximation factors for the BMUMD achievable through OPMs. Here k is the rank of $\mathcal{S}$. . . . 69
4.2 A summary of approximation factors for the BMUMD achievable through SPMs. . . . 69

List of Figures

5.1 An example from Thanassoulis (2004) contrasting the optimal item and lottery pricings. The regions R_1, R_2, and R_{Lot} denote the sets of valuations at which the agent buys item 1, item 2, and the $(1/2, 1/2)$ lottery, respectively. 71
5.2 The allocation function for the lottery pricing $\mathcal{L}$. 99

Abstract

We study the classic mathematical economics problem of *Bayesian mechanism design* where a principal aims to optimize an objective, typically expected revenue or welfare, when allocating resources to self-interested agents with preferences drawn from a known distribution. In single-parameter settings where each agent's preference is given by a single private value, this problem is solved (Myerson, 1981). Unfortunately, these single-parameter optimal mechanisms are impractical and rarely employed (Ausubel and Milgrom, 2006), and furthermore the underlying economic theory fails to generalize to the important, relevant, and unsolved multi-parameter setting, where each agent's preferences require multiple values to describe (Manelli and Vincent, 2007).

The main theme of this work is that we can solve multi-parameter mechanism design problems by analogy to properly chosen single-parameter ones. In order to implement this approach, however, the mechanisms we design must be robust to changes in the setting they are used in. Thus, in contrast to the theory of optimal mechanisms, we develop a theory of sequential posted-price mechanisms, where agents are offered take-it-or-leave-it prices in sequence. We prove that these mechanisms are approximately optimal in single-parameter settings. Further, these posted-price mechanisms avoid many of the properties of optimal mechanisms that make the latter impractical, allowing them to generalize naturally to multi-parameter settings.

We consider two types of multi-parameter settings in this thesis. First, we consider settings with multiple services and unit-demand agents. In such settings agents value each service differently, but desire at most one. In particular, we achieve constant-factor approximations for the multi-parameter multi-unit revenue-maximizing auction problem, and bound the economic benefit of randomization in such settings. Second, we consider settings where agents face budget constraints. An agent with a hard budget constraint cannot make payments exceeding it, even when the agent's value for service greatly exceeds the budget. We show how unconstrained Bayesian mechanism design can guide the design of mechanisms for budget-constrained agents, and achieve approximations to both revenue and welfare. Some of our results extend to settings where budgets are private information of the agents.

1 Introduction

Over the history of computing, the very nature of how, when, and why we compute has changed dramatically. While computing began as a costly and centralized resource, reserved for important optimization tasks, the advent of smartphones and cloud computing has made it a near-constantly-available part of daily life for many people. Whereas in the early days of computer hardware, processing time was often reserved for business or mission-critical optimization problems, in the modern day most people think little of the having the ability to check the quickest route to a coffee shop on a smartphone they can carry in their pocket. Additionally, the frontier of tractability for large-scale optimization problems has also greatly expanded. This fundamental shift in the scale and nature of how we compute has led to the examination both of what role computation plays in other fields, and how other fields can help us understand the role computation plays in our lives. The topic of this thesis lies in the area of algorithmic game theory, at the intersection of the fields of computer science and economics.

On one hand, as computing resources become ever more ubiquitous and distributed, questions of coordination and resource allocation become ever more important. The reliability of collaborative efforts such as Wikipedia or Yelp for acquiring information and guiding decision critically depends on the ability to identify and filter out any harmful behaviors that can arise among users. Another example is trying to understand how shared resources such as network bandwidth and cloud computing services can be allocated: how should we sell such services to satisfy user demands, when those demands are themselves shaped by our choice of a selling strategy? Questions of this sort are both natural and important as we face environments where informational and computing resources are shared among self-interested agents, and these questions are a natural fit for game theory.

On the other hand, as the availability of computing power increases and the cost decreases, it becomes important to understand what role computation plays in economic settings. When considering settings such as complex markets or spectrum auctions, the underlying optimization problem facing any sort of coordinator becomes complicated, and understanding how algorithmic and computational complexity considerations impact the optimization in question is critical. As the core optimization problem becomes more difficult, so too do the individual optimization problems participants must solve if they are to take actions that give them their best possible outcomes; if we want to predict what actions individuals will take, we must understand their ability to solve the optimization problem they face as well. Further, as computer science looks at economic mechanisms to solve allocation problems such as selling Internet advertising or cloud resources, many new use cases with novel aspects arise. Thus, this change in how computing impacts our daily lives also impacts how we think about game theory.

Within the area of algorithmic game theory, the particular focus of this thesis is on problems in mechanism design. In particular, we study approximation in the context of Bayesian multi-parameter mechanism design settings. Before describing our results in Section 1.2, we first make each part of this work's title precise.

1.1 Approximation in Bayesian multi-parameter mechanism design

In this section, we examine the topic of this thesis, approximation in Bayesian multi-parameter mechanism design in more detail. Specifically, we describe each component of the topic in turn.

Mechanism design. Mechanism design has a similar core goal to that of algorithm design: an optimizer wishes to map inputs to outputs in order to achieve some global objective. For example, we might want to find a max-weight matching in a bipartite graph. Typically, the goal in algorithm design is to find a good solution in a computationally efficient manner. In mechanism design, however, the pieces of the input are private information held by a collection of agents; these agents each have preferences over the outcomes, and may lie about the information they hold if they believe it will produce an outcome that is better for them. So, in the matching example, we might have that each edge is owned by a different individual, who experiences some benefit (equal to the weight) for having their edge included in the matching. A designer who wants to maximize the total value received by all individuals (i.e. maximize social welfare) has the same goal as the algorithm designer in our original problem statement. Note, however, that as currently phrased, we cannot hope to solve our problem optimally – every agent should simply claim the weight on their edge is the largest number they can think of, since they only care about getting their own edge selected.

In order to address this issue, mechanism design gives the designer extra power: in addition to specifying an outcome, the designer also specifies monetary transfers to or from each agent. For example, the mechanism designer could specify that they will select a maximum-weight matching, but will charge each individual included in the matching an amount equal to the weight they reported for their edge. While this will remove the issue of agents reporting absurdly high values – since they would be required to make absurdly high payments – it creates a new incentive to lie, namely to under-report values in the hopes of reducing the amount paid. The goal of mechanism design is to find mappings from inputs to (outcome, payment) pairs in order to achieve some objective in the presence of individuals who act in their own best interest. A mechanism design problem has several key features:

- the objective of the designer, for example social welfare (total value received by individuals), revenue (total payments received by the mechanism), or fairness;
- the feasibility constraints faced by the designer, both in terms of what outcomes are possible and what payments are possible;
- how individuals interact with the mechanism; and
- how we can predict the behavior of individuals, and therefore the outcome of running the mechanism.

The last two points are especially important. If we want to be able to predict outcomes of a mechanism, we must be able to understand how individuals will interact with it. In this work, we focus on direct revelation mechanisms where individuals interact with the mechanism by making reports in the form of values; further, we focus on mechanisms that are incentive compatible and individually rational (this is, in fact, without loss of generality by the Revelation Principle; see Nisan (2007) for details). The first says that every individual has a *dominant strategy* of being truthful about their value, while the second says that participating in the mechanism is no worse for any individual than not participating. An action is said to constitute a dominant strategy if there is no situation where, considering the actions of others, it is strictly better to deviate and take another action. Thus, it is quite robust – it says that an individual can settle on this strategy based only on their own information and not regret their choice later, once the actions of others are revealed. While we relax the former in some cases, we still rely on dominant strategies to characterize individual behavior.

Multi-parameter. In the context of mechanism design, whether a setting is multi-parameter or single-parameter is determined by the underlying mathematical complexity of participants' preferences over outcomes. Formally, an agent's type specifies what value they place on every possible outcome; however, we can typically find a description that is far more compact than explicitly listing the value for each possible outcome. For example, consider how an agent might value the outcomes in a single-item auction. In the simplest case, the agent may simply care whether or not they win the item, receiving some fixed value from winning it and no value from losing. Since we can summarize such an agent's preferences with just the value the agent places on winning, we term the agent single-parameter. In a more complicated scenario, an agent may further differentiate outcomes. Perhaps if the agent does not win, they care whether or not the item was won by a friend of theirs; they may obtain some value from knowing the winner. For such an agent, we cannot hope to summarize their preferences over outcomes with a single number, and so would describe such an agent as multi-parameter.

In this work, our focus is on understanding multi-parameter settings. While single-parameter settings are generally well-understood, there appears to be a fundamental difficulty in extending techniques from single-parameter settings to multi-parameter ones. Some success has been seen in the case of low-dimensional cases, but many of these results depend on a transformation of the input space so that either it becomes a single-parameter problem, or it decomposes nicely into two independent single-parameter problems. While multi-parameter settings have typically resisted attempts to characterize optimal solutions, we believe that by building analogies to appropriate single-parameter settings, we can gain insight into the structure of solutions that are nearly optimal.

Bayesian. When designing either algorithms or mechanisms, a critical question is how to anticipate the performance of a proposed solution. Typically, our goal is to design a solution that is robust enough to handle a wide variety of different problem instances, often in the face of uncertainty about which instances will be seen in practice. After all, if we only want to solve a single, known instance of a problem, it is often easier and cheaper to just solve that particular instance in an ad hoc fashion. A natural question, then, is how we measure the performance of a solution when we do not know the instances it will be run on in advance, and how much robustness to future uncertainty it will have. Historically, this question has been approached quite differently by the fields of economics and computer science.

The field of computer science has long focused on worst-case analysis, which seeks performance guarantees that hold for every possible instance, or every possible instance in some natural class of inputs. While such a guarantee is incredibly robust – we need no information about the input instances we will see – it is often hard to achieve such a strong guarantee. In economics, however, it is common to work in a Bayesian setting. In such a setting, the goal is not to ensure an outcome with certain qualities for every possible input, but instead to achieve those qualities on average. In particular, in a Bayesian setting we assume that we face an instance that is drawn from a known distribution, and our performance is measured in expectation over the distribution. Furthermore, we assume that we know the distribution in advance, and can tune our solution to take said distribution into account.

While it is clear that a Bayesian setting gives the designer extra power, an important point to note is that in the context of mechanism design it may be impossible to even come up with a performance benchmark without such an assumption. For many mechanism design objectives, such as revenue, there is no single well-defined optimal mechanism we can compare against when seeking a worst-case guarantee. This arises because truthfulness constraints typically bind *across* possible inputs, in the sense that the outcome chosen for one input instance of a problem places constraints on the outcomes that can be chosen for other input instances. This means the mechanism designer is frequently forced to make tradeoffs in performance between inputs, that is, it is often the case that one cannot increase performance on a particular input without decreasing performance on some other inputs. This makes worst-case benchmarks unrealistic: any mechanism a designer might propose will inevitably be outperformed on a given input instance by a mechanism that optimizes for that particular input at the cost of poor performance on all other inputs. This hints at the core issue, namely that since the designer must make performance tradeoffs between input instances, there is no way to guarantee good performance overall while remaining entirely oblivious to what inputs are likely to occur. In fact, without such information, even the idea of an optimal mechanism becomes ill-defined. A Bayesian assumption on inputs provides a natural resolution to this problem. When inputs are drawn according to a (known) distribution, the performance of a mechanism can be measured by its expected performance under that distribution, and the optimal mechanism becomes a well-defined benchmark that gives a clear measure of performance for any proposed mechanism.

Approximation. The theory of approximation in the field of computer science has arisen as a response to the fact that many important, natural problems have been shown to be computationally hard to optimize exactly. In an approximation algorithm, computational efficiency is traded for quality of solution: an algorithm is said to be an α-approximation (for some $\alpha \geqslant 1$) for a problem if it guarantees a solution within an α-factor of optimal for every instance of the problem. Formally, this means that if we let $\mathrm{ALG}(\mathcal{I})$ and $\mathrm{OPT}(\mathcal{I})$ denote the objectives values of an algorithm's solution and the optimal solution for an instance $\mathcal{I}$ of a problem, then we say that ALG provides an α-approximation if we have that

$$\max_{\mathcal{I}} \left(\frac{\mathrm{OPT}(\mathcal{I})}{\mathrm{ALG}(\mathcal{I})} \right) \leqslant \alpha \qquad \text{or} \qquad \max_{\mathcal{I}} \left(\frac{\mathrm{ALG}(\mathcal{I})}{\mathrm{OPT}(\mathcal{I})} \right) \leqslant \alpha,$$

when the problem is a maximization or minimization problem, respectively, and $\mathcal{I}$ varies over all possible instances of the problem. The goal is to achieve the best possible tradeoff between improvements in tractability or runtime and losses in the objective value; for example, providing a polynomial-time constant-approximation for a problem where exact optimization is not believed to be polynomial-time solvable is typically considered a good result when the constant is small.

In the context of Bayesian settings, where the instance $\mathcal{I}$ is drawn from a (known) distribution, we measure the performance of an algorithm as the ratio of expectations rather than the expectation of the ratio, that is we say ALG provides an α-approximation if

$$\frac{\mathrm{E}_{\mathcal{I}}[\mathrm{OPT}(\mathcal{I})]}{\mathrm{E}_{\mathcal{I}}[\mathrm{ALG}(\mathcal{I})]} \leqslant \alpha \qquad \text{or} \qquad \frac{\mathrm{E}_{\mathcal{I}}[\mathrm{ALG}(\mathcal{I})]}{\mathrm{E}_{\mathcal{I}}[\mathrm{OPT}(\mathcal{I})]} \leqslant \alpha$$

when the problem is a maximization or minimization problem, respectively.

One natural criticism of approximation theory is that in practical settings, even losing half of the optimal objective value might be intolerable. It is important to note, however, that

- since the guarantee must apply for *every* problem instance $\mathcal{I}$, it may be reasonable to hope that, in fact, for practical instances the performance is much closer to optimal;
- when developing ad hoc heuristics for a problem, approximation algorithms can provide a principled baseline from which to start, since unlike an arbitrarily chosen starting point we have some guarantee that they provide reasonable answers; and
- most importantly, a good approximation algorithm gives insight into the structure that both occurs in relatively good solutions and can be recognized in a computationally efficient manner, thus guiding the development of specialized algorithms for practical settings.

Thus, a key aspect of approximation algorithms is that even when such an algorithm does not promise a practical performance guarantee, it can provide critical insight into developing algorithms that do provide good performance in practical settings.

1.2 Contributions

Our goal in this work is to gain insight into the structure of near-optimal mechanisms for multi-parameter settings through the lens of approximation. While extending techniques for optimal mechanism design from single-parameter settings to multi-parameter ones appears to be challenging or even impossible in general, we believe that single-parameter mechanism design can give strong insight into the structure of multi-parameter mechanisms. We propose that while single-parameter techniques may not immediately yield mechanisms that are *exactly* optimal for multi-parameter settings, that they can help us get close to this goal and design mechanisms that are *nearly* optimal. Currently, the design of optimal mechanisms for multi-parameter settings seems largely intractable, with only sparse successes (especially when compared to results in single-parameter settings). Thus, it is very appealing to find a way to leverage our understanding of single-parameter mechanism design in order to gain insight into the structure of good mechanisms for multi-parameter settings.

The main theme of our results in this work is that we can understand how to design mechanisms for a multi-parameter setting by drawing an analogy to a related single-parameter setting. At a high level, our goal is to solve a multi-parameter problem by deriving a single-parameter problem from it, leveraging the many mechanism design techniques at our disposal in single-parameter settings, and then using the mechanism we design for the single-parameter setting to generate a good mechanism for the original multi-parameter setting. Our claim is that by carefully aligning a single-parameter setting with a given multi-parameter one, we can find mechanisms that approximate the optimal revenue in both.

As we shall see, however, both directions of the analogy in such a result require care: picking the correct single-parameter setting is critical since it must capture the essential structure of the multi-parameter instance; furthermore, we may need to impose extra constraints or design criteria in the single-parameter setting if we want the analogy to allow us to translate mechanisms from the single-parameter setting back to the multi-parameter one. Failing to adequately address either of these issues can lead to either poor approximation ratios, or a complete breakdown of the desired analogy. In the remaining chapters of this work, we both demonstrate single-parameter mechanisms that achieve relevant robustness properties, and show how to implement this approach in several multi-dimensional settings. We briefly sketch the focus of each chapter below.

Chapter 3: Sequential pricings for single-parameter settings. In this chapter, we focus on single-parameter mechanism design problems. While designing optimal mechanisms for both welfare and revenue is well-understood for single-parameter settings, we study a class of mechanisms that trade optimality (in terms of the objective) for other desirable properties. In particular, we focus on a class of mechanisms called sequential posted pricings. Such mechanisms approach agents one-by-one in turn, making each agent a take-it-or-leave-it offer of service at a (pre-computed) price. We consider settings both where the order of offers is in control of the posted-pricing's designer and settings where the designer is *oblivious* to the order agents will arrive in; we denote these variants of sequential posted pricings as SPMs and OPMs, respectively.

What posted-price mechanisms lose in their objective values, they make up for with several other desirable properties. In particular, they are both easy for the designer to run, and easy for the agents to understand. Furthermore, their simple structure makes them extremely robust to collusion. And as we show, while such mechanisms do not achieve exact optimality, we are able to give strong approximation guarantees for a variety of feasibility constraints. Our results are summarized in Table 1.1.

Chapter 4: Multi-service settings with unit-demand agents. In this chapter, we give our first reduction from a multi-parameter setting to a single-parameter one. We focus on multi-service settings with unit-demand agents. In such a setting, a single seller has several services or items for sale, and they are serving a collection of unit-demand agents. Such an agent may be interested in several different options, but only wants to be allocated at most one of them: for example, think of an individual purchasing a television set, a car, or a plane ticket. Even though each agent only wants a single service, the fact that they can value different options at differing levels means that the setting is fundamentally multi-parameter.

We show, however, that we can build an analogy to a specific single-parameter setting. In particular, we show that we can relate each agent in our setting to a collection of representatives, one per service, who represent the agent's interest in the corresponding service. By assuming that the representatives for a particular agent act independently of (and possibly to the detriment of) each other, we get a standard single-parameter setting. Furthermore, the competition between representatives can only increase revenue, and as long as we require strong properties of collusion-robustness of the mechanisms we design for the single-parameter setting, they can be easily translated back to our original multi-parameter setting. We summarize our results for this setting in Table 1.2.

Chapter 5: The power of randomization in multi-service settings. In this chapter, we examine the power of randomization for multi-service settings with unit-demand agents. In single-parameter settings, randomization gives the mechanism designer no extra power – the revenue-optimal mechanism is, in fact, deterministic. In the multi-parameter settings we consider in Chapters 4 and 5, however, this no longer holds: even very simple examples give a gap in revenue between randomized and deterministic mechanisms (see the introduction to Chapter 5 for a concrete example).

In Chapter 4, we saw how to build an analogy between multi-service settings with unit-demand agents, and carefully chosen single-parameter settings. Unfortunately, while that reduction did produce deterministic mechanisms, it relied on an upper bound that only works for the optimal *deterministic* mechanism for the multi-parameter settings we considered. In particular, it crucially relied on the observation that since different representatives capture an agent's interest in different items, we could only increase revenue by letting them compete with each other. Unfortunately, this intuition only holds for deterministic mechanisms. The key issue is this: when allocations are deterministic, a higher value for one item can only make it less likely other items are allocated; when allocations can mix multiple items, however, a higher value for one item may make an agent interested in mixtures that happen to give higher allocations for both that item *and* other items at the same time. In other words, when an agent is forced to choose between items, they compete for attention, but when mixtures of the items are possible it can introduce synergies between those items. We show, however, that with proper modifications the same analogy can also allow us to achieve approximations to the optimal randomized mechanism; we summarize our results in Table 1.3.

Chapter 6: Settings with budget-constrained agents. In this chapter, we consider the problem of designing mechanisms for agents who are budget-constrained. The amount such an agent can and will pay for service is based both on how much they value having the service, and on how much money they have available to spend on the service. To get intuition for such settings, consider how a company might decide whether or not to invest in a project or significant capital purchase: when deciding whether or not to make the investment, they will weigh both what resources they can spare, and how much projected benefit the investment will produce. Note that both place a limit on how much cost the company can or will accept for the investment, but do so for very different reasons.

In this chapter, our goal is to understand how to design mechanisms for settings where agents face such absolute budget constraints; we consider the revenue and welfare objectives, both in settings where the agents' budget constraints are public common knowledge, and in settings where budget constraints are part of an agent's private information. As we see, budget constraints have a significant impact on the designer's problem. Once again, however, we will see that we can solve our multi-parameter problem by proper analogy to a single-parameter setting. We summarize some of our results in Table 1.4.

1.3 Related works

Single-parameter mechanism design. Myerson (1981) was the first to characterize (revenue) optimal single-parameter mechanisms. Unfortunately, the optimal mechanism is frequently quite complicated, and rarely implemented in practice. The question of whether simple mechanisms can achieve near-optimal revenue was considered recently by Hartline and Roughgarden (2009). Except for their result on single-item auctions with anonymous reserve prices, their VCG-based mechanisms are likely to suffer the same impracticality criticisms as the optimal mechanism. The essay "The Lovely but Lonely Vickrey Auction" by Ausubel and Milgrom (2006) discusses why this is the case. As a consequence of the near-optimality of sequential posted prices, we answer one of their open questions in the positive, namely, that the gap between the revenue optimal mechanism and a VCG mechanism with appropriate reserve prices is a constant (i.e., 2) in matroid settings but with arbitrary valuation distributions. This bound matches their result for regular distributions.

Sequential posted price mechanisms have been considered previously in single-dimensional settings. Sandholm and Gilpin (2006) show experimentally that these mechanisms compare favorably to Myerson's optimal mechanisms. Blumrosen and Holenstein (2008) show how to compute the optimal posted prices in the special case where agents' values are distributed identically, and also show that in this case the revenue of these mechanisms approaches the optimal revenue asymptotically. Several papers study revenue maximization through online posted pricings in the context of adversarial values, albeit in the simpler context of multi-unit auctions (Blum et al., 2004; Kleinberg and Leighton, 2003; Blum and Hartline, 2005).

Multi-service settings. Revenue-optimal mechanisms in multi-parameter settings are poorly understood. Following Myerson's characterization (1981) of optimal single-parameter mechanisms, there were a number of attempts to obtain simple characterizations of optimal mechanisms in the multi-parameter setting (McAfee and McMillan, 1988; Rochet and Chone, 1998; Manelli and Vincent, 2007), however no general-purpose characterization of such mechanisms is known (Manelli and Vincent, 2007). For further work in economics on optimal multi-dimensional mechanism design, see Manelli and Vincent (2007) and references therein. For further work in computer science on multi-dimensional pricing for a single agent, see Chawla et al. (2007) and references therein. Our single-agent setting is most closely related to the work of Chawla et al. (2007) who gave a 3-approximation to the optimal deterministic mechanism for single-agent product-distribution instances, and builds upon techniques developed in that work. We extend the settings in that work to multiple agents and improve their approximation for a single agent from 3 to 2.

The power of randomness. Randomness is a useful resource in mechanism design. In settings involving uncertainty, it allows the designer to hedge against adversarial input: in prior-free mechanism design where the designer has no information about buyers' values, (anonymous) deterministic mechanisms provably cannot obtain any guarantees on revenue and randomness is crucial (see, e.g., Hartline and Karlin, 2007, and references therein). Randomness is also useful when computation is a costly resource and the underlying optimization problem is computationally intractable (e.g. Dobzinski and Dughmi, 2009; Dughmi and Roughgarden, 2010). In the settings that we consider, neither of these effects are present: the designer knows the distribution from which agent types are drawn and we ignore computational issues. In our settings randomness helps for purely economic reasons—it gives the seller more latitude to price discriminate among buyers with different preferences.

Riley and Zeckhauser (1983) were the first to study the question of whether lotteries offer more revenue than item pricing; they showed that for a variety of single-parameter settings the optimal mechanism is deterministic. Subsequently Thanassoulis (2004) noted that there exist multi-parameter instances with valuations drawn from product distributions where randomness helps increase the revenue by about 8-10%. Manelli and Vincent (2006) and Pavlov (2006) presented other examples with small gaps. Briest et al. (2010) were the first to uncover the extent of the benefit of randomization, as well as to study the hardness of finding the optimal randomized mechanism in multi-parameter settings. They showed that lottery pricings can be arbitrarily better than item pricings in terms of revenue even for the case of 4 items offered to a single agent.

Our mechanism design setting with unit-demand agents is closely related to the standard setting for envy-free pricing problems considered in literature (Guruswami et al., 2005; Balcan and Blum, 2006; Balcan et al., 2008; Briest, 2006; Chawla et al., 2007); those works study the single-agent problem with a correlated value distribution and aim to approximate the optimal deterministic mechanism (item pricing).

Settings with budgets. Several works in economics have studied characterizations of optimal BIC IIR budget-feasible mechanisms (e.g., Pai and Vohra, 2008; Laffont and Robert, 1996; Che and Gale, 2000; Maskin, 2000). However, these works are generally weak in the kinds of settings they consider (typically just single-item auctions) and the kinds of value distributions they allow[1]. Laffont and Robert (1996) considered single item settings where bidders have a private value and public common budget. Che and Gale (2000) considered the setting with a single item and a single buyer, but allowed both the value and the budget to be private. Pai and Vohra (2008) gave a more general result in which they designed an optimal auction for a single item and multiple buyers with private i.i.d. values and private budgets.

Bhattacharya et al. (2010) were the first to study settings beyond single-item auctions and focused on revenue maximization. They considered a setting with heterogeneous items and additive values, and exhibited a (large) constant factor DSIC approximation mechanism as well as an all-pay auction which admits truthtelling as a BNE and in that BNE obtains a 4-approximation. However, these results required the value distributions to satisfy the MHR condition. The mechanisms are LP-based. In contrast most of our mechanisms are easy to compute, work for general distributions, enforce EPIR, and achieve small approximation factors.

In prior-free settings few results are known for revenue maximization. Borgs et al. (2005) looked at multi-unit auctions for homogeneous goods where agents have private values and budgets and considered worst case competitive ratio (see also Abrams, 2006). They designed a mechanism based on random sampling that maximizes revenue when the number of bidders is large.

Social welfare maximization has also been considered under budget constraints. Maskin (2000) considered the setting of a single item and multiple buyers with public budgets. He defined and showed how to compute the constrained efficient mechanism, the truthful feasible mechanism under budget constraints that maximizes the expected social welfare (however, the result holds only for some distribution

[1] E.g., Pai and Vohra (2008) and Maskin (2000) make the assumption that value distributions have a monotone hazard rate as well as a nondecreasing density function, unnatural conditions that few distributions satisfy simultaneously.

functions (Pai and Vohra, 2008)). In prior-free settings for multi unit homogeneous items, Dobzinski et al. (2008) studied Pareto efficient DSIC mechanisms with budget constraints. They showed that if the budgets are private there is no Pareto optimal incentive compatible mechanism; for public budgets they showed that there exists a unique mechanism based on the *clinching auction*. Chen et al. (2010) considered a setting with multiple goods and unit demand buyers and showed how to compute competitive prices that enforce truthfulness under budget constraints if such prices exist. Finally, the work of Alaei et al. (2010) stands out in their study of "soft" budgets constraints, where buyers pay an increasing interest rate for payments made above their budgets. They showed how to exactly compute the smallest competitive prices in this setting that result in an incentive compatible mechanism with an outcome in the core.

Relevant Techniques. Our setting of sequential posted pricing with a matroid constraint is very closely related to the so-called matroid secretary problem (Babaioff et al., 2007, 2009; Korula and Pál, 2009), but there are two important differences: (a) they assume that agents' values are adversarial, whereas in our setting they are drawn from known distributions, and (b) in their setting agents arrive in random order, whereas we consider optimized and adversarial orderings. Some of our results are reminiscent of that work, but our techniques are necessarily different.

Finally, our results for OPMs in the multi-unit auction setting are based on work on prophet inequalities from optimal stopping theory. While that work applies directly to the analysis of OPMs in the single-item auction setting, we show that it extends to k-unit auctions with no loss in approximation factor.

Further work. Some of these results have been improved upon by Yan (2011) and Alaei (2011). Devanur et al. (2011) and Roughgarden et al. (2012) extend these approaches to give "prior-independent" multi-item auctions for unit-demand agents; these auctions give constant approximations to the (prior-dependent) revenue optimal auction. Kleinberg and Weinberg (2012) extend the prophet inequalities to matroids, and achieve a constant approximation for general matroids via an adaptive posted-pricing mechanism for which truthful reporting is a dominant strategy. Recently, a series of papers (Cai and Daskalakis, 2011; Daskalakis and Weinberg, 2012; Cai et al., 2012) presented a different approach for computationally finding arbitrarily good approximations to the optimal multi-item auction for agents with unit-demand or additive values.

1.4 Bibliographic notes

The work presented in this thesis is based on the following joint works:

- The results in Chapters 3 and 4 are based on joint work with Shuchi Chawla, Jason Hartline, and Balasubramanian Sivan, previously appearing in STOC 2010 as "Multi-parameter mechanism design and sequential posted pricing"(see Chawla et al., 2010);

- The results in Chapter 5 are based on joint work with Shuchi Chawla and Balasubramanian Sivan, previously appearing in EC 2010 and in *Games and Economic Behavior* as "The power of randomness in Bayesian optimal mechanism design"(see Chawla et al., 2012); and

- The results in Chapter 6 are based on joint work with Shuchi Chawla and Azarakhsh Malekian, previously appearing in EC 2011 as "Bayesian mechanism design for budget-constrained agents"(see Chawla et al., 2011).

Feasibility constraint $\mathcal{S}$	Type of posted pricing	approximation
Uniform matroid, Partition matroid	OPM	2
Graphical matroid	OPM	3
General matroid	SPM	2

Table 1.1: A selection of approximation factors for single-dimensional settings through posted pricings

Feasibility constraint $\mathcal{S}$	Solution concept	approximation
Intersection of two part. matroids	DSIC	5.83
Matching with i.i.d. agents	DSIC	$2e/(e-1) \approx 3.17$
Graphical matroid $\cap$ partition matroid	DSIC	7.47
Intersection of two matroids	PDSE	8

Table 1.2: A selection of approximation factors for multi-dimensional unit-demand settings.

Feasibility constraint $\mathcal{S}$	Solution concept	approximation
Single agent	DSIC	4
Multi-item multi-agent auction	DSIC	29.15
General matroid $\cap$ unit-demand constraint	PDSE	40

Table 1.3: A selection of approximation factors between deterministic and randomized mechanisms for multi-dimensional unit-demand settings.

Feasibility constraint $\mathcal{S}$	objective	budget	approximation
General	revenue	public	2
Downwards-closed	revenue	private	$3(1+e)$
General	welfare	private	$2(1+e)$

Table 1.4: A selection of approximation factors for single-service settings with budget-constrained agents.

2 Definitions and notations

2.1 Bayesian settings for mechanism design

In this section, we give formal definitions for each type of mechanism design instance we consider throughout this work. In all of these settings, a seller offers one or more services to a collection of agents, who place (possibly different) values on the services they might receive. We focus on Bayesian settings, where the seller has distributional information about the agents' values, and wants to choose feasible assignments of services to agents in order to maximize some prespecified objective in expectation with respect to these distributions. At a high level, for every setting we consider, an instance is specified by four characteristics:

- the set of agents interested in receiving service;
- the set of services the seller can provide;
- a description of how the agents' values are distributed; and
- a description of what allocations of services to agents are simultaneously feasible.

We now make this characterization formal for the various settings considered in this thesis; at the end of the section, we consider the last two of these criteria in more detail.

BSMD: Bayesian Single-parameter Mechanism Design problem. The Bayesian single-parameter mechanism design problem (BSMD for short) is an abstraction of the setting where each agent has a single private value for any "good outcome" of the mechanism. In this setting, the mechanism can produce a good outcome for agent i in which case their valuation is v_i, or a bad outcome in which case their valuation is zero. We denote an instance of BSMD by the tuple $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$, where:

- $I = [n]$ is a set of n agents.
- $\mathcal{S} \subset 2^I$ is a feasibility constraint. It specifies the sets of agents the seller can simultaneously serve. We assume $\mathcal{S}$ is downward closed, i.e., for $S \in \mathcal{S}$ all subsets $S' \subset S$ are in $\mathcal{S}$ as well.

- $\mathbf{F} = F_1 \times \cdots \times F_n$ is the joint product distribution on agent values for being served. I.e., v_i is drawn independently from distribution F_i with density function f_i.

Bayesian multi-parameter unit-demand mechanism design. The Bayesian multi-parameter unit-demand mechanism design problem (BMUMD for short) is an abstraction of the setting where a seller can provide a number of different services to agents, where each agent desires at most one service. An agent values each service differently, but the values for each service are drawn independently from known distributions. Formally we denote an instance of this problem by the tuple $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ where:

- $I = [n]$ is a set of n agents.
- $J = [m]$ is a set of m services, which is partitioned as $\Pi = (J_1, \ldots, J_n)$ among the n agents.[1] The services in J_i are the ones being targeted at agent i. These agents are *unit-demand* in that they each desire at most one service from their partition.
- $\mathcal{S} \subset 2^J$ is a feasibility constraint. It specifies the sets of services the seller can simultaneously provide. As with the BSMD, we assume $\mathcal{S}$ is downward closed, i.e., for $S \in \mathcal{S}$ all subsets $S' \subset S$ are in $\mathcal{S}$. Further, we assume $\mathcal{S}$ respects the partitioning Π and the unit-demand constraint, i.e., $S \in \mathcal{S}$ and $i \in [n]$ implies $|S \cap J_i| \leqslant 1$.
- $\mathbf{F} = F_1 \times \cdots \times F_m$ is the joint product distribution on agent values over the m services. I.e., v_j is drawn from distribution F_j with density function f_j.

Note that any instance of the BSMD can be represented as a special case of the multi-dimensional setting where there is exactly one service available to each agent, i.e., $n = m$ and $J_i = \{i\}$.

Mechanism design problems with budgets. In Bayesian mechanism design problems with budgets, agents' preferences are guided not only by how much value they place on services, but also by how much they can afford to pay to receive

[1] This partitioning is for notational convenience only. Since we allow for an arbitrary feasibility constraint over the set J, the assumption that the sets J_i are disjoint is without loss of generality.

those services. We model this with a hard budget constraint, which places a precise upper limit on the amount an agent can pay for receiving service. We consider both settings where the budget constraint is public and common knowledge to both the mechanism designer and the other agents; and settings where an agent's budget is private information of the respective agent.

An instance of the BSMD with public budgets is given by the tuple $\mathcal{I} = (I, \mathcal{S}, \mathbf{F}, \mathbf{B})$; an instance of the BMUMD with public budgets is given by the tuple $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F}, \mathbf{B})$. In both cases, the only difference from the corresponding problems without budgets is the addition of the parameter $\mathbf{B}$; all other parameters retain the same definitions and meanings. The new parameter $\mathbf{B}$ is a vector with B_i being the budget of agent i. In the case of private budgets, we simply replace $\mathbf{B}$ with a (product) distribution $\mathbf{G} = \prod_i G_i$ and agent i has a budget B_i drawn independently from distribution G_i.

2.1.1 The distribution F

Throughout most of this work, we assume agents' values are all independently distributed, that is, drawn from a product distribution. In particular, we assume both that values of different agents are independent of each other, and that when there are multiple services available, the values a given agent places on them are also independent. While this assumption is necessary to many of our techniques, it can be quite a strong assumption; in particular, in BMUMD settings where an agent wants to receive only one among many alternatives, it seems unnatural to assert that the values of these options are wholly independent.

We seek to address the above concerns in Section 5.3 of Chapter 5, where we consider correlated values. In particular, we consider the *common base value* model, where an agent's values for the various services available are split into a portion that is common to all the services (introducing a natural form of correlation for BMUMD settings) and a portion specific to the particular service. More formally, each agent i is assumed to have an $(m+1)$-dimensional type $\{t_{i0}, \cdots, t_{im}\}$ with each t_{ij} being distributed independently according to a known distribution. Agent i's value for service j is then given by $v_{ij} = t_{i0} + t_{ij}$. Here, t_{i0} represents the agent's base value for receiving any service and the remaining t_{ij}'s capture the agent's preferences among the different services.

2.1.2 The feasibility constraint $\mathcal{S}$

In this work, all of our results require that the feasibility constraint $\mathcal{S}$ be *downward-closed*. A set system $\mathcal{S}$ is said to be downward-closed if we have that for any $T \in \mathcal{S}$, $S \subset T$ implies $S \in \mathcal{S}$. While in some instances this restriction on $\mathcal{S}$ is sufficient to enable our results, in many cases we find that stronger conditions are required to achieve positive results. One particular family of set systems that prove important to our results are *matroids*, which we describe in the remainder of this subsection.

Matroids and related set systems. Many of our techniques work for feasibility constraints $\mathcal{S}$ that are matroids or close to matroids. We define these set systems here. The set system $(X, \mathcal{S})$ over a universe X with $\mathcal{S} \subseteq 2^X$ is called a matroid if it satisfies the following conditions:

1. (**h**eredity) For every $A \in \mathcal{S}$, $B \subset A$ implies $B \in \mathcal{S}$.
2. (**a**ugmentation) For every $A, B \in \mathcal{S}$ with $|A| > |B|$, there exists $e \in A \setminus B$ such that $B \cup \{e\} \in \mathcal{S}$.

Sets in $\mathcal{S}$ are called independent, and maximal independent sets are called *bases*. A simple consequence of the above properties is that all bases are equal in size. The rank of a set $S \subseteq X$ is defined to be the size of any maximal independent subset of S. The *span* of a set $S \subseteq X$, $\text{span}(S)$, is the maximal set $T \supseteq S$ with $\text{rank}(T) = \text{rank}(S)$.

Here we give some examples of special matroid constraints. A k-uniform matroid on the universe X is a matroid where every subset of X of size at most k is independent. A partition matroid $(X, \mathcal{S})$ is a union of two or more uniform matroids $\{(X_i, \mathcal{S}_i)\}_i$, where $\{X_i\}_i$ is a partition of X and $\mathcal{S} = \{\cup_i A_i : A_i \subseteq \mathcal{S}_i \ \forall i\}$. A transversal matroid $(X, \mathcal{S})$ is defined by a bipartite graph $G = (V, E)$, where V is partitioned as $X \cup Y$ for some Y. The independent sets $\mathcal{S}$ are precisely those subsets of X that can be matched one-to-one to a subset of Y in G.[2] A set system $(X, \mathcal{S})$ is called a matroid intersection if there are two matroids $(X, \mathcal{S}_1)$ and $(X, \mathcal{S}_2)$, such that $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2$. An example of a matroid intersection is a matching in a bipartite graph.

