

NORTHWESTERN UNIVERSITY

Optimal Multi-parameter Auction Design

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Computer Science

By

Nima Haghpanah

EVANSTON, ILLINOIS

August 2014

© Copyright by Nima Haghpanah 2014

All Rights Reserved

ABSTRACT

Optimal Multi-parameter Auction Design

Nima Haghpanah

This thesis studies the design of Bayesian revenue-optimal auctions for a class of problems in which buyers have general (non-linear and multi-parameter) preferences. This class includes the classical *linear single-parameter* problem considered by Myerson (1981), for which he provided a simple characterization of the optimal mechanism, leading to numerous applications in theory and practice. However, for fully general preferences no generic and practical solution is known (various negative computational or structural results exist for special cases), even for the problem of designing a mechanism for a single agent.

With general preferences, the optimal mechanism can be complex and impractical. This thesis identifies key conditions implying that the optimal mechanism is practical. Our main results are different in that they identify different conditions implying different notions of practicality, but are all similar in adopting a *modular view* to the problem that separates the task of designing a solution for the single-agent problem as the main module, from the task of combining these modules to form an optimal multi-agent mechanism. First, for multi-parameter linear settings, we specify a large class of distributions over values that implies that the optimal single-agent mechanism is posted pricing, and the optimal multi-agent mechanism maximizes *virtual values* for players' favorite items. When agents are identical, the mechanism becomes the second price

auction with reserve for favorite items. Second, and more generally, we specify a condition called *revenue-linearity*, defined beyond multi-parameter linear settings, that implies that optimizing agents' *marginal revenue* maximizes revenue. Finally, adopting efficient computability as the notion of practicality, we show that for any setting in which single-agent solutions are efficiently computable, multi-agent solutions are also computable.

Acknowledgements

First of all, I would like to thank my advisor Jason Hartline for his guidance and encouragement. I truly admire his passion for research, and appreciate his selfless devotion to the success of his students. Jason is always available to give me advice and share his deep and unique intuition on research. He taught me the importance of clear communication in research, be it in technical and general writing, discussing work with colleagues, or giving talks. My experience of being his student for five years has been, indeed, fantastic.

I am grateful to my other committee members, Eddie Dekel, Robert Kleinberg, and David Parkes for their valuable feedback before, during, and after my defense. Their feedback greatly improved the material of this thesis, and provided me with a more clear vision about my research. I am very fortunate have such a strong group of researchers with diverse backgrounds as my committee.

In addition to being on my committee, Bobby and David generously hosted me for extended periods of times when I visited Cornell and Harvard Universities in the Spring semesters of 2013 and 2014, and for that I am thankful to them. Their groups at Cornell and Harvard are two of the best in computer science, each with their own unique strengths and focuses, and I was lucky to be a part of their group meetings and discussions.

My great coauthors made the process of doing research much more fun for me, and for that I thank all of them. In particular, I thank Saeed Alaei, Hu Fu, Jason, and Azarakhsh Malekian, with whom I coauthored some of the papers that are the basis of this thesis.

I would also like to thank my fellow grad students at the Theory and Economics group at Northwestern University, for being both great colleagues and fun friends. Together we held student-organized meetings to read papers, practiced talks, participated in many sports event, and played board games. Having a group of great colleagues was definitely an important part of my student life.

I would like to specially thank Nicole Immorlica for her support, encouragement, and friendship. She helped make this journey much more delightful by sharing research and work stories, and when I needed a break from work, a cat picture or two.

Last but not least, I am forever grateful to my parents Maryam and Ali. I dedicate this thesis to them, even though I can never thank them enough for their unconditional love and limitless support.

Dedication

to *maman-o-baba*

Table of Contents

ABSTRACT	3
Acknowledgements	5
Dedication	7
List of Figures	11
Chapter 1. Introduction	13
1.1. Optimal Auction Design: Background	13
1.2. Our Contributions	17
1.3. Practicality and Optimality	21
1.4. Multi-Parameter Models	23
1.5. Organization	24
Chapter 2. Models, Background, and Warmup	26
2.1. Models and the Method of Virtual Values	27
2.2. The Marginal Revenue Mechanism	47
2.3. Interim Feasibility and Computation	59
Chapter 3. The Single Agent Problem and Incentives	63
3.1. Introduction	64
3.2. Preliminaries	74
3.3. Amortization of Revenue	77

3.4. Optimality of Uniform Pricing for Unit Demand Preferences	81
3.5. Optimality of Bundle Pricing for Additive Preferences	91
3.6. Multi Agent Extensions	98
Chapter 4. Single-Agent to Multi-Agent Solutions: The Structure	103
4.1. Introduction	104
4.2. Multi-dimensional and Nonlinear Preferences	113
4.3. Implementation	125
4.4. Approximation	131
4.5. Single Dimensional Extension Theorems	143
Chapter 5. Single-Agent to Multi-Agent Solutions: Computation	150
5.1. Introduction	151
5.2. Preliminaries	158
5.3. The Single-agent Problem	162
5.4. Multi- to Single-agent Reductions	166
5.5. Optimization and Implementation of Interim Allocation Rules	168
5.6. Conclusions and Extensions	187
Bibliography	189
Appendix A. The Single Agent Problem and Incentives	196
A.1. Extension for $m > 2$ Items	196
A.2. Missing Proofs	198
Appendix B. Single-Agent to Multi-Agent Solutions: The Structure	206
B.1. Proofs from Subsection 4.3.2	206
B.2. Proofs from Subsection 4.3.3	209