The following proposition is a simple consequence of the above conditions on matroid set systems and will be useful in our analysis.

---

[2]In fact, we can require that the subsets of X in $\mathcal{S}$ have matchings to independent sets in some matroid $(Y, \mathcal{S}')$, and the result is still a transversal matroid. See, e.g., Aigner (1997) for more details.

Proposition 2.1. *Let* B_1 *and* B_2 *be arbitrary independent sets in some matroid set system* $\mathcal{S}$*. Then there exists a set* $B_2' \subseteq B_2$ *and a one to one function* $g : B_2' \to B_1$ *such that for all* $e \in B_2'$, $B_1 \setminus \{g(e)\} \cup \{e\}$ *is independent in* $\mathcal{S}$*, and for all* $e \in B_2 \setminus B_2'$, $B_1 \cup \{e\}$ *is independent in* $\mathcal{S}$*.*

Proof. Our proof relies on the following well known theorem (see, for example, Brualdi, 1969).

Theorem 2.1. *If* B_1 *and* B_2 *are two bases of a matroid, then there exists a one to one function* $g : B_2 \to B_1$ *such that* $(B_1 \setminus \{g(e)\}) \cup \{e\}$ *is a basis for all* $e \in B_2$*.*

Now, in order to apply Theorem 2.1 we need two bases. Let B be a basis of $\mathcal{S}$. Repeatedly apply the augmentation property to B_2 and B to produce a basis $\bar{B_2} \supset B_2$, and then do the same with B_1 and $\bar{B_2}$ to produce a basis $\bar{B_1} \supset B_1$. Now, Theorem 2.1 guarantees us a one to one function $g : \bar{B_2} \to \bar{B_1}$ such that for all $e \in \bar{B_2}$, $\bar{B_1} \setminus \{g(e)\} \cup \{e\}$ is independent. Note that for all $e \in \bar{B_2} \cap \bar{B_1}$, we must have that $g(e) = e$, since otherwise

$$\bar{B_1} \setminus \{g(e)\} \cup \{e\} = \bar{B_1} \setminus \{g(e)\} \subsetneq \bar{B_1},$$

and so is not a basis. Since g is one to one, this means that $g(e) \in \bar{B_1} \setminus \bar{B_2}$ for all $e \in \bar{B_2} \setminus \bar{B_1}$.

Set ${B_2}' = B_2 \setminus \bar{B_1} \subset \bar{B_2} \setminus \bar{B_1}$. Since $\bar{B_1} \setminus \bar{B_2} \subset B_1$, we may view g as a one to one function $g : {B_2}' \to B_1$. It has the first specified property, since for any $e \in {B_2}'$, $B_1 \setminus \{g(e)\} \cup \{e\} \subset \bar{B_1} \setminus \{g(e)\} \cup \{e\}$ is independent. Furthermore, $e \in B_2 \setminus {B_2}' \subset \bar{B_1}$ implies $B_1 \cup \{e\} \subset \bar{B_1}$ is independent, and so the second specified property holds as well. □

While we do not need the following proposition until Section 5.3, we present it now in the interest of collecting the results we need on matroids in one place. We present the proposition without proof (see, e.g. Korte and Hausmann, 1978, Theorems 1.7 and 3.2).

Proposition 2.2. *Let* $\mathcal{S}_1$ *and* $\mathcal{S}_2$ *be arbitrary matroid set systems over a common ground set. Assume that each element* e *in the ground set has a weight* w_e*. Then if we construct an independent set* $\mathcal{G} \in \mathcal{S}_1 \cap \mathcal{S}_2$ *in a greedy fashion, i.e. by repeatedly selecting the highest* *weight element we can feasibly add to $\mathcal{G}$ until none are left, then*

$$\sum_{e \in \mathcal{G}} w_e \geqslant 1/2 \sum_{e \in S} w_e$$

for any $S \in \mathcal{S}_1 \cap \mathcal{S}_2$; furthermore, this holds no matter how we break ties (if any) when constructing $\mathcal{G}$.

2.2 Mechanism design desiderata

A deterministic mechanism for the settings we consider in this work maps any set of bids $\mathbf{b}$ made by the agents to an allocation $M(\mathbf{b}) \in \mathcal{S}$ and a pricing $\pi(\mathbf{b})$ with a price π_i to be paid by agent i. A randomized mechanism maps a set of bids to a distribution over $\mathcal{S}$; we use $M(\mathbf{b})$ to denote this distribution.

The participants in a setting evaluate their outcomes in terms of *utility*. The utility an agent i receives with a valuation v_i under a bid profile $\mathbf{b}$ is the value they receive minus the payment they must make: $v_i \cdot M_i(\mathbf{b}) - \pi_i(\mathbf{b})$. In the case of randomized mechanisms, we assume that agents are *risk-neutral*, and experience their expected utility: $E[v_i \cdot M_i(\mathbf{b}) - \pi_i(\mathbf{b})]$, where the expectation is over any randomization the mechanism performs.

We use $\mathcal{R}_{\mathcal{J}}^M(\mathbf{v})$ to denote the revenue of a mechanism M for instance $\mathcal{J}$ at valuation vector $\mathbf{v}$: $\mathcal{R}_{\mathcal{J}}^M(\mathbf{v}) = \sum_{i \in I} \pi_i(\mathbf{v})$. We drop the subscript $\mathcal{J}$ when it is clear from the context. To aid disambiguation, we sometimes use $\mathcal{R}_i^M(\mathbf{v})$ to denote $\pi_i(\mathbf{v})$ for M. The expected revenue of a mechanism is $\mathcal{R}^M = E_{\mathbf{v}}[\mathcal{R}^M(\mathbf{v})]$.

We now list desiderata for the mechanism designer.

Objective. We consider two objectives:

- Social welfare, or the sum of utilities of all of the participants. Since payments act as a transfer of utility from the buyers to the seller, they cancel out and social welfare can be written as simply the total value all buyers receive from their respective allocations.
- Revenue of the seller, or the total sum of all payments made by buyers to the seller.

Mechanisms that maximize social welfare are called (economically) efficient, while those that maximize the seller's revenue are called optimal.

Compatibility with agents' incentives. Since agents act in their own best interests, any mechanism we propose must respect their incentives. In the settings we consider, this has two components:

- Voluntary participation, or individual rationality. This property requires that every agent be better off participating in the mechanism than not, i.e. every agent receives nonnegative utility from the mechanism.
- Truthfulness. This property requires that it is always utility-maximizing for an agent to report their value truthfully. While this may appear to be a strong property, in fact the Revelation Principle (see Nisan (2007) for a formal statement and details) ensures that we may assume truthfulness without loss of generality.

Depending on the context, we may require the above properties hold in either dominant strategies or as a Bayes-Nash equilibrium, as appropriate.

In some settings, it is challenging to obtain truthfulness in dominant strategies. In such cases, we relax our goals to a weaker solution concept called *partial dominant strategy equilibrium* (PDSE). An outcome is in PDSE if all agents that have dominant strategies follow said strategies and other agents follow (arbitrary) undominated strategies.

2.3 Optimal single-parameter mechanism design

Myerson (1981) describes the revenue maximizing mechanism for the Bayesian single-parameter mechanism design problem. *Virtual valuations* are given by the formula $\phi(v) = v - \frac{1-F(v)}{f(v)}$. When the value distributions F_i are *regular*, i.e., virtual valuations are monotone in valuations, the optimal mechanism first computes virtual values for each agent, and then allocates to a feasible subset of agents that maximizes the "virtual surplus"—the sum of the virtual values of agents in the set minus the cost of serving that set of agents (Myerson, 1981). For a single agent, this mechanism allocates to the agent as long as his value is above the threshold $\phi^{-1}(0)$; we call this threshold the *monopoly reserve price* corresponding to the value distribution.

When the distributions F_i are irregular, that is, virtual valuations are not monotone in valuations, Myerson's mechanism as described above will no longer be truthful. Myerson addressed this case by "ironing" the virtual valuation function and converting it into a monotone non-decreasing function called the ironed virtual value function denoted by $\bar{\phi}(v)$. We skip the description of this procedure (the reader is referred to Bulow and Roberts, 1989; Chawla et al., 2007, for details).

Some of our results require a stronger condition on distributions called the monotone hazard rate condition, a common assumption in mechanism design literature. This condition is satisfied by many common distributions such as the uniform, Gaussian, exponential, and power law distributions.

Definition 2.2. *A distribution F with density f is said to have a **m**onotone hazard rate if the function $h(v) = f(v)/(1 - F(v))$ is non-decreasing in v. Distributions satisfying MHR are regular.*

We use $\mathcal{R}^{\mathcal{M}}_{\mathcal{I}}$ to denote the expected revenue of Myerson's mechanism on a single-parameter instance $\mathcal{I}$. For our analyses, we primarily require the following three characterizations of incentive compatible mechanisms, all due to Myerson (1981). The first tells us that we have Bayesian incentive compatibility if and only if we have a monotone allocation rule, and furthermore that the payment rule induced by a monotone allocation rule is unique up to additive "shifts". The last two formalize our earlier claim that revenue maximization is equivalent to (ironed) virtual value maximization.

Theorem 2.3. *(Myerson, 1981) A single-parameter mechanism with allocation rule $\mathbf{x}(\cdot)$ and payment rule $\mathbf{p}(\cdot)$ is Bayesian incentive compatible if and only if for all i:*

1. *$x_i(v_i)$ is monotone nondecreasing in v_i and*
2. *$p_i(v_i) = v_i x_i(v_i) - \int_0^{v_i} x_i(z)dz + p_i(0)$.*

Proposition 2.3. *(Myerson, 1981) When the distributions F_i are regular, the expected revenue of any incentive compatible single-parameter mechanism M is equal to its expected virtual surplus.*

Proposition 2.4. *(Myerson, 1981) The expected revenue of any incentive compatible single-parameter mechanism M is no more than its expected ironed virtual surplus. If the probability with which the mechanism serves agent i, as a function of v_i, is constant over any valuation* *range in which the ironed virtual value of* i *is constant, the expected revenue is equal to expected ironed virtual surplus.*

3 Sequential pricings for single-parameter settings

Consider a single-parameter setting where each agent has a private value for service and there is a combinatorial feasibility constraint on the set of agents that can be simultaneously served. For this setting a *sequential posted pricing* (SPM) is a mechanism defined by a price for each agent, a sequence on agents, and the semantics that each agent is offered their corresponding price in sequence as a take-it-or-leave-it while-supplies-last offer. Meaning: if it is possible to serve the agent given the set of agents already being served then the agent is offered the price. A rational agent will accept if and only if the price is no more than their private value for service. That prices are associated with the agents and not the sequence reflects the possibility that agents may play asymmetric roles for a given feasibility constraint or value distribution.

Consider the following hotel rooms example with one room, two attendees, and attendee values independently and identically distributed uniformly between $100 and $200. The optimal mechanism is the Vickrey auction and its expected revenue is $133. The optimal sequential posted pricing is for the organizers to offer the room to attendee 1 at a price of $150. If the attendee accepts, then the room is sold, otherwise it is offered to attendee 2 for $100. The expected revenue of this SPM is $125.

We are interested in comparing the optimal mechanism to the optimal posted pricing in general settings. A special class of SPMs is one where mechanisms have provable performance guarantees for any sequence of the agents. These *order-oblivious posted pricings* (OPM) are mechanisms defined by a price for each agent and the semantics that each agent is offered their corresponding price in some arbitrary sequence as a take-it-or-leave-it while-supplies-last offer.

In single-parameter settings, the advantages of sequential posted pricings speak to the many reasons optimal auctions are rarely seen in practice (Ausubel and Milgrom, 2006), and explain why posted pricings are ubiquitous (Holahan, 2008). First, take-it-or-leave-it offers result in trivial game dynamics: truthful responding is a dominant strategy. Second, SPMs satisfy strong notions of collusion resistance, e.g., *group strategyproofness* (see Goldberg and Hartline, 2005): the only way in which an agent can "help" another agent is to decline an offer that he could have accepted, thereby hurting his own utility. Third, agents do not need to precisely know or report their value, they must only be able to evaluate their offer; therefore, they risk minimal exposure of their private information. Fourth, agents learn immediately whether they will be served or not. In conclusion, the robustness of SPMs in single-parameter settings makes their approximation of optimal mechanisms independently worthy of study.

The final robustness property of SPMs, which will be of paramount importance when we study multi-parameter settings in later chapters, is that they minimize the role of agent competition, implying that single-parameter SPMs can be used "as-is" in multi-parameter settings with only a constant factor loss in performance. In our translation from the multi-parameter setting to the single-parameter setting, each multi-parameter agent has many single-parameter representatives. A good OPM for the single-parameter setting can be viewed as an OPM for the multi-parameter setting by grouping all representatives of an agent together and making their offers simultaneously to the agent. The agent will of course accept the offer that maximizes their utility. The resulting mechanism is incentive compatible and achieves the same performance guarantee as the single-parameter OPM. For SPMs where we are not free to group each multi-parameter agent's single-parameter representatives together, an agent possibly faces a strategic dilemma of whether to accept an offer (e.g., for one hotel room) early on or wait for a later offer (e.g., another hotel room) which may or may not still be available. Our guarantee is robust to the actions of any agent with such a strategic option; if all agents with dominant strategies follow said strategies then our performance is a constant fraction of the original SPM's performance. (This is a non-standard notion of *dominant strategy implementation*.) Given the advantages of SPMs over standard dominant strategy mechanisms, these partially dominant strategy mechanisms may be more practically relevant.

Finally, we note that most of our results for posted pricings are constructive and there are efficient algorithms for them. A posted price mechanism has two components where computation is necessary: an offline computation of the prices to post (and for SPMs, the sequence of agents) and an online while-supplies-last offering of said prices.[1] The agents are only present for the online part where the mechanism is trivial. All of the computational burden for an SPM is in the

[1] This is similar, for example, to nearest neighbor algorithms, where one distinguishes the time taken to construct a database, and the time taken to compute nearest neighbors over that database given a query.

offline part. The offline computation of our posted price mechanisms is based on a subroutine that repeatedly samples the distribution of agent values and simulates Myerson's mechanism on the sample. This clearly requires more computation than just running Myerson's mechanism on the real agents in the first place; however, we benefit from the robustness that comes from the trivial online implementation of posted pricings.

3.1 Preliminaries

In this section, we formally define sequential posted-price mechanisms, and discuss technical issues related to their implementation. While the current chapter is focused on developing mechanisms for the single-parameter setting, in the interests of completeness we discuss posted-pricing mechanisms in the context of the more general multi-parameter setting, since we implement them in such settings in later chapters.

3.1.1 Posted-price mechanisms

We will consider sequential posted-price mechanisms based on the following high-level protocol that is parameterized by $\mathbf{p}$, a vector of prices, one for each service, and σ, an ordering over the services.

The generic sequential posted pricing protocol for $(\mathbf{p}, \sigma)$ is as follows:

1. Initialize $A \leftarrow \emptyset$.
2. For $j = 1$ through m, do:
   a) If $A \cup \{\sigma(j)\} \in \mathcal{S}$, offer service $\sigma(j)$ at price p_j.
   b) If the agent accepts, $A \leftarrow A \cup \{\sigma(j)\}$.
3. Provide the services in A to the corresponding agents.

We denote the revenue of this mechanism on valuation profile $\mathbf{v}$ by $\mathcal{R}^{(\mathbf{p},\sigma)}_{(I \times J, \mathcal{S}, \mathbf{F})}(\mathbf{v})$. In the strategically-simple single-parameter setting this revenue can be calculated with the assumption that service j is accepted by the agent when offered if $v_j \geqslant p_j$. We defer the discussion of incentives in the more complicated multi-parameter setting to Section 3.1.2.

It is clear that it is always better from the designer's point of view to be able to choose the ordering σ. Unfortunately, this may not always be possible. We therefore distinguish between the following two kinds of posted-price mechanisms.

Sequenced Posted-Price Mechanisms. An SPM is given by $(\mathbf{p}, \sigma)$. Its expected revenue is:

$$\mathcal{R}^{(\mathbf{p},\sigma)}_{(I,\mathcal{S},\mathbf{F})} = \mathrm{E}_{\mathbf{v}\sim\mathbf{F}}[\mathcal{R}^{(\mathbf{p},\sigma)}_{(I,\mathcal{S},\mathbf{F})}(\mathbf{v})]$$

Order-oblivious Posted-pricing Mechanisms. An OPM is given simply by the pricing $\mathbf{p}$ where we allow the order over the services to be picked adversarially after the valuations of the agents are drawn. This pessimistically bounds the worst possible revenue for a given pricing. Formally:

$$\mathcal{R}^{\mathbf{p}}_{(I,\mathcal{S},\mathbf{F})} = \mathrm{E}_{\mathbf{v}\sim\mathbf{F}}[\min_{\sigma} \mathcal{R}^{(\mathbf{p},\sigma)}_{(I,\mathcal{S},\mathbf{F})}(\mathbf{v})]$$

When it is clear from the context we will omit the subscript $(I \times J, \mathcal{S}, \mathbf{F})$ or $(I, \mathcal{S}, \mathbf{F})$.

In some settings we consider randomized versions of SPMs and OPMs where the pricing $\mathbf{p}$ is picked randomly. In this case, we assume that the prices are drawn first and then the order σ is determined based on the prices (adversarially or by the designer).

3.1.2 Incentives

Most of the literature on mechanism design (especially in computer science) focuses on sealed-bid single-round direct-revelation mechanisms. These are mechanisms that consist of two steps: first agents report bids, to be interpreted as their preferences over possible outcomes of the mechanisms, and second the mechanism selects an outcome and agent payments. In this context a mechanism is *incentive compatible* if each agent has a (weakly) dominant strategy of truthful reporting. It is assumed that agents report their true preferences in an incentive compatible mechanism.

Our posted price mechanisms do not take this general single-round form. Instead our mechanism will offer each agent a sequence of prices (and these offers may be arbitrarily interleaved among agents). Strategically, a bidder i when offered price p_j for $j \in J_i$ has two options. They can accept or reject the offer. An agent with value v_j for service j is *sincere* if they accept offers $p_j \leqslant v_j$ and reject offers $p_j > v_j$.

Sincere bidding is a dominant strategy for an agent only when the ordering σ respects the agent's incentives. Formally, we say that an ordering σ is J_i-respecting if for all $j_1, j_2 \in J_i$, $v_{j_1} - p_{j_1} > v_{j_2} - p_{j_2} \geqslant 0$ implies $\sigma^{-1}(j_1) < \sigma^{-1}(j_2)$. That is, the offers made to agent i are ordered by decreasing utility for the agent (although they may be interleaved arbitrarily with offers for other agents). An ordering is Π-respecting if it is J_i-respecting for all i. The following lemma formalizes the connection between sincere bidding and Π-respecting orderings.

Lemma 3.1. *For $\mathbf{v}_{J_i}$ (the values of agent i) sincere bidding is a (weakly) dominant strategy for i in sequential posted pricing $(\mathbf{p}, \sigma)$ if and only if σ is J_i-respecting.*

It is easy to see that if the condition on σ is met then an agent will have no reason not to respond sincerely. If the condition is not met, an agent may strategize in the following way. When offered item j' at a desirable price $p_{j'}$, the agent might reject the offer in hopes of later being offered item j for which the agent has even higher utility.

The condition of the lemma is met in three special cases of interest:

1. When the agent is single dimensional, i.e., $|J_i| = 1$.
2. When the agent has positive utility for at most one service, i.e., $|\{j \in J_i : v_j - p_j > 0\}| \leqslant 1$
3. When the agent can choose the relative order of σ on J_i.

Given the last point, for an OPM in the multi-parameter setting we assume that orderings of interest are Π-respecting and define the worst-case revenue of the mechanism accordingly:

$$\mathcal{R}^{\mathbf{p}}_{(I \times J, \mathcal{S}, \mathbf{F})} = \mathrm{E}_{\mathbf{v} \sim \mathbf{F}} \left[\min_{\sigma : \sigma \text{ is } \Pi\text{-respecting}} \mathcal{R}^{(\mathbf{p}, \sigma)}_{(I \times J, \mathcal{S}, \mathbf{F})}(\mathbf{v}) \right]$$

In SPMs in multi-parameter settings, the ordering is not necessarily Π-respecting. We assume that all bidders for whom sincere bidding is a (weakly) dominant strategy indeed bid sincerely. We derive robust bounds on our mechanism performance in the presence of arbitrary manipulations of agents that do not have dominant strategies. We term this solution concept partial dominant strategy equilibriumimplementation (see Section 2.2 for details).

As is standard in mechanism design, we assume that the agents understand the mechanism. As our mechanisms are parameterized by prices $\mathbf{p}$, it is assumed that agents know these prices in advance. This assumption is only necessary for SPMs when we implement them in multi-parameter settings where an agent i must know whether there is a future offer p_j for $j \in J_i$ such that $v_j - p_j > 0$.

Lemma 3.2. *If F_i is regular for each i, the revenue of any incentive compatible mechanism M over the n agents is bounded from above by $\sum_i p_i^M q_i^M$ where q_i^M is the probability (over $v_1, \cdots, v_n$) with which M allocates to agent i and $p_i^M = F_i^{-1}(1 - q_i^M)$.*

Furthermore for every i (with a regular or non-regular value distribution), there exist two prices $\underline{p_i}$ and $\overline{p_i}$ with corresponding probabilities $\underline{q_i}$ and $\overline{q_i}$, and a number $x_i \leqslant 1$, such that $x_i \underline{q_i} + (1 - x_i)\overline{q_i} = q_i^M$, and the expected revenue of M is no more than $\sum_i x_i \underline{p_i q_i} + (1 - x_i)\overline{p_i q_i}$.

Proof. We prove the regular case first. Consider the revenue that M draws from serving agent i. This is clearly bounded above by the optimal mechanism that sells to only i, but with probability at most q_i^M. By Proposition 2.3, such a mechanism should sell to agent i with probability 1 whenever the value of the agent is above $F_i^{-1}(1 - q_i^M)$ and with probability 0 otherwise. The revenue of the optimal such mechanism is therefore $p_i^M q_i^M$.

In the non-regular case, note that the value p_i^M may fall in a valuation range that has constant ironed virtual value. Let $\underline{p_i}$ denote the infimum $\inf\{v : \bar{\phi}_i(v) = \bar{\phi}_i(p_i^M)\}$ of this range and $\overline{p_i}$ denote the supremum $\sup\{v : \bar{\phi}_i(v) = \bar{\phi}_i(p_i^M)\}$. Let $\underline{q_i} = 1 - F_i(\underline{p_i})$ and $\overline{q_i} = 1 - F_i(\overline{p_i})$. Then, $\overline{q_i} \leqslant q_i^M \leqslant \underline{q_i}$, and there exists an x_i such that $x_i \underline{q_i} + (1 - x_i)\overline{q_i} = q_i^M$. Now an easy consequence of Proposition 2.4 is that the optimal mechanism with selling probability q_i^M sells to the agent with probability x_i if the agent's value is between $\underline{p_i}$ and $\overline{p_i}$, and with probability 1 if the value is above $\overline{p_i}$. The revenue of this mechanism is exactly $x_i \underline{q_i p_i} + (1 - x_i)\overline{q_i p_i}$. ☐

3.1.3 Computing the posted prices

For all but one of the approximately-optimal posted-price mechanisms that we present, prices and orderings can be computed efficiently in a computational model where we have black box access to the distribution $\mathbf{F}$. Please see Section 3.4 for details.

3.2 Order-oblivious sequential pricings

In this section we instantiate the reduction described in Section 4.1 for several different settings. We begin by developing our approach in the context of a "k item auction" setting, obtaining a (tight) 2-approximation in that setting. We then expand this approach to obtain constant-factor approximations for other matroids as well as matroid intersection constraints. For example, we show that in a setting with heterogeneous items and multiple unit-demand agents with identically distributed values (that we call the "supermarket setting"), there exists a pricing of the items that obtains a $2e/(e-1)$ approximation to the optimal mechanism.

Our most general positive result is a non-constructive $O(\log k)$-approximation for arbitrary matroids, where k is the rank of the matroid. On the negative side, we show that the gap between OPMs and the optimal achievable revenue can be as large as $\frac{\log n}{\log \log n}$ for non-matroid feasibility constraints, where n is the number of agents. Table 3.1 summarizes the approximation factors for all the settings considered.

k-uniform matroids

To illustrate our main technique for designing good OPMs we begin by presenting a simple 4-approximation for the case where $\mathcal{S}$ is a k-uniform matroid and $\mathbf{F}$ is regular. Our approach is to determine the probabilities q_i^M with which Myerson's mechanism serves each agent. We then offer to each agent a price that is accepted by the agent with probability roughly the same as q_i^M (or smaller by a constant factor). We then exploit properties of the feasibility constraint to show that with these prices the probability that an agent gets "blocked" by other agents and does not receive an offer is small for every agent.

Theorem 3.3. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD with $\mathcal{S}$ being a k-uniform matroid and $\mathbf{F}$ being regular. Then, there exists a set of prices $\mathbf{p}$ such that $\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geqslant \frac{1}{4}\mathcal{R}_{\mathcal{I}}^{\mathcal{M}}$.*

Proof. For each $j \in J$, let q_j^M be the probability with which an optimal mechanism would sell service j, and $p_j^M = F_j{}^{-1}(1 - q_j^M)$. Set $q_j = q_j^M/2$ and $p_j = F_j{}^{-1}(1 - q_j) \geqslant p_j^M$. We show that these prices $\mathbf{p}$ achieve the claimed approximation. Define c_j to be the probability the OPM $\mathbf{p}$ may still offer service j when it is reached in some arbitrary ordering σ. Then we can see that for all $j \in J$,

$c_j \geqslant \Pr[\text{The OPM } \mathbf{p} \text{ sells less than } k \text{ services}]$. On the other hand,

$$\begin{aligned} &\mathrm{E}[\text{Number of services sold by OPM } \mathbf{p}] \\ &\qquad \leqslant \mathrm{E}[\text{Number of services desired at prices } \mathbf{p}] \\ &\qquad \leqslant \sum_{j \in J} q_j \leqslant k/2, \end{aligned}$$

by our choice of q_j, and the fact that no mechanism can sell more than k services in setting $\mathcal{I}$. Thus, applying Markov's inequality we can see that $c_j \geqslant 1/2 \, \forall j$, and,

$$\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} = \sum_{j \in J} c_j q_j p_j \geqslant \frac{1}{4} \sum_{j \in J} p_j^M q_j^M \geqslant \frac{1}{4} \mathcal{R}_{\mathcal{I}}^{\mathcal{M}},$$

where the last inequality follows from Lemma 3.2. □

We now demonstrate an improved (and tight) 2-approximation. Our analysis follows an approach developed in the context of prophet inequalities by Samuel-Cahn (1984). The following lemma encapsulates our use of techniques from that literature.

Lemma 3.4. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD with a k-uniform matroid constraint. Given any set of prices $\mathbf{p}$, define $q_j = 1 - F_j(p_j)$ and $q = \sum_{j \in J} q_j$. Then there exist prices $\mathbf{p}'$ such that ${p_j}' \geqslant p_j$ for all $j \in J$ and*

$$\mathcal{R}_{\mathcal{I}}^{\mathbf{p}'} \geqslant \frac{1}{1 + q/k} \sum_{j \in J} p_j q_j.$$

Proof. We define the desired prices $\mathbf{p}'$ as follows. For a random variable X, let $(X)^+$ denote the positive portion of X, i.e. $(X)^+ = \max(0, X)$. Let r^* be the unique solution to the equation

$$kr = \sum_{j \in J} q_j (p_j - r)^+.$$

Thus, we have that

$$\begin{aligned}\sum_{j\in J} q_j p_j &\leqslant \sum_{j\in J} q_j \left(r^* + (p_j - r^*)^+\right)\\ &= \left(r^* \sum_{j\in J} q_j + \sum_{j\in J} q_j (p_j - r^*)^+\right)\\ &\leqslant kr^* \left(1 + q/k\right).\end{aligned}$$

For all $j \in J$, define the prices (and associated probabilities)

$$\begin{aligned}&{p_j}' = \max\{p_j, r^*\}, \text{ and}\\ &{q_j}' = 1 - F_j({p_j}').\end{aligned}$$

Then r^* is also the unique solution to the equation

$$kr = \sum_{j\in J} {q_j}'({p_j}' - r).$$

Now, as in the proof of Theorem 3.3, define c_j to be the probability that OPM $\mathbf{p}'$ may still offer service j when it is reached; then we can see that for all $j \in J$, $c_j \geqslant \Pr[\text{The OPM } \mathbf{p}' \text{ sells less than k services}]$. So we can lower bound the revenue from $\mathbf{p}'$ as

$$\begin{aligned}\mathcal{R}_{\mathcal{J}}^{\mathbf{p}'} &= \sum_{j\in J} c_j {q_j}' {p_j}' = \sum_{j\in J} c_j {q_j}' r^* + \sum_{j\in J} c_j {q_j}'({p_j}' - r^*)\\ &= r^* \cdot \mathrm{E}[\text{No. of services sold}] + \sum_{j\in J} c_j {q_j}'({p_j}' - r^*)\\ &\geqslant kr^* \cdot \Pr[\text{k services sold}]\\ &\qquad + \left(\sum_{j\in J} {q_j}'({p_j}' - r^*)\right) \Pr[\text{less than k services sold}]\\ &= kr^*.\end{aligned}$$

Combining these gives us the claimed lower bound. □

We now proceed to prove the improved approximation mentioned earlier.

Theorem 3.5. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD where the constraint $\mathcal{S}$ is a k-uniform matroid. Then there exist prices $\mathbf{p}$ such that $\mathcal{R}^{\mathbf{p}}_{\mathcal{I}} \geqslant \frac{1}{2}\mathcal{R}^{\mathcal{M}}_{\mathcal{I}}$.*

Proof. We first consider the setting for regular distributions. Recall from Lemma 3.2 that for regular distributions the revenue of the optimal truthful mechanism for $\mathcal{I}$ is bounded by $\sum_{j \in J} p^M_j q^M_j$, where q^M_j is the probability with which the optimal mechanism allocates service j and $p^M_j = F_j{}^{-1}(1 - q^M_j)$. Since any mechanism for $\mathcal{I}$ can sell no more than k items, we know that $\sum_{j \in J} q^M_j \leqslant k$. The claimed result then follows immediately from applying Lemma 3.4 to $\mathcal{I}$ with prices $\mathbf{p}^M$. □

By observing that in a partition matroid there is no interaction between the constituent uniform matroid constraints we immediately get the following corollary.

Corollary 3.6. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD where $\mathcal{S}$ is a partition matroid. Then there exist prices $\mathbf{p}$ such that $\mathcal{R}^{\mathbf{p}}_{\mathcal{I}} \geqslant \frac{1}{2}\mathcal{R}^{\mathcal{M}}_{\mathcal{I}}$.*

A lower bound of 2. We now show that OPMs cannot approximate the optimal revenue to within a factor better than 2 even in the single-item setting. Consider a seller with one item and two agents. The first agent has a fixed value of 1. The second has a value of $1/\epsilon$ with probability ϵ and 0 otherwise, for some small constant $\epsilon > 0$. Then, the optimal mechanism can obtain a revenue of $2 - \epsilon$ by first offering a price of $1/\epsilon$ to the second agent, and then a price of 1 to the first if the second declines the item. On the other hand, if the mechanism is forced to offer the item to the first agent first, then it has two choices: (1) offer the item at price 1 to agent 1; the agent always accepts, and (2) skip agent 1 and offer the item at price $1/\epsilon$ to agent 2; the agent accepts with probability ϵ. In either case, the expected revenue of the mechanism is 1.

Graphical matroids

While we do not know how to obtain constant factor approximations through OPMs for general matroid feasibility constraints, we now demonstrate that OPMs are constant factor optimal for a large class of matroids, namely graphical matroids. The ground set for a graphical matroid is the set of all edges of an undirected graph; A subset of edges is independent if it forms a forest (that is, it contains no cycles).

In order to obtain an approximation, however, we need to extend our definition of OPMs to allow the mechanism to be more restrictive in enforcing feasibility.

Specifically, a **c**onstrained order-oblivious posted-price mechanism (COPM, for short) is given by the tuple $(\mathbf{p}, \mathcal{S}')$ where $\mathcal{S}' \subseteq \mathcal{S}$, and (as for OPMs) we allow the order σ over the agents to be picked adversarially, after the valuations of the agents are drawn. The selling protocol for a COPM offers a service j if the service along with previously allocated services is feasible in the set system $\mathcal{S}'$, and not merely in $\mathcal{S}$.

For graphical matroids, Babaioff et al. (2009) and Korula and Pál (2009) develop approaches for reducing this case to a partition matroid that in our setting yield a 8-approximation to the optimal revenue; we use a similar approach but exploit the connection between prophet inequalities and partition matroids to obtain a 3-approximation.

Theorem 3.7. *Let $\mathcal{I}$ be an instance of the BSMD with a graphical matroid feasibility constraint. Then there is a COPM $(\mathbf{p}, \mathcal{S}')$, where $\mathcal{S}'$ is a partition matroid, that 3-approximates $\mathcal{R}^{\mathcal{M}}$ for $\mathcal{I}$.*

Proof. Our technique here is to partition the elements of the matroid such that we may treat each part as a 1-uniform matroid yet still respect the original feasibility constraint, and achieve good revenue while doing so.

Let $G = (V, J)$ be the graph defining our matroid constraint, where J is the set of services/edges. As before, let q_j^M denote the probability with which edge (service) j is allocated by the optimal mechanism. Let $\delta(v)$ denote the set of edges incident on a vertex v, and for each $v \in V$ define $q_v = \sum_{j \in \delta(v)} q_j^M$. Now, we can see that

$$\sum_{v \in V} q_v = \sum_{j \in E} 2q_j^M \leqslant 2(|V| - 1),$$

This implies that there exists a v for which $q_v \leqslant 2$; Let $\delta(v)$ be one of our partitions. Note that the edge set $\delta(v)$ forms a cut in G, and so given an independent set of edges from $J \setminus \delta(v)$ we may add any single edge from $\delta(v)$ while retaining independence. We apply this argument recursively to $(V \setminus \{v\}, J \setminus \delta(v))$ to form the rest of our partition. At the end, we have a partition of J such that each part has total mass no more than 2 and any collection of edges using no more than one edge from each part is independent.

We first note how to obtain a simple 8-approximation and then describe the changes needed to obtain a 3-approximation. We define $\mathcal{S}'$ to be the union of 1-

uniform matroids, each over the different parts of J defined above. The prices $\mathbf{p}$ for the 8-approximation are defined as follows: $q_j = q_j^M/4$ and $p_j = F_j^{-1}(1 - q_j)$. Then, the optimal revenue is at most $4\sum_j p_j q_j$, whereas, our mechanism offers each service with probability at least 1/2, and therefore, obtains a revenue of $1/2\sum_j p_j q_j$.

To obtain a 3-approximation, we use the same constraint $\mathcal{S}'$ a before, but here we invoke Lemma 3.4. Consider a single part $X \subset J$ within $\mathcal{S}'$. Since X is a 1-uniform matroid, and we chose it such that $\sum_{j\in X} q_i^M \leqslant 2$, we can see that applying Lemma 3.4 to each part in turn will yield a set of prices $\mathbf{p}$ such that the revenue of the COPM $(\mathbf{p}, \mathcal{S}')$ is at least $\frac{1}{3}\sum_{j\in J} p_j^M q_j^M$. □

General matroids

We now prove the $O(\log k)$ approximation for general matroids, where k is the rank of matroid. We remark that a similar result was obtained by Babaioff et al. (2007) for the related matroid secretary problem. However, we show though an example at the end of this subsection that their approach cannot give a non-trivial approximation in our setting.

Theorem 3.8. *Let $\mathcal{I}$ be an instance of the BSMD with a matroid feasibility constraint. Then, there exists a set of prices $\mathbf{p}$ such that $\mathcal{R}_{\mathcal{I}}^{\mathbf{p}}$ $O(\log k)$-approximates the optimal mechanism's (Myerson's) revenue $\mathcal{R}_{\mathcal{I}}^{M}$.*

Proof. We first consider the setting for regular distributions. Recall from Lemma 3.2 that for regular distributions the revenue of the optimal truthful mechanism for $\mathcal{I}$ is bounded by $\sum_j p_j^M q_j^M$, where q_j^M is the probability with which the optimal mechanism allocates service j and $p_j^M = F_j^{-1}(1 - q_j^M)$. The OPM we describe sets the same prices as the optimal mechanism, i.e. $p_j = p_j^M$. Note that since the feasibility constraint is a matroid, for any instantiation of values, the worst (least revenue) allocation is achieved when agents arrive in the order of increasing prices. Hereafter we assume that agents always arrive in that order. Let c_i be the probability that service i gets offered in this ordering. Note that the expected revenue may be expressed as $\sum_i c_i p_i q_i$.

Now consider a hypothetical situation where the prices are all equal to 1 but the probabilities with which the agents accept the offered prices are still q_i. Then, the expected revenue of this hypothetical mechanism would be given by $\sum_i c_i q_i$ which is at least $1/2\sum_i q_i$ by the argument in Theorem 3.12. In other words, the weighted average of the c_is is at least 1/2, weighted by the q_is. We get that

$$(1/2)\sum_i q_i \leqslant \sum_i c_j q_i \leqslant \sum_{i:c_j<1/4} q_i/4 + \sum_{i:c_j\geqslant 1/4} q_i$$
$$= (1/4)\sum_i q_i + (3/4)\sum_{i:c_j\geqslant 1/4} q_i$$

and hence

$$\sum_{i:c_j\geqslant 1/4} q_i \geqslant (1/3)\sum_i q_i.$$

This means that the probability mass of elements having $c_j \geqslant 1/4$ is at least a third of the total. Let $G = \{i|c_j \geqslant 1/4\}$; the revenue obtained from serving only the agents in G is

$$\sum_{i\in G} c_i p_i q_i \geqslant 1/4 \sum_{i\in G} p_i^M q_i^M. \tag{3.1}$$

Consider recursively applying the above argument to the elements outside G. At step j, let G_j be the newly found G, and let E_j be the set of agents still under consideration, defined as $E_1 = [n]$ and $E_j = E_{j-1} - G_{j-1}$ for $j > 1$. Now, at each stage, G_j contains at least one third of the total probability mass of the remaining elements; thus, at stage $\ell = \lceil 1 + \log_{3/2} k \rceil$, we would have reduced the total probability mass to less than 3/4; by noting that any singleton set is independent in a matroid and applying Markov's inequality we may see that $G_\ell = E_\ell$. Since the collection of G_j's form a size $O(\log k)$ partition of $[n]$, and summing (3.1) over the collection gives a total expected revenue of at least $\mathcal{R}^{\mathcal{M}}/4$, we may conclude that there is some G_j which gives a $\Omega(1/\log k)$-fraction of $\mathcal{R}^{\mathcal{M}}$ regardless of ordering. □

Proof of Theorem 3.8 for non-regular distributions. We now prove Theorem 3.8 when distributions are non-regular. For doing this, we first give a different proof for Theorem 3.8 for regular distributions. While in Section 3.2 we proved that for any G_j, when the services in G_j are offered in the "worst" ordering, the probability c_i of a service $i \in G_j$ being offered is at least 1/4, we now prove that for any $i \in G_j$, even if i is placed last in the ordering among services in G_j, i has a probability 1/4 of being offered.

The mechanism sets $q_i = q_i^M/2$ and $p_i = F_i^{-1}(1 - q_i)$. For any service i consider placing it last. Let D be the set of services that are "desired", that is, services for which $v_i \geqslant p_i$. Note that the mechanism offers to provide service i only if the services sold till i don't span it. Thus, at each step, the served set has the same span as the portion of D seen so far. We then have that

$$\begin{aligned} \Pr_D[\text{We offer service } i \text{ if placed last}] &= \Pr_D[i \notin \mathrm{span}(D - \backslash\{i\})] \\ &\geqslant \Pr_D[i \notin \mathrm{span}(D)]; \end{aligned} \tag{3.2}$$

Let c_i denote the probability (3.2), which is over the desired set D. Fix any basis B. For any desired set D, at least $(k - \mathrm{rank}(D))$ elements of B are not spanned by D. Thus, we have

$$\begin{aligned} \sum_{i \in B} c_i &= \sum_{i \in B} \sum_{D : i \notin \mathrm{span}(D)} \Pr[D \text{ is desired}] \\ &= \sum_D (k - \mathrm{rank}(D)) \Pr[D \text{ is desired}] \\ &= k - \mathop{E}_D[\mathrm{rank}(D)] \geqslant k/2, \end{aligned}$$

where the inequality follows from the fact that

$$\mathop{E}_D[\mathrm{rank}\, D] \leqslant \mathop{E}_D[|D|] = \sum_{i \in [n]} q_i \leqslant k/2.$$

Thus, $\sum_{i \in B} c_i \geqslant k/2$ which implies that at least one third of the agents in B must have $c_i \geqslant 1/4$. Let $G = \{i | c_i \geqslant 1/4\}$. The revenue obtained from serving only the agents in G is

$$\sum_{i \in G} c_i p_i q_i \geqslant 1/8 \sum_{i \in G} p_i^M q_i^M.$$

As before, we can recursively apply the above argument to the elements outside G. At step j, let G_j be the newly found G, and let E_j be the set of services still under consideration, defined as $E_1 = [n]$ and $E_j = E_{j-1} - G_{j-1}$ for $j > 1$. Now, at each stage, G_j contains at least one third of the the remaining services. Thus we have $O(\log k)$ G_j's, and one of these must contain $\Omega(1/\log k)$-fraction of the revenue.

Now, when the distributions F_j are non-regular, we pick the prices p_j randomly as suggested by Lemma 3.2, such that the probability that service j is accepted if offered a randomized price is exactly $q_j^M/2$, and $\mathcal{R}_J^{\mathcal{M}}$ is bounded from above by $2\sum_j \mathrm{E}[p_j q_j]$. We assume that for each instantiation of the prices, the mechanism faces the "worst" ordering for that instantiation, i.e., the ascending order of prices for every intantiation. To bound the expected revenue that the OPM obtains from allocating service $j \in G_i$, we note that in any instantiation of the prices (and corresponding "worst" ordering over services), we can pessimistically defer offering service j until all other services in G_i have been offered. Then, following the above analysis, the probability that j is offered is at least $c_j \geqslant 1/4$. Then, the expected revenue from j is at least $1/4\,\mathrm{E}[p_j q_j]$, which is 1/8-th of the revenue that this service contributes to $\mathcal{R}^{\mathcal{M}}$. Therefore, as before, one of the G_i's contain $\Omega(1/\log k)$ fraction of the revenue. □

We remark that while all the other posted price mechanisms we give can be computed efficiently(see Section 3.4), we do not know of an efficient algorithm for computing an $O(\log k)$-approximate OPM.

As was mentioned earlier, Babaioff et al. (2007) obtain a similar result for the related matroid secretary problem. In Babaioff et al.'s setting agents arrive in a random order but their values are adversarial. They present an $O(\log k)$ approximation by picking a price uniformly at random in the set $\{h/k, 2h/k, \cdots, h\}$ and charging it to every agent; here h is the largest among all values. In our setting such an approach does not work: the example below shows that no uniform pricing can achieve an $o(\log h)$ approximation even for $k = 1$.

Example 3.9. *Let $k = 1$ and consider a group of h agents where agent i has a value of i with probability $1/2i^2$ and zero otherwise. Then an SPM that sets a price of i for agent i obtains an expected revenue of $\Omega(\log h)$. On the other hand, an SPM that uses a uniform price of c only obtains expected revenue $\sum_{i\in[c,h]} c/2i^2 < c/2c = 1/2$.*

Non-matroid constraints

We now show that the approximations described above cannot extend to general non-matroid set systems. In particular, the example below describes a family of instances with i.i.d. agents and a symmetric non-matroid constraint for which the ratio between the expected revenue of Myerson's mechanism and that of the optimal

OPM is $\Omega(\log n/\log\log n)$ where n is the number of agents. In fact the same lower bounds holds even for SPMs, that is, when we are able to choose the best ordering over offers.

Example 3.10. *For a given* k, *set* $n = k^{k+1}$. *Partition* $[n]$ *into* k^k *groups* $G_1, \cdots, G_{k^k}$ *of size* k *each, with* $G_i \cap G_j = \emptyset$ *for all* $i \neq j$. *The set system* $\mathcal{S}$ *contains all subsets of the groups* G_i, *that is,* $\mathcal{S} = \cup_i 2^{G_i}$ *Each agent has a value of* 1 *with probability* $1 - 1/k$ *and* k *with probability* $1/k$.

For any given valuation profile, let us call the agents with a value of k *to be good agents and the rest to be bad agents. The probability that a group contains* k *good agents is* k^{-k}. *Therefore in expectation one group has* k *good agents and Myerson's mechanism can obtain revenue* k^2 *from such a group:* $\mathcal{R}^{\mathcal{M}} = \Omega(k^2)$.

Next consider any SPM. Once the mechanism commits to serving an agent, it can only serve agents within the same group in the future. These have a total expected value less than $2k$. *Therefore, the revenue of any SPM is at most* k *from the first agent it serves and* $2k$ *in expectation from subsequent agents, for a total of* $3k$. *We get a gap of* $\Omega(k) = \Omega(\log n/\log\log n)$.

The above example also shows that while in many single-parameter pricing problems when the values are distributed in the range $[1, h]$ it is possible to obtain a $\log h$ approximation to social welfare, the same does not hold in our general setting. In the example we have $h = m$ and the gap between the expected revenue of the optimal SPM and that of Myerson's mechanism is $\Omega(h)$. On the other hand, the gap is always bounded by $O(h)$ and is achieved by an SPM that charges each agent a uniform price of 1.

While the above example illustrates that OPMs cannot obtain a constant-factor approximation in non-matroid settings, we now present an example with a non-matroid constraint for which the revenue obtained by ordering the agents in the optimal way is a factor of $\Omega(\log n/\log\log n)$ larger than that obtained by ordering the agents in the least optimal way.

Lemma 3.11. *There exists an instance of the single-parameter mechanism design problem with a non-matroid feasibility constraint, along with two orderings* σ_1 *and* σ_2 *such that the revenue of the optimal SPM using ordering* σ_1 *is a factor of* $\Omega(\log n/\log\log n)$ *larger than that of the optimal SPM using ordering* σ_2.

Proof. Consider the following example. Construct a complete m-ary tree of height $m + 1$, and place a single agent at each node other than the root. The agents' valuations are i.i.d. , where any agent has a valuation of m with probability $1/m$, and a valuation of 0 otherwise. Our constraint on serving the agents is that we may serve any set of agents that lie along a single path from the root of the tree to some leaf – it is easy to verify that this is downward-closed.

Consider what happens when we may serve the agents in order by level from the root of the tree to the leaves. At each level of the tree, we may offer to serve at least m different agents, regardless of the outcome on previous levels. Since we may never sell to more than one agent per level, our revenue is either 0 or m on each level. We get a revenue of 0 if and only if every agent has a valuation of 0; this occurs with probability at most

$$(1 - 1/m)^m \leqslant 1/e,$$

and thus our expected revenue overall is at least

$$m^2 \cdot (1 - 1/e) = \Omega(m^2).$$

On the other hand, if we must serve the agents in order by level from the leaves of the tree to the root, then the first agent we serve commits us to a specific path. So we cannot hope to achieve revenue better than m for this specific node, plus the revenue expected revenue for an arbitrarily chosen path. Since each agent has an expected valuation of 1, this is bounded by

$$m + (m - 1) \cdot 1 = O(m).$$

Thus, the difference in revenue between the described orderings is $\Omega(m)$; since the total number of agents is $n = O(m^m)$, in terms of n this gap is $\Omega(\log n / \log \log n)$. □

3.3 Order-specifying sequential pricings

In this section we show that we can achieve much better approximations for some settings by picking the right ordering σ over offers, that is, through sequenced posted-price mechanisms. While there is no direct reduction from SPMs in multi-parameter settings to SPMs in single-parameter settings analogous to Theorem 4.3, we show that for matroid and matroid intersection settings our results carry over in an approximation preserving way to multi-parameter instances as well.

We begin with a 2 approximation to single-parameter instances with a general matroid feasibility constraint, and show an improved 1.58 approximation for the special cases of uniform and partition matriods. We then describe a 3 approximation for general matroid intersection. We conclude this section by proving an 8 approximation in PDSE to multi-parameter instances with a general matroid intersection constraint. Table 3.2 summarizes the approximation factors for all the settings considered.

3.3.1 A 2-approximation for matroids

Consider the instance $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ where $\mathcal{S}$ is a matroid with rank k. Assume first that all the distributions F_j are regular. Let $q_j = q_j^M$ be the probability with which the optimal truthful mechanism (Myerson's mechanism) allocates service j. Let $p_j = p_j^M = {F_j}^{-1}(1 - q_j^M)$. Let σ be the order of decreasing prices p_j over the services. Our approximately optimal SPM is given by $(\mathbf{p}, \sigma)$. When the distributions F_j are non-regular, we define the prices randomly as suggested by Lemma 3.2, and for each instantiation of the prices, pick the greedy ordering over services in order of decreasing prices.

Theorem 3.12. *Let $\mathcal{I}$ be an instance of the BSMD with a matroid feasibility constraint. Then, the mechanism $(\mathbf{p}, \sigma)$ described above 2-approximates the revenue of Myerson's mechanism for $\mathcal{I}$.*

Proof. We show that the mechanism $(\mathbf{p}, \sigma)$ achieves an expected revenue of at least $\frac{1}{2}\sum_i p_i q_i$. Once again we start with the assumption that all the distributions F_j are regular. Note that if the mechanism ignored the feasibility constraint $\mathcal{S}$, and offered the prices $\mathbf{p}$ to all agents, serving any agent that accepted its offered price, then its expected revenue would be exactly $\sum_i p_i q_i$. So our proof accounts for the total revenue lost due to agents "blocked" from getting an offer by previously served agents.

Formally, let $S = \{i_1 < i_2 < \cdots < i_\ell\}$ be the set of agents served, and let S_j denote the first j elements of S. Let $\text{span}(S)$ denote the span of set S. Define the sets

$B_j = (\mathrm{span}(S_j) \setminus \mathrm{span}(S_{j-1})) \cap \{i : i > i_j\}$. Note that the sets B_j partition the set of agents blocked by those previously served. Moreover, $p_i \leqslant p_{i_j}$ for all $i \in B_j$, since $B_j \subseteq \{i : i > i_j\}$.

Denote the price offered to agent i_j by p^j. Then, the expected revenue lost given that S is served is

$$\begin{aligned}
&\sum_{j=1}^{\ell} \sum_{i \in B_j} p_i q_i \\
&\qquad \leqslant p^1 \left(\sum_{i \in sp(S_1)} q_i \right) + \sum_{j=2}^{\ell} p^j \left(\sum_{i \in sp(S_j)} q_i - \sum_{i \in sp(S_{j-1})} q_i \right) \\
&\qquad = \sum_{j=1}^{\ell-1} \left((p^j - p^{j+1}) \sum_{i \in sp(S_j)} q_i \right) + p^{\ell} \left(\sum_{i \in sp(S_\ell)} q_i \right) \\
&\qquad \leqslant \sum_{j=1}^{\ell-1} (p^j - p^{j+1}) \cdot j + p^{\ell} \cdot \ell = \sum_{j=1}^{\ell-1} p^j,
\end{aligned}$$

which is the revenue obtained by serving S. Here we used the fact that $\sum_{i \in sp(S_j)} q_i \leqslant \mathrm{rank}(S_j) = |S_j| = j$. Therefore,

$$\mathrm{E}[\text{revenue lost}] \leqslant \sum_{S} \sum_{j \in S} p^j \cdot \Pr[S \text{ is served}] = \mathcal{R}^{(\mathbf{p},\sigma)},$$

and so it follows that $\sum_i p_i q_i \leqslant 2\mathcal{R}^{(\mathbf{p},\sigma)}$.

Next consider the case of non-regular distributions. As mentioned earlier, we pick prices randomly as suggested by Lemma 3.2. Let p_j be the average price offered for service j if j is the first service offered. Consider a hypothetical posted-price mechanism that orders the services in decreasing order of p_j, and then defers the instantiation of the prices to be offered to just before the service is offered. Then, the acceptance probabilities for the services are exactly q_j, and the previous analysis continues to work in this case. Our SPM, that picks the greedy ordering for every instantiation of prices performs no worse than this hypothetical mechanism and we obtain the same approximation factor as before. □

We note that this approximation factor is not known to be tight. Blumrosen and Holenstein (2008) show that the gap between the optimal SPM and Myerson's mechanism can be as large as $\sqrt{\pi/2} \approx 1.253$ even in the single item auction case with i.i.d. agents. We reproduce the example here for completeness. There are n agents, each with a value distributed independently according to function $F(v) = 1 - 1/v^2$. The seller has one item to sell. Then, the expected revenue of Myerson's mechanism is $\Gamma(1/2)\sqrt{n}/2$, where $\Gamma()$ is the Gamma function. On the other hand, the expected revenue of the optimal SPM can be computed to be $\sqrt{n/2}$. Therefore, we get a gap of $\Gamma(1/2)/\sqrt{2} = \sqrt{\pi/2} \approx 1.253$.

Improved approximations for partition matroids

We now show an improved $e/e-1 = 1.58$ approximation for uniform and partition matroids. We begin by proving it for 1-uniform matroids, then extend it k-uniform matroids and partition matroids.

Let $q_j = q_j^M$ be the probability with which the Myerson's mechanism allocates service j. The SPM sets a price of $p_j = p_j^M = F_j{}^{-1}(1 - q_j^M)$ for service j. Let σ be the order of decreasing prices p_j over the services. and let c_j denote the probability of offering service j in this ordering. Note that for uniform matroids, the offer probabilities are always in descending order, i.e. if $i < j$, then $c_{\sigma(i)} > c_{\sigma(j)}$. The expected revenue $\mathcal{R}_\mathbf{p}^\sigma$ of this SPM is exactly $\sum_{i=1}^n c_i p_i q_i$. Let p be the price satisfying the equation

$$\sum_i p_i q_i = p \sum_i q_i. \tag{3.3}$$

Now consider a hypothetical situation where the prices are all equal to p but the probability with which agents accepted the items are still q_i. Let $\bar{\mathcal{R}}$ (=$\sum_i c_i p q_i$) denote the revenue obtained in this hypothetical setting. We first prove that $\mathcal{R}_\mathbf{p}^\sigma \geqslant \bar{\mathcal{R}}$.

Lemma 3.13. $\mathcal{R}_\mathbf{p}^\sigma \geqslant \bar{\mathcal{R}}$

Proof. For notational convenience, let us assume w.l.g. that for all i, $\sigma(i) = i$. Recall that for uniform matroids, for any given order σ, the $c_{\sigma(i)}$'s are in descending order. Since we assume $\sigma(i) = i$, the c_i's are in descending order. Let $\delta_i = q_i(p_i - p)$. So we have

$$\mathcal{R}_\mathbf{p}^\sigma = \sum_{i=1}^n c_i p_i q_i = \sum_{i=1}^n c_i (p q_i + \delta_i) = \bar{\mathcal{R}} + \sum_{i=1}^n c_i \delta_i \geqslant \bar{\mathcal{R}},$$

where the inequality follows from observing that:

- c_i's are in descending order;
- $\exists j$ such that δ_i is non-negative for all $i \leqslant j$ and negative otherwise; and
- $\sum_i \delta_i = 0$ (By (3.3)) .

□

Theorem 3.14. *Let $\mathcal{I}$ be an instance of the BSMD with a 1-uniform matroid feasibility constraint. Then, the mechanism $(\mathbf{p}, \sigma)$ achieves an $(e/e-1)$-approximation to the revenue of Myerson's mechanism for $\mathcal{I}$.*

Proof. Let $\sum_i q_i = s$. Note that $s \leqslant 1$, as we are in a 1-uniform matroid setting. We first prove that $\bar{\mathcal{R}} \geqslant (1 - \frac{1}{e})ps$.

$$
\begin{aligned}
\bar{\mathcal{R}} &= p(\Pr[\text{Some agent is served}]) = p(1 - \Pr[\text{No agent is served}]) \\
&= p\left(1 - \prod_{i=1}^{n}(1 - q_i)\right) \\
&\geqslant p\left(1 - \prod_{i=1}^{n}(1 - s/n)\right) \qquad (3.4) \\
&\geqslant (1 - 1/e)ps,
\end{aligned}
$$

where (3.4) follows since the product is maximized when the q_i's are all equal. Now, from (3.3) we have $ps = \sum_i p_i q_i$. This, along with Lemma 3.13 proves the theorem when distributions are regular.

For non-regular distributions, if p_i is defined as

$$
\frac{x_i \underline{p_i q_i} + (1 - x_i)\overline{p_i q_i}}{q_i},
$$

from Lemma 3.2, we know that the optimal revenue achievable by selling an item with probability q_i to agent i is no more than $p_i q_i$, and thus expected revenue of Myerson's mechanism is upper bounded by $\sum_i p_i q_i$. On the other hand, the revenue of the SPM, which picks prices randomly, is $\sum_i c_i p_i q_i$. The rest of analysis is identical to that for regular distributions. □

Next we consider the k-uniform case.

Theorem 3.15. *Let $\mathcal{I}$ be an instance of the BSMD with a k-uniform matroid feasibility constraint. Then, the mechanism $(\mathbf{p}, \sigma)$ achieves an $(e/e-1)$-approximation to the revenue of Myerson's mechanism for $\mathcal{I}$.*

Proof. Our proof technique is closely related to the proof for 1-uniform matroids. Let $p_i = p_i^M$ and let p be the common price for all agents which satisfies (3.3). If we define $\bar{\mathcal{R}}$ as defined for the 1-uniform case, then the proof of Lemma 3.13 extends to k-uniform matroids also. Thus it would be enough to argue that $\bar{\mathcal{R}} \geqslant (1-1/e)p \sum_i q_i$.

For any set of probabilities $\{q_i\}$ in the k-item case, let us define $q_i' = q_i/k$. Note that the probabilities $\{q_i'\}$ form a valid set of probabilities for a 1-item case because

$$\sum_i q_i' = \sum_i q_i/k \leqslant 1$$

We can come up with distributions F_i' for the 1-item case such that the price $F_i'^{-1}(1-q_i') = p$ for all services. Let c_i' denote the probability that service i is offered in this derived 1-item case. By Theorem 3.14, we know that the revenue in this 1-item case, given by $\sum_i c_i' p q_i'$, is at least $(1-1/e)\sum_i p q_i'$. If we can prove that the revenue $\bar{\mathcal{R}}$ for the k-item case, given by $\sum_i c_i p q_i$, is at least k times $\sum_i c_i' p q_i'$, then we are done. We prove this by the following induction. We assume that for $j-1 \leqslant n$, the revenue R_{j-1} from the first $j-1$ items is at least k times the revenue R'_{j-1} from the first $j-1$ items in the corresponding 1-item case i.e. $\sum_{i=1}^{j-1} c_i p q_i \geqslant k \cdot \sum_{i=1}^{j-1} c_i' p q_i'$. The base case is trivially true. We prove the same for j through two cases.

1. If $c_j \geqslant c_j'$, then we are done, because we know that revenue R_j from the first j items can be written as

   $$R_j = R_{j-1} + c_j p q_j \geqslant k(R'_{j-1}) + k c_j' p q_j' = k R_j'.$$

   The inequality uses the induction hypothesis.

2. If $c_j < c_j'$, we show that the revenue obtained is better than when $c_j = c_j'$ and then we will be done. To see this observe that the revenue R_j can be written as being conditioned on whether or not k items were sold in the first $j-1$ items.

So we have

$$R_j = (1 - c_j)kp + c_j\big(pq_j + E[R_{j-1} | \text{ at most } k-1 \text{ of first } j-1 \text{ served}]\big);$$

since we have that

$$kp \geqslant \big(pq_j + E[R_{j-1} | \text{at most } k-1 \text{ of first } j-1 \text{ served}]\big),$$

we conclude that the revenue only decreases by increasing c_j to c'_j.

Thus in either case, the k-item case has a better revenue, guaranteeing an approximation factor of $\frac{e}{e-1}$. □

Corollary 3.16. *Let $\mathcal{I}$ be an instance of the BSMD with a partition matroid feasibility constraint. Then, the mechanism $(\mathbf{p}, \sigma)$ achieves an $(e/e-1)$-approximation to the revenue of Myerson's mechanism for $\mathcal{I}$.*

BSMD with a matroid intersection constraint

The SPM $(\mathbf{p}, \sigma)$ we use here has prices $\mathbf{p}$ and the ordering σ picked in a manner similar to the one employed in Section 3.3.1.

Theorem 3.17. *Let $\mathcal{I}$ be an instance of the BSMD with a feasibility constraint that is an intersection of m matroids. Then, the mechanism $(\mathbf{p}, \sigma)$ $(m+1)$-approximates the revenue of Myerson's mechanism for $\mathcal{I}$.*

Proof. Let the m matroids be denoted by $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_m$. Let $\text{rank}_a(S)$ and $\text{span}_a(S)$ denote respectively the rank and span of set S in the matroid $\mathcal{M}_a$. Note that for any subset S and any $a \in [m]$, we have $\sum_{i \in S} q_i \leqslant \text{rank}_a(S)$. Once again, let $S = \{i_1 < i_2 < \cdots < i_\ell\}$ denote the set of agents served. We prove the theorem by showing that the expected revenue of $\mathcal{S}$ is at least $1/(m+1) \sum_i p_i q_i$, by arguing that the total price paid by agents in S is at least $1/m$ times the expected revenue from agents that are "blocked" by S.

Let S_j denote the first j elements of S. For each $1 \leqslant a \leqslant m$, define sets B_j^a with respect to matroid $\mathcal{M}_a$ as in the proof of Theorem 3.12. That is, $B_j^a = \text{span}_a(S_j) \setminus$

$\text{span}_a(S_{j-1})$. Denote the price of item i_j by p^j. Then, if we let $B_j = \cup_{a=1}^{m} B_j^a$, we can upper bound the expected revenue lost when S is served by

$$\sum_{1 \leqslant j \leqslant \ell} \sum_{i \in B_j} p_i q_i \leqslant \sum_{a=1}^{m} \sum_{1 \leqslant j \leqslant \ell} \sum_{i \in B_j^a} p_i q_i \leqslant m \sum_{1 \leqslant j < \ell} p^j.$$

Here we used the same algebraic transformation as in the proof of Theorem 3.12 along with the fact that $\sum_{i \in B_j^a} q_i \leqslant \sum_{i \in \text{span}_a(S_j)} q_i \leqslant j$. Therefore as before we get $\sum_i p_i q_i \leqslant (m+1)\mathcal{R}_{\mathbf{p}}^{\sigma}$ proving the theorem for regular distributions.

For non-regular distributions, we pick the prices p_j randomly as suggested by Lemma 3.2, and thus $\mathcal{R}_{\mathcal{J}\text{reps}}^{\mathcal{M}}$ is bounded from above by $\sum_j \mathrm{E}[p_j q_j]$. The ordering σ in one where the services are offered in the decreasing order of their expected prices. The instantiation of price for a service is deferred till the service is offered. Then along the lines of the proof of regular distributions, the expected total price in B (expectation over the values as well as the randomization over prices), conditioned on S, is at most m times the expected total price contained in S (expectation over the randomization over prices). Thus, the expected total price contained in S is at least $\frac{1}{m+1} \sum_j \mathrm{E}[p_j q_j]$. □

Combinatorial auctions with small bundles

Consider a situation where the seller has multiple copies of a number of items on sale, and each agent is interested in some (commonly known) bundles over items and has a common value for all of these bundles . When each desired bundle is of size at most m, we call this setting a single-parameter combinatorial auction with known bundles of size m. In this case the SPM $(\mathbf{p}, \sigma)$, similar to the one described in Section 3.3.1, achieves an $m+1$ approximation.

Theorem 3.18. *Let $\mathcal{I}$ be an instance of a single-parameter combinatorial auction with known bundles of size m. Then, the mechanism $(\mathbf{p}, \sigma)$ $(m+1)$-approximates the revenue of Myerson's mechanism for $\mathcal{I}$.*

Proof. Let A denote the set of items available to the seller, each with some multiplicity. First suppose that each agent is single-minded, that is, each agent is interested in only one bundle of items, the bundle being of size at most m. Then, the feasibility constraint is an intersection over $|A|$ uniform matroids, one corresponding to each item, with each agent participating in only m of the matroids. Now it is easy to adapt the proof of Theorem 3.17 to obtain an $m + 1$ approximation.

More generally suppose that every agent is interested in a collection of bundles, each of size at most m, and modify the mechanism $\mathcal{S}$ so that in addition to deciding whether or not to serve an agent, it also arbitrarily allocates any available desired bundle to every agent it serves. Then we can argue that for any set S, and set B blocked by the agents in S, the sum of the probabilities q_i over the set B is no more than m times the size of S. Therefore, once again following along the proof of Theorem 3.17, we get an $m + 1$ approximation.

The extension to non-regular distributions is identical to that described in Theorem 3.17.

□

3.4 Computing the near-optimal posted-price mechanisms

We now describe how to compute the approximately optimal OPMs and SPMs designed in Sections 3.2 and 3.3. We assume that we are given access to the following oracles and algorithms:

- An algorithm to compute the optimal price to charge to a single-parameter agent given the agent's value distribution. Note that given such an algorithm and some value x, we can modify it to return the optimal price in the range $[x, \infty)$ to charge the agent.

- An oracle that given a value v and index i returns $F_i(v)$ and $f_i(v)$, as well as, given a probability α returns $F_i^{-1}(\alpha)$. Note that the oracle can be used to compute the virtual value $\phi_i(v)$.

- An oracle for computing ironed virtual values in order to compute the approximately optimal SPM for non-regular distributions.

- An algorithm to maximize social welfare over the given feasibility constraint in order to be able to compute the outcome of Myerson's mechanism.

All of the mechanisms designed by us require computing the probabilities q_i^M. We first show how to estimate these probabilities within small constant factors:

1. Let $\epsilon = 1/3n$. Sample $N = 4n^4 \log n/\epsilon^2$ value profiles from $F_1 \times F_2 \times \cdots \times F_n$. For each sample, compute the (ironed) virtual value for each agent, and use these to compute the outcome of Myerson's mechanism for that value profile.

2. Estimate the probabilities q_i^M using the samples. Call the estimates $\widehat{q_i^M}$.

3. If $\widehat{q_i^M} < 1/n^2$, set $\widehat{q_i} = 1/n^2$, else set $\widehat{q_i} = \widehat{q_i^M}/(1-\epsilon)$. Compute for each i the value $\widehat{p_i} = {F_i}^{-1}(1 - \widehat{q_i})$.

4. Find the optimal price in the range $[\widehat{p_i}, \infty)$ to charge to agent i. Call it p_i. Let $q_i = 1 - F_i(p_i)$.

5. Output the prices computed in the last step and order the agents in order of decreasing prices.

In order to analyse the performance of this approach, we compare it to a mechanism that charges agent i the price $p_i^M = {F_i}^{-1}(1 - q_i^M)$ but uses the same ordering as the mechanism above. We first show that the probabilities q_i closely estimate the probabilities q_i^M.

Lemma 3.19. *With probability at least* $1 - 2/n$, *we have* $\widehat{q_i} \in [q_i^M, (1+3\epsilon)q_i^M + 2/n^2]$.

Proof. First, for any i with $q_i^M \geqslant 1/n^4$, using Chernoff bounds we get that

$$\mathbf{Pr}\left[\left|\widehat{q_i^M} - q_i^M\right| \geqslant \epsilon q_i^M\right] \leqslant 2e^{-\epsilon^2 q_i^M N/2} \leqslant 2/n^2$$

$\widehat{q_i^M} \in (1 \pm \epsilon)q_i^M$ in turn implies by definition that $q_i^M \leqslant \widehat{q_i} \leqslant (1+\epsilon)/(1-\epsilon)q_i^M \leqslant (1+3\epsilon)q_i^M$. Therefore we have $\widehat{q_i} \in [q_i^M, (1+3\epsilon)q_i^M]$. On the other hand, for $q_i^M < 1/n^4$, by Markov's inequality, with probability $1 - 1/n^2$, $\widehat{q_i^M} < 1/n^2$, and so $\widehat{q_i} \in [q_i^M, 1/n^2]$. The lemma now follows by employing the union bound. □

Furthermore, conditioned on the event defined in the statement of the above lemma (call it $\mathcal{E}$), since p_i^M lies in the range $[\widehat{p_i}, \infty)$, we have that $q_i^M p_i^M \leqslant q_i p_i$. This implies that the prices p_i give a good estimate on the revenue of Myerson's mechanism.

Next, we compare the real mechanism $\mathcal{S}$ with prices p_i to the theoretically good mechanism $\mathcal{S}'$ that charges prices p_i^M. Let S be the set of services for which $\widehat{q_i^M} < 1/n^2$. The probability that any of these services is offered in $\mathcal{S}$ is at most $1/n$. Conditioned on this event not happening, the probability that a service is offered in $\mathcal{S}$ is no smaller than its counterpart in $\mathcal{S}'$. Moreover, conditioned on being made an offer, the revenue from service i is $q_i p_i \geqslant q_i^M p_i^M$.

Therefore, conditioned on the event $\mathcal{E}$, the expected revenue of $\mathcal{S}$ is at least a $(1 - 1/n)$ fraction of the expected revenue of $\mathcal{S}'$. But the event $\mathcal{E}$ happens with probability $1 - 2/n$, therefore, we get a $(1 - o(1))$ approximation to the expected revenue of $\mathcal{S}'$.

3.5 Revenue maximization through VCG mechanisms

A consequence of our constant-factor approximation to revenue through SPMs is that in matroid settings VCG mechanisms with appropriate reserve prices are near-optimal in terms of revenue. This follows from noting, as we show below, that VCG mechanisms perform no worse in terms of expected revenue than SPMs with the same reserve prices. Although VCG mechanisms aim to maximize the social welfare of the outcome, setting high enough reserve prices allows them to also obtain good revenue.

Formally, a Vickrey-Clarke-Groves (VCG) mechanism $\mathcal{V}^{\mathbf{p}}$ with reserve prices $\mathbf{p}$ sells the set S of services, with $v_i \geqslant p_i$ for all $i \in S$, that maximizes $\sum_{i \in S} v_i$.

Hartline and Roughgarden (2009) show that in several single-parameter settings the VCG mechanism with monopoly reserve prices gives a constant factor approximation to revenue. This result holds when all the value distributions satisfy the so-called monotone hazard rate condition, or with a matroid feasibility constraint when all the value distributions are regular. Their result does not extend to the case of matroids with general (non-regular) value distributions. One of the main questions left open by their work is whether there is some set of reserve prices (not necessarily equal to the monopoly reserve prices) for which the VCG mechanism gives a constant factor approximation to revenue in the matroid setting with general value distributions. We answer this question in the positive. We use the following fact about matroids.

Proposition 3.1. *Let* B_1 *and* B_2 *be any two independent sets of equal size in a matroid set system* $\mathcal{S}$. *Then there is a bijective function* $g : B_1 \setminus B_2 \to B_2 \setminus B_1$ *such that for all* $e \in B_1 \setminus B_2$, $B_1 \setminus \{e\} \cup \{g(e)\}$ *is independent in* $\mathcal{S}$.

Theorem 3.20. *For any instance of the single-parameter Bayesian mechanism design problem with a matroid feasibility constraint, there exists a set of reserve prices such that the expected revenue of the VCG mechanism with those reserve prices is at least half of the expected revenue of Myerson's mechanism.*

Proof. We prove that when the set system $\mathcal{S}$ is a matroid, for any collection of prices $\mathbf{p}$, the revenue of the SPM $\mathcal{S}^{\mathbf{p}}$ is no more than the revenue of the VCG mechanism $\mathcal{V}^{\mathbf{p}}$. The result then follows from Theorem 3.12.

Fix a value vector $\mathbf{v}$ and let A denote the set sold by $\mathcal{S}^{\mathbf{p}}$ and B denote the set sold by $\mathcal{V}^{\mathbf{p}}$. Then, since both mechanisms sell a maximal independent set among the set of services with $v_i \geqslant p_i$, we have $|A| = |B|$. Proposition 3.1 then implies the existence of a bijection g such that for all $e \in B \setminus A$, $B \setminus \{e\} \cup \{g(e)\}$ is independent. This implies that $\mathcal{V}^{\mathbf{p}}$ sells e at a price of at least the value of $g(e)$, which is at least the reserve price $p_{g(e)}$. On the other hand, by definition, the price charged to any $e \in B \cap A$ is at least p_e. Therefore, the revenue of $\mathcal{V}^{\mathbf{p}}$ in this case is at least $\sum_{e \in B \cap A} p_e + \sum_{e \in B \setminus A} p_{g(e)} = \sum_{e \in A} p_e$ which is equal to the revenue of $\mathcal{S}^{\mathbf{p}}$. □

3.6 Discussion

We designed approximately-optimal posted price mechanisms for a variety of single-parameter settings. The approximation factors we obtain depend on the kind of feasibility constraint that the seller faces. The exact constants are summarized in Tables 3.1 and 3.2.

Our approach does not extend beyond matroid and matroid-like settings. However, it is possible that there is some other class of simple near-optimal mechanisms for non-matroid single-parameter settings that do not exploit competition among agents. Such mechanisms may lead to approximately-optimal multi-parameter mechanisms for a broader class of feasibility constraints.

Feasibility constraint $\mathcal{S}$	Gap from optimal	
	upper bound	lower bound
Uniform matroid, Partition matroid	2	2
Graphical matroid	3	2
Intersection of two part. matroids	5.83	2
Matching with i.i.d. agents	$2e/(e-1) \approx 3.17$	2
Graphical matroid $\cap$ partition matroid	7.47	2
General matroid	$O(\log k)$	2
Non-matroid downward closed	-	$\Omega(\frac{\log n}{\log \log n})$

Table 3.1: A summary of approximation factors for the BSMD achievable through OPMs. Here k is the rank of $\mathcal{S}$.

Feasibility constraint $\mathcal{S}$	Gap from optimal	
	upper bound	lower bound
General matroid	2	$\sqrt{\pi/2} \approx 1.25$
Uniform matroid, Partition matroid	$e/(e-1) \approx 1.58$	1.25
Intersection of two matroids (BSMD)	3	1.25
Non-matroid downward closed	-	$\Omega(\log n / \log \log n)$

Table 3.2: A summary of approximation factors for the BSMD achievable through SPMs.

4 Multi-service settings with unit-demand agents

Suppose that the local organizers for a prominent symposium on computer science need to arrange for suitable hotel accommodations in the Boston area for the attendees of the conference. There are a number of hotel rooms available with different features and attendees have preferences over the rooms. The organizers need a mechanism for soliciting preferences, assigning rooms, and calculating payments. Fortunately, they have distributional knowledge over the participants' preferences (e.g., from similar conferences). This is a stereotypical multi-parameter setting for mechanism design that, for instance, also arises in most resource allocation problems in the Internet. What mechanism should the organizers employ to maximize their objective (e.g., revenue)?

The economic theory of optimal mechanism design is elegant and predictive in single-parameter settings. Here Myerson's theory (1981) of virtual valuations and characterizations of incentive constraints via monotonicity guide the design of optimal incentive-compatible mechanisms with practical (often non-incentive-compatible) implementations (Ausubel and Milgrom, 2006). The challenge posed by multi-parameter settings (e.g., in the likely case that conference attendees, i.e., agents, have different values for different hotel rooms) is two-fold. First, multi-parameter settings are unlikely to permit succinct descriptions of optimal mechanisms (McAfee and McMillan, 1988; Rochet and Chone, 1998; Manelli and Vincent, 2007). Second, in multi-parameter settings optimal mechanisms are unlikely to have practical implementations – even asking agents to report their true types across the many possible outcomes of the mechanism may be impractical. In summary, theory and practical considerations from optimal mechanism design in single-parameter settings fail to generalize to multi-parameter settings.

This work approaches these issues through the lens of approximation. Our main results are simple, practical, approximately optimal mechanisms for a large class of multi-parameter settings. We consider the multi-parameter setting through a single parameter analogy wherein each multi-parameter agent is represented by many independent single-parameter agents (e.g., one for each hotel room). The optimal revenue for this single-parameter setting is well understood and, due to increased competition among agents, upper-bounds that of the original multi-parameter setting. We leverage the fact that the sequential posted-price mechanisms described in the previous chapter achieve their approximations without inter-agent competition. This gives a robustness to deviations in modeling assumptions and, for instance, the same mechanism continues to be approximately optimal in the original multi-parameter setting. Therefore, our theory for approximately optimal single-parameter mechanisms generalizes to multi-parameter settings.