	10
B.3. Proofs for unit-demand approximation	213
B.4. Revenue Linearity for Unit Demand Valuations Uniform on Hypercubes	218
B.5. Reverse Auctions	221
Appendix C. Single-Agent to Multi-Agent Solutions: Computation	224
C.1. Proofs from Section 5.5.1	224
C.2. Proofs from Section 5.5.2	229
C.3. Proofs from Section 5.5.3	231

List of Figures

2.1	Monotonicity of the allocation function and truthfulness of the mechanism.	30
2.2	The allocation of the posted pricing mechanism optimizes expected virtual value pointwise, and therefore in expectation.	33
2.3	The ironing process.	40
2.4	The ironed revenue curve.	50
2.5	The connection between the two constrained lottery pricing problems.	52
2.6	Decomposing an allocation function using the basis of step functions.	53
2.7	The quantile shift operator.	56
3.1	Distributions satisfying the conditions of the chapter's main theorems.	70
3.2	The trapezoidal set defined by the maximum-value and ratio-of-values parameters.	84
3.3	The trapezoidal set parameterized by the sum-of-values and ratio-of-values parameters.	97
B.1	The quantile mapping for the two dimensional uniform distribution.	221
C.1	The flow network representing the stochastic transition probabilities in the stochastic sequential allocation mechanism.	224

- C.2 An application of the REROUTE routine to the flow graph of the stochastic sequential allocation mechanism. 227
- C.3 The bipartite graph used in the max-flow/min-cut argument of the proof of Theorem 21. 229

CHAPTER 1

Introduction**1.1. Optimal Auction Design: Background**

Finding the best way to distribute scarce resources and goods has long been a central question studied by researchers in computer science, economics, and operations research. How should bandwidth be allocated to Internet users? Who should be hired in a job market when there is competition between candidates and available spots? What company's advertisement should be displayed together with the content when a user visits a webpage? The rise of the Internet has reformed the way scientists study these problems. The input to an economic system (the cost of routing a packet on the Internet; the qualifications and the preferences of the job candidate; the amount an advertiser is willing to pay to receive the ad slot) is now provided by millions of users with possibly conflicting interests who will use the system in order to improve their own well-being. As a result, we must now design systems that work well when their inputs are provided by self-interested users. This brings forward a category of interesting yet demanding problems in which economic design and engineering challenges converge.

Auction theory provides a framework to study such problems, using the guiding principle that "people respond to incentives." Self-interested agents will choose actions, e.g. how much to bid, depending on their preferences and the actions of other players in the auction. Equilibrium notions are used as models to predict the outcome of the auction when rational agents choose their actions to maximize their own well-being. The goal of the auction designer is to set the rules of the mechanism such that performance, typically welfare or revenue, is maximized in the outcome prescribed by the in-equilibrium behavior.

Optimal auction design is a branch of auction theory which aims to specify the auction that maximizes the revenue of the seller of goods or services. Auctions for display advertising and sponsored search, and eBay and Priceline auctions are applications of this theory. A key feature of the optimal auction design theory is *Bayesian representation of imperfect information*. That is, instead of assuming that the seller knows *everything* about the problem, including anything that is known to any of the players, we only assume that the seller has a *belief* about the unknown parameters. The belief is represented by a probability distribution, and the seller wishes to maximize the *expected revenue* when the unknown parameters are drawn at random. These unknown parameters typically include the preferences of the buyers over *outcomes* (the way goods, services, and money are assigned), and affect the actions of the buyers in the mechanism.

As a starting example, consider the design of *display advertising* auctions as a simple optimal auction design scenario. The seller is a content generating website, which wishes to display some advertisement, say a single one for simplicity, next to its content to a user viewing its webpage, in exchange for money. The buyers are advertisers wishing to *purchase the ad slot*, that is, to show their ad next to the content. Each advertiser has a utility over outcomes (whether or not it receives the ad slot and the payment), in a *linear single-parameter* form: the *value* for the ad slot minus the payment (hence the term “linear”). The value of each advertiser is unknown to the seller (in fact, that is the only unknown, hence the term “single parameter”), but is believed to come from a distribution, independently from the value of other advertisers (possibly from different distributions). The seller wishes to maximize its revenue in expectation.

In his seminal paper, Myerson (1981) solved for the optimal auction in the above setting. For a single buyer, the optimal mechanism is to simply post a price for the ad slot (among the class of all mechanisms, which could involve for example randomization and arbitrary message passing or negotiation protocols). For *identical* buyers, that is, when the distributions of agents are the same, and under a natural *regularity* assumption on the distribution, the optimal auction

is a second price auction with a reserve price. For non-identical buyers, the optimal auction is to give the item to the agent who has the maximum positive *virtual value*. Myerson described the virtual value as a transformation of the buyers value based on its distribution that encapsulates how much the buyer is *valued by the seller* for the purpose of revenue. In particular, Myerson showed that the revenue of the seller is exactly equal to the virtual value of the winner, and therefore revenue is maximized by selecting the buyer with the highest virtual value.

Myerson's analysis via virtual values, and the simple characterization of the optimal solution has since found numerous applications from theory to practice. For example, Bulow and Klemperer (1996) use Myerson's result to show that attracting one extra buyer is more valuable than any amount of bargaining power for the seller. Ostrovsky and Schwarz (2011) confirm via a field experiment the significance of using the reserve prices suggested by theory. Hartline and Roughgarden (2009) use Myerson's analysis to show that second price auctions with reserve remain approximately optimal even with non-identical agents. However, Myerson's setting does not accommodate features natural to many real world problems, when the utility of the players has more complicated form than linear, or an agent has more unknown parameters than a single value.

As an example of a more realistic scenario, let us revisit the ad auctions problem. The seller