In the context of computer science literature this work is an extension of *algorithmic pricing* (see, e.g., Guruswami et al., 2005) to settings with multiple agents; it is unrelated to the standard computational questions of *algorithmic mechanism design* (see, e.g., Lehmann et al., 1999; Nisan and Ronen, 1999). The central problem in algorithmic pricing can be viewed (for the most part) as Bayesian revenue maximization in a single agent setting (see, e.g., Guruswami et al., 2005). Algorithmic pricing is hard to approximate when the agent's values for different outcomes are generally correlated (Briest, 2006); however, when the values are independent there is a 3-approximation (Chawla et al., 2007). In this context, our results improve and extend the independent case to settings with multiple agents and combinatorial feasibility constraints. Notice that the challenge in these problems is one imposed by the multi-parameter incentive constraints and not one from an inherent complexity of an underlying non-game-theoretic optimization problem. (E.g., in the hotel example the underlying optimization problem is simply maximum weighted matching.) In contrast, most work in algorithmic mechanism design addresses settings where economic incentives are well understood but the underlying optimization problem is computationally intractable (e.g., combinatorial auctions (Lehmann et al., 1999)).

While our exposition focuses on revenue maximization, all of our techniques and results apply equally well to *social welfare*. Social welfare is unique among objectives in that designing optimal mechanisms in multi-parameter settings is solved (by the VCG mechanism). Therefore, the interesting implication of our work on social welfare maximization is that sequential posted pricing approximates the welfare of the VCG mechanism and may be more practical.

4.1 Reducing multi-service instances to single-service instances

We begin by presenting our general reduction from the multi-parameter unit-demand mechanism design problem to the single-parameter problem. Using this reduction we can argue that if there exists an approximately optimal sequential posted-price mechanism in the single-parameter setting, there also exists one in the original multi-parameter setting. Understanding the properties of optimal mechanisms in multi-parameter settings is tricky so our approach is based on upper and lower bounds for single-parameter settings.

There are four main steps to give and instantiate our reduction. They are:

1. (Analogy) Define a single-parameter analog for any multi-parameter setting.

2. (Lower bound) Show that the revenue of the optimal single-dimensional analog is at least the revenue of the optimal multi-dimensional mechanism.

3. (Reduction) Show that if we had a sequential posted pricing for the single-dimensional analog, we can convert it into a posted pricing for the multi-dimensional setting without much loss in performance.

4. (Instantiation) Show for a given multi-dimensional setting that there exist sequential pricings that approximate the optimal mechanism for the single-parameter analog.

We describe the analogy and state the lower bound as well as a reduction for OPMs here. Section 4.2 instantiates the reduction for OPMs in various settings of interest. While we do not obtain a general purpose reduction from SPMs, as we do for OPMs, in Section 4.3 we describe SPMs that obtain approximate optimality in PDSE in many single as well as multi-parameter settings.

The analogy

The main concept behind our reduction is a single-parameter analogy. Consider an instance $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ of the BMUMD with $n = |\Pi|$ agents and $m = |J|$ services. The *single-parameter analog* is the setting we get when we assume that each service is demanded by a distinct agent, i.e., $\mathcal{I}^{\text{reps}} = (I, \mathcal{S}, \mathbf{F})$. Formulaically, this analogy is trivial; intuitively, it replaces each agent i with $|J_i|$ distinct agents (called representatives or "reps" hereafter). Each rep is interested in a single service $j \in J_i$ and behaves independently of (and potentially to the detriment of) other reps. Notice that $\mathcal{I}^{\text{reps}}$ has $m = |J|$ agents and services.

Lower bound

Notice that $\mathcal{I}^{\text{reps}}$ is similar to $\mathcal{I}$ except that it involves more competition (among different reps of the same multi-parameter agent). Therefore it is natural to expect that a seller can obtain more revenue in the instance $\mathcal{I}^{\text{reps}}$ than in $\mathcal{I}$. The following lemma formalizes this.

Lemma 4.1. *Let $\mathcal{A}$ be any individually rational and incentive compatible deterministic mechanism for instance $\mathcal{I}$ of the BMUMD. Then the expected revenue of $\mathcal{A}$, $\mathcal{R}^{\mathcal{A}}$, is no more than the expected revenue of Myerson's mechanism for the single-parameter instance $\mathcal{I}^{\text{reps}}$.*

Proof. Truthful mechanisms in multi-parameter settings satisfy the weak monotonicity condition defined below. For a Mechanism M, and a value vector $\mathbf{v}$, let $M_j(\mathbf{v})$ denote the probability with which M provides service j at value vector $\mathbf{v}$.

Definition 4.2. *A mechanism M satisfies weak monotonicity if for any agent i and any two types $\mathbf{v}^1$ and $\mathbf{v}^2$ with $v^1_j = v^2_j$ for all $j \in J \setminus J_i$, the following holds:*

$$\sum_{j \in J_i} \left(M_j(\mathbf{v}^1) v^1_j + M_j(\mathbf{v}^2) v^2_j \right) \geqslant \sum_{j \in J_i} \left(M_j(\mathbf{v}^2) v^1_j + M_j(\mathbf{v}^1) v^2_j \right)$$

We show that we can construct a truthful mechanism $\mathcal{A}^{\text{reps}}$ for the instance $\mathcal{I}^{\text{reps}}$ with revenue no less than that of $\mathcal{A}$. The lemma then follows from the optimality of Myerson's mechanism. Given a vector of values $\mathbf{v}$, the mechanism $\mathcal{A}^{\text{reps}}$ allocates to the set that $\mathcal{A}$ allocates to in $\mathcal{I}$ for the same value vector $\mathbf{v}$. We first claim that the allocation rule of $\mathcal{A}^{\text{reps}}$ is monotone non-decreasing in any v_j, implying that there exists a payment rule that makes the mechanism truthful. To prove the claim, fix any agent i and $j \in J_i$, and consider two value vectors $\mathbf{v}^1$ and $\mathbf{v}^2$ with $v^1_j = x$, $v^2_j = y$, and $v^1_{j'} = v^2_{j'}$ for $j' \neq j$. Let α_x and α_y denote the probabilities of serving agent i with service j under the two value vectors respectively, and let β_x and β_y denote the total value that agent i obtains from other services $j' \in J_i$, $j' \neq j$, in the two cases respectively. Then the weak-monotonicity (Definition 4.2) of $\mathcal{A}$ implies that

$$(x\alpha_x + \beta_x) + (y\alpha_y + \beta_y) \geqslant (x\alpha_y + \beta_y) + (y\alpha_x + \beta_x).$$

This is equivalent to $(x - y)(\alpha_x - \alpha_y) \geqslant 0$ and so the claim holds.

It remains to prove that the expected revenue of $\mathcal{A}^{\text{reps}}$ given $\mathcal{I}^{\text{reps}}$ is no less than the expected revenue of $\mathcal{A}$ given $\mathcal{I}$. Note that any deterministic multi-parameter mechanism can be interpreted as offering a price menu with one price for each item or service to each agent as a function of other agents' bids (Wilson, 1997). The agent then chooses the item or service that brings her the most utility. Given this characterization, suppose that for a fixed set $\mathbf{v}$ of values, mechanism $\mathcal{A}$ offers a price menu with prices $\mathbf{p}$ to agent i. Then, it draws a revenue of p_j from i whenever service j is offered. On the other hand, mechanism $\mathcal{A}^{\text{reps}}$ charges the agent j the minimum amount it needs to bid to be served, which is no less than p_j, as $\mathcal{A}$ is individually rational. □

A reduction for OPMs

The main advantage of a seller in the single-parameter analog is increased competition. Intuitively, if we can design mechanisms for the single-parameter setting that do not exploit competition, then it is reasonable to expect them to obtain similar performance in the multi-dimensional setting. Here, sequential posted pricings are exactly what is needed.

Theorem 4.3. *If OPM* $\mathbf{p}$ *is an* α*-approximation to the optimal mechanism for the single-parameter setting* $\mathcal{I}^{\text{reps}}$ *then it is an* α*-approximation in PDSE to the optimal mechanism for the multi-parameter setting* $\mathcal{I}$*.*

Proof. Let $\mathbf{p}$ be an α-approximate OPM for $\mathcal{I}^{\text{reps}}$. Consider its performance on I. For a fixed instantiation $\mathbf{v}$ of values let σ be any Π-respecting ordering that minimizes the revenue of the mechanism. Note that whenever the mechanism $(\mathbf{p}, \sigma)$ offers a service to agent i it is a dominant strategy for the agent to accept the service if and only if the agent gets non-negative value from the service. This is because any future offers that the agent gets can only bring him lower utility. Therefore, the sequence of offers and outcome of $(\mathbf{p}, \sigma)$ is identical under $\mathcal{I}$ and $\mathcal{I}^{\text{reps}}$, and, $\mathcal{R}_{\mathcal{I}}^{(\mathbf{p},\sigma)}(\mathbf{v}) = \mathcal{R}_{\mathcal{I}^{\text{reps}}}^{(\mathbf{p},\sigma)}(\mathbf{v}) \geqslant \mathcal{R}_{\mathcal{I}^{\text{reps}}}^{(\mathbf{p})}(\mathbf{v})$.

Therefore, the revenue of $\mathbf{p}$ in the multi-parameter setting is no less than its revenue in the single-parameter setting. Then the result follows from Lemma 4.1. $\square$

4.2 Approximation through order-oblivious sequential pricings

In this section we instantiate the reduction described in Section 4.1 for several different settings. We begin by developing our approach in the context of a "k item auction" setting, obtaining a (tight) 2-approximation in that setting. We then expand this approach to obtain constant-factor approximations for other matroids as well as matroid intersection constraints. For example, we show that in a setting with heterogeneous items and multiple unit-demand agents with identically distributed values (that we call the "supermarket setting"), there exists a pricing of the items that obtains a $2e/(e-1)$ approximation to the optimal mechanism.

Our most general positive result is a non-constructive $O(\log k)$-approximation for arbitrary matroids, where k is the rank of the matroid. On the negative side, we show that the gap between OPMs and the optimal achievable revenue can be as large as $\frac{\log n}{\log \log n}$ for non-matroid feasibility constraints, where n is the number of agents. Table 4.1 summarizes the approximation factors for all the settings considered.

4.2.1 Intersections of partition matroids

In this section, as well as in Sections 4.2.2 and 4.2.3, we look at instances $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ where $\mathcal{S}$ is given by the intersection $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2$ of a partition matroid $\mathcal{S}_1$ with some other matroid constraint $\mathcal{S}_2$. These instances arise as single-parameter analogs of instances $(I \times J, \mathcal{S}, \mathbf{F})$ with $\mathcal{S}_1$ formalizing the unit-demand constraint over the agents and $\mathcal{S}_2$ representing an unrelated matroid constraint over services (e.g. a supply constraint).

Our approach is to first use Lemma 3.4 to show that we can achieve good revenue under just the constraint $\mathcal{S}_1$, and then bound the impact of applying the additional constraint $\mathcal{S}_2$ on this revenue. We develop this approach in the settings where $\mathcal{S}_2$ is also a partition matroid. In Section 4.2.3 we extend this approach to more general matroids $\mathcal{S}_2$.

We begin by extending Corollary 3.6 to show that for the feasibility constraint $\mathcal{S}_1$, for any constant $\alpha \leqslant 1$, we can construct prices $\mathbf{p}^\alpha$ such that the OPM $\mathbf{p}^\alpha$ is approximately optimal for $(J, \mathcal{S}_1, \mathbf{F})$ and provides service to an agent with probability a factor of α smaller than the corresponding probability in Myerson's mechanism. By chosing an appropriate α we can then ensure that the additional constraint $\mathcal{S}_2$ minimally effects the OPM's revenue.

Lemma 4.4. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD where $\mathcal{S} = \mathcal{S}_1 \cap \mathcal{S}_2$ and $\mathcal{S}_1$ is a partition matroid. Consider a modified instance $\mathcal{I}' = (J, \mathcal{S}_1, \mathbf{F})$. Then, for each $\alpha \in (0, 1]$ there exist prices $\mathbf{p}$ such that $q_j = 1 - F_j(p_j) \leqslant \alpha q_j^M$ for all $j \in J$, and,*

$$\mathcal{R}^{\mathbf{p}}_{\mathcal{J}'} \geqslant \left(\frac{\alpha}{1+\alpha}\right) \mathcal{R}^{\mathcal{M}}_{\mathcal{J}}.$$

Proof. Let q_j^M, p_j^M be the probabilities and associated prices with which Myerson's mechanism provides service $j \in J$ under the full constraint $\mathcal{S}$. Fix $\alpha \in (0, 1]$, and consider the modified set of probabilities and prices

$$\begin{aligned} q_j{}' &= \alpha q_j^M \\ p_j{}' &= F_j{}^{-1}(1 - q_j{}') \geqslant p_j^M. \end{aligned}$$

Let $X \subset J$ be one of the uniform matroids making up $\mathcal{S}_1$, and let k be its rank. Then we know that $\sum_{j \in X} q_j{}' \leqslant \alpha k$, and so applying Lemma 3.4 to each such X in turn with prices $\mathbf{p}'$ will produce prices $\mathbf{p}$ such that

$$\mathcal{R}^{\mathbf{p}}_{\mathcal{J}'} \geqslant \frac{1}{1+\alpha} \sum_{j \in J} p_j{}' q_j{}' \geqslant \frac{\alpha}{1+\alpha} \sum_{j \in J} p_j^M q_j^M.$$

Lemma 3.2 then implies the claimed bound. □

Next, we show that when we consider the full constraint $\mathcal{S}$, expected revenue is not much less than we would get under just $\mathcal{S}_1$.

Theorem 4.5. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD with $\mathcal{S}$ being an intersection of two partition matroids. Then, there exists a set of prices $\mathbf{p}$ such that $\mathcal{R}^{\mathbf{p}}_{\mathcal{J}} \geqslant \frac{1}{5.83} \mathcal{R}^{\mathcal{M}}_{\mathcal{J}}$.*

Proof. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be the two partition matroids making up $\mathcal{S}$. Let $\mathcal{J}' = (J, \mathcal{S}_1, \mathbf{F})$ and let $\mathbf{p}$ be the prices obtained by applying Lemma 4.4 to $\mathcal{J}$ for some $\alpha \in (0, 1]$ to be picked later. Let σ be the arbitrary order in which the OPM offers services to agents. For each service $j \in J$ we define the following events:

- X_j is the event that service j would be sold by the OPM $\mathbf{p}$ in the setting $\mathcal{J}'$; and
- Y_j is the event that service j is not blocked in matroid $\mathcal{S}_2$ when it is reached by the OPM $\mathbf{p}$ in the ordering σ, that is, the set of services sold so far along with j is independent in $\mathcal{S}_2$.

Then, we have that

$$\mathcal{R}_{\mathcal{J}}^{\mathbf{p}} \geqslant \mathrm{E}\left[\sum_{j \in J} p_j \Pr[X_j | Y_j] \Pr[Y_j]\right] \geqslant \mathrm{E}\left[\sum_{j \in J} p_j \Pr[X_j] \Pr[Y_j]\right],$$

where the last inequality follows from the fact that conditioning on previous offers being rejected can only increase the probability we may make the current offer.

Now, as in the proof of Theorem 3.3, we can apply Markov's Inequality to see that for all $j \in J$ we have that $\Pr[Y_j] \geqslant (1-\alpha)$. Thus, we have that

$$\mathcal{R}_{\mathcal{J}}^{\mathbf{p}} \geqslant (1-\alpha)\,\mathrm{E}\left[\sum_{j \in J} p_j \Pr[X_j]\right] = (1-\alpha)\mathcal{R}_{\mathcal{J}'}^{\mathbf{p}} \geqslant (1-\alpha)\left(\frac{\alpha}{1+\alpha}\right)\mathcal{R}_{\mathcal{J}}^{\mathcal{M}}.$$

Choosing $\alpha = \sqrt{2}-1$ yields an approximation factor of $(\sqrt{2}-1)^{-2} \approx 5.83$. □

4.2.2 Supermarket setting

We now consider a special case of the matroid intersection setting studied above, where agents have identically distributed types. Specifically, we consider the setting of a supermarket offering a number of different items indexed by set T and with different multiplicities (supply limits). The set J of services is simply $I \times T$. Note that each service $j \in J$ is a tuple (i, t) corresponding to agent i and item t. We use v_{it} to denote the value of agent i for item $t \in T$. The feasibility constraint $\mathcal{S}$ is an intersection of two partition matroids $\mathcal{S}_1$ and $\mathcal{S}_2$, where $\mathcal{S}_1$ represents the unit-demand constraint over agents (i.e. every agent is allocated at most one item), and $\mathcal{S}_2$ represents the supply limits over items (i.e. every item is allocated to at most as many agents as the number of its units available). We show that when the values

v_{it} for different agents and the same item t are distributed identically, we can obtain an improved $2e/(e-1) \approx 3.17$ approximation.

The improvement in approximation factor arises from the following improvement to Theorem 3.5.

Lemma 4.6. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD where $\mathcal{S}$ is a k-uniform matroid and all the distributions F_i are identical and independent. Then the set of prices $\mathbf{p}$ with $p_i = p_i^M$ obtains $\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geqslant \frac{e-1}{e}\mathcal{R}_{\mathcal{I}}^{\mathcal{M}}$.*

Proof. Note that since the distributions are all identical, the probabilities q_i^M and prices p_i^M can be assumed to be identical. In the rest of this proof we will use p to denote the common price, but will not use the fact that the probabilities q_i^M are identical.

We first prove the lemma for the special case of $k = 1$. Let $\sum_i q_i = s \leqslant 1$. Then

$$\begin{aligned}
\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} &= p(\text{Pr[Some agent is served]}) \\
&= p(1 - \text{Pr[No agent is served]}) \\
&= p\left(1 - \prod_{i=1}^{n}(1 - q_i)\right) \\
&\geqslant p\left(1 - \prod_{i=1}^{n}(1 - s/n)\right) \qquad (4.1) \\
&\geqslant (1 - 1/e)ps = (1 - 1/e)\sum_i p_i q_i,
\end{aligned}$$

where (4.1) follows since the product is maximized when the q_i's are all equal. This proves the theorem when distributions are regular.

To prove the lemma for general k, we note that $\sum_i q_i = s \leqslant k$. Let $q' = q_i/k$ and $p' = F^{-1}(1 - q_i{}')$. Consider the OPM $\mathbf{p}'$ on the given instance of the BSMD but with a 1-uniform matroid constraint instead of a k-uniform one. Then the analysis in the preceding paragraph shows that the expected revenue of this BSMD is at least $(1 - 1/e)\sum_i p'q'$. Therefore, the probability with which the OPM sells an item is at least $(1 - 1/e)\sum_i q_i/k$. We now claim that the OPM $\mathbf{p}$ with the k-uniform matroid constraint sells at least k times as many items in expectation, which implies the result. □

We now prove the main result of this section.

Theorem 4.7. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD where $\mathcal{S}$ represents an instance of the matching setting described above. Then, there exists a set of prices $\mathbf{p}$ such that $\mathcal{R}^{\mathbf{p}}_{\mathcal{I}} \geqslant \frac{e-1}{2e}\mathcal{R}^{\mathcal{M}}_{\mathcal{I}}$.*

Proof. First, note that by the assumption that agents are identical, we have that prices and probabilities in Myerson's are constant across agents. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ by the matroids arising from customers being unit-demand and items having limited supply, respectively. Furthermore, if we apply Lemma 4.4 to $\mathcal{I}$ with some fixed α and $\mathcal{S}_1$, the prices $\mathbf{p}$ it provides will also be constant across agents. So for all i, j we have $p_{ij} = p_j$ and $q_{ij} = q_j$, some p_j, q_j. Let c_{ij} be the probability we could offer item j to agent i under just $\mathcal{S}_2$. Since there is no interaction between the uniform matroids making up $\mathcal{S}_2$, and prices are constant within each uniform matroid, we can apply Theorem 3.15 to get that for each item j

$$\sum_i q_{ij}p_{ij}c_{ij} \geqslant (1 - 1/e)\sum_i q_{ij}p_{ij} = q_j p_j (1 - 1/e)n,$$

implying that

$$\sum_i c_{ij} \geqslant (1 - 1/e)n.$$

Then, if we define X_{ij} and Y_{ij} as in the proof of Theorem 4.5, we can see that

$$\begin{aligned}
\mathcal{R}^{\mathbf{p}} &\geqslant \sum_i \sum_j \mathrm{E}[X_{ij}p_{ij}] \Pr[Y_{ij} = 1] \\
&= \sum_j \mathrm{E}[X_{ij}p_{ij}] \sum_i \Pr[Y_{ij} = 1] \\
&\geqslant \sum_j \mathrm{E}[X_{ij}p_{ij}] \sum_i c_{ij} \qquad (4.2) \\
&\geqslant \sum_j \mathrm{E}[X_{ij}p_{ij}] n(1 - 1/e) \\
&= (1 - 1/e) \sum_i \sum_j \mathrm{E}[X_{ij}p_{ij}] \\
&= (1 - 1/e)\mathcal{R}^{\mathbf{p}}_{\mathcal{I}'}.
\end{aligned}$$

Note that (4.2) follows because we know that the c_{ij}'s decrease with i, and failing to make some offers because of the additional constraint $\mathcal{S}_1$ can only improve the probability we can offer an item to a later agent.

Combining the above with Lemma 4.4 yields

$$\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geqslant \frac{(e-1)\alpha}{e(1+\alpha)} \mathcal{R}_{\mathcal{I}}^{\mathcal{M}};$$

choosing $\alpha = 1$ yields the claimed approximation factor of $2e/(e-1) \approx 3.164$. □

4.2.3 Intersection of graphical matroid and a partition matroid

Consider the instance $(I, \mathcal{S}, \mathbf{F})$ of the BSMD, where $\mathcal{S}$ is an intersection of a graphical matroid and a partition matroid Such an instance arises as the single-parameter analog $\mathcal{I}^{\text{reps}}$ of the following BMUMD instance: consider a graph $G = (V, E)$ where agents have independent values for different edges and are interested in buying one edge each; the seller can allocate any forest in the graph.

Theorem 4.8. *Let $\mathcal{I} = (I, \mathcal{S}, \mathbf{F})$ be an instance of the BSMD with $\mathcal{S}$ being an intersection of graphical matroid and a partition matroid. Then there is a COPM $(\mathbf{p}, \mathcal{S}')$, where $\mathcal{S}'$ is a partition matroid, such that $\mathcal{R}_{\mathcal{I}}^{\mathbf{p},\mathcal{S}'} \geqslant \frac{1}{7.47}\mathcal{R}_{\mathcal{I}}^{\mathcal{M}}$.*

Proof. Note that though the feasibility constraint we are facing is the intersection of a graphical matroid and partition matroid (from the unit demand constraint), we can view the feasibility constraint as if it was an intersection of two partition matroids. This follows from the proof of Theorem 3.7, where we showed that graphical matroids admit a COPM where the constrained set system is a union of 1-uniform matroids i.e., a partition matroid. The total probability mass of the elements of each 1-uniform matroid is at most 2. Fix $\alpha \in (0, 1]$, and consider trying to apply Lemma 4.4 to $\mathcal{I}$ with α and this partition matroid. The only change would be in the invocation of Lemma 3.4, where instead of bounding the sum of q_j within a part as the rank of the part, we now can only bound it by twice that. Thus, the proof will go through exactly as before, but with the modified final bound of

$$\mathcal{R}_{\mathcal{I}'}^{\mathbf{p}} \geqslant \left(\frac{\alpha}{1+2\alpha}\right) \mathcal{R}_{\mathcal{I}}^{\mathcal{M}}.$$

Proceeding as in the proof of Theorem 4.5 then gives us that

$$\mathcal{R}_{\mathcal{I}}^{\mathbf{p}} \geqslant (1-\alpha) \left(\frac{\alpha}{1+2\alpha}\right) \mathcal{R}_{\mathcal{I}}^{\mathcal{M}};$$

choosing $\alpha = (\sqrt{3}-1)/2$ yields the claimed approximation factor of $(1-\sqrt{3}/2)^{-1} \approx 7.47$. □

4.3 Approximation through order-specifying sequential pricings

In this section we show that we can achieve much better approximations for some settings by picking the right ordering σ over offers, that is, through sequenced posted-price mechanisms. While there is no direct reduction from SPMs in multi-parameter settings to SPMs in single-parameter settings analogous to Theorem 4.3, we show that for matroid and matroid intersection settings our results carry over in an approximation preserving way to multi-parameter instances as well.

We begin with a 2 approximation to single-parameter instances with a general matroid feasibility constraint, and show an improved 1.58 approximation for the special cases of uniform and partition matriods. We then describe a 3 approximation for general matroid intersection. We conclude this section by proving an 8 approximation in PDSE to multi-parameter instances with a general matroid intersection constraint. Table 4.2 summarizes the approximation factors for all the settings considered.

4.3.1 BMUMD with a matroid intersection constraint

Recall that Theorems 4.3 and 4.5 together show that we can achieve a constant factor approximation through OPMs to instances of the BMUMD with a feasibility constraint that is an intersection of two partition matroids. We now extend this result to general matroid intersections and combinatorial auctions with small bundles, albeit through SPMs in a slightly weaker solution concept — partial dominant-strategy implementation.

The prices and ordering in our approximately optimal SPM is picked in a manner similar to the one employed in Section 3.3.1 for matroids. Specifically, let $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ be the instance of BMUMD that we are interested in. Assume, to begin with, that all the distributions F_j are regular. Consider the instance $\mathcal{I}^{\text{reps}} = (I, \mathcal{S}, \mathbf{F})$, and let q_j^M be the probability with which the optimal truthful mechanism (Myerson's mechanism) allocates service j in that setting. Let $p_j^M = {F_j}^{-1}(1 - q_j^M)$. We define q_j

to be $q_j^M/2$ and $p_j = F_j^{-1}(1 - q_j)$. Let σ be the order of decreasing prices p_j over the services. Our approximately optimal SPM for $\mathcal{I}$ is $(\mathbf{p}, \sigma)$.

Theorem 4.9. *Given any instance $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ of the BMUMD, if $\mathcal{S}$ is a matroid intersection set system, then the SPM $(\mathbf{p}, \sigma)$ described above is an 8-approximation in PDSE to the revenue of the optimal incentive compatible mechanism for $\mathcal{I}$. Given an instance of a multi-parameter combinatorial auction with known bundles of size 2, the SPM $(\mathbf{p}, \sigma)$ implements an 8-approximation in PDSE.*

Proof. Let the pricing $\mathbf{p}$ and ordering σ be as defined above. Recall that Lemmas 3.2 and 4.1 together imply that the revenue of any incentive compatible mechanism for $\mathcal{I}$ is bounded above by $\sum_j p_j^M q_j^M \leqslant 2 \sum_j p_j q_j$.

Now consider the SPM $(\mathbf{p}, \sigma)$. We say that an agent i desires a service $j \in J_i$ if $v_j > p_j$, and i uniquely desires j if j is the only service in J_i with that property. As noted in Section 3.1.2, for an agent that uniquely desires a service, sincere bidding is a (weakly) dominant strategy. We first note that for every agent i, with probability 1/2, i bids sincerely. This follows from Markov's inequality by noting that $\sum_{j \in J_i} \Pr[i \text{ desires } j] = \sum_{j \in J_i} q_j = 1/2 \sum_{j \in J_i} q_j^M \leqslant 1/2$.

Now divide the set of all services into three groups—S, the set of *sold* services, B the set of services that are desired by their corresponding agents but "blocked" by services in S, and U the set of services that are desired by their corresponding agents and not in sets S or B. Note that these sets S, B, and U are random variables depending on the instantiation of agents' values. Then our observation above implies that services in U are not uniquely desired. Now, the expected total price in the union of the sets S, B and U is exactly $\sum_j p_j q_j$. Moreover, every desired service is uniquely desired with probability at least 1/2; therefore, the expected total price in U is at most half the total price of all desired services, that is, at most $1/2 \sum_j p_j q_j$. Finally, following the proof of Theorem 3.12, the expected total price in B conditioned on S is at most the total price contained in S. Therefore, putting everything together we get that the expected total price obtained from S is at least $1/4 \sum_j p_j q_j$. By our choice of $\mathbf{p}$ and $\mathbf{q}$, this is an 8-approximation.

When the distributions F_j are non-regular, we pick the prices p_j randomly as suggested by Lemma 3.2, such that the probability that service j is accepted if offered a randomized price is exactly $q_j^M/2$, and $\mathcal{R}^{\mathcal{M}}_{\mathcal{J}\text{reps}}$ is bounded from above by $2 \sum_j \mathrm{E}[p_j q_j]$. The order σ is one where services are offered in the decreasing order of their expected prices. The instantiation of price for a service is deferred till the service is offered. Then along the lines of the proof of Theorem 3.12, the expected total price in U is at most $1/2 \sum_j \mathbf{E}[p_j q_j]$, and, the expected total price in B (expectation over the values as well as the randomization over prices), conditioned on S, is at most the expected total price contained in S (expectation over the randomization over prices). Thus, the expected total price contained in S is at least $1/4 \sum_j \mathbf{E}[p_j q_j]$.

The argument for the combinatorial auction setting is identical and based on Theorem 3.18. We omit it for brevity. □

4.4 Discussion

We presented constant factor approximations to revenue for several classes of multi-parameter mechanism design problems by leveraging the approximately-optimal posted price mechanisms of Chapter 3 for single-parameter settings. The approximation factors we obtain depend on the kind of feasibility constraint that the seller faces. The exact constants are summarized in Tables 4.1 and 4.2. Note that the first is, in fact, identical to Table 3.1 since OPMs translate from single-parameter to multi-parameter settings with no loss in approximation factor.

While these approximation factors are with respect to the optimal deterministic incentive compatible mechanism, in the next chapter we show that (slightly worse) constant-factor approximation guarantees can be obtained against the optimal randomized incentive compatible mechanism as well.

More generally, two important assumptions underlie our work: (1) agents are unit-demand, and (2) their values for different services are distributed independently. In the absence of either of these assumptions the upper bound on the optimal revenue based on the single-parameter setting with representatives does not remain valid. An important open question is to design a reasonably tight upper bound in those cases, and use it to approximate the optimal mechanism.

Feasibility constraint $\mathcal{S}$	Gap from optimal	
	upper bound	lower bound
Uniform matroid, Partition matroid	2	2
Graphical matroid	3	2
Intersection of two part. matroids	5.83	2
Matching with i.i.d. agents	$2e/(e-1) \approx 3.17$	2
Graphical matroid $\cap$ partition matroid	7.47	2
General matroid	$O(\log k)$	2
Non-matroid downward closed	-	$\Omega(\frac{\log n}{\log \log n})$

Table 4.1: A summary of approximation factors for the BMUMD achievable through OPMs. Here k is the rank of $\mathcal{S}$.

Feasibility constraint $\mathcal{S}$	Gap from optimal	
	upper bound	lower bound
General matroid	2	$\sqrt{\pi/2} \approx 1.25$
Uniform matroid, Partition matroid	$e/(e-1) \approx 1.58$	1.25
Intersection of two matroids (BMUMD)	8	1.25
Non-matroid downward closed	-	$\Omega(\log n / \log \log n)$

Table 4.2: A summary of approximation factors for the BMUMD achievable through SPMs.

5 The power of randomization in multi-service settings

A fundamental problem in mechanism design is to characterize the optimal selling strategy for a monopolist trying to maximize his revenue. One of the big successes in this area is Myerson's characterization (1981) of the optimal mechanism in settings where buyers' types are single-dimensional. The optimal mechanism can be described in simple terms: it is a maximizer of ironed virtual values, weakly monotone transformations of agents' values. Remarkably, the optimal mechanism is deterministic. For example, the optimal selling strategy for a single-good monopolist facing a single buyer is to offer the buyer a take-it-or-leave-it price that depends on the value distribution of the buyer. Unfortunately, as we move away from the single-parameter setting, the design of optimal mechanisms becomes far more complex. Manelli and Vincent (2007) noted that the class of all mechanisms that are optimal for some distribution of agent types includes nearly all mechanisms[1]; so no nontrivial characterization of optimal mechanisms is possible. In particular, this class includes mechanisms that use randomness. Thanassoulis (2004) gave specific examples of settings where randomized mechanisms are strictly better than deterministic ones. This begs the question: *to what extent is randomness useful for revenue maximization? Are deterministic mechanisms near-optimal in multi-parameter settings as well?*

In this chapter we investigate the power of randomness in the context of the Bayesian mulit-parameter unit-demand mechanism design problm. To answer our questions in the BMUMD we must first understand the structure of randomized mechanisms in such multi-dimensional settings. In the context of a single unit-demand agent and a seller offering multiple items, any deterministic mechanism is simply a price for each of the items with the agent picking the one that maximizes her utility (her value for the item minus its price). Likewise, randomized mechanisms can be thought of as pricings for distributions or convex combinations over items. These convex combinations are called *lotteries*. A risk-neutral buyer with a quasilinear utility function buys the lottery that maximizes his expected value minus the price of the lottery.

[1]This holds even if we restrict the type distributions to be independent across agents and items, the setting we consider in this chapter.

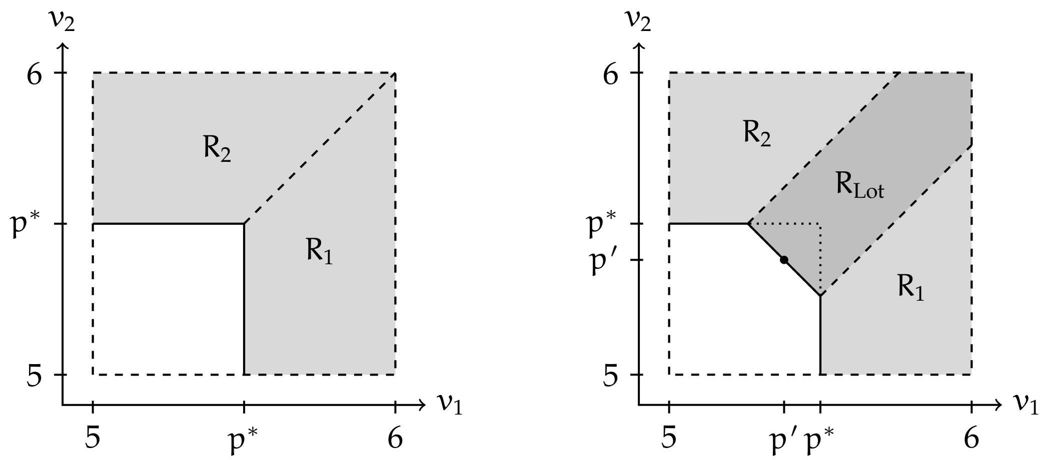

Figure 5.1: An example from Thanassoulis (2004) contrasting the optimal item and lottery pricings. The regions R_1, R_2, and R_{Lot} denote the sets of valuations at which the agent buys item 1, item 2, and the $(1/2, 1/2)$ lottery, respectively.

The following example due to Thanassoulis (2004) illustrates how lotteries work. Suppose that a seller offers two items for sale to a single buyer, and that the buyer's value for each of the items is independently and uniformly distributed in the interval $[5, 6]$. The optimal deterministic mechanism for the seller is to simply price each of the items at $p^* = \$5.097$ (see Figure 5.1). In a randomized mechanism, the seller may in addition price a $(1/2, 1/2)$ distribution over the two items at a slightly lower price of $p' = \$5.057$. If the buyer purchases this lottery, the seller tosses a coin and allocates the first item to her with probability 1/2 and the second with probability 1/2. A buyer that is nearly indifferent between the two items would choose to buy the lottery because of its lower cost, rather than either one of the items. While the seller loses some revenue by selling the lower priced lottery with some probability, he gains by selling to a larger segment of the market (those that cannot afford either of the individual items but can afford the lower priced lottery). In this example the gain is more than the loss, so that introducing the lottery improves the seller's revenue.

Lotteries are thus a mechanism for sellers to screen buyers on their relative preferences for different items. For example, car dealerships often offer a discount on new cars when buyers do not have a strict preference for color. Likewise many travel agencies offer discounted vacation packages in which the vendor providing the service is unknown and depends on the demand at the time that the vacation is undertaken. This partitions the market into buyers that are indifferent among the different colors or vendors and buyers that greatly prefer some of the options over others.

In general, a randomized mechanism can offer to the buyer a menu of prices for arbitrarily many lotteries. We call such a menu a *lottery pricing*, and likewise a deterministic pricing an *item pricing*. In multiple-agent settings randomized mechanisms can be more complicated, but when viewed from the perspective of a single buyer they behave like a lottery pricing.

From the point of view of an optimizing designer, knowing the quantitative benefit from using lotteries can be crucial. On the one hand, optimizing revenue over the space of all lottery pricings is easier than optimizing over the space of all item pricings—Briest et al. (2010) show that the former can be done by solving a linear program. On the other hand, in general the optimal lottery pricing can contain as many different options on the menu as the number of different buyer types in the market. This may be infeasible to implement from a practical point of view. However, optimizing revenue over menus with few options (e.g., a single lottery in addition to item prices) appears to be harder than finding the optimal item pricing. If lotteries offer only a marginal improvement in revenue over item pricings, then the seller may be better off just using item prices.

Until recently, the largest gap known between item pricings and lottery pricings for a single agent was a gap of 3/2 due to Pavlov (2006); for the special case where values for different items are independent, Thanassoulis (2004) gave the best gap example with a gap of 1.1. Briest et al. (2010) showed that in single-agent settings the gap between lottery pricings and item pricings can in fact be unbounded even with only 4 items. Specifically, they construct a discrete distribution over the agent's values, with each pair of value vectors having a large "angle" between them (and therefore representing different segments of the market). Then a lottery pricing, by offering different convex combinations over items to different segments, can obtain nearly the entire social value in the system, whereas a deterministic item pricing cannot price discriminate as effectively. Briest et al. show that when the number of items is at least 4, an unlimited number of such value vectors can be packed into the distribution, leading to an unbounded gap between the revenues of the optimal item and lottery pricings. The value distributions they construct, however, are quite unnatural with the values of different items being correlated in a specific way.

In this work we show that in single-agent problems with independent values (the setting considered by Thanassoulis) the revenue obtained by lottery pricings is no more than four times the revenue obtained by item pricings. We extend this result to settings involving multiple agents and general supply constraints, as well as a limited kind of correlation between values, again obtaining constant factor bounds on the benefit of randomness. While the constant factors we achieve are large in some cases, our results are in sharp contrast to the findings of Briest et al. where the improvement offered by lottery pricings is nearly as large as the number of different agent types in the market. Informally this implies that (when values are uncorrelated) randomized mechanisms cannot offer arbitrarily large improvements over deterministic ones. We believe that the factors we achieve can be improved considerably and the gap between randomized and deterministic mechanisms is much smaller in practice.

5.1 Warm up: Single-agent setting

In this section, we introduce our approach through the following fundamental BMUMD setting: we have a seller offering n items for sale to a single unit-demand agent. Recall, such an agent is interested in buying at most one item. For intuition, think of a buyer trying to choose which car or which laptop to buy. Let n be the number of available options, and $v_1, \ldots, v_n$ be the values the agent places on them, which are independently drawn from distributions $F_1, \ldots, F_n$, respectively.

Our goal is to study the gap between the revenue achievable from an optimal randomized truthful mechanism and that achievable from an optimal deterministic truthful mechanism. In this simple single-agent setting, deterministic truthful mechanisms are *item pricings*: the seller picks a vector of prices $\mathbf{p}$ for the items, and the buyer purchases the item i that maximizes her utility, $v_i - p_i$, or buys nothing if all items provide negative utility.

On the other hand, randomized truthful mechanisms can be interpreted as *lottery pricings* that are defined as follows. An n-dimensional *lottery* is a vector $\ell = (q_1, \cdots, q_n, p)$ where p is the price of the lottery and $(q_1, \cdots, q_n)$ is a probability distribution over the n items, $\sum_{i \in [n]} q_i \leqslant 1$. A lottery pricing $\mathcal{L} = \{\ell_1, \ell_2, \cdots\}$ is a menu of (an arbitrary number of) lotteries offered by the seller to the buyer. As for item pricings, the buyer purchases the lottery that maximizes her utility. We assume that the agent is risk-neutral, and so her utility from a lottery $\ell = (q_1, \cdots, q_n, p)$ is given by $\sum_{i\in[n]} q_i v_i - p$.

Characterizing and understanding optimal item and lottery pricings proves quite difficult. In order to relate them, we follow the same approach as in Chapter 4: we relate both to the (revenue) optimal mechanism for a related single-parameter instance, which we *do* understand well. In the simple single-agent case we are currently considering, the appropriate single-parameter problem to consider is a single-item, n-agent auction, where agent i has a value v_i for the item under sale. As before, v_i is drawn independently from F_i. We use the same terminology as in Chapter 4, and call agent i in this new setting a representative of the multi-parameter agent in the original setting (rep i for short). We call an instance of the original problem $\mathcal{I}$ and the corresponding instance of the new single-item auction problem $\mathcal{I}^{\text{reps}}$.

Item pricings for $\mathcal{I}$ and deterministic mechanisms for $\mathcal{I}^{\text{reps}}$ are closely related, and in particular, their input-output behavior is identical. Both accept a value vector $\mathbf{v}$ drawn from the product distribution $\mathbf{F} = \prod_i F_i$ as input and return as output an index i and a price to be paid. Not only do the two have similar structure, but their performance is closely related as well: Theorems 4.1 and 3.5 show that the revenue of the optimal (deterministic) mechanism for $\mathcal{I}^{\text{reps}}$ lies between the revenue of the optimal item pricing for $\mathcal{I}$ and twice this value, respectively. Note, however, that in single-parameter settings such as $\mathcal{I}^{\text{reps}}$ randomization gives no increase in revenue; the overall optimal mechanism *is* deterministic.

Recall that our goal is to understand how much the extra power lottery pricings give the designer can increase revenue relative to item pricings. As already mentioned, however, we know that the revenues of item pricings for $\mathcal{I}$ and mechanisms for $\mathcal{I}^{\text{reps}}$ are good approximations of each other, and so it is sufficient to relate the revenue of lottery pricings for $\mathcal{I}$ to mechanisms for $\mathcal{I}^{\text{reps}}$. While it remains the case that the input-output behavior of lottery pricings for $\mathcal{I}$ and randomized mechanisms for $\mathcal{I}^{\text{reps}}$ are the same, it becomes much harder to relate their revenues than it was in the deterministic case. The main challenge is to apportion the price of a lottery across reps in such a way that a large fraction of the original payment is recovered, while at the same time ensuring the mechanism is incentive compatible for the reps.

In order to make the issue of relating revenues between lottery pricings for $\mathcal{I}$ and randomized mechanisms for $\mathcal{I}^{\text{reps}}$, we begin by considering the most natural candidate mechanism to relate to a given lottery pricing. For any lottery pricing

$\mathcal{L}$, a natural candidate mechanism for $\mathcal{I}^{\text{reps}}$ based on $\mathcal{L}$ is one that implements the same allocation rule as $\mathcal{L}$: at a value vector $\mathbf{v}$, if $\mathcal{L}$ allocates an item i to the multi-parameter agent with probability q_i, then the mechanism for $\mathcal{I}^{\text{reps}}$ would allocate the item to rep i with probability q_i. We note that this allocation rule is monotone nondecreasing in the values v_i: if the multi-parameter agent increases his value for item i keeping other values the same, then he may switch to a different lottery that allocates item i with a higher probability, but never to one with a lower probability for item i. Therefore, by Theorem 2.3 there exists a truthful payment rule for this allocation rule in the $\mathcal{I}^{\text{reps}}$ setting, and furthermore this pricing rule is unique up to additive shifts. Call the resulting mechanism $\mathcal{A}^{\mathcal{L}}$. We normalize payments in this mechanism so that a rep with value 0 always pays 0.

Ideally we would like to claim that the revenue of $\mathcal{A}^{\mathcal{L}}$ is close to that of $\mathcal{L}$. As the following simple example shows, however, this is not always the case. Let $n = 2$ and consider a lottery pricing consisting of a single lottery $\ell = (1/2, 1/2, 1)$. Whenever the agent has a value of 2 for any one of the items, he buys the lottery and pays 1; this holds in particular when the agent values both items at 2. Now consider the corresponding mechanism $\mathcal{A}^{\mathcal{L}}$. Whenever any one of the reps values the item at 2, both reps get the item with a 1/2 probability each; in particular the other rep gets the item with a 1/2 probability regardless of her report. So when both of the reps value the item at 2, neither pays anything! Now, if the values are $(2, 2)$ with probability 1, the lottery pricing $\mathcal{L}$ gets a revenue of 1, whereas $\mathcal{A}^{\mathcal{L}}$ gets a revenue of 0. A key feature of this example is that this disparity in revenues occurs only at value vectors where both reps have sufficiently high values. But these are precisely the value vectors where another mechanism for $\mathcal{I}^{\text{reps}}$ obtains good revenue, namely the Vickrey auction. Intuitively, we should be able to bound the loss in revenue from $\mathcal{A}^{\mathcal{L}}$ (relative to $\mathcal{L}$) by the revenue of the Vickrey auction $\mathcal{V}$. As we shall next show, this intuition continues to hold even for more complicated instances of $\mathcal{I}$ than the simple example we discussed here.

We now formally prove that the combined revenues of $\mathcal{A}^{\mathcal{L}}$ and the Vickrey auction can be related to the revenue of $\mathcal{L}$. Let $\mathcal{R}^{M}(\mathbf{v})$ denote the revenue of a mechanism M at value vector $\mathbf{v}$ and $\mathcal{R}^{M}_{i}(\mathbf{v})$ denote the contribution of rep i to this revenue. Then we claim for all $\mathbf{v}$:

$$\mathcal{R}^{\mathcal{L}}(\mathbf{v}) \leqslant \mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) + \mathcal{R}^{\mathcal{V}}(\mathbf{v}).$$

As an intermediate step towards proving this claim, it will help to first study a related mechanism $\tilde{\mathcal{A}}^{\mathcal{L}}$ that has the same allocation rule as $\mathcal{A}^{\mathcal{L}}$ but a different (shifted) pricing rule that follows the lottery pricing more closely. Given the lottery pricing $\mathcal{L}$, the mechanism $\tilde{\mathcal{A}}^{\mathcal{L}}$ forms a one-dimensional lottery pricing for each of the n reps in $\mathcal{I}^{\text{reps}}$. Each rep then selects her utility maximizing lottery and purchases it (or elects to buy nothing, if no option yields nonnegative utility). We denote the lottery pricing offered to rep i as $\mathcal{L}_i$, and construct it as follows. For a given $\mathbf{v}_{-i}$ and $\ell = (q_1, q_2, \ldots, q_n, p) \in \mathcal{L}$, we add a lottery $\ell_i = (q', p')$ to $\mathcal{L}_i$, where

$$\begin{aligned} q' &= q_i; \text{ and} \\ p' &= p - \sum_{j \neq i} q_j v_j. \end{aligned}$$

Note that since $\tilde{\mathcal{A}}^{\mathcal{L}}$ offers each rep i a menu of options that does not depend on that rep's reported value v_i, we may immediately conclude that $\tilde{\mathcal{A}}^{\mathcal{L}}$ is truthful. Since a rep always has the option to reject all of the offered lotteries, receiving and paying nothing, we may also conclude that every rep always gets nonnegative utility from $\tilde{\mathcal{A}}^{\mathcal{L}}$ (i.e. it satisfies individual rationality). The following lemma shows that $\tilde{\mathcal{A}}^{\mathcal{L}}$ implements the same allocation rule as $\mathcal{L}$ (and thus $\mathcal{A}^{\mathcal{L}}$).

Lemma 5.1. *For any valuation vector $\mathbf{v}$, if the agent in $\mathcal{I}$ chooses lottery ℓ from $\mathcal{L}$, then each rep i in $\mathcal{I}^{reps}$ will select the lottery ℓ_i corresponding to ℓ in $\tilde{\mathcal{A}}^{\mathcal{L}}$.*

Proof. Fix some $\mathbf{v}$ and some i. Now, for any $\ell_i \in \mathcal{L}_i$, we can write the utility that rep i receives from it as

$$q' v_i - p' = q_i v_i - \left(p - \sum_{j \neq i} q_j v_j \right) = \sum_j q_j v_j - p,$$

precisely the utility the original agent received from the lottery ℓ that ℓ_i was derived from. Since both the original agent and the rep are utility maximizers, the result follows. □

We can now relate the revenue obtained by $\tilde{\mathcal{A}}^{\mathcal{L}}$ to that of $\mathcal{L}$.

Lemma 5.2. *For any fixed valuation vector $\mathbf{v}$, there exists an agent i^* such that $\mathcal{R}^{\mathcal{L}}(\mathbf{v}) \leqslant \mathcal{R}^{\tilde{\mathcal{A}}^{\mathcal{L}}}_{i^*}(\mathbf{v}) + \mathcal{R}^{\mathcal{V}}(\mathbf{v})$, where $\mathcal{V}$ denotes Vickrey's auction.*

Proof. Fix a valuation vector $\mathbf{v}$. Let $i^* = \text{argmax}_i v_i$. Let $\ell = (q_1, \ldots, q_n, p)$ be the lottery an agent with valuation vector $\mathbf{v}$ would select from $\mathcal{L}$ in $\mathcal{I}$. Then, $\mathcal{R}^{\mathcal{L}}(\mathbf{v})$ is precisely p, and so we can see that

$$\begin{aligned}
\mathcal{R}^{\mathcal{L}}(\mathbf{v}) &= \left(p - \sum_{i \neq i^*} q_i v_i\right) + \sum_{i \neq i^*} q_i v_i \\
&= \mathcal{R}_{i^*}^{\tilde{\mathcal{A}}^{\mathcal{L}}}(\mathbf{v}) + \sum_{i \neq i^*} q_i v_i \\
&\leqslant \mathcal{R}_{i^*}^{\tilde{\mathcal{A}}^{\mathcal{L}}}(\mathbf{v}) + \max_{i \neq i^*} v_i,
\end{aligned}$$

where the second equality follows from Lemma 5.1 and the definition of $\tilde{\mathcal{A}}^{\mathcal{L}}$, and the inequality follows from the fact that $\sum_i q_i \leqslant 1$ always. Since the revenue of Vickrey's auction for the setting $\mathcal{I}^{\text{reps}}$ on valuation vector $\mathbf{v}$ is precisely $\max_{i \neq i^*} v_i$, the claim follows. □

Note that while the payment rule of $\tilde{\mathcal{A}}^{\mathcal{L}}$ makes it easy to relate its revenue from some agent to the total revenue of $\mathcal{L}$, we cannot conclude that $\mathcal{R}^{\mathcal{L}}$ is at most the total revenue of $\tilde{\mathcal{A}}^{\mathcal{L}}$ plus $\mathcal{R}^{\mathcal{V}}$, because $\tilde{\mathcal{A}}^{\mathcal{L}}$ may charge some reps negative prices, i.e. compensate them, to align their preferences with those of the multi-parameter agent. Instead, we use the fact that $\tilde{\mathcal{A}}^{\mathcal{L}}$ and $\mathcal{A}^{\mathcal{L}}$ have identical allocation rules, and so have closely related payment rules.

In particular, Lemma 5.1 and Theorem 2.3 imply that the mechanisms $\tilde{\mathcal{A}}^{\mathcal{L}}$ and $\mathcal{A}^{\mathcal{L}}$ differ only by the payments they charge the reps at value 0. Since $\tilde{\mathcal{A}}^{\mathcal{L}}$ is individually rational, it must charge a nonpositive payment to rep i at $v_i = 0$. Therefore, $\mathcal{A}^{\mathcal{L}}$ charges payments that are no smaller than the payments in $\tilde{\mathcal{A}}^{\mathcal{L}}$, and obtains more revenue than the latter. Noting that $\mathcal{A}^{\mathcal{L}}$ always charges nonnegative payments, we conclude the following.

Lemma 5.3. *For any fixed valuation vector* $\mathbf{v}$ *and rep* i, *we have that* $\mathcal{R}_i^{\tilde{\mathcal{A}}^{\mathcal{L}}}(\mathbf{v}) \leqslant \mathcal{R}_i^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v}) \leqslant \mathcal{R}^{\mathcal{A}^{\mathcal{L}}}(\mathbf{v})$.

Lemmas 5.2 and 5.3 imply that $\mathcal{R}^{\mathcal{L}} \leqslant \mathcal{R}^{\mathcal{A}^{\mathcal{L}}} + \mathcal{R}^{\mathcal{V}} \leqslant 2\mathcal{R}^{\mathcal{M}}$ where $\mathcal{M}$ is the optimal mechanism for $\mathcal{I}^{\text{reps}}$. We can combine this with Theorems 4.3 and 3.5 from the previous chapters (which give an improvement on a corresponding theorem of Chawla et al. (2007)) to get the main result of this section: when values are distributed according to a product distribution, for any lottery pricing there exists an item pricing whose expected revenue is at least one fourth that of the lottery pricing.

Theorem 5.4. *Given a setting* $\mathcal{I}$ *where* $\mathbf{F}$ *is a product distribution, there exists an item pricing* $\mathbf{p}$ *for the setting* $\mathcal{I}$ *whose expected revenue is at least* $1/4$ *that of any lottery pricing for the setting* $\mathcal{I}$*.*

Pricings with a single lottery

As we noted earlier, optimal lottery pricings can contain as many lotteries as the number of different buyer types in the market. In many practical settings, offering such large menus is unreasonable, and the seller may instead want to construct a menu with a single lottery in addition to item prices. A natural question is whether we can improve the bound in Theorem 5.4 by restricting our attention to lottery pricings with only one lottery. The following simple argument shows that the bound improves from a factor of 4 to a factor of just 2.

For a lottery pricing of the given form, consider offering an agent either just the item pricings it contains, or just the single lottery it contains. Note that reducing the options in a lottery system only causes an agent to change their behavior if we remove their favorite option; thus, the combined revenue from offering these parts is at least the revenue of the original lottery pricing. Furthermore, an agent buying a lottery must value at least one of the items it randomizes over at the price of the lottery or higher. This implies that the revenue from offering a single lottery in isolation is no more than the revenue from offering each item in its support at the same price as the lottery. Since this gives two item pricings with combined revenue at least that of the original lottery pricing, one of the item pricings must give at least half of this amount.

5.2 Multi-agent settings

We will now prove the main result of the chapter, namely that the increase in a seller's revenue from using randomization in a multi-agent multi-parameter setting can be bounded by a small constant factor. We will extend the approach outlined in Section 5.1 for the special case of a single buyer: we bound the revenue of any BIC, IR randomized mechanism for an instance $\mathcal{I}$ of the BMUMD by those of three truthful deterministic mechanisms for the corresponding single-parameter instance with representatives, $\mathcal{I}^{\text{reps}}$.

Given a BIC, IR randomized mechanism M for $\mathcal{I}$, we first study the properties of the mechanism $\mathcal{A}^M$ for $\mathcal{I}^{\text{reps}}$ that has the same allocation rule as M (Section 5.2.1). Then, in an argument similar to the one in Section 5.1, we show that the revenue of M can be bounded by the revenue of $\mathcal{A}^M$ plus the revenues of two VCG-style mechanisms for $\mathcal{I}^{\text{reps}}$ (Section 5.2.2). This argument requires us to use some properties of matroid set systems that we describe in Section 2.1.2.

5.2.1 A mechanism for $\mathcal{I}^{\text{reps}}$

Consider an instance $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ of the BMUMD. Given a randomized BIC, IR mechanism M for $\mathcal{I}$, we define a mechanism $\mathcal{A}^M$ for the instance $\mathcal{I}^{\text{reps}}$ that also satisfies BIC and IR. As in the single-agent case, our goal is to relate the revenue of M to the revenue of the mechanism $\mathcal{A}^M$ implementing the same allocation rule in $\mathcal{I}^{\text{reps}}$. To this end, we also need to ensure $\mathcal{A}^M$ makes no positive transfers to the reps in $\mathcal{I}^{\text{reps}}$. Let $M(\mathbf{v})$ and $\pi(\mathbf{v})$ denote the allocation and payment rules, respectively, of the mechanism M. In an effort to unify our notation and theorem statements with Section 5.1, we define the quantities

$$q_{ij}(\mathbf{v}_i) = \mathop{\mathrm{E}}_{\mathbf{v}_{-i}}[M_{ij}(\mathbf{v}_{-i}, \mathbf{v}_i)]; \text{ and}$$
$$p_i(\mathbf{v}_i) = \mathop{\mathrm{E}}_{\mathbf{v}_{-i}}[\pi_i(\mathbf{v}_{-i}, \mathbf{v}_i)].$$

We are now ready to define the mechanism $\mathcal{A}^M$. For a given valuation vector $\mathbf{v}$, $\mathcal{A}^M$ will simulate the original mechanism M on $\mathbf{v}$. It makes an allocation of $M_{ij}(\mathbf{v})$ to rep (i, j), and charges a price of

$$p_i(\mathbf{v}_i) - \sum_{k \neq j} q_{ik}(\mathbf{v}_i) v_{ik} + u_{ij}(\mathbf{v}_{i,-j}).$$

Note that this payment rule is similar to the one defined for mechanism $\tilde{\mathcal{A}}^{\mathcal{L}}$ in Section 5.1 except for the additive term $u_{ij}(\mathbf{v}_{i,-j})$. The terms $u_{ij}(\mathbf{v}_{i,-j})$ are normalization factors that ensure that $\mathcal{A}^M$ never makes positive transfers to the agents. In the proof of Lemma 5.5 below we describe how to set these terms so that the resulting mechanism is BIC, IR, and makes no positive transfers. For the sake of continuity we defer the proof of this lemma to the end of this subsection.

Lemma 5.5. *If* M *satisfies BIC and IR, then for appropriately defined* $u_{ij}(\mathbf{v}_{i,-j})$*'s* $\mathcal{A}^M$ *satisfies BIC and IR, and makes no positive transfers.*

Following Section 5.1 as a guide, we would now like to upper bound the revenue of M point-wise, i.e. at every value vector $\mathbf{v}$, by the revenue of $\mathcal{A}^M$ at $\mathbf{v}$ and a few other terms. Note, however, that we defined the payments of $\mathcal{A}^M$ not in terms of the payments M actually charges, but in terms of their expectations (with respect to $\mathbf{v}_{-i}$). Thus, in order to achieve the sort of point-wise guarantee we want, we need to change how we account for the revenue of M. Specifically, we define the quantities

$$\bar{\mathcal{R}}_i^M(\mathbf{v}_i) = \mathop{\mathrm{E}}_{\mathbf{v}_{-i}}[\mathcal{R}_i^M(\mathbf{v}_{-i}, \mathbf{v}_i)]; \qquad \text{and} \qquad \bar{\mathcal{R}}^M(\mathbf{v}) = \sum_{i \in I} \bar{\mathcal{R}}_i^M(\mathbf{v}_i).$$

Note that $\mathrm{E}_{\mathbf{v}}[\mathcal{R}^M(\mathbf{v})] = \mathrm{E}_{\mathbf{v}}[\bar{\mathcal{R}}^M(\mathbf{v})]$, and so it suffices to get a point-wise bound for $\bar{\mathcal{R}}^M(\mathbf{v})$. Furthermore, $\bar{\mathcal{R}}^M(\mathbf{v})$ is a more natural candidate for a point-wise bound, since $\bar{\mathcal{R}}_i^M(\mathbf{v}_i) = p_i(\mathbf{v}_i)$.

While the mechanism $\mathcal{A}^M$ collects revenue from a total of mn reps, the proof of our bound only relies on the revenue extracted from a small subset of them. We denote this subset by $a(\mathbf{v})$. We choose a function $a(\mathbf{v})$ with the property that it includes at most one rep (i, j) for any given i. Formally, a *unit-demand function* $a(\cdot)$ is a function mapping valuation vectors to sets of reps that respect the unit-demand constraint, i.e. for any valuation vector $\mathbf{v}$ and $i \in I$, $|a(\mathbf{v}) \cap J_i| \leqslant 1$.

The following generalization of Lemma 5.2 is our main characterization of $\mathcal{A}^M$.

Lemma 5.6. *For any unit-demand function* $a(\cdot)$ *and any valuation vector* $\mathbf{v}$*, we have*

$$\begin{aligned}\bar{\mathcal{R}}^M(\mathbf{v}) &\leqslant \sum_{(i,j) \in a(\mathbf{v})} \mathcal{R}_{ij}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v}_i) v_{ij} \\ &\leqslant \mathcal{R}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{(i,j) \notin a(\mathbf{v})} q_{ij}(\mathbf{v}_i) v_{ij}.\end{aligned}$$

Proof. We bound the contribution of each agent to the revenue term $\bar{\mathcal{R}}^M(\mathbf{v})$ independently. Fix some agent i, and let $\mathbf{q}_i(\mathbf{v}_i)$ and $p_i(\mathbf{v}_i)$ denote the expected allocation and payment for agent i under mechanism M (where the expectation is over the other agents' values $\mathbf{v}_{-i}$). Recalling the definition of $\mathcal{A}^M$, we can see that

$$\begin{aligned}\bar{\mathcal{R}}_i^M(\mathbf{v}_i) = p_i(\mathbf{v}_i) &= \left(p_i(\mathbf{v}_i) - \sum_{k\neq j} q_{ik}(\mathbf{v}_i)v_{ik}\right) + \sum_{k\neq j} q_{ik}(\mathbf{v}_i)v_{ik} \\ &\leqslant \mathcal{R}_{ij}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{k\neq j} q_{ik}(\mathbf{v}_i)v_{ik},\end{aligned} \tag{5.1}$$

for any j, where $\mathcal{R}_{ij}^{\mathcal{A}^M}(\mathbf{v})$ is the revenue of mechanism $\mathcal{A}^M$ from the rep (i, j). The inequality follows from the fact that the normalization terms $u_{ij}(\mathbf{v}_{i,-j})$ we chose are always nonnegative. Note that for bounding the portion of $\bar{\mathcal{R}}^M(\mathbf{v})$ corresponding to agent i, we have used the portion of $\mathcal{R}^{\mathcal{A}^M}$ obtained from just a single rep (i, j). In particular, for a given agent i, we use the revenue obtained from rep $(i, j) \in a(\mathbf{v})$ when bounding $\bar{\mathcal{R}}_i^M(\mathbf{v}_i)$. Not every agent is guaranteed to have a rep in $a(\mathbf{v})$, however; for such an agent i, we instead use a formulation of the IR constraint for i under M: $p_i(\mathbf{v}_i) \leqslant \sum_k q_{ik}(\mathbf{v}_i)v_{ik}$. Adding the inequality (5.1) or the IR constraint (as appropriate) over all i, we get that for any unit-demand function $a(\cdot)$,

$$\begin{aligned}\bar{\mathcal{R}}^M(\mathbf{v}) &\leqslant \sum_{(i,j)\in a(\mathbf{v})} \mathcal{R}_{ij}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{(i,j)\notin a(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij} \\ &\leqslant \mathcal{R}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{(i,j)\notin a(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij},\end{aligned}$$

since we chose normalization terms $u_{ij}(\mathbf{v}_{i,-j})$ in the payments of $\mathcal{A}^M$ so that the mechanism never made positive transfers to agents. This is precisely our claimed bound. □

We now present the proof of Lemma 5.5.

Proof of Lemma 5.5. We first prove that $\mathcal{A}^M$ is BIC. Consider a rep (i, j) with value v_{ij}. Our key observation is that the expected utility that the rep gets in $\mathcal{A}^M$ from reporting a value v'_{ij} can be related to the expected utility that the corresponding agent i gets in M from reporting a value vector $(v'_{ij}, \mathbf{v}_{i,-j})$. In particular, we can write Rep (i,j)'s expected utility from reporting v'_{ij} in $\mathcal{A}^M$ as

$$\mathop{\mathrm{E}}_{\mathbf{v}_{-ij}}\left[M_{ij}(\mathbf{v}_{-ij},v'_{ij})v_{ij} - \left(p_i(\mathbf{v}_{i,-j},v'_{ij}) - \sum_{k\neq j} q_{ik}(\mathbf{v}_{i,-j},v'_{ij})v_{ik} + u_{ij}(\mathbf{v}_{i,-j})\right)\right]$$
$$= \mathop{\mathrm{E}}_{\mathbf{v}_{i,-j}}[\underbrace{\textstyle\sum_{k\in J} q_{ik}(\mathbf{v}_{i,-j},v'_{ij})v_{ik} - p_i(\mathbf{v}_{i,-j},v'_{ij})}_{\text{Agent } i\text{'s expected utility from reporting } (\mathbf{v}_{i,-j},v'_{ij}) \text{ in } M}] - \mathop{\mathrm{E}}_{\mathbf{v}_{-ij}}[u_{ij}(\mathbf{v}_{i,-j})]$$

Note that the first term inside the expectation above is the utility that agent i receives in M from reporting $(\mathbf{v}_{i,-j},v'_{ij})$; Since M is BIC, for every $\mathbf{v}_{i,-j}$ this term is maximized when agent i reports $(\mathbf{v}_{i,-j},v_{ij})$. The second term, on the other hand, is independent of v_{ij}. So the entire expression is maximized with $v'_{ij}=v_{ij}$, and $\mathcal{A}^M$ is BIC.

We now proceed to show that $\mathcal{A}^M$ satisfies IR, that is, all reps get nonnegative expected utility. First, consider omitting the $u_{ij}(\mathbf{v}_{i,-j})$ terms from the payments in $\mathcal{A}^M$. Then the equation above implies that rep (i,j) gets expected utility from $\mathcal{A}^M$ that is exactly equal to the expected utility that agent i gets from M given $\mathbf{v}_i$. Thus, if the $u_{ij}(\mathbf{v}_{i,-j})$ terms were all zero, then $\mathcal{A}^M$ would satisfy IR because M satisfies IR. Note, however, that the $u_{ij}(\mathbf{v}_{i,-j})$ terms cannot all be zero—an agent i may value some item j at $v_{ij}=0$, yet receive positive utility under M via the allocation of some other item; the only way for $\mathcal{A}^M$ to match the utility of rep (i,j) to that of agent i is by making a positive transfer. To cancel this positive transfer, we choose

$$u_{ij}(\mathbf{v}_{i,-j}) = \sum_{k\neq j} q_{ik}(\mathbf{v}_{i,-j},0)v_{ik} - p_i(\mathbf{v}_{i,-j},0),$$

For a rep with value $v_{ij}=0$ this makes the payment as well as utility of the rep equal to 0. To complete the proof we claim that (1) the rep gets positive expected utility at all v_{ij}, implying interim IR, and, (2) the rep makes nonnegative payments at all value vectors $\mathbf{v}$, implying no positive transfers. The first claim follows by noting that the mechanism is BIC and therefore the expected utility of any rep is a nondecreasing function of his value.

For the second claim, we use the fact that M is BIC to note that for any rep (i,j) and values $\mathbf{v}_{i,-j}$, the utility of the rep in expectation over $\mathbf{v}_{-i}$ is maximized when the rep reports his value truthfully. This is slightly stronger than the BIC condition proved above because it holds for all $\mathbf{v}_{i,-j}$ and not just in expectation over those values. This implies that the payment made by rep (i,j) in expectation over

$\mathbf{v}_{-i}$ is a nondecreasing function of his value v_{ij}. However, observe that rep (i, j)'s payment is independent of $\mathbf{v}_{-i}$. Therefore, the ex post payment function is in fact a nondecreasing function of v_{ij}. Now, using the fact that payments are zero at $v_{ij} = 0$, we conclude that payments are always nonnegative. □

5.2.2 Main theorem

As noted earlier, the motivation for this setting arises in the context of multi-unit multi-item auctions. Consider, in particular, a seller with m different items and k_j copies of item j for all j. Each of the n unit-demand buyers have independently distributed values for each item. The seller's constraint is to allocate item j to no more than k_j agents, and to allocate at most one item to each agent. Note that the unit-demand constraint and the item supply constraints are each instances of partition matroids. Thus the system $\mathcal{S}$ in this setting is an intersection of two partition matroids.

More generally, in this section we consider set systems $\mathcal{S}$ that are intersections of the partition matroid corresponding to the unit-demand constraint over agents (call it $\mathcal{U}$) and an arbitrary other matroid over $I \times J$ (call it $\mathfrak{M}$).

As in the single-agent case, we want to bound the revenue of a randomized mechanism M for $\mathcal{I}$ by the revenues of $\mathcal{A}^M$ and a suitably defined Vickrey-style auction. We can apply Lemma 5.6 to each value vector to bound the revenue. If we want to achieve an analog of Theorem 5.4 for our current setting, however, we need to relate the second term in the bound Lemma 5.6 provides to the revenue of a feasible mechanism for $\mathcal{I}^{\text{reps}}$. As in the single-agent case, we may bound the second term by the value of the second best item for each agent. However, this may be far larger than the value of *any* feasible allocation. For instance, if the seller has a single copy of some item j and j is every agent's highest-valued item, then we cannot give every agent their highest-valued item without violating the supply constraint. Instead, we will use the fact that the q_{ij}'s arise from distributions over feasible allocations, and the corresponding second term can therefore be bounded by revenue obtained from feasible allocations. We now present the details.

Lemma 5.7. *Let $\mathcal{I}$ be an instance of the BMUMD. The revenue from any BIC, IR randomized mechanism M for $\mathcal{I}$ is at most five times the expected revenue of the optimal mechanism for $\mathcal{I}^{reps}$.*

Proof. Consider any mechanism M for $\mathcal{I}$ and recall the characterization of the corresponding mechanism $\mathcal{A}^M$ from Lemma 5.6:

$$\bar{\mathcal{R}}^M(\mathbf{v}) \leqslant \mathcal{R}^{\mathcal{A}^M}(\mathbf{v}) + \sum_{(i,j)\notin a(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij}, \tag{5.2}$$

where $a(\cdot)$ is any unit-demand function.

Recall that in the single-agent setting, we bound the second term in this characterization by the second highest value in $\mathbf{v}$. This essentially corresponds to taking $a(\mathbf{v}) = \text{argmax}_i v_i$. Likewise, here we will pick $a(\mathbf{v})$ to be the maximum value feasible set of reps; note that this is a unit-demand function. Let $A_1(\mathbf{v})$ denote this set (we drop the argument $\mathbf{v}$ wherever it is obvious). Summing Equation (5.2) over all value vectors we get

$$\bar{\mathcal{R}}^M(\mathbf{v}) \leqslant \mathcal{R}^{\mathcal{A}^M}(\mathbf{v}) + \underbrace{\sum_{(i,j)\notin A_1(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij}}_{T}. \tag{5.3}$$

Our goal in the rest of the proof becomes to bound the second term above, labeled T. Informally, as in Section 5.1, we bound T by the second best feasible set. To this end, we define the set A_2 as the maximum valued feasible set over the remaining $(I \times J) \setminus A_1$ reps:

$$A_2(\mathbf{v}) = \underset{S\in\mathcal{S}; S\cap A_1(\mathbf{v})=\emptyset}{\text{argmax}} \; v(S).$$

Note that our definition of $q_{ij}(\mathbf{v}_i)$ here does not exactly match up with the corresponding definition in Section 5.1. In that setting, the q_i values represented overall expected allocations; here, they represent the allocation agent i expects to receive knowing his or her own value, but not other agents' values. In particular, if an agent knows they have a relatively high value for a service with only small supply, it is reasonable for them to expect to receive it; but if *every* agent happens to value that service highly for a particular $\mathbf{v}$, most of the agents will necessarily be disappointed in *any* feasible allocation. The point is that for a particular value vector, the sum T might give a value much larger than any feasible set, so we cannot hope to get a point-wise bound. We know, however, that $\mathbf{q}_i(\mathbf{v}_i)$ is defined in terms of a feasible allocation, and so we can relate it to the values in A_2 in expectation. We formalize this in the following claim.

Claim 5.8. *The sum of values of all reps in* A_2 *is no less than* T *in expectation:*

$$\mathop{\mathrm{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\in A_2(\mathbf{v})} v_{ij}\Big] \geqslant \mathop{\mathrm{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\notin A_1(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij}\Big].$$

Proof. This follows simply by computing

$$\mathop{\mathrm{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\notin A_1(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij}\Big] = \mathop{\mathrm{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\notin A_1(\mathbf{v})} M_{ij}(\mathbf{v})v_{ij}\Big] \leqslant \mathop{\mathrm{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\in A_2(\mathbf{v})} v_{ij}\Big];$$

the final inequality is a result of the fact that M always allocates a feasible set, and that A_2 has the largest value among all feasible sets in the complement of A_1. □

Next, we want to claim that a VCG-style mechanism can extract the value of the set A_2. At a high level, we can do so by constructing a mechanism that serves reps in A_1 and charges them prices that are at least as large as the values of the reps they displace in A_2. We need to ensure that each rep in A_2 becomes a price setter for at least one rep in A_1. If we faced either of the constraints $\mathcal{U}$ or $\mathfrak{M}$ alone, we could simply run a VCG mechanism and know that each displaced rep in A_2 would set the price for some rep in A_1. When we consider the intersection $\mathcal{U} \cap \mathfrak{M}$, however, it might be the case that a rep in A_2 is displaced by two different reps from A_1, each with respect to a different constraint, and so fails to set a price for either one (since neither can be assigned sole responsibility for the displacement). Our approach is to design two VCG-style mechanisms M_1 and M_2, which focus on extracting revenue related to displacements arising from the constraints $\mathcal{U}$ and $\mathfrak{M}$ respectively.

To formalize this, we use Proposition 2.1 to construct two maps from A_2 to A_1, one for each of the matroid constraints:

$$\begin{aligned} g_1 : A_2 \to A_1 \text{ s.t. } \forall e \in A_2 : \quad & g_1(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathcal{U}, \text{ or} \\ & g_1(e) \text{ is defined and } A_1 \setminus \{g_1(e)\} \cup \{e\} \in \mathcal{U}; \text{ and} \\ g_2 : A_2 \to A_1 \text{ s.t. } \forall e \in A_2 : \quad & g_2(e) \text{ is undefined and } A_1 \cup \{e\} \in \mathfrak{M}, \text{ or} \\ & g_2(e) \text{ is defined and } A_1 \setminus \{g_2(e)\} \cup \{e\} \in \mathfrak{M}. \end{aligned}$$

It follows that for any (i,j) in A_2, the set $A_1 \cup \{(i,j)\} \setminus \{g_1(i,j), g_2(i,j)\}$ is a feasible set. Furthermore, by the optimality of A_1, we have $v_{ij} \leqslant v_{g_1(i,j)} + v_{g_2(i,j)}$. The maximality of A_1 implies that every element of A_2 has an image under either g_1 or g_2 or both.

We are now ready to define the mechanisms M_1 and M_2 by specifying their allocation rules. Given a valuation vector $\mathbf{v}$, the mechanism M_1 serves only those reps (i, j) that belong to A_1 and for which $v_{ij} \geqslant v_{g_1^{-1}(i,j)}/2$ (if g_1^{-1} is defined at that point). Likewise, mechanism M_2 serves only those reps $(i, j) \in A_1$ that have $v_{ij} \geqslant v_{g_2^{-1}(i,j)}/2$ (again, if it is defined). We note that M_1 and M_2 have monotone allocation rules, and are therefore IC. Truthful payments satisfying IR can be defined appropriately. Note that by our choice of allocation rule, whenever mechanism M_1 or mechanism M_2 serves rep (i, j), it charges a payment of at least $v_{g_1^{-1}(i,j)}/2$ or $v_{g_2^{-1}(i,j)}/2$, respectively. Since both M_1 and M_2 serve subsets of A_1, they are both feasible under $\mathcal{S}$ as well.

The following claim lower bounds the combined revenue of M_1 and M_2.

Claim 5.9. *Twice the combined revenue of mechanisms M_1 and M_1 is no less than the sum of values of all reps in A_2 in expectation, i.e.,*

$$2\left(\mathcal{R}^{M_1} + \mathcal{R}^{M_2}\right) \geqslant \mathop{\mathbf{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\in A_2} v_{ij}\Big].$$

Proof. Consider any rep $(i, j) \in A_2$, and the reps $g_1(i, j)$ and $g_2(i, j) \in A_1$ (if defined). Note that $A_1' = A_1 \cup (i, j) \setminus \{g_1(i, j), g_2(i, j)\}$ is feasible. Suppose both $v_{g_1(i,j)}$ and $v_{g_2(i,j)}$ are less than $v_{ij}/2$; then the set A_1' is a feasible set and $v(A_1') > v(A_1)$ which is a contradiction to the optimality of A_1. Thus one of $v_{g_1(i,j)}$ or $v_{g_2(i,j)}$ must be at least $v_{ij}/2$ and so M_1 or M_2 charges that rep this amount, respectively. So we get that

$$2\left(\mathcal{R}^{M_1}(\mathbf{v}) + \mathcal{R}^{M_2}(\mathbf{v})\right) \geqslant \sum_{(i,j)\in A_2} v_{ij},$$

for any $\mathbf{v}$; taking expectation over $\mathbf{v}$ yields the claim. □

From equation (5.3) we can see that

$$\mathcal{R}^M = \mathop{\mathbf{E}}_{\mathbf{v}}[\bar{\mathcal{R}}^M(\mathbf{v})] \leqslant \mathcal{R}^{\mathcal{A}^M} + \mathop{\mathbf{E}}_{\mathbf{v}}\Big[\sum_{(i,j)\notin A_1(\mathbf{v})} q_{ij}(\mathbf{v}_i)v_{ij}\Big];$$

combining the above with Claims 5.8 and 5.9 completes the proof of Lemma 5.7. □

Combining the lemma with Theorems 4.9 and 4.5 we obtain the following theorem.

Theorem 5.10. *Given an instance $\mathcal{I}$ of the BMUMD with unit-demand agents and a matroid feasibility constraint, there exists a deterministic mechanism for $\mathcal{I}$ that obtains in PDSE at least a* $1/40$ *fraction of the revenue of the optimal BIC, IR randomized mechanism for $\mathcal{I}$. In the special case of multi-unit multi-item auctions, the revenue of any BIC, IR randomized mechanism is at most* 33.75 *times the revenue of the optimal deterministic mechanism for $\mathcal{I}$.*

5.3 Common base value correlation

In the previous sections we considered settings where buyers have independent values for the different services offered. We now consider a more general model for agent types. Since the agents are unit-demand, we can think of the services being offered as perfect substitutes. A natural form of correlation, then, is for the agent to have some "base" value for being served (regardless of which service is received), plus an additive deviation specific to the particular service received.

Formally, in the *common base value* setting, agents' types consist of $(m+1)$ independently distributed values $\{t_0, t_1, \ldots, t_m\}$, with t_0 being the base value for getting served and t_i being the additional benefit of obtaining service i; the agent's value for service i becomes $v_i = t_i + t_0$. Henceforth, we use the abbreviation CBV to refer to the version of BMUMD with this type of common base value correlation.

5.3.1 Warm up: single-agent setting

Once again we introduce our techniques through the single-agent setting. At a high level, our approach is to "reduce" the CBV setting to BMUMD with independent values. In particular, given an instance $\mathcal{I}$ of the former, we construct an instance $\widehat{\mathcal{I}}$ of the latter such that a lottery pricing for the former can be converted into one for the latter without much loss in revenue, and conversely an item pricing for the latter can be converted into one for the former. Then we can just apply Lemma 5.7 to $\widehat{\mathcal{I}}$ to obtain a bound on the benefit of randomness for $\mathcal{I}$. The transformation from $\mathcal{I}$ to $\widehat{\mathcal{I}}$ is straightforward except that $\widehat{\mathcal{I}}$ does not satisfy the unit-demand constraint, and we need to modify the proof of Lemma 5.7 appropriately. We now present the details.

Theorem 5.11. *Given an instance $\mathcal{I}$ of the CBV, there exists an item pricing* $\mathbf{p}$ *such that the revenue of any lottery menu $\mathcal{L}$ for $\mathcal{I}$ satisfies $\mathcal{R}^{\mathcal{L}} \leqslant 8\mathcal{R}^{\mathbf{p}}$.*

Proof. We begin by proving a bound with a weaker multiplicative factor of 9 and then show how to improve it to a factor of 8. We first define an uncorrelated instance $\widehat{\mathcal{I}}$ of the BMUMD. $\widehat{\mathcal{I}}$ is a single-agent setting with $(m+1)$ items; we interpret the tuple $\{t_0, \dots, t_m\}$ making up an agent's type in $\mathcal{I}$ as being the values of the agent in setting $\widehat{\mathcal{I}}$ for the $(m+1)$ items. In keeping with $\mathcal{I}$, the feasibility constraint we associate with $\widehat{\mathcal{I}}$ is that we may sell item 0, and at most one additional item from among items $1, \dots, m$. Note that the agent in $\widehat{\mathcal{I}}$ is not a unit-demand agent.

We first show how to convert a lottery menu $\mathcal{L}$ for $\mathcal{I}$ into a lottery menu $\widehat{\mathcal{L}}$ for $\widehat{\mathcal{I}}$ with no loss in revenue. For a lottery $\ell = (q_1, \dots, q_m, p)$ in $\mathcal{L}$, we define $q_0 = \sum_{i=1}^m q_i$, and add the lottery $\widehat{\ell} = (q_0, \dots, q_m, p)$ to $\widehat{\mathcal{L}}$. Note that $\widehat{\ell}$ does not necessarily satisfy the requirement that the q_i's sum to at most one; it does, however, satisfy the feasibility constraint indicated for $\widehat{\mathcal{I}}$. Furthermore, fixing a type $t_0, \dots, t_m$, the utility an agent in $\mathcal{I}$ receives from a lottery $\ell \in \mathcal{L}$ is

$$\sum_{i=1}^{m} q_i v_i - p = \sum_{i=1}^{m} q_i(t_i + t_0) - p = \sum_{i=0}^{m} q_i t_i - p,$$

which is precisely the utility a corresponding agent in $\widehat{\mathcal{I}}$ would receive from the corresponding $\widehat{\ell} \in \widehat{\mathcal{L}}$. We thus have $\mathcal{R}^{\mathcal{L}} = \mathcal{R}^{\widehat{\mathcal{L}}}$.

Next we will prove an analog of Theorem 5.4 for $\widehat{\mathcal{I}}$. Note that we cannot apply that theorem directly because the instance does not satisfy the unit-demand constraint.

Consider the setting $\widehat{\mathcal{I}}^{\text{reps}}$ obtained by replacing the multi-parameter agent in $\widehat{\mathcal{I}}$ by $m+1$ single-parameter reps. Let $\widehat{\mathcal{M}}$ be the optimal mechanism for this instance. Likewise, let $\widehat{\mathcal{V}}$ be the Vickrey auction and $\mathcal{A}^{\widehat{\mathcal{L}}}$ be the mechanism with the same allocation rule as $\widehat{\mathcal{L}}$ for the instance $\widehat{\mathcal{I}}^{\text{reps}}$. Then, by following the proof of Lemma 5.2, due to the less restrictive feasibility constraint ($\sum_{i=0}^m q_i \leqslant 2$) we get

$$\mathcal{R}^{\widehat{\mathcal{L}}} \leqslant \mathcal{R}^{\mathcal{A}^{\widehat{\mathcal{L}}}} + 2\mathcal{R}^{\widehat{\mathcal{V}}} \leqslant 3\mathcal{R}^{\widehat{\mathcal{M}}}.$$

To complete the argument, we need to relate the revenue of the mechanism $\widehat{\mathcal{M}}$ for $\widehat{\mathcal{I}}^{\text{reps}}$ to that of a deterministic pricing for $\mathcal{I}$. To this end, a key observation is that our feasibility constraint for $\widehat{\mathcal{I}}^{\text{reps}}$ (carried over from $\widehat{\mathcal{I}}$) implies that $\widehat{\mathcal{M}}$ may make decisions about allocations and prices for rep 0 separately from those for reps $1, \ldots, m$; as such, $\widehat{\mathcal{M}}$ effectively consists of two mechanisms, one serving rep 0 and another serving reps $1, \ldots, m$, both under a unit-demand constraint. Now, the optimal mechanism for serving rep 0 is a pricing with a single price; applying this price to all of the items in $\mathcal{I}$ gives us a pricing with revenue at least as large as that of the optimal mechanism to serve rep 0. For the optimal mechanism serving reps $1, \ldots, m$, Theorem 3.5 implies that there exists a pricing in $\widehat{\mathcal{I}}$ that obtains half the revenue of the mechanism; we apply the same pricing to $\mathcal{I}$ and note that the agent continues to select the same item under this pricing. Therefore, we can see that

$$\mathcal{R}_0^{\widehat{\mathcal{M}}} \leqslant \mathcal{R}^{\mathbf{p}} \quad \text{and} \quad \mathcal{R}_{-0}^{\widehat{\mathcal{M}}} \leqslant 2\mathcal{R}^{\mathbf{p}},$$

where $\mathbf{p}$ is the optimal pricing for $\mathcal{I}$, $\mathcal{R}_0^{\widehat{\mathcal{M}}}$ is the revenue of $\widehat{\mathcal{M}}$ from rep 0 in $\widehat{\mathcal{I}}^{\text{reps}}$, and $\mathcal{R}_{-0}^{\widehat{\mathcal{M}}}$ is the revenue of $\widehat{\mathcal{M}}$ from the other reps in $\widehat{\mathcal{I}}^{\text{reps}}$. Putting everything together, we get

$$\mathcal{R}^{\mathcal{L}} \leqslant 3\mathcal{R}^{\widehat{\mathcal{M}}} = 3(\mathcal{R}_0^{\widehat{\mathcal{M}}} + \mathcal{R}_{-0}^{\widehat{\mathcal{M}}}) \leqslant 3(\mathcal{R}^{\mathbf{p}} + 2\mathcal{R}^{\mathbf{p}}) \leqslant 9\mathcal{R}^{\mathbf{p}}.$$

In order to improve the factor from 9 to 8, we employ better bounds on the contribution of t_0 and the contribution of other values to the revenue of a mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$, depending on the type vector at which they are evaluated. Denote these quantities as $\mathcal{R}_0^M(\mathbf{t})$ and $\mathcal{R}_{-0}^M(\mathbf{t})$, respectively, for a mechanism M at a particular valuation vector $\mathbf{t}$. As previously noted, the optimal mechanism $\mathcal{M}$ in $\widehat{\mathcal{I}}^{\text{reps}}$ treats rep 0 independently from reps $1, \ldots, m$; thus, we have that any mechanism M in this setting must satisfy both $\mathcal{R}_0^M \leqslant \mathcal{R}_0^{\widehat{\mathcal{M}}}$ and $\mathcal{R}_{-0}^M \leqslant \mathcal{R}_{-0}^{\widehat{\mathcal{M}}}$.

Since we know that $\sum_{i=1}^m q_i \leqslant 1$, when t_0 is the maximum among all the t_i, Lemma 5.6 implies that

$$\mathcal{R}^{\mathcal{L}}(\mathbf{t}) = \mathcal{R}^{\widehat{\mathcal{L}}}(\mathbf{t}) \leqslant \mathcal{R}_0^{\mathcal{A}^{\widehat{\mathcal{L}}}}(\mathbf{t}) + \mathcal{R}_0^{\widehat{\mathcal{V}}}(\mathbf{t});$$

on the other hand, when one of $t_1, \ldots, t_m$ takes on the maximum value, we end up with, for some i,

$$\mathcal{R}^{\mathcal{L}}(\mathbf{t}) = \mathcal{R}^{\widehat{\mathcal{L}}}(\mathbf{t}) \leqslant \mathcal{R}_i^{\mathcal{A}^{\widehat{\mathcal{L}}}}(\mathbf{t}) + 2\mathcal{R}_i^{\widehat{\mathcal{V}}}(\mathbf{t}).$$

Combining these two gives us a point-wise guarantee of

$$\mathcal{R}^{\mathcal{L}}(\mathbf{t}) \leqslant \mathcal{R}_0^{\mathcal{A}^{\widehat{\mathcal{L}}}}(\mathbf{t}) + \mathcal{R}_0^{\widehat{\mathcal{V}}}(\mathbf{t}) + \mathcal{R}_{-0}^{\mathcal{A}^{\widehat{\mathcal{L}}}}(\mathbf{t}) + 2\mathcal{R}_{-0}^{\widehat{\mathcal{V}}}(\mathbf{t}).$$

Therefore, if we let $\mathbf{p}$ be the optimal pricing for $\mathcal{I}$, we get

$$\mathcal{R}^{\mathcal{L}} \leqslant \mathcal{R}_0^{\mathcal{A}^{\hat{\mathcal{L}}}} + \mathcal{R}_0^{\hat{\mathcal{V}}} + \mathcal{R}_{-0}^{\mathcal{A}^{\hat{\mathcal{L}}}} + 2\mathcal{R}_{-0}^{\hat{\mathcal{V}}} \leqslant 2\mathcal{R}_0^{\widehat{\mathcal{M}}} + 3\mathcal{R}_{-0}^{\widehat{\mathcal{M}}} \leqslant 2\mathcal{R}^{\mathbf{p}} + 6\mathcal{R}^{\mathbf{p}},$$

implying the claimed bound of 8. □

5.3.2 Multi-agent settings

We now consider the common base value model for multi-agent settings. As with the single-agent case, given an instance $\mathcal{I} = (I \times J, \mathcal{S}, \mathbf{F})$ of our problem, we construct a related setting $\widehat{\mathcal{I}}$ with independent values in such a way that any randomized mechanism M for $\mathcal{I}$ corresponds naturally to a randomized mechanism $\widehat{M}$ for $\widehat{\mathcal{I}}$ that achieves the same expected revenue. Our argument has three main steps:

1. We define the related setting $\widehat{\mathcal{I}}$ and mechanism $\widehat{M}$.

2. We extend Lemma 5.7 to bound the revenue of the optimal mechanism for $\mathcal{I}$ in terms of the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$.

3. We apply results from Chapters 3 and 4 to obtain good deterministic mechanisms for $\mathcal{I}$ that approximate the revenue of the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$.

Step 1: Defining the setting $\widehat{\mathcal{I}}$ and mechanism $\widehat{M}$

The main idea behind our construction of the modified instance $\widehat{\mathcal{I}}$ remains unchanged from the single-agent case, namely, to create a new service that explicitly captures agents' base values for being served. We assign this extra service index 0, and map an agent i with values $(t_{i0} + t_{i1}, \dots, t_{i0} + t_{im})$ in $\mathcal{I}$ to an agent with values $(t_{i0}, t_{i1}, \dots, t_{im})$ in $\widehat{\mathcal{I}}$. We construct the mechanism $\widehat{M}$ from M by extending the allocation rule so that whenever M allocates service j to agent i, $\widehat{M}$ allocates both service j and service 0 to agent i. We use the payment rule of M without any changes for $\widehat{M}$. It is easy to see that $\widehat{M}$ is BIC and IR.

All that remains is to define the feasibility constraint $\widehat{\mathcal{S}}$ for the setting $\widehat{\mathcal{I}}$. In single-agent CBV settings, we used the constraint "service 0 plus at most one other service." In multi-agent CBV settings, we generalize this to "the *projection* of any feasible allocation for $\mathcal{I}$ onto service 0, plus any feasible allocation for $\mathcal{I}$." To make this formal, write the feasibility constraint for $\mathcal{I}$ as $\mathcal{S} = \mathcal{U} \cap \mathfrak{M}$, where $\mathcal{U}$ represents the unit-demand constraint and $\mathfrak{M}$ is the matroid associated with the instance $\mathcal{I}$. Define constraint $\mathcal{S}_0$ on the reps for service 0 as

$$\mathcal{S}_0 = \{S \subseteq I \times \{0\} : \exists S' \in \mathfrak{M} \text{ such that } (i, 0) \in S \Leftrightarrow (i, j) \in S' \text{ for some } j\}.$$

Our overall feasibility constraint for the instance $\widehat{\mathcal{I}}$ is then the direct product $\widehat{\mathcal{S}} = \mathcal{S}_0 \times \mathcal{S}$.The proposed feasibility constraint has several nice properties:

A. $\widehat{\mathcal{S}}$ is compatible with the modified mechanism $\widehat{M}$;

B. $\widehat{\mathcal{S}}$ is "close" to a unit-demand constraint (in particular, $\mathcal{S}_0$ and $\mathcal{S}$ are both unit-demand, and so every agent in $\widehat{\mathcal{I}}$ is allocated at most two services); and

C. $\widehat{\mathcal{S}}$ ensures that feasible allocations of service 0 and of services $1, \ldots, m$ are both related to feasible allocations in $\mathcal{S}$, but, importantly, allows decisions about these two to be made *independently*.

The second property above is what allows us to extend Lemma 5.7 to our current setting, and we discuss it further in the next subsection. The third property is critical to allowing us to find mechanisms for the original setting $\mathcal{I}$ that approximate the revenue of the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$; it means that a mechanism for $\widehat{\mathcal{S}}$ can be "split" into two mechanisms, one for service 0 and one for services $1, \ldots, m$, which respect $\mathcal{S}$ (or a projection of it).

Step 2: Bounding the optimal revenue for $\mathcal{I}$ via the optimal revenue for $\widehat{\mathcal{I}}^{\text{reps}}$

We now prove the following analogue to Lemma 5.7 for CBV instances. As in the proof of Theorem 5.11, we will use the fact that while the setting $\widehat{\mathcal{I}}$ is not unit-demand, it is "close" to unit-demand in that every feasible allocation can be covered by at most two unit-demand feasible allocations.

Lemma 5.12. *Consider an instance $\mathcal{I}$ of the CBV. The revenue from any BIC, IR mechanism M for $\mathcal{I}$ is at most nine times the expected revenue of the optimal mechanism for the instance $\widehat{\mathcal{I}}^{reps}$.*

Proof. We follow the same framework as in the proof of Lemma 5.7; in fact, much of the proof goes through exactly as written there, and we refer the reader to that proof in such instances.

In order to be able to leverage Lemma 5.6 in our proof, we need a unit-demand restriction of the feasibility constraint $\widehat{\mathcal{S}}$. We define $\widehat{\mathcal{S}}_{\mathcal{U}} = \mathcal{U} \cap (\mathcal{S}_0 \times \mathfrak{M})$ for this purpose. Note that $\widehat{\mathcal{S}}_{\mathcal{U}} \subseteq \widehat{\mathcal{S}}$, and further, for any $S \in \widehat{\mathcal{S}}$ there exist two sets $S_1, S_2 \in \widehat{\mathcal{S}}_{\mathcal{U}}$ such that $S \subseteq S_1 \cup S_2$.

We begin by considering the instance $\widehat{\mathcal{I}} = (I \times \{0\} \cup J, \widehat{\mathcal{S}}, \mathbf{F})$ and define, as before, a "best" set A_1 and a "second best" set A_2, and two VCG-style mechanisms M_1 and M_2. In this case, however, we define them with respect to the unit-demand feasibility constraint $\widehat{\mathcal{S}}_{\mathcal{U}}$ rather than the true feasibility constraint $\widehat{\mathcal{S}}$. We set

$$A_1(\mathbf{v}) = \operatorname*{argmax}_{S \in \widehat{\mathcal{S}}_{\mathcal{U}}} v(S) \qquad \text{and} \qquad A_2(\mathbf{v}) = \operatorname*{argmax}_{S \in \widehat{\mathcal{S}}_{\mathcal{U}}; S \cap A_1(\mathbf{v}) = \emptyset} v(S),$$

respectively. Given a mechanism M for $\mathcal{I}$, we can apply Lemma 5.6 to the mechanism $\widehat{M}$ for $\widehat{\mathcal{I}}$ as defined above and take an expectation over $\mathbf{v}$ to get

$$\mathcal{R}^{M} = \mathcal{R}^{\widehat{M}} \leqslant \mathcal{R}^{\mathcal{A}^{\widehat{M}}} + \mathop{\mathbf{E}}_{\mathbf{v}}\Big[\sum_{(i,j) \notin A_1(\mathbf{v})} q_{ij}(\mathbf{v}_i) v_{ij}\Big].$$

Recall that $\widehat{M}$ always allocates a set in $\widehat{\mathcal{S}}$, and that every such set can be covered by two sets in $\widehat{\mathcal{S}}_{\mathcal{U}}$. So, for any $\mathbf{v}$, we have

$$\sum_{(i,j) \notin A_1(\mathbf{v})} \widehat{M}_{ij}(\mathbf{v}) v_{ij} \leqslant 2 \sum_{(i,j) \in A_2(\mathbf{v})} v_{ij}.$$

Then, we can apply the rest of the proof of Lemma 5.7 unchanged (in particular, Claim 5.9) with respect to the feasibility constraint $\widehat{\mathcal{S}}_{\mathcal{U}}$ to get that

$$\mathcal{R}^{M} \leqslant \mathcal{R}^{\mathcal{A}^{\widehat{M}}} + 2 \mathop{\mathbf{E}}_{\mathbf{v}}\Big[\sum_{(i,j) \in A_2(\mathbf{v})} v_{ij}\Big] \leqslant \mathcal{R}^{\mathcal{A}^{\widehat{M}}} + 4(\mathcal{R}^{M_1} + \mathcal{R}^{M_2}).$$

Since $\widehat{\mathcal{S}}_{\mathcal{U}} \subseteq \widehat{\mathcal{S}}$, all three of these mechanisms are feasible for the instance $\widehat{\mathcal{I}}^{\text{reps}}$, and the result follows. □

Step 3: Bounding the optimal revenue for $\widehat{\mathcal{I}}^{\text{reps}}$

We now show that the revenue obtained by the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$ is no more than a constant factor larger than that obtained by a deterministic mechanism for $\mathcal{I}$. Note that as in the single-agent case, the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$ treats service 0 independently from services $1, \ldots, m$, and so we may bound the revenue from each of these separately. In the single-agent setting, it is straightforward to apply the techniques of Chapter 4 to get approximately optimal pricings for the original setting $\mathcal{I}$. In the multi-agent setting, however, agents can affect each other's outcomes in nontrivial ways, and our approximation arguments necessarily become more complex. Dealing with these issues for service 0 and for services $1, \ldots, m$ is substantially different, and so we address these cases separately.

Step 3a: Revenue from the base value

Consider a mechanism M for allocating service 0 for the instance $\widehat{\mathcal{I}}^{\text{reps}}$, that is, when the feasibility constraint is given by $\mathcal{S}_0$ alone. We can extend this mechanism to one for $\mathcal{I}$ as follows. For every agent that gets served at a certain price in M, we offer the agent any of the services 1 through m at the same price. Then, there is some feasible allocation to the agents of services that obtains the same expected revenue as the mechanism M for $\widehat{\mathcal{I}}^{\text{reps}}$. Of course, the multi-parameter agents in $\mathcal{I}$ may not choose to buy the services prescribed by this feasible allocation—each agent would instead buy the service with the most value. While an agent changing which service they buy has no effect on the revenue we receive from that agent, the choices of one agent affect what offers we can make to another, and hence when an agent changes what service they buy it can have a significant impact on *total* revenue. In order to show that we do not lose too much revenue in this manner, we need to leverage the fact that the approximations of Chapter 4 were via posted-price mechanisms.

We note that the feasibility constraint $\mathcal{S}_0$ is a matroid.[2] Therefore, we can employ the following theorem of Yan (2011), which improves upon the corresponding theorem in Chapter 4.

Theorem 5.13 (Yan 2011, Theorem 3.1 and Lemma 4.1)**.** *Given an instance of the BSMD with a matroid feasibility constraint, there exists a PPM with expected revenue at least a* $(1 - 1/e)$ *fraction of that of the optimal mechanism for the instance.*

We are now ready to bound the optimal revenue from service 0 for $\widehat{\mathcal{I}}^{\text{reps}}$.

[2] Precisely, $\mathcal{S}_0$ can be seen to be a transversal matroid via a bipartite graph with $I \times \{0\}$ on one side and $I \times J$ on the other, and an edge between $(i, 0)$ and (i, j) for every i and j.

Lemma 5.14. *Given an instance $\mathcal{I}$ of the CBV, there exists a truthful deterministic mechanism for $\mathcal{I}$ whose revenue is at least a $1/2(1-1/e)$ fraction of the revenue that an optimal mechanism for $\widehat{\mathcal{I}}^{reps}$ extracts from reps for service 0.*

Proof. Consider the instance $\widehat{\mathcal{I}}^{\text{reps}}$ with just the constraint $\mathcal{S}_0$. We first apply Theorem 5.13 to obtain an SPM P for this setting yielding a $(1-1/e)$ fraction of the optimal expected revenue. We now construct a mechanism for $\mathcal{I}$ as follows. Consider agents in decreasing order of the prices at which P offers service 0 to them in $\widehat{\mathcal{I}}^{\text{reps}}$. To each agent, offer any services that can be feasibly allocated at a uniform price equal to that in P, and allocate the service (if any) chosen by the agent.

Let us bound the revenue of this mechanism. Note that the revenue of P is no more than the sum of prices corresponding to a maximum-total-price independent set in $\mathcal{S}_0$, or equivalently, a maximum-total-price independent set in $\mathcal{S}$. On the other hand, the constructed mechanism for $\mathcal{I}$ can be thought of as a greedy algorithm for $\mathcal{S}$, a matroid intersection constraint, that breaks ties in some arbitrary manner (namely according to agents' utilities). Proposition 2.2 then implies that this greedy algorithm obtains at least half of the revenue of the best independent set in $\mathcal{S}$. We therefore obtain the lemma. □

Step 3b: Revenue from the service-specific values

Finally, we bound the revenue that the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$ extracts from reps for services $1, \ldots, m$. At a high level, our approach is similar to that followed in the previous subsection to bound the revenue from reps for service 0. We again invoke a theorem of Chapter 4 to get an SPM for $\widehat{\mathcal{I}}$, and then analyze its performance when run in $\mathcal{I}$ in the natural way – offering services to agents in the same order and for the same prices.

Relating the outcome and revenue of an SPM applied to $\mathcal{I}$ and $\widehat{\mathcal{I}}$ presents two challenges: (1) agents in $\mathcal{I}$ may make different decisions from those in $\widehat{\mathcal{I}}$; (2) the offers made to any single agent are potentially interleaved by offers to other agents. The latter complicates an agent's decision to either accept the current offer, or wait for a future offer that brings in more utility but may get preempted by a sale to another agent. To deal with these challenges, we leverage a stronger version of the revenue guarantees that SPMs provide for instances of the BMUMD that is implicit in the proof techniques of Chapter 4. Let us define the favorite service of an agent to be the one that brings him the most utility, $j^*(i) = \text{argmax}_j(v_{ij} - p_{ij})$ for an agent i.[3] Note that when an agent is offered his favorite service, it is a dominant strategy for the agent to accept this service. In particular, the agent need not strategize about what offers he may receive in the future. The SPMs designed in Chapter 4 get large expected revenue from agents who are only interested in buying their favorite service at the announced prices because the other services bring them negative utility. Focusing on this set of agents makes it easier for us to relate the agents' behavior in $\mathcal{I}$ and $\widehat{\mathcal{I}}$; in particular, each agent's favorite (utillity maximizing) service is the same in both settings.

The following stronger version of Theorem 4.9 is implicit from the proof given in Chapter 4.

Theorem 5.15. *Given an instance $\widehat{\mathcal{I}}$ of the BMUMD with unit-demand agents and a matroid feasibility constraint, there exists a PPM P for $\widehat{\mathcal{I}}$, such that the expected revenue that it obtains in PDSE from agents that derive positive utility from only one service (by definition, their favorite one) at the given prices, is at least a $1/8$ fraction of the revenue of any truthful mechanism for the instance $\widehat{\mathcal{I}}^{reps}$.*

We use the above theorem to approximate the revenue that the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$ extracts from reps for services $1, \ldots, m$.

Lemma 5.16. *Given an instance $\mathcal{I}$ of the CBV, there exists a deterministic mechanism for $\mathcal{I}$ achieving revenue in PDSE that is at least one sixteenth of that which the optimal mechanism for $\widehat{\mathcal{I}}^{reps}$ extracts from reps for services $1, \ldots, m$.*

Proof. Consider the instance $\widehat{\mathcal{I}}^{\text{reps}}$ with just the services 1 through m and the constraint $\mathcal{S}$. Applying Theorem 5.15 to the optimal mechanism for this instance, we get a SPM P whose expected revenue over $\widehat{\mathcal{I}}$ from agents that only desire their favorite service is at least an eighth of the expected revenue that the optimal mechanism for $\widehat{\mathcal{I}}^{\text{reps}}$ extracts from reps for services $1, \ldots, m$. We apply the mechanism P to setting $\mathcal{I}$ in the natural way: we offer services to agents in the same order and at the same prices as the ones used by P for the setting $\widehat{\mathcal{I}}$.

We will now relate the revenue of P from $\mathcal{I}$ to the revenue it obtains in $\widehat{\mathcal{I}}$ from agents that desire only their favorite service. Fix a type vector $\mathbf{t}$. Consider an agent

[3] If there are multiple utility-maximizing services, we break ties in favor of the one with the lowest price, i.e., in favor of the one occuring latest in the ordering of offers.

i that in mechanism P in the setting $\widehat{\mathcal{I}}$ desires only his favorite service $j^*(i)$ at the given prices. Recall that $j^*(i)$ is still the agent's favorite service in the mechanism P applied to setting $\mathcal{I}$. So, whenever the agent is offered $j^*(i)$, it is a dominant strategy for the agent to accept the service.

Formally, let $F(\mathbf{t})$ denote the set of (agent, service) pairs (i, j) where j is i's favorite service and the only service that the agent desires in the setting $\widehat{\mathcal{I}}$. Let $A(\mathbf{t})$ denote the set of (agent, service) pairs (i, j) where i accepts (purchases) j in the setting $\mathcal{I}$. For any pair (i, j) that belongs to $F(\mathbf{t})$ but not to $A(\mathbf{t})$ it must be the case that P does not offer service j to i in $\mathcal{I}$ because it is blocked by a previous sale, i.e., it cannot be allocated feasibly. Accordingly, let $B(\mathbf{t})$ denote the set of (agent, service) pairs (i, j) that are "blocked" in $\mathcal{I}$. Then we have $F(\mathbf{t}) \subseteq A(\mathbf{t}) \cup B(\mathbf{t})$.

Now we will focus on the set $A(\mathbf{t}) \cup B(\mathbf{t})$. Note that P's run on this set of (agent, service) pairs is essentially a greedy algorithm in which any pair that is not blocked is purchased. In particular, any pair that is not in $A(\mathbf{t}) \cup B(\mathbf{t})$ is offered but not accepted, and so does not effect P's run over this set of pairs. Since $\mathcal{S}$ is a matroid intersection constraint, we can apply Proposition 2.2 to conclude that the total revenue obtained by this run of the mechanism P in $\mathcal{I}$, namely the sum of the prices corresponding to the pairs in $A(\mathbf{t})$, is at least half of the total price in any feasible subset of $A(\mathbf{t}) \cup B(\mathbf{t})$.

On the other hand, by the definition of $F(\mathbf{t})$, the expected revenue of P in $\widehat{\mathcal{I}}$ from agents that only desire their favorite service is equal to the total price in $F(\mathbf{t})$. Since $F(\mathbf{t})$ is a feasible subset of $A(\mathbf{t}) \cup B(\mathbf{t})$, we may conclude that the revenue of P from $\mathcal{I}$ when the type vector is $\mathbf{t}$ is at least half of the revenue that P gets in $\widehat{\mathcal{I}}$ from agents desiring only their favorite service when the type vector is $\mathbf{t}$.

Taking expectations over $\mathbf{t}$ and applying Theorem 5.15 we get the lemma. □

Final approximation

We now combine the lemmas from each of the above subsections to get the main theorem of this section.

Theorem 5.17. *The revenue obtained by any BIC, IR randomized mechanism for an instance $\mathcal{I}$ of the CBV is at most $(162 + 18/(e-1)) \approx 172.5$ times the revenue of the optimal IR deterministic mechanism for $\mathcal{I}$ implemented in partial dominant strategies.*

Proof. Consider any BIC, IR mechanism M for $\mathcal{I}$. We apply Lemma 5.12 to M to get that $\mathcal{R}^{M}_{\mathcal{I}} \leqslant 9\mathcal{R}^{\mathcal{M}}_{\widehat{\mathcal{I}}^{\text{reps}}}$ where $\mathcal{M}$ is the optimal mechanism for the setting $\widehat{\mathcal{I}}^{\text{reps}}$.

Now, applying Lemmas 5.14 and 5.16 give us deterministic mechanisms for $\mathcal{I}$ that guarantee a $(e-1)/2e$ fraction and a $1/16$ fraction of the revenue that $\mathcal{M}$ extracts from reps for service 0 and services $1, \ldots, m$, respectively, in $\widehat{\mathcal{I}}^{\text{reps}}$. Thus, the better of these mechanisms guarantees revenue of at least $\mathcal{R}^{\mathcal{M}}_{\widehat{\mathcal{I}}^{\text{reps}}}/(16 + 2e/(e-1)) \geqslant \mathcal{R}^{M}_{\mathcal{I}}/(162 + 18/(e-1))$. □

5.4 A gap example

We showed in Sections 5.1 and 5.2 that the revenue of the optimal mechanism for $\mathcal{I}^{\text{reps}}$ gives an upper bound within constant factors to the revenue of an optimal randomized mechanism for $\mathcal{I}$. A natural question is whether it is possible to tighten our analysis to reduce the factor to 1, i.e., whether the optimal revenue for $\mathcal{I}^{\text{reps}}$ is a true upper bound on the revenue of an optimal randomized mechanism for $\mathcal{I}$. In this section we give a simple example where the revenue of a lottery pricing for a single-agent BMUMD instance $\mathcal{I}$ is strictly larger than the optimal revenue for the instance $\mathcal{I}^{\text{reps}}$; we then describe a generalization of this example where the former exceeds the latter by a factor of 1.13, the largest gap we know of.

The first instance we consider is defined as follows. There is a single agent with i.i.d. valuations for two items, each of which is drawn according to a discrete distribution with three point masses. Specifically, for $i = 1, 2$ we have that the valuation for item i is distributed as

$$v_i \sim \begin{cases} 1 & \text{with probability} 1/2; \\ 2 & \text{with probability} 1/2 - 1/H; \text{ and} \\ H & \text{with probability} 1/H, \end{cases}$$

where $H > 2$ is a parameter to be fixed later.

Before we describe the optimal pricing and lottery, let us note some properties of this instance. Agents who value both items highly are so rare that their contribution to revenue will be negligible. Thus, we are concerned with extracting revenue from two groups of agents: those who place a low value on both items, and those who place a high value on exactly one item. Ideally, we would like to offer a low price to the former and a high price to the latter, but item pricings do not give us the ability to target agents in this way. In the setting with reps, we expect that increased competition leads to increased revenue. However, the rarity of the event that both reps have high values means that competition does little to improve things. On the other hand, lotteries are perfectly suited to this example since they allow us to screen agents based on the strength of their preferences between the two items. Essentially, randomization lets us add a "new item" to our offerings, which is really a $(1/2, 1/2)$ chance on the two items. Agents with low values place a low value on this item; agents with exactly one high value place an *intermediate* value on this item. Thus, we can place a low price on the lottery to serve the former agents, and yet still sell items outright to the latter agents at an increased price. This improved market segmentation leads to the revenue increase that we shall shortly see.

Consider the optimal revenue in the setting $\mathcal{I}^{\text{reps}}$. If we relax the feasibility constraint so that we may serve both reps at once, we can only improve the optimal revenue. On the other hand, this relaxation means that allocation decisions for the two reps can be made independently, so the optimal mechanism just offers each rep a fixed price for service. Note, however, that by our choice of distribution, any price achieves expected revenue of at most 1. Thus, we can upper bound the revenue of the optimal mechanism for the setting $\mathcal{I}^{\text{reps}}$ by 2. The same argument also shows that the revenue from an optimal item pricing in the setting $\mathcal{I}$ is also upper bounded by 2.

We now proceed to give a lottery system that achieves expected revenue strictly exceeding 2. Consider the following lottery pricing $\mathcal{L}$ for $\mathcal{I}$:

$$\mathcal{L} = \left\{ \left(\frac{1}{2}, \frac{1}{2}, \frac{3}{2}\right), \left(1, 0, \frac{H+1}{2}\right), \left(0, 1, \frac{H+1}{2}\right) \right\}.$$

The first two coordinates in every lottery denote the probabilities with which items 1 and 2 are offered by that lottery and the third coordinate is the price. So this corresponds to putting a price of $(H+1)/2$ on each item, as well as offering a lottery that allocates each item with probability $1/2$ at a price of $3/2$. Then we can see that if the agent values exactly one of the items at H, they will buy that item outright; if they value both items at 1, they will buy nothing; and otherwise, they will buy the lottery. Thus, with probability $3/4$, the agent makes a purchase; and with probability $2/H(1 - 1/H)$ they choose to buy an item outright. The expected revenue will be

$$\frac{3}{4} \cdot \frac{3}{2} + \left(\frac{H+1}{2} - \frac{3}{2}\right) \frac{2}{H} \left(1 - \frac{1}{H}\right) = \frac{17}{8} - \frac{3}{H} + \frac{2}{H^2},$$

which is strictly greater than 2 for $H = 24$.

We now give a generalization of the above example to continuous distributions, and show that the gap in expected revenues increases to a factor of 1.13. The new instance $\mathcal{I}$ is defined as follows. There is still a single agent with i.i.d. valuations for two items; now the items' values are distributed according to a continuous *equal-revenue* distribution bounded at H. Formally, the valuations v_1 and v_2 for items 1 and 2 have distributions F_1 and F_2 such that

$$F_1(x) = F_2(x) = \begin{cases} 1 - 1/x & 1 \leqslant x < H; \text{ and} \\ 1 & x = H. \end{cases}$$

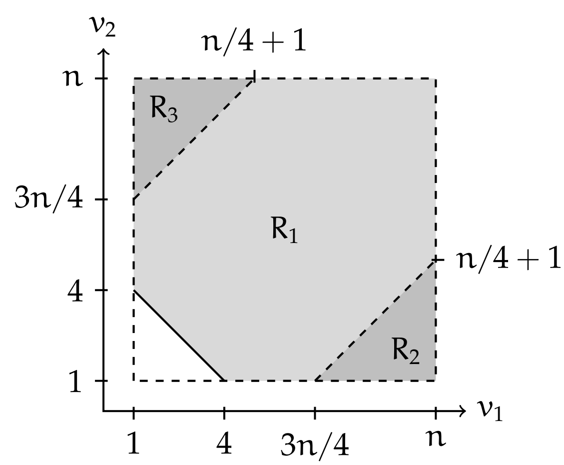

Figure 5.2: The allocation function for the lottery pricing $\mathcal{L}$.

Note that the distributions are regular. They also have the property that every fixed price yields expected revenue of at most 1. Thus, as in the previous example, we can think of relaxing the feasibility constraint to allow simultaneous allocation to both reps (or of both items) to get an upper bound of 2 on the expected revenue of any mechanism for the single-parameter setting $\mathcal{I}^{\text{reps}}$ (or for any item pricing for the setting $\mathcal{I}$).

The lottery system we consider for this new setting again consists of a price for each item, plus a lottery that gives an equal chance between the items:

$$\mathcal{L} = \left\{ \left(\frac{1}{2}, \frac{1}{2}, \frac{5}{2}\right), \left(1, 0, 2 + \frac{3H}{8}\right), \left(0, 1, 2 + \frac{3H}{8}\right) \right\}.$$

Figure 5.2 shows the allocation function of this lottery pricing. In particular, R_i for $i = 1, 2, 3$ is the set of valuations where lottery i is bought. The probability mass of regions R_2 and R_3 together can be computed to be $8/3H - o(1/H)$. The probability mass of region R_1 is $2/5 + (\ln 16)/25 - o(1)$. Therefore, the revenue of $\mathcal{L}$ can be computed to be $2 + (\ln 4)/5 - o(1) \approx 2.277 - o(1)$. This is a factor of 1.13 higher than the optimal revenue for $\mathcal{I}^{\text{reps}}$, or the revenue of any item pricing for $\mathcal{I}$.

5.5 Conclusions and open problems

In this work, we studied the gap between the revenues of the optimal deterministic and randomized mechanisms for multi-parameter settings with unit-demand agents, and showed that this gap is small when agents' values for different items are independent. Our results extend to a limited form of positive correlation between item values. Several open problems remain. Our bounds on the benefit of randomness are in some cases quite large (although always independent of the parameters of the instance) and we believe they can be improved. For example, in the case of a single-agent, we bound the gap by 4, while the largest gap example known only shows an improvement of approximately 1.13. It would also be interesting to extend our bounds to more general forms of correlation between item values, while avoiding the unbounded gap that Briest et al. (2010) achieve. Another direction of interest is to develop an understanding of special classes of randomized mechanisms, such as those in which the number or the kinds of lotteries offered is restricted.

6 Settings with budget-constrained agents

Auction and mechanism design have for the most part focused on agents with quasilinear utility functions: each agent is described by a function that assigns values to possible outcomes, and the agent's utility from an outcome is her value minus any payment that she makes to the mechanism. This implies, for example, that an agent offered an outcome at a price below her value for the outcome should in the absence of better alternatives immediately accept that outcome. This simple model fails to capture a basic practical issue—agents may not necessarily be able to afford outcomes that they value highly. For example, most people would value a large precious stone such as the Kohinoor diamond at several millions of dollars (for its resale value, if not for personal reasons), but few can afford to pay even a fraction of that amount. Many real-world mechanism design scenarios involve financially constrained agents and values alone fail to capture agents' preferences. Budget constraints have frequently been observed in FCC spectrum auctions Brusco and Lopomo (2009); Che and Gale (1998), Google's auction for TV ads Nisan et al. (2009), and sponsored search auctions, to take a few examples.

From a theoretical viewpoint, the introduction of budget constraints presents a challenge in mechanism design because they make the utility of an agent nonlinear and discontinuous as a function of the agent's payment—the utility decreases linearly with payment while payment stays below the budget, but drops to negative infinity when the payment crosses the budget. The assumption of linearity in payments (i.e. quasilinearity of utility) underlies much of the theoretical framework for mechanism design. Consequently, standard mechanisms such as the VCG mechanism can no longer be employed in settings involving budgets.

The goal of this chapter is to develop connections between budget-constrained mechanism design and the well-developed theory of unconstrained mechanism design. Specifically we ask "when can budget-constrained mechanism design be reduced to unconstrained mechanism design with some small loss in performance?" We consider this question in the context of the two most well-studied objectives in mechanism design—social welfare and revenue. Some of our results assume that the mechanism knows the budgets of the agents, but others hold even when budgets are private and agents need to be incentivized to reveal them truthfully.

Recent work in computer science has begun exploring a theory of mechanism design for budget-constrained agents (see, for example, Abrams (2006); Borgs et al. (2005); Dobzinski et al. (2008); Chen et al. (2010); Alaei et al. (2010)). Most of this work has focused on *prior-free* or worst-case settings, where the mechanism designer has no information about agents' preferences. Unsurprisingly, the mechanism designer has very little power in such settings, and numerous impossibility results hold. For example, in the worst-case setting no truthful mechanism can obtain a non-trivial approximation to social welfare Borgs et al. (2005). The goal of achieving good social welfare has therefore been abandoned in favor of weaker notions such as Pareto optimality Dobzinski et al. (2008). For the revenue objective while approximations can be achieved in simple enough settings, e.g. multi-unit auctions Borgs et al. (2005), hardness results hold for more general feasibility constraints even in the absence of budgets. In this work, we sidestep these impossibility results by considering Bayesian settings where the mechanism designer has prior information about the distributions from which agents' private values and private budgets are drawn.

We restrict our attention to direct revelation truthful mechanisms. Our mechanisms are allowed to randomize, and agents' utilities are computed in expectation over the randomness used by the mechanism. As is standard, we assume that both the mechanism and the agents possess a common prior from which values are drawn. While we optimize over the class of Bayesian incentive compatible (BIC) mechanisms, all of the mechanisms we develop are dominant strategy incentive compatible (DSIC) (see, for example, Nisan (2007) for definitions of these solution concepts).

In addition, we require that our mechanisms satisfy the ex-post individual rationality (EPIR) constraint, namely that the payment of any agent never exceeds her value for the mechanism's outcome. This implies, in particular, that the mechanism cannot charge any agent to whom no item or service is allocated. In contrast, most previous work has enforced the individual rationality constraint only in expectation over the mechanism's randomness as well as the randomness in other agents' values (i.e. interim IR).

It is worth noting here that the EPIR constraint is not without loss in performance. Consider the following example: suppose we are selling a single item to one of n agents, each with a value of v with probability 1 (that is publicly known) and a public budget of v/n with probability 1. Now, under the IIR constraint, the optimal auction asks agents to pay what they bid and offers each agent that pays at least v/n a fair chance at winning the item. Each agent pays v/n, the item is allocated to a random agent, and the mechanism's revenue is v. Under the EPIR constraint, however, a mechanism can only charge the agent that wins the item and can charge this agent no more than v/n. As we can see, the revenue gap between the optimal IIR and the optimal EPIR mechanism gets larger and larger as n grows.

It is well known that over the class of BIC IIR mechanisms, the revenue-optimal as well as welfare-optimal mechanisms are both so-called "all-pay" auctions Maskin (2000); Pai and Vohra (2008). In all-pay auctions agents pay the mechanism a certain (distribution dependent) function of their value regardless of the allocation that the mechanism makes. The optimality of all-pay auctions follows by noting that any allocation rule that admits some BIC budget-feasible payment function can be implemented with an all-pay payment rule with worst-case payments that are no larger than those in any other truthful payment rule and are therefore budget-feasible. Unfortunately all-pay auctions have many undesirable properties. In many settings it is simply not feasible to force the agents to pay upfront without knowing the outcome of the mechanism. Moreover all-pay auctions may admit many Bayes-Nash equilibria (BNE), truthtelling being merely one of them. Then the fact that a certain objective is achieved when all the agents report their true types does not necessarily imply that the objective will be achieved in practice if a different BNE gets played out. Therefore, in a departure from previous work, we choose to enforce ex-post individual rationality.

6.1 Maximizing revenue

We first consider the revenue objective, and begin by characterizing the optimal budget feasible mechanism for a single agent setting. The characterization relies on describing the mechanism as a collection of so-called lotteries or randomized pricings. We then consider settings with public budgets. Our general approach towards budget-constrained mechanism design in these settings is to approximate the optimal revenue in two parts: the contribution to optimal revenue by agents whose budget is binding (i.e. their budget is less than their value), and the contribution by agents whose budget is not binding (i.e. their budget is above their value).

We present different mechanisms for approximating these two benchmarks. We demonstrate this approach first in the simple setting of single-parameter agents with public budgets and an arbitrary downwards closed feasibility constraint. Then in subsequent sections we extend the approach to settings involving more complicated incentive constraints—multi-dimensional values and private budgets. In private budget settings, instead of asking agents to reveal budgets directly, our mechanism once again relies on collections of lotteries to motivate agents to pay a good fraction of their budgets when their values are high enough.

6.1.1 Single agent settings with public budgets

Before presenting our general approach, we first consider the most basic version of this problem—namely a setting with one single-parameter agent and a public budget constraint. Even this simple setting, however, reveals the challenges budget constraints introduce to the problem of mechanism design. Without the budget constraint, the optimal mechanism is to offer the item at a fixed price. With budgets, however, the following example shows that a single fixed price can be a factor of 2 from optimal. After the example we proceed to characterize the optimal mechanism.

Example 6.1. *Fix $n > 1$. Consider an agent whose value for receiving an item is $v = 1$ with probability $1 - 1/n$, and is $v = n^2$ with probability $1/n$. Let the agent have a budget of $B = n$. Any single fixed price that respects the budget in this setting receives a revenue of at most 1.*

We now describe the optimal mechanism. The mechanism offers two options to the agent: either buy the item at price n, or receive the item with a probability of $n/(n+1)$ at a price of $n/(n+1)$. This generates an expected revenue of $2n/(n+1) = 2 - o(1)$.

The optimal mechanism in the above example is what we call a *lottery menu* mechanism. A lottery is a pair (q, p) and offers to the agent at a price p a probability q of winning. A lottery menu is a collection of lotteries that an agent is free to choose from in order to maximize his expected utility. We will now show that for any single agent setting with a public budget, the optimal mechanism is a lottery menu mechanism with at most two options.

Consider a setting $\mathcal{I}$ with a single agent with private value $v \sim F$ and a public budget B. Let ϕ be the virtual value function corresponding to F. For ease of exposition, throughout the following discussion we will assume that F is regular and ϕ is non-decreasing; when F is non-regular, we can merely replace ϕ by $\bar{\phi}$, the ironed virtual value, in the following discussion.

We first note that if $B \geqslant \phi^{-1}(0)$ then the unconstrained optimal mechanism is already budget feasible. Therefore, for the rest of this section we assume that $B < \phi^{-1}(0)$. Using Myerson's theorem relating revenue to virtual values (see Proposition 2.3 in Section 2.3), our goal is to solve the following optimization problem.

$$\begin{aligned} \max_q \int q(v)\phi(v)f(v)dv & \qquad \text{subject to} \\ \int (q_{max} - q(v))dv \leqslant B \cdot q_{max}, & \qquad \text{and,} \\ q(v) \text{ is a non-decreasing function.} & \end{aligned}$$

Here $q_{max} \leqslant 1$ is the probability of allocation at the upper end of the support of the value distribution. The first constraint encodes the budget constraint. In particular, the left hand side of the inequality is the expected payment made by the agent at his highest value; the right side is an upper bound on the expected payment under EPIR because the agent can pay a maximum of B when he gets allocated, and 0 otherwise.

Let q^* be the optimal solution to the above optimization problem. We make the following observations (proofs are deferred to the end of this subsection in the interests of readability). In the following, we denote the inverse virtual value of 0 as $v^* = \phi^{-1}(0)$ to simplify notation.

Claim 6.2. *Without loss of generality, we may assume* $q^*_{max} = 1$.

Claim 6.3. *Without loss of generality, we may assume that for all* $v \geqslant v^*$, $q^*(v) = 1$.

Following these claims, our optimization problem changes to the following (the monotonicity constraint on q is implicit).

$$\begin{aligned} \max_q \int q(v)\phi(v)f(v)dv \quad & \text{subject to} \\ \int (1 - q(v))dv \leqslant B & \\ q(v) = 1 \ \ \forall v \geqslant v^* & \end{aligned}$$

This can be simplified to:

$$\min_{q} \int_0^{v^*} q(v)(-\phi(v)f(v))dv \text{ subject to}$$
$$\int_0^{v^*} (1-q(v))dv = B$$

Note that we replace the inequality in the budget constraint with an equality. This is because if the constraint is not tight, we can feasibly reduce $q(v)$ and thereby reduce the objective function value. For the sake of brevity, we define $B' = v^* - B$, and $g(v) = -\phi(v)f(v)$. The budget constraint then changes to $\int_0^{v^*} q(v)dv = B'$. Note that $v^* \geqslant B' \geqslant 0$, and g is nonnegative on $[0, v^*]$. Finally, we define the set of allocations

$$\mathcal{A} = \left\{ \begin{array}{c} \text{increasing } q : [0, v^*] \to [0,1] \\ \text{such that } \int_0^{v^*} q(v)dv = B' \end{array} \right\}$$

Then, we can express our objective as

$$\min_{q \in \mathcal{A}} \int_0^{v^*} q(v)g(v)dv.$$

If g is non-increasing on $[0, v^*]$, then we immediately have that the optimal solution is to set $q(v) = 1$ if $v \geqslant v^* - B'$ $(= B)$ and 0 otherwise.

If g is not non-increasing, we "iron" the function g to produce a non-increasing function $\hat{g}$ with the property that any non-decreasing function q that is constant over intervals where $\hat{g}$ is constant has the same integral with respect to $\hat{g}$ as with respect to g. Let $\widetilde{\mathcal{A}}$ be the subset of $\mathcal{A}$ containing all functions q that are constant over intervals where $\hat{g}$ is constant. We obtain the following lemma. (The details of the ironing procedure and the proof of the following lemma can be found at the end of this subsection, along with deferred proofs.)

Lemma 6.4. *For all $q \in \mathcal{A}$, there exists $\tilde{q} \in \widetilde{\mathcal{A}}$ such that $\int_0^{v^*} q(v)g(v)dv \geqslant \int_0^{v^*} \tilde{q}(v)g(v)dv$.*

The lemma lets us confine our optimization to the set $\widetilde{\mathcal{A}}$:

$$\min_{q \in \mathcal{A}} \int_0^{v^*} q(v)g(v)dv = \min_{q \in \widetilde{\mathcal{A}}} \int_0^{v^*} q(v)g(v)dv$$

Finally, we define $\mathcal{A}^*$ to be a subset of $\widetilde{\mathcal{A}}$ in which functions q take on at most three different values – 0, 1, and an intermediate value. The final part of our proof is to show that the optimal solution lies in this set.

Theorem 6.5. *For any single agent setting $\mathcal{I} = (F, B)$, there is an optimal mechanism with allocation rule in the set $\mathcal{A}^*$.*

Proof. Recall that the optimal solution q^* lies in the set $\widetilde{\mathcal{A}}$. Suppose for contradiction that this function takes on two different intermediate values, $q^*(v_1) = y$ and $q^*(v_2) = z$, between 0 and 1 with $y < z$. Then, since $\hat{g}$ is non-increasing and q^* is non-decreasing, we must have $\hat{g}(v_1) > \hat{g}(v_2)$. Now we can improve our objective function value by increasing q^* between v_2 and the value at which it becomes 1, and decreasing q^* between the value at which it becomes strictly positive and v_1, while maintaining the budget constraint. This contradicts the optimality of q^*. □

Deferred proofs and the ironing procedure. Here we present the proofs and ironing procedure omitted from earlier in this section. We first give the proofs of two claims regarding the revenue-optimal allocation rule q^* for this setting.

Proof of Claim 6.2. Suppose that $q^*_{\max} < 1$, and consider setting $\hat{q}(v) = \frac{q^*(v)}{q^*_{\max}}$. Then we have that $\hat{q}_{\max} = 1$ and

$$\int (\hat{q}_{\max} - \hat{q}(v))dv = \frac{1}{q^*_{\max}} \int (q^*_{\max} - q^*(v))dv \leqslant B.$$

Furthermore q^* must yield a nonnegative objective value (since $q(v) = 0$ is also valid solution). Thus, we get that

$$\int \hat{q}(v)\phi(v)f(v)dv \geqslant \int q^*(v)\phi(v)f(v)dv,$$

and so $\hat{q}$ only improves upon q^*. □

Proof of Claim 6.3. Suppose that $q^*(v^*) < 1$, and consider setting $\hat{q}(v) = q^*(v)$ for $v < v^*$ and $\hat{q}(v) = 1$ otherwise. Then we have

$$\int (1 - \hat{q}(v))dv \leqslant \int (1 - q^*(v))dv.$$

Furthermore,

$$\begin{aligned}\int \hat{q}(v)\phi(v)f(v)dv &= \int_{v\leqslant v^*} q^*(v)\phi(v)f(v)dv + \int_{v\geqslant v^*} \phi(v)f(v)dv \\ &\geqslant \int_{v\leqslant v^*} q^*(v)\phi(v)f(v)dv + \int_{v\geqslant v^*} q^*(v)\phi(v)f(v)dv \\ &= \int q^*(v)\phi(v)f(v)dv,\end{aligned}$$

since $\phi(v) \geqslant 0$ for $v \geqslant v^*$. This contradicts the optimality of q^*. □

We now proceed to give the ironing procedure previously mentioned in greater detail. Recall that our goal is to find an optimal solution to

$$\min_{q\in\mathcal{A}} \int_0^{v^*} q(v)g(v)dv, \text{ where}$$

$$\mathcal{A} = \left\{ \begin{array}{c} \text{increasing } q : [0, v^*] \to [0, 1] \\ \text{such that } \int_0^{v^*} q(v)dv = B' \end{array} \right\}.$$

The solution to the above is a simple step function if g is non-increasing; however, if g is not non-increasing, we need to "iron" the function g to produce a non-increasing function $\hat{g}$. Let $G(v) = \int_0^v g(t)dt$; $\hat{G}(v)$ be the convex upper envelope of G; and $\hat{g} = \frac{d}{dv}\hat{G}(v)$. We will find it useful to focus on a subset of allocation functions that are compatible with the ironed $\hat{g}$ in that they are constant on ironed regions. Given any v such that $G(v) \neq \hat{G}(v)$, define

$$\begin{aligned}\underline{v} &= \sup\{v' \leqslant v : G(v') = \hat{G}(v')\}, \text{ and} \\ \bar{v} &= \inf\{v' \geqslant v : G(v') = \hat{G}(v')\}.\end{aligned}$$

Then $[\underline{v}, \bar{v}]$ is the ironed region containing v. Now, given any $q \in \mathcal{A}$, we define the modified allocation function $\tilde{q}$ by

$$\tilde{q}(v) = \begin{cases} q(v) & \text{if } G(v) = \hat{G}(v)\text{; and} \\ \frac{1}{\bar{v}-\underline{v}} \int_{\underline{v}}^{\bar{v}} q(t)dt & \text{if } G(v) \neq \hat{G}(v), \end{cases}$$

Note that $\tilde{q} \in \widetilde{\mathcal{A}}$.

We break the proof of Lemma 6.4 into the chain of comparisons

$$\int_0^{v^*} q(v)g(v)dv \geqslant \int_0^{v^*} q(v)\hat{g}(v)dv = \int_0^{v^*} \tilde{q}(v)\hat{g}(v)dv = \int_0^{v^*} \tilde{q}(v)g(v)dv.$$

We formalize each of these relations in a lemma; first, however, we state a fact that will be useful in their proofs.

Fact 6.6. *For any fixed* $q \in \mathcal{A}$*, and for any interval* (a, b) *such that* $G(v) = \hat{G}(v)$ *for all* $v \in (a, b)$*, we have that*

$$\int_a^b q(v)g(v)dv = \int_a^b q(v)\hat{g}(v)dv.$$

We now proceed with proving the lemmas.

Lemma 6.7. *For any* $q \in \mathcal{A}$*, we have that* $\int_0^{v^*} q(v)g(v)dv \geqslant \int_0^{v^*} q(v)\hat{g}(v)dv$*.*

Proof. By Fact 6.6, we know we need only focus on regions where $\hat{g}$ is ironed. So let v be such that $G(v) \neq \hat{G}(v)$, and consider the interval $[\underline{v}, \bar{v}]$.

By our definition of $\hat{G}(v)$, we know that for any $v' \in [\underline{v}, \bar{v}]$ we have

$$\int_{\underline{v}}^{v'} g(t)dt \leqslant \int_{\underline{v}}^{v'} \hat{g}(t)dt,$$

with equality if and only if $v' = \underline{v}$ or $v' = \bar{v}$. This implies the following. Let

$$\gamma = \int_{\underline{v}}^{\bar{v}} g(t)dt = \int_{\underline{v}}^{\bar{v}} \hat{g}(t)dt;$$

then both $G(\cdot)/\gamma$ and $\hat{G}(\cdot)/\gamma$ are distributions on $(\underline{v}, \bar{v})$, and the former stochastically dominates the latter. As such, since $q(\cdot)$ is monotone increasing, we must have that the expectations of $q(\cdot)$ under these two distributions are related as

$$\int_{\underline{v}}^{\bar{v}} q(t)\frac{g(t)}{\gamma}dt \geqslant \int_{\underline{v}}^{\bar{v}} q(t)\frac{\hat{g}(t)}{\gamma}dt \iff \int_{\underline{v}}^{\bar{v}} q(t)g(t)dt \geqslant \int_{\underline{v}}^{\bar{v}} q(t)\hat{g}(t)dt.$$

Since the claimed relation holds on ironed intervals as well, we may conclude that it holds overall. $\square$

Lemma 6.8. *For any* $q \in \mathcal{A}$*, we have that* $\int_0^{v^*} q(v)\hat{g}(v)dv = \int_0^{v^*} \tilde{q}(v)\hat{g}(v)dv$*.*

Proof. First, note that by our definition of $\tilde{q}$, we know that $\tilde{q}(v) = q(v)$ whenever $G(v) = \hat{G}(v)$. This means we once again need only consider ironed regions. Let v be such that $G(v) \neq \hat{G}(v)$, and consider the interval $[\underline{v}, \bar{v}]$. Now, we know that on $(\underline{v}, \bar{v})$, $\hat{g}(\cdot)$ takes on the constant value $\hat{g}(v)$. So we have that

$$\int_{\underline{v}}^{\bar{v}} q(t)\hat{g}(t)dt = \hat{g}(v)\int_{\underline{v}}^{\bar{v}} q(t)dt = \hat{g}(v)(\bar{v}-\underline{v})\tilde{q}(t) = \int_{\underline{v}}^{\bar{v}} \tilde{q}(t)\hat{g}(t)dt,$$

and the claim follows. □

Lemma 6.9. *For any* $\tilde{q} \in \widetilde{\mathcal{A}}$*, we have that* $\int_0^{v^*} \tilde{q}(v)\hat{g}(v)dv = \int_0^{v^*} \tilde{q}(v)g(v)dv$.

Proof. Once again, Fact 6.6 implies we know we need only focus on regions where $\hat{g}$ is ironed. So let v be such that $G(v) \neq \hat{G}(v)$, and consider the interval $[\underline{v}, \bar{v}]$. Then we know that $\tilde{q}(\cdot)$ is constant on the interval $(\underline{v}, \bar{v})$. Thus,

$$\int_{\underline{v}}^{\bar{v}} \tilde{q}(t)\hat{g}(t)dt = \tilde{q}(v)\int_{\underline{v}}^{\bar{v}} \hat{g}(t)dt = \int_{\underline{v}}^{\bar{v}} \tilde{q}(t)g(t)dt,$$

and the result follows. □

6.1.2 Single parameter setting with public budgets

We now consider single parameter settings with multiple agents. Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance of single-parameter budget-constrained revenue maximization. Define the truncated distributions $\widehat{F}_i$ as follows.

$$\widehat{F}_i(v) = \begin{cases} F_i(v) & \text{if } v < B_i\text{; and} \\ 1 & \text{if } v \geqslant B_i. \end{cases} \tag{6.1}$$

Let $\widehat{\mathcal{I}} = (\widehat{\mathbf{F}}, \mathcal{S})$ be the modified setting where we replace $\mathbf{F}$ with $\widehat{\mathbf{F}}$ — note that for each i, the support of $\widehat{F}_i$ ends at or before B_i, and so we may remove the budgets since they place no constraint on the instance $\widehat{\mathcal{I}}$.

A mechanism for $\widehat{\mathcal{I}}$ naturally extends to $\mathcal{I}$, while satisfying budget feasibility and obtaining the same revenue. Our general technique will be to relate the revenue of a mechanism for $\mathcal{I}$ to that of a mechanism for $\widehat{\mathcal{I}}$. In general, the latter can be quite small, and so we introduce the following quantity to bound this loss. Define the set

$\mathcal{B}$ as

$$\mathcal{B} = \operatorname*{argmax}_{S \in \mathcal{S}} \left\{ \sum_{i \in S} B_i \middle| \, \forall i \in S, v_i \geqslant B_i \right\}. \tag{6.2}$$

Our basic approach is to design a BIC mechanism $\widehat{M}$ for the setting $\widehat{\mathcal{I}}$ based on the original mechanism M such that we have

$$\mathcal{R}^M \leqslant \mathcal{R}^{\widehat{M}} + \mathbf{E}\left[\sum_{i \in \mathcal{B}} B_i\right]. \tag{6.3}$$

Then, the first term on the right is bounded above by the revenue of the optimal mechanism for $\widehat{\mathcal{I}}$. We further demonstrate in each case that we can bound the expectation $\mathbf{E}\left[\sum_{i \in \mathcal{B}} B_i\right]$ by another mechanism for $\widehat{\mathcal{I}}$.

We define the mechanism $\widehat{M}$ in terms of its expected allocation and payment. Let $\mathbf{x}(\mathbf{v})$ and $\mathbf{p}(\mathbf{v})$ be the expected allocations and expected payments for M, respectively. Define the expected allocation and expected payment rules for $\widehat{M}$ as follows. For each agent i in the setting $\widehat{\mathcal{I}}$ with valuation $\hat{v}_i$, draw a corresponding v_i consistent with $\hat{v}_i = \min(v_i, B_i)$; in this case that simply means $v_i = \hat{v}_i$ if $\hat{v}_i < B_i$, and $v_i \sim F_i(v \mid v \geqslant B_i)$ otherwise. Then $\widehat{M}$'s expected allocation and payment are given by

$$\hat{x}_i(\hat{\mathbf{v}}) = x_i(\mathbf{v}_{-i}, \hat{v}_i); \text{ and } \qquad \hat{p}_i(\hat{\mathbf{v}}) = p_i(\mathbf{v}_{-i}, \hat{v}_i),$$

respectively.

Lemma 6.10. *$\widehat{M}$ is a feasible BIC mechanism for $\widehat{\mathcal{I}}$.*

Proof. We first note that from the point of view of a single agent i, the expected allocation and price function of $\widehat{M}$ behave as though other agents' values are the same as before. Therefore, the expected allocation is still an increasing function of value and the payments satisfy BIC. We will now argue that the expected allocation function can be implemented in a way that the resulting outcome is a randomization over feasible outcomes. To do so, we first compute $x_i(\mathbf{v})$, as well as $\hat{x}_i(\hat{\mathbf{v}})$ for all i. Starting with the allocation returned by $\mathbf{x}(\mathbf{v})$, for every agent i in this allocation, with probability $\hat{x}_i(\hat{\mathbf{v}})/x_i(\mathbf{v})$, we serve this agent, and with the remaining probability we remove her from the allocated set. Since $\mathcal{S}$ is a downward closed feasibility constraint, feasibility is maintained, and we achieve the target allocation probabilities. We remark here that our goal is to merely exhibit that $\widehat{M}$ is feasible and not to actually compute it. □

We now prove the bound (6.3) on $\mathcal{R}^M$.

Lemma 6.11. *Given any mechanism* M *for* $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$*, where* $\mathcal{S}$ *is downward-closed, if we define the mechanism* $\widehat{M}$ *for* $\widehat{\mathcal{I}}$ *as above, then* (6.3) *holds.*

Proof. In order to prove the statement, we couple the values $\mathbf{v}$ that $\widehat{M}$ draws for each $\hat{\mathbf{v}}$ with the $\mathbf{v}$ in the other expectations. So fix some corresponding pair of value vectors $\mathbf{v}$ and $\hat{\mathbf{v}}$; consider the contribution of each agent i to the revenue of M. Split the agents into two sets L and H, defined by

$$L = \{i | v_i \leqslant B_i\}; \text{ and } \qquad H = \{i | v_i \geqslant B_i\}.$$

Recall that for all $i \in L$, we have that $v_i = \hat{v}_i$, and so $\hat{p}_i(\hat{\mathbf{v}}) = p_i(\mathbf{v}_{-i}, \hat{v}) = p_i(\mathbf{v})$. Furthermore, since M faces the downward-closed feasibility constraint $\mathcal{S}$, any subset of H that M serves is one of the sets $\mathcal{B}$ maximizes over. Since M can never charge any agent more than their budget, we can see that

$$\mathcal{R}^M(\mathbf{v}) = \sum_{i \in L} \mathcal{R}_i^M(\mathbf{v}) + \sum_{i \in H} \mathcal{R}_i^M(\mathbf{v}) \leqslant \sum_{i \in L} \mathcal{R}_i^{\widehat{M}}(\hat{\mathbf{v}}) + \sum_{i \in \mathcal{B}} B_i \leqslant \mathcal{R}^{\widehat{M}}(\hat{\mathbf{v}}) + \sum_{i \in \mathcal{B}} B_i.$$

Taking expectations on both sides (according to the previously mentioned coupling) proves the claim. □

Note that $\mathcal{R}^{\widehat{M}}$ can be easily achieved by simply running the (unconstrained) revenue-optimal mechanism over $\widehat{\mathcal{I}}$. It remains to be shown that we can, in fact, upper bound $E[\sum_{i \in \mathcal{B}} B_i]$ also by the revenue of the same (unconstrained) revenue-optimal mechanism over $\widehat{\mathcal{I}}$.

Lemma 6.12. *There exists a mechanism* $M_{\mathcal{B}}$ *for the setting* $\widehat{\mathcal{I}}$ *such that* $E_{\mathbf{v} \sim \mathbf{F}}\left[\sum_{i \in \mathcal{B}} B_i\right] \leqslant \mathcal{R}^{M_{\mathcal{B}}}$.

Proof. We define the mechanism $M_{\mathcal{B}}$ as implementing the allocation rule $\mathcal{B}$. Note that membership of i in $\mathcal{B}$ is monotone in v_i, and that the truthful payment for $i \in \mathcal{B}$ is precisely B_i, since this is the minimum value required for allocation. Thus, we can immediately see that

$$\mathcal{R}^{M_{\mathcal{B}}} = \mathop{\mathrm{E}}_{\hat{\mathbf{v}} \sim \widehat{\mathbf{F}}}\left[\sum_{i \in \mathcal{B}} B_i\right] = \mathop{\mathrm{E}}_{\mathbf{v} \sim \mathbf{F}}\left[\sum_{i \in \mathcal{B}} B_i\right],$$

as desired. □

By combining the results of Lemmas 6.11 and 6.12, we get the following theorem.

Theorem 6.13. *Given a single parameter setting $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, the optimal mechanism $\mathcal{M}$ for the modified setting $\widehat{\mathcal{I}} = (\widehat{\mathbf{F}}, \mathcal{S})$ gives a 2-approximation to the optimal revenue for $\mathcal{I}$.*

6.1.3 Multi-parameter setting with public budgets

We next consider settings where a seller offers multiple kinds of service and agents have different preferences over them. Agents are unit-demand and want any one of the services; the seller faces a general downward closed feasibility constraint. As before, we use the tuple $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ to denote an instance of this problem; throughout, i indexes agents and j indexes services. Let $\mathcal{S}$ be a downward-closed feasibility constraint over (i, j) pairs, and furthermore assume each agent i has a budget B_i.

Ideally, we would like to follow the same approach as in the previous section. We use the same basic benchmark, defining $\widehat{\mathbf{F}}$ and $\mathcal{B}$ analogously to (6.1) and (6.2) for the instance $\mathcal{I}$. Note that here, $\mathcal{B}$ is a collection of (i, j) pairs; since agents are unit-demand, however, we can also think of $\mathcal{B}$ as a set of agents. We can't apply the same reduction from M to $\widehat{M}$ directly, however, because truncating each of a multi-parameter agent's values to their budget affects the agents' preferences across different items, a concern we did not have before.

Instead, we make use of the reduction from multi-parameter Bayesian MD to single-parameter Bayesian MD we presented in Chapter 4. We now describe how to modify the reduction to accomodate budgets. Specifically, starting with a budget feasible mechanism M for $\mathcal{I}$, we first follow the approach of Chapter 4 and convert it into a mechanism M' for $\mathcal{I}^{\text{reps}}$ with the same allocation rule as for M. Unfortunately, M' is not necessarily budget feasible, since the allocation rule of M may induce larger payments in the single-parameter setting $\mathcal{I}^{\text{reps}}$. We therefore modify the mechanism so that any representative in $\mathcal{I}^{\text{reps}}$ offered service at a price larger than his budget is dropped from the allocated set. This makes M' budget feasible. We then apply the approach of the previous section to construct a mechanism $\widehat{M}$ for the instance $\widehat{\mathcal{I}}^{\text{reps}}$ based on the modified M'.

Finally, we note that any sequential posted price mechanism for the setting $\widehat{\mathcal{I}}^{\text{reps}}$ is necessarily budget-feasible for the setting $\mathcal{I}$ because in $\widehat{\mathcal{I}}^{\text{reps}}$ values do not exceed budgets. We thus can apply the approximately-optimal posted-pricing we obtain via the techniques of Chapters 3 and 4 unchanged to the original setting $\mathcal{I}$. We call this final mechanism $\mathcal{S}$.

Lemma 6.14. *Consider a multi-parameter setting $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, and let α denote the approximation to revenue that an appropriately chosen $\mathcal{S}$ achieves in this setting. Then for any deterministic budget feasible mechanism M, the mechanism $\mathcal{S}$ defined above satisfies*

$$\mathcal{R}^{\mathcal{S}} \geqslant 1/\alpha \mathop{\mathbf{E}}_{\mathbf{v}}\left[\sum_{i \in L} \mathcal{R}_i^M(\mathbf{v})\right]$$

and consequently,

$$\mathcal{R}^M \leqslant \alpha \mathcal{R}^{\mathcal{S}} + \mathop{\mathbf{E}}_{\mathbf{v} \sim \mathbf{F}}\left[\sum_{i \in \mathcal{B}} B_i\right].$$

Proof. We first note that the revenue that M' derives from the agents in L is no smaller than the revenue that M derives from these agents because payments for these agents never exceed their budgets. Following the analysis of Lemma 6.11, we further conclude that the revenue of $\widehat{M}$ is an upper bound on the contribution of agents in L to $\mathcal{R}^M$. Finally, the statement of the lemma implies that the revenue of the mechanism $\mathcal{S}$ is at least as large as a $1/\alpha$ fraction of $\mathcal{R}^{\widehat{M}}$. This proves the first part of the lemma. The second statement follows along the lines of the proof of Lemma 6.11. □

Finally, we show how to approximate the benchmark $\mathbf{E}[\sum_{i \in \mathcal{B}} B_i]$ in the multi-parameter setting.

Lemma 6.15. *If $\mathcal{S}$ is a matroid set system, there exists a mechanism $M_{\mathcal{B}}$ for the setting $\widehat{\mathcal{I}}$ such that $\mathbf{E}_{\mathbf{v} \sim \mathbf{F}}\left[\sum_{i \in \mathcal{B}} B_i\right] \leqslant 2\mathcal{R}^{M_{\mathcal{B}}}$.*

Proof. We begin by noting that the unit-demand constraint on agents is precisely a partition matroid; furthermore, taking a subset of a matroid induces a matroid.

Hence for any fixed $\mathbf{v}$, our objective $\sum_{i \in \mathcal{B}} B_i$ is precisely a maximum weighted set in the intersection of two matroids. Note that each (i, j) that is a valid element of $\mathcal{B}$ corresponds to an agent i who would be willing to pay a price of B_i for service j.

Our proposed mechanism $M_{\mathcal{B}}$ sequentially approaches each agent in order by decreasing B_i, and offers them all services that are still feasible under $\mathcal{S}$ (based on previous decisions by agents), at a price of B_i. Then our revenue is the weight of a greedily selected independent set in the matroid intersection $\mathcal{B}$ is optimal over; by Theorems 1.1 and 3.2 of (Korte and Hausmann, 1978), this set's weight is at least 1/2 that of $\mathcal{B}$. Note that as a sequential posted pricing, the mechanism is immediately truthful; taking expectations over $\mathbf{v}$ yields the claimed revenue bound. □

Combining Lemmas 6.14 and 6.15 immediately gives us the following theorem.

Theorem 6.16. *Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance with multi-parameter, unit-demand agents and $\mathcal{S}$ being a matroid or simpler feasibility constraint. Then, there exists a polynomial time computable mechanism for $\widehat{\mathcal{I}}$ that is budget feasible and DSIC for $\mathcal{I}$ and obtains a constant fraction of the revenue of the optimal budget-feasible mechanism for $\mathcal{I}$.*

6.1.4 Private budgets

We next consider settings where budgets are part of agents' private types, but where the mechanism designer knows the distributions from which budgets are drawn. We assume that values and budgets are drawn from independent distributions. We focus on settings where agents' values are single-dimensional.

Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ denote an instance of this setting. We follow a similar approach as for public budgets. Switching from public to private budgets, however, adds new complexity; in particular it becomes tricky to achieve our benchmark $E[\sum_{i \in \mathcal{B}} B_i]$ in general. In this section, we present an approximately optimal mechanism for the case when each distribution in $\mathbf{F}$ satisfies the MHR condition (see Definition 2.2 in Section 2.3). As with the public budgets case, we begin by considering the single agent case, and then show how to extend our results to more general multi-agent cases.

Given a pair of value and budget vectors, we consider the "extractable value" of an agent i to be $\min(v_i, B_i)$; we modify our definition of $\mathcal{B}$ to reflect this:

$$\mathcal{B} = \operatorname*{argmax}_{S \in \mathcal{S}} \sum_{i \in S} \min(v_i, B_i).$$

Similarly to the public budgets case, our approach is to split the revenue of an arbitrary mechanism into two terms, which (loosely speaking) look like revenue in a truncated value setting, and the sum of the budgets in $\mathcal{B}$; we then demonstrate a lottery menu mechanism whose revenue upper bounds both of these terms.

6.1.5 Single agent settings with private budget

We begin by considering the case of a single agent with a private budget. A key idea will be the "extractable value" of the agent, which we will define to be $\min(v, B)$. The idea behind focusing on this is that value and budget function similarly in terms of how they limit the revenue that can be achieved – both upper-bound it. Recall that, in settings where value are distributed according to an MHR distribution, we know that the revenue and welfare objectives are quite close, in the sense that the we can design mechanisms whose revenue gives a constant approximation to the optimal welfare. We begin by recalling the following two theorems about MHR distributions (both special cases of Theorem 3.11, Dhangwatnotai et al., 2010).

Theorem 6.17. *If distribution F satisfies the MHR condition, then $E_{v\sim F}[v] \leqslant e \cdot r^*(1 - F(r^*))$, where r^* is the monopoly reserve price for F.*

Theorem 6.18. *If distribution F satisfies the MHR condition, then $1 - F(r^*) \geqslant 1/e$, where r^* is the monopoly reserve price for F.*

The first theorem, in fact, immediately captures our claim about revenue and welfare for the case of a single agent, since $r(1 - F(r))$ is precisely the expected revenue from offering a price r to an agent. The second theorem captures the fact that MHR distributions can't have overly long tails.

In the case where agents have budgets, we can no longer hope to obtain expected revenue that approximates the optimal welfare in every case; for example, if the budget happens to be deterministically 0, we cannot achieve *any* revenue no matter how high the agent's value for service is. As we shall see, however, we *can* extend the above theorem to our setting by altering it to compare revenue with extractable value. In particular, by truncating both the value and the monopoly reserve price to the budget, we get the following analogue of Theorem 6.17.

Lemma 6.19. *If distribution* F *satisfies the MHR condition, then for any* $B \geqslant 0$ *we have* $E_{v \sim F}[\min(v, B)] \leqslant e \cdot \min(r^*, B)(1 - F(r^*))$, *where* r^* *is the monopoly reserve price for* F.

Proof. Observing that the expectation of the minimum of two terms can never exceed the minimum of their respective expectations, we can see that

$$\underset{v \sim F}{E}[\min(v, B)] \leqslant \min(\underset{v \sim F}{E}[v], B) \leqslant \min(e \cdot r^*(1 - F(r^*)), B),$$

where the second inequality follows from Theorem 6.17. Recall, however, that by Theorem 6.18 we have that $e(1 - F(r^*)) \leqslant 1$. Thus, we may conclude that

$$\underset{v \sim F}{E}[\min(v, B)] \leqslant e \cdot (1 - F(r^*)) \min(r^*, B),$$

exactly as claimed. □

The above echoes what we saw in the case of public budgets. If the budget is above the monopoly price, then our optimal mechanism remains unchanged; otherwise, we need to use the budget as a guide to forming our offers. As there, lotteries will play a key role in extracting good revenue. Our proposed mechanisms will be lottery pricings $\mathcal{L}(p)$ that are parameterized by a "target price" p. $\mathcal{L}(p)$ contains, for all $0 \leqslant \alpha \leqslant 2/3$, a lottery that with probability $(1/3 + \alpha)$ allows the agent to purchase service at a price of $\alpha p/2$.

Note that the probability of allocation in the above lottery system rises faster as a function of α than the price of the lottery does. This ensures that the agent is willing to buy the most expensive lottery that he can afford. So, in particular, if all lotteries bring positive utility, then the agent spends his entire budget – yielding the maximum revenue that any mechanism can hope to achieve from the agent. This powerful idea is what enables our approximation. We formalize this in the following lemma.

Lemma 6.20. *When an agent with* $v \geqslant p$ *is offered the menu* $\mathcal{L}(p)$, *he purchases an option yielding expected revenue at least* $\min(p, B)/3$.

Proof. Consider the utility of an agent with value v when purchasing the lottery with parameter α, which we denote $u_\alpha(v)$. We can see that

$$u_\alpha(v) = (1/3+\alpha)(v - \alpha p/2), \text{ and so}$$
$$\frac{\partial u_\alpha(v)}{\partial \alpha} = v - (\alpha + 1/6)p \geqslant 0$$

when $v \geqslant p$. So we can see that an agent with value $v \geqslant p$ will purchase the lottery with the highest α value they can afford; since the lottery for $\alpha = 2/3$ assigns service with probability 1 at a price of $p/3$, and every lottery provides service with probability at least one third, we can see that an agent will purchase a lottery yielding revenue at least $\min(p, B)/3$. □

The key intuition behind the above lemma is that it says that we can think of the lottery pricing $\mathcal{L}(p)$ as a "truncated" version of the price p; whenever an agent without budget constraints would accept the price p, a budget-constrained agent will choose an option from $\mathcal{L}(p)$ that gives revenue of $\min(p/3, B)$ when allocation occurs. Note that this revenue guarantee closely matches the upper bound we derived in Lemma 6.19, assuming we set $p = r^*$. We formalize this in the following theorem, which is our main result of the section.

Theorem 6.21. *Consider a single-agent instance $\mathcal{I} = (F, G)$ of the BSMD problem with private budgets, where F satisfies the monotone hazard rate condition. The lottery pricing $\mathcal{L}(r^*)$ provides a $3e$-approximation to the optimal revenue for $\mathcal{I}$.*

Proof. Recall that any EPIR, budget-feasible mechanism for $\mathcal{I}$ can never charge the agent more than $\min(v, B)$. Thus, $E_{v,B}[\min(v, B)]$ upper bounds the optimal revenue. Now, consider fixing some budget B. By combining Lemmas 6.19 and 6.20, we get that

$$\mathop{E}_{v}[\mathcal{R}^{\mathcal{L}(r^*)}(v, B)] \geqslant (1/3)\min(r^*, B)(1 - F(r^*)) \geqslant (1/3e)\mathop{E}_{v}[\min(v, B)].$$

Taking expectations on both sides with respect to B gives exactly the claimed bound. □

6.1.6 Multi-agent settings with private budgets

We focus on settings $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ where $\mathcal{S}$ is a downwards-closed set system, and each distribution in $\mathbf{F}$ satisfies the MHR condition. As we shall see, the approach we used for settings with a single agent will apply here as well with only minor adjustments.

We begin with some definitions. Once again, we observe that no EPIR mechanism can ever charge an agent a price for service that exceeds either the agent's value or their budget. Thus, we rewrite $\mathcal{B}$ in terms of the "extractable value" to reflect this upper bound:

$$\mathcal{B} = \operatorname*{argmax}_{S \in \mathcal{S}} \sum_{i \in S} \min(v_i, B_i).$$

Recall that in the previous section, we showed how to construct a lottery system for a single agent that achieved revenue approximating the truncated welfare. In this section, we extend the approach to multiple agents. Our general approach has two parts. First, we observe that if we run the VCG mechanism for the truncated setting in our original setting, it remains truthful and achieves optimal truncated welfare. Second, we show that we can modify our lottery systems slightly so they take the place of threshold payments in the VCG mechanisms; by doing so, we achieve a mechanism whose revenue approximates the truncated welfare of the VCG mechanism.

We begin by describing how we modify the lottery pricings. Here, our lottery pricings $\mathcal{L}(\underline{p}, \bar{p})$ are parameterized by both a minimum price $\underline{p}$ and a maximum price $\bar{p}$. There are two cases: if $\underline{p} \geqslant \bar{p}/3$, then $\mathcal{L}(\underline{p}, \bar{p})$ contains the single fixed price of $\underline{p}$; otherwise, it contains, for all $2\underline{p}/\bar{p} \leqslant \alpha \leqslant 2/3$, a lottery that with probability $(1/3 + \alpha)$ allows the agent to purchase service at a price of $\alpha\bar{p}/2$. These lottery systems are a simple modification of those we used before. Starting with one of those lotteries, we just remove every option with (ex post) price strictly below $\underline{p}$; if this removes *every* option in the lottery pricing, we switch to offering a fixed price of $\underline{p}$. The following lemma follows immediately from the above observations.

Lemma 6.22. *An agent receives nonnegative utility from some option in $\mathcal{L}(\underline{p}, \bar{p})$ if and only if* $\min(v, B) \geqslant \underline{p}$.

Furthermore, we get the following extension of Lemma 6.20

Lemma 6.23. *Let* $p = \max(\underline{p}, \bar{p})$*. When an agent with* $v \geqslant p$ *is offered the menu* $\mathcal{L}(\underline{p}, \bar{p})$*, the agent purchases an option yielding expected revenue at least* $\min(p, B)/3$*.*

We begin by noting the following more general form of Theorem 6.17 (which is a restatement of Theorem 3.11, Dhangwatnotai et al., 2010).

Theorem 6.24. *Given any* $t \geqslant 0$*, if* F *is an MHR distribution, we have that* $\mathrm{E}[v|v \geqslant t] \Pr[v \geqslant t] \leqslant e \cdot p(1 - F(p))$*, where* $p = \max(t, r^*)$*.*

As in the previous section, the above has a natural extension to settings with budgets

Lemma 6.25. *For any* $t \geqslant 0$ *and any* $B \geqslant t$*, if* F *is an MHR distribution, we have that*

$$\mathop{\mathrm{E}}_{v}[\min(v, B)|v \geqslant t] \Pr[v \geqslant t] \leqslant e \cdot \min(p, B)(1 - F(p)),$$

where $p = \max(t, r^*)$*.*

Proof. Since moving an expectation inside of the $\min(\cdot)$ operator can only cause an increase in an expression's value, we can see that

$$\begin{aligned}\mathop{\mathrm{E}}_{v}[\min(v, B)|v \geqslant t] \Pr[v \geqslant t] &\leqslant \min(\mathop{\mathrm{E}}_{v}[v|v \geqslant t] \Pr[v \geqslant t], B \Pr[v \geqslant t]) \\ &\leqslant \min(e \cdot \min(p, B)(1 - F(p)), B(1 - F(t)),\end{aligned}$$

by applying Theorem 6.24. Now, either $p = t$ or $1 - F(p) = 1 - F(r^*) \geqslant 1/e$; in either case, however, we have that $1 - F(t) \leqslant e(1 - F(p))$ and so the claim follows. □

The mechanism $M^{\mathcal{L}}$ we propose serves (a subset of) the set $\mathcal{B}$. For each i, let T_i be the threshold corresponding to inclusion in $\mathcal{B}$, i.e. $T_i = \min\{v' : i \in \mathcal{B} \text{ for } ((\mathbf{v}_{-i}, v'), (\mathbf{B}_{-i}, v'))\}$. Our mechanism offers the lottery system $\mathcal{L}(T_i, {\phi_i}^{-1}(0))$ to agent i, where $\phi_i^{-1}(0)$ is the monopoly price for i. We get the following theorem.

Theorem 6.26. *For any setting* $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ *where* $\mathcal{S}$ *is a downwards-closed set constraint and each distribution in* $\mathbf{F}$ *satisfies MHR, the mechanism* $M^{\mathcal{L}}$ *is a* $3e$*-approximation to the optimal revenue.*

Proof. We shall prove a stronger claim than that of the theorem statement, namely that $M^{\mathcal{L}}$ gives a $3e$-approximation to our benchmark of $\mathrm{E}_{\mathbf{v},\mathbf{B}}[\sum_{i \in \mathcal{B}} \min(v_i, B_i)]$.

Consider some agent i, and fix the values $\mathbf{v}_{-i}$ and budgets $\mathbf{B}_{-i}$ of all other agents. Note that this fixes the threshold T_i as well. Now, for any $B_i \geqslant T_i$ we may combine Lemma 6.25 and 6.23 to get that

$$\begin{aligned}\mathop{E}_{v_i}[\min(v_i, B_i)|v_i \geqslant T_i]\Pr[v_i \geqslant T_i] &\leqslant e\min(p, B_i)(1 - F_i(p)) \\ &\leqslant 3e\mathop{E}_{v_i}\left[\mathcal{R}^{\mathcal{L}(T_i, {\phi_i}^{-1}(0))}(v_i, B_i)\right],\end{aligned}$$

where $p = \max(T_i, {\phi_i}^{-1}(0))$. Noting that our lottery systems never make positive transfers to the agents, we may take expectations with respect to T_i and B_i and conclude that

$$\mathop{E}_{v_i, B_i, T_i}[\min(v_i, B_i)|v_i, B_i \geqslant T_i]\Pr[v_i, B_i \geqslant T_i] \leqslant 3e\mathop{E}_{v_i, B_i, T_i}\left[\mathcal{R}^{\mathcal{L}(T_i, {\phi_i}^{-1}(0))}(v_i, B_i)\right].$$

Recall, however, that $i \in \mathcal{B}$ if and only if $v_i, B_i \geqslant T_i$ by Lemma6.22, and that our mechanism $M^{\mathcal{L}}$ offers each agent i the lottery system $\mathcal{L}(T_i, {\phi_i}^{-1}(0))$. Thus, if we sum the above inequality over i, we get

$$\mathop{E}_{\mathbf{v},\mathbf{B}}[\sum_{i \in \mathcal{B}} \min(v_i, B_i)] \leqslant 3e\mathcal{R}^{M^{\mathcal{L}}},$$

exactly as desired. □

6.2 Maximizing welfare

In this section we focus on the welfare objective. In particular, the seller's goal is to maximize the total value of the allocation in expectation. Once again we assume that budgets are known publicly.

We first note that we cannot use the approach of the previous section as a roadmap. Even with public budgets, truncating values to the corresponding budgets does not work for the social welfare objective. In particular, the following example shows it is possible for a budget feasible mechanism to distinguish between values above the budget without exceeding the budget in payments.

Example 6.27. *Consider an n agent single-item auction, where agents have i.i.d. values for the item. Each agent has a budget of 1. Each agent's value is 1 with probability $1 - 1/n$ and*

n with probability $1/n$. Then a mechanism that simply truncates values to budgets cannot distinguish between the agents and gets a social welfare of at most 2. On the other hand, consider a mechanism that orders agents in an arbitrary order and offers two options to each agent in turn while the item is unallocated: getting the item for free with probability $1/n$ and nothing otherwise, or purchasing the item at a price of 1. Then, an agent picks the first option if and only if her value is below $n/(n-1)$, and otherwise picks the second option. In particular, an agent with value n always picks the second option, and an agent with value 1 always picks the first option. For large n, with probability approaching $1-1/e$ at least one agent has value n, and with probability at least $1/e$ the item is unsold before the first agent with value n is made an offer. The mechanism's expected welfare is therefore at least $1/e(1-1/e)n = \Omega(n)$.

Note that the precise choice of budgets in the above example was critical: if budgets were any lower, the proposed mechanism would have been infeasible; and if they were any higher, truncation would have still allowed for distinguishing between agents with low and high values. This suggests considering bicriteria approximations where we compete against an optimal mechanism that faces smaller budgets. We first demonstrate a mechanism achieving an approximation of this sort; we then show that our mechanism also gives a good approximation if we relax the EPIR constraint to an IIR constraint, instead of relaxing budgets. Of course, our ultimate goal is to provide a good approximation for the social welfare objective via an EPIR budget feasible mechanism. While we are unable to do so in general, the final section presents a constant factor approximation for settings where the distributions F_i for every agent i satisfy the MHR condition (Definition 2.2 in Section 2.3).

6.2.1 A bicriteria approximation

Consider a setting $\mathcal{I}$ with budgets $\mathbf{B}$. Let OPT$'$ denote the EPIR mechanism that is welfare-maximizing and feasible for budgets $(1-\epsilon)\mathbf{B}$ (i.e. where each budget is scaled down by a factor of $1-\epsilon$). We claim that we can approximate the welfare of this mechanism while maintaining budget feasibility with respect to the original budgets $\mathbf{B}$.

Theorem 6.28. *For a given instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, let $\mathcal{I}'$ be the instance $(\mathbf{F}, \mathcal{S}, (1-\epsilon)\mathbf{B})$ where each agent's budget is scaled down by a factor of $1-\epsilon$. Let OPT$'$ denote the welfare optimal budget feasible mechanism for $\mathcal{I}'$. Then, there exists an easy to compute ex post IR* *mechanism (namely, the VCG mechanism over a modified instance) that is budget feasible for $\mathcal{I}$ and obtains at least an ϵ fraction of the social welfare of OPT$'$.*

Proof. We first use OPT$'$ to construct a new mechanism M. M proceeds as follows. It elicits values from agents. For all agents i with $v_i \geqslant B_i$, it resamples agent i's value from the distribution F_i restricted to the set $[B_i, \infty)$. Other values are left unchanged. It then runs the mechanism OPT$'$ on the resampled values. It is easy to see that M is budget feasible valid for $\mathcal{I}'$.

We claim that the social welfare of M is at least ϵ times the social welfare of OPT$'$. To prove the claim, consider a single agent i, and let ${q_i}^1$ denote the probability of allocation for this agent in OPT$'$ when her value is B_i, and ${q_i}^2$ denote the probability of allocation for this agent in OPT$'$ when her value is $v_i^{\max}$ (the agent's maximum possible value). Note that the expected payment that the agent makes at $v_i^{\max}$ is at least $({q_i}^2 - {q_i}^1)B_i$ plus the payment she makes at B_i. Then, EPIR implies that $({q_i}^2 - {q_i}^1)B_i$ is at most the budget $(1-\epsilon)B_i$ times ${q_i}^2$. This implies ${q_i}^2 < {q_i}^1/\epsilon$. Now, noting that the value distributions for agents are unaltered by resampling, it holds that for $v_i \geqslant B_i$, the probability of allocation for agent i at v_i under M is equal to the expected probability of allocation for the agent under OPT$'$ conditioned on the agent's value being in the range $[B_i, \infty)$. Since the probability of allocation under OPT$'$ for this range is always between ${q_i}^1$ and ${q_i}^2$, the expected probability of allocation is at least ${q_i}^1 \geqslant \epsilon {q_i}^2$. So compared to those under OPT', the probabilities of allocation under M are at most a factor of ϵ smaller. Therefore, the expected social welfare of M is also at most a factor of ϵ smaller than that of OPT$'$.

Our goal will then be to approximate the social welfare of M. Note that for every agent i, M treats values above B_i identically. We can therefore consider the following optimization problem: for an instance $\mathcal{I}$, construct a DSIC EPIR mechanism that maximizes social welfare subject to the additional constraint that for every agent i the mechanism's (distribution over) allocation should be identical across value vectors that differ only in agent i's value and where agent i's value is $\geqslant B_i$. For any such mechanism, agent i's expected contribution to social welfare from value vectors with $v_i \geqslant B_i$ conditioned on being allocated is $\bar{v}_i$, where $\bar{v}_i = E[v_i | v_i \geqslant B_i]$. Therefore, the following mechanism maximizes welfare over the above class of mechanisms: for every agent i with $v_i \geqslant B_i$, replace v_i by $\bar{v}_i$; other values remain unmodified; run the VCG mechanism over the modified value vector; charge every agent the minimum of the payment returned by the VCG mechanism and their budget. It is easy to verify that this mechanism is DSIC, ex post IR, budget feasible for the original budgets B_i, and obtains expected social welfare at least that of M. Therefore, it satisfies the claim in the theorem. □

6.2.2 An interim IR mechanism

Next we note that it is in fact easy to remove the approximation on budget in the above theorem if we are willing to give up on EPIR. In particular, consider an optimal mechanism OPT on the instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$. Then the above theorem implies the existence of a mechanism V that is budget feasible for $\mathcal{I}' = (\mathbf{F}, \mathcal{S}, 2\mathbf{B})$ and obtains half the welfare of OPT (taking $\epsilon = 1/2$). Now consider the mechanism V' described as follows. V' simulates V on the given value vector. Then for every agent i it charges i half the payment charged by V and with probability 1/2 makes an allocation to i if V makes an allocation to i. Agent i's expected utility from any strategy under V' is exactly half her expected utility from the same strategy under V. Therefore, V' is DSIC. Moreover, it is budget feasible for the original budgets $\mathbf{B}$ since it always charges half the payments in V. Its expected social welfare is exactly half that of V. We therefore get the following theorem.

Theorem 6.29. *For a given instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, let OPT denote the welfare optimal EPIR budget feasible mechanism for $\mathcal{I}$. Then, there exists an easy to compute IIR mechanism that is budget feasible for $\mathcal{I}$ and obtains at least a quarter of the social welfare of OPT.*

6.2.3 An ex-post IR mechanism for MHR distributions

As previously remarked, our ultimate goal is to provide a good approximation for the social welfare objective via an EPIR budget feasible mechanism. We now show that under an MHR condition on distributions, we can achieve precisely this goal. In particular, we present a constant factor approximation for settings where the distributions F_i for every agent i satisfy the MHR condition (Definition 2.2 in Section 2.3).

Under the MHR condition, we can exhibit a close relationship between the welfare and revenue of any mechanism. Using this relationship along with results from the previous section, we can come up with a budget feasible approximately-revenue-maximizing mechanism that also provides an approximation to social welfare. The MHR condition is quite crucial to our approach. In fact our solution consists of two mechanisms, one of which charges no payments, and the other of which truncates values to their corresponding budgets – approaches that don't work for the example we considered above.

Let $v_i^* = {\phi_i}^{-1}(0)$ denote the monopoly price for the distribution F_i. We then get the following bound on social welfare, which we are able to approximate.

Lemma 6.30. *For any instance $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$, if all the distributions F_i satisfy the MHR condition, then for any non-decreasing allocation function $x(v)$, we have that*

$$\int_{\mathbf{v}} \left(\sum_i v_i x_i(\mathbf{v}) \right) d\mathbf{F}(\mathbf{v}) \leqslant \int_{\mathbf{v}} \left(\sum_i (\phi_i(v_i) + 2v_i^*) x_i(\mathbf{v}) \right) d\mathbf{F}(\mathbf{v}).$$

In order to prove the Lemma, we require some new definitions and claims. Consider a single agent with MHR distribution F, virtual value function ϕ, and monopoly price v^*. Let ϕ^+ and ϕ^- be the positive and negative portions of ϕ respectively; i.e. for all v, $\phi^+(v), \phi^-(v) \geqslant 0$ and $\phi(v) = \phi^+(v) - \phi^-(v)$. We can then claim the following (the first is a restatement of Hartline and Roughgarden 2009, Lemma 3.1).

Lemma 6.31. *For a distribution F satisfying the MHR, all values v satisfy $v \leqslant v^* + \phi^+(v) = v^* + \phi(v) + \phi^-(v)$.*

Lemma 6.32. *For any monotone allocation function $x(\cdot)$,*

$$\int \phi^-(v) x(v) dF(v) \leqslant \int v^* x(v) dF(v).$$

Proof. We begin by recalling that the expected revenue of any BIC mechanism is equal to its expected virtual surplus (see Proposition 2.3 in Section 2.3). Now consider a mechanism for a single agent with value distribution F that always serves the agent. Clearly the revenue of this mechanism is 0. So we get

$$\int_v (\phi^+(v) - \phi^-(v)) dF(v) = 0,$$

which implies that

$$\int_v \phi^+(v) dF(v) = \int_v \phi^-(v) dF(v).$$

Second, the revenue from offering the agent the monopoly price v^* is precisely $v^*(1 - F(v^*))$. Therefore,

$$\int \phi^+(v)dF(v) = \int_{v^*}^{\infty} (\phi^+(v) - \phi^-(v))dF(v) = v^* \int_{v^*}^{\infty} dF(v),$$

where the first equality follows from regularity of F.

Note that the regularity of F implies that ϕ^+ and ϕ^- are identically 0 below and above v^*, respectively. Now, from the above two equalities, we can see that if x is monotone non-decreasing, then

$$\begin{aligned}\int \phi^-(v)x(v)dF(v) &\leqslant \int \phi^-(v)x(v^*)dF(v) \\ &= x(v^*)\int \phi^+(v)dF(v) \\ &= x(v^*)v^* \int_{v^*}^{\infty} dF(v) \\ &\leqslant \int v^* x(v)dF(v);\end{aligned}$$

the claim follows. □

The proof of Lemma 6.30 follows immediately by combining Lemmas 6.31 and 6.32. The Lemma gives us the following approximation.

Theorem 6.33. *Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{B})$ be an instance where all distributions F_i satisfy the MHR condition. Then, one of the following mechanisms obtains a $2(1+e)$-approximation to the social welfare of a welfare-optimal budget-feasible mechanism for $\mathcal{I}$. Both of these mechanisms are DSIC, EPIR and budget feasible.*

- Mechanism 1: *Always allocate to the set S_1^* and charge zero payments, where* $S_1^* = \text{argmax}_{S \in \mathcal{S}} \sum_{i \in S} v_i^*$.
- Mechanism 2: *Elicit values from agents; for all i with $v_i > B_i$, replace v_i by B_i; run Myerson's mechanism on the resulting instance.*

Proof. We begin by noting that an immediate consequence of Lemma 6.18 is that for a distribution F satisfying the MHR, $E_F[v] \geqslant v^*/e$.

Now, consider some budget feasible mechanism M for the instance $\mathcal{I}$. Then by Theorem 6.13, the optimal mechanism for the truncated distributions (6.1) obtains revenue, and therefore also social welfare, no less than a 1/2 fraction of the expected revenue $\int \sum_i \phi_i(v_i)x_i(\mathbf{v})d\mathbf{F}(\mathbf{v})$.

Moreover, we can see that $\int_{\mathbf{v}} \sum_i v_i^* x(\mathbf{v})d\mathbf{F}(\mathbf{v})$ is upper bounded by $\sum_{i \in S^*} v_i^*$, where $S^* = \text{argmax}_{S \in \mathcal{S}} \sum_{i \in S} v_i^*$. Then a mechanism which always allocates to the set S^* and charges no payments is budget feasible, DSIC, and obtains welfare $\sum_{i \in S^*} E_{F_i}[v_i] \geqslant 1/e \sum_{i \in S^*} v_i^*$, where the inequality follows from Lemma 6.18. The original claim follows. □

6.3 Incentive compatibility for budgets

Let $\mathcal{I} = (\mathbf{F}, \mathcal{S}, \mathbf{G})$ be a mechanism design instance with single-parameter agents with private budgets, and let M be a truthful mechanism for this setting. We get the following characterization for the allocation and payment functions of M.

Lemma 6.34. *If M is truthful, then for each i there exists some monotone function $\tilde{x}_i(v)$ such that for each B, $x_i(v, B)$ has the form*

$$x_i(v, B) = \begin{cases} \tilde{x}_i(v) & \text{if } v < v_B \\ \tilde{x}_i(v_B) & \text{if } v \geqslant v_B, \end{cases}$$

where v_B is a monotone non-decreasing function of B, and the payment has the form

$$p_i(v, B) = x_i(v, B)v - \int_0^v x_i(t, B)dt.$$

Proof. For now, assume that $x_i(v, B)$ and $p_i(v, B)$ are continuous. Begin by fixing some arbitrary budget B, and considering the truthfulness constraints on v. Just as in the case where we have no budgets, truthfulness in reporting valuation implies that $x_i(v, B)$ is monotone increasing, and payments are of the form

$$p_i(v, B) = x_i(v, B)v - \int_0^v x_i(t, B)dt;$$

so we need only show the claimed relation between allocation curves for different budgets. Denote the utility i receives by $u_i(v, B)$. Note by the above, we get that

$$\begin{aligned} u_i(v, B) &= x_i(v, B)v - p_i(v, B) \\ &= x_i(v, B)v - \left(x_i(v, B)v - \int_0^v x_i(t, B)dt\right) \\ &= \int_0^v x_i(t, B)dt. \end{aligned}$$

Begin by fixing some v, and considering two different budgets B and B' (without loss of generality $B < B'$). Since an agent with budget B' must be able to afford any payment that an agent with budget B can, we know that $u_i(v, B') \geqslant u_i(v, B)$. Assume that $u_i(v, B') = u_i(v, B)$; then we have that $x_i(v, B') = x_i(v, B)$ as well. Suppose not; without loss of generality, say $x_i(v, B') > x_i(v, B)$. Then by continuity, this holds on some neighborhood $(v - \delta, v + \delta)$ of v. But then we get that

$$\begin{aligned} u_i(v - \delta, B') &= \int_0^{v-\delta} x_i(t, B')dt \\ &= \int_0^{v-\delta} x_i(t, B)dt + \int_{v-\delta}^{v} x_i(t, B) - x_i(t, B')dt \\ &< u_i(v - \delta, B), \end{aligned}$$

a contradiction. A symmetric argument gives us that if $x_i(v, B') > x_i(v, B)$, there is some δ such that $u_i(v + \delta, B') < u_i(v + \delta, B)$.

On the other hand, if $u_i(v, B') > u_i(v, B)$, then we may conclude that $p_i(v, B') > B$ (since otherwise an agent with budget B would have incentive to lie and report B').

We can combine the above two observations to get the claimed characterization as follows. Consider the revenue curve for the maximum possible budget $\widehat{B}$, and for any other budget B. Then we know that at each v, either they are identical, or the payment for budget $\widehat{B}$ exceeds that for budget B. By continuity, the last point where they were identical must have had a corresponding price of B. After that point, the allocation curve for B must remain constant, since any increase would imply an increase in price.

While the above assumed that $x_i(v, B)$ was continuous, note that the function is monotone and bounded. As such, it is continuous a.e., and so our characterization holds a.e. for general $x_i(v, B)$. □

We now give an example showing that when budgets are private the $E[\sum_{i \in \mathcal{B}} B_i]$ benchmark cannot be approximated merely by mechanisms for $\widehat{\mathcal{I}}$.

Consider trying to serve a single agent with a fixed value for service of $v = 3n+1$, and a budget B distributed according to the distribution $G(B) = 1-1/B$ for $B \in [1, n)$ and $G(n) = 1$. If we consider the truncated value distribution, since we have a single agent offering a fixed price is optimal; but any fixed price p gives revenue of $p(1 - F(p)) = 1$. On the other hand, consider offering the agent the menu

$$\{(1/4 + \alpha, 2\alpha^2 n) : \alpha \in [0, 3/4]\}.$$

Note that differentiating the agent's utility as a function of α gives us

$$\frac{d}{d\alpha}\left((1/4 + \alpha)(3n + 1) - 2\alpha^2 n\right) = 3n + 1 - 4\alpha n \geqslant 1$$

for $\alpha \in [0, 3/4]$, and so the agent always buys the most expensive lottery he or she can. Prices on the menu range from 0 to $9n/8$, and each offer is made with probability at least $1/4$; so this menu extracts revenue of at least $E[B]/4 = \Theta(\log n)$.

REFERENCES

Abrams, Zoë. 2006. Revenue maximization when bidders have budgets. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithms*, 1074–1082. SODA '06, New York, NY, USA: ACM.

Aigner, M. 1997. *Combinatorial theory*. Classics in mathematics, Springer-Verlag.

Alaei, S., K. Jain, and A. Malekian. 2010. Competitive equilibrium for unit-demand buyers with non quasi-linear utilities. In *CORR*.

Alaei, Saeed. 2011. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, 512–521. Washington, DC, USA: IEEE Computer Society.

Ausubel, Lawrence M., and Paul Milgrom. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions*, ed. P. Cramton, Y. Shoham, and R. Steinberg, chap. 1, 17–40. MIT Press, Cambridge, MA.

Babaioff, M., N. Immorlica, and R. Kleinberg. 2007. Matroids, secretary problems, and online mechanisms. In *Proc. 19th ACM Symp. on Discrete Algorithms*.

Babaioff, Moshe, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar. 2009. Secretary problems: weights and discounts. In *SODA '09: Proceedings of the Nineteenth Annual ACM -SIAM Symposium on Discrete Algorithms*, 1245–1254. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.

Balcan, Maria-Florina, and Avrim Blum. 2006. Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 29–35. New York, NY, USA: ACM.

Balcan, Maria-Florina, Avrim Blum, and Yishay Mansour. 2008. Item pricing for revenue maximization. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, 50–59. New York, NY, USA: ACM.

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

Blum, A., and J. Hartline. 2005. Near-optimal online auctions. In *Proc. 16th ACM Symp. on Discrete Algorithms*. ACM/SIAM.

Blum, A., V. Kumar, A. Rudra, and F. Wu. 2004. Online learning in online auctions. *Theoretical Computer Science* 324:137–146.

Blumrosen, L., and T. Holenstein. 2008. Posted prices vs. negotiations: an asymptotic analysis. In *Proc. 10th ACM Conf. on Electronic Commerce*.

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

Briest, P. 2006. Towards hardness of envy-free pricing. Tech. Rep. TR06-150, Electronic Colloquium on Computational Complexity (ECCC).

Briest, Patrick, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. 2010. Pricing randomized allocations. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, 585–597. Philadelphia, PA, USA: SIAM.

Brualdi, Richard A. 1969. Comments on bases in dependence structures. *Bulletin of the Australian Mathematical Society* 1(02):161–167.

Brusco, Sandro, and Giuseppe Lopomo. 2009. Simultaneous ascending auctions with complementarities and known budget constraints. *Economic Theory* 38(1): 105–124.

Bulow, J., and J. Roberts. 1989. The simple economics of optimal auctions. *The Journal of Political Economy* 97:1060–90.

Cai, Yang, and Constantinos Daskalakis. 2011. Extreme-value theorems for optimal multidimensional pricing. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, 522–531. Washington, DC, USA: IEEE Computer Society.

Cai, Yang, Constantinos Daskalakis, and S. Matthew Weinberg. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, 459–478.

Chawla, S., J. Hartline, and R. Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, 243–251.

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

Chawla, Shuchi, David Malec, and Balasubramanian Sivan. 2012. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*.

Chawla, Shuchi, David L Malec, and Azarakhsh Malekian. 2011. Bayesian mechanism design for budget-constrained agents. In *Proceedings of the 12th ACM conference on Electronic commerce*, 253–262. ACM.

Che, Yeon-Koo, and Ian Gale. 1998. Standard auctions with financially constrained bidders. *Review of Economic Studies* 65(1):1–21.

———. 2000. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory* 92(2):198–233.

Chen, Ning, Xiaotie Deng, and Arpita Ghosh. 2010. Competitive equilibria in matching markets with budgets. *CoRR*.

Daskalakis, Constantinos, and S. Matthew Weinberg. 2012. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 370–387.

Devanur, Nikhil, Jason Hartline, Anna Karlin, and Thach Nguyen. 2011. Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*, ed. Ning Chen, Edith Elkind, and Elias Koutsoupias, vol. 7090 of *Lecture Notes in Computer Science*, 122–133. Springer Berlin / Heidelberg.

Dhangwatnotai, Peerapong, Tim Roughgarden, and Qiqi Yan. 2010. Revenue maximization with a single sample. In *Proceedings of the 11th ACM conference on Electronic commerce*, 129–138. ACM.

Dobzinski, Shahar, and Shaddin Dughmi. 2009. On the power of randomization in algorithmic mechanism design. In *Proceedings of the 50th Annual IEEE Symposium* *on Foundations of Computer Science*, 505–514. Washington, DC, USA: IEEE Computer Society.

Dobzinski, Shahar, Noam Nisan, and Ron Lavi. 2008. Multi-unit auctions with budget limits. In *Proc. of the 49th Annual Symposium on Foundations of Computer Science (FOCS)*.

Dughmi, Shaddin, and Tim Roughgarden. 2010. Black-box randomized reductions in algorithmic mechanism design. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, 775–784. Washington, DC, USA: IEEE Computer Society.

Goldberg, A., and J. Hartline. 2005. Collusion-resistant mechanisms for single-parameter agents. In *Proc. 16th ACM Symp. on Discrete Algorithms*.

Guruswami, Venkatesan, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. 2005. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1164–1173. Philadelphia, PA, USA: SIAM.

Hartline, J., and T. Roughgarden. 2009. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*.

Hartline, Jason, and Anna Karlin. 2007. Profit maximization in mechanism design. In *Algorithmic Game Theory*, ed. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, chap. 13, 331–362. Cambridge Press.

Holahan, Catherine. 2008. Auctions on eBay: A Dying Breed. http://www.businessweek.com/technology/content/jun2008/tc2008062_112762.htm.

Kleinberg, R., and T. Leighton. 2003. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *Proc. 44th IEEE Symp. on Foundations of Computer Science*.

Kleinberg, Robert, and S. Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, 123–136.

Korte, Bernhard, and Dirk Hausmann. 1978. An analysis of the greedy heuristic for independence systems. In *Algorithmic Aspects of Combinatorics*, ed. P. Hell B. Alspach and D.J. Miller, vol. 2 of *Annals of Discrete Mathematics*, 65 – 74. Elsevier.

Korula, Nitish, and Martin Pál. 2009. Algorithms for secretary problems on graphs and hypergraphs. In *ICALP '09: Proceedings of the 36th Internatilonal Collogquium on Automata, Languages and Programming*, 508–520. Berlin, Heidelberg: Springer-Verlag.

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

Lehmann, D., L. I. O'Callaghan, and Y. Shoham. 1999. Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*, 96–102. ACM Press.

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

———. 2007. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory* 137(1):153–185.

Maskin, Eric S. 2000. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review* 44(4-6):667–681.

McAfee, R.P., and J. McMillan. 1988. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory* 46(2):335–354.

Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6: 58–73.

Nisan, N., and A. Ronen. 1999. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, 129–140. ACM Press.

Nisan, Noam. 2007. Introduction to mechanism design. In *Algorithmic Game Theory*, ed. Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, chap. 9, 209–242. Cambridge Press.

Nisan, Noam, Jason Bayer, Deepak Chandra, Tal Franji, Robert Gardner, Yossi Matias, Neil Rhodes, Misha Seltzer, Danny Tom, Hal Varian, and Dan Zigmond. 2009.

Google's auction for tv ads. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming: Part II*, 309–327. ICALP '09, Berlin, Heidelberg: Springer-Verlag.

Pai, Mallesh M., and Rakesh Vohra. 2008. Optimal auctions with financially constrained bidders. Discussion Papers, Northwestern University.

Pavlov, Gregory. 2006. Optimal mechanism for selling substitutes. Boston University - Department of Economics - Working Papers Series WP2006-014, Boston University - Department of Economics.

Riley, John, and Richard Zeckhauser. 1983. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics* 98(2):267–289.

Rochet, Jean-Charles, and Philippe Chone. 1998. Ironing, Sweeping, and Multidimensional Screening. *Econometrica* 66(4):783–826.

Roughgarden, Tim, Inbal Talgam-Cohen, and Qiqi Yan. 2012. Supply-limiting mechanisms. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 844–861. New York, NY, USA: ACM.

Samuel-Cahn, Ester. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability* 12(4): 1213–1216.

Sandholm, Tuomas, and Andrew Gilpin. 2006. Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 1127–1134. New York, NY, USA: ACM.

Thanassoulis, John E. 2004. Haggling Over Substitutes. *Journal of Economic Theory* 117(2):217–245.

Wilson, Robert B. 1997. *Nonlinear pricing*. Oxford University Press.

Yan, Qiqi. 2011. Mechanism design via correlation gap. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, 710–719. Philadelphia, PA, USA: SIAM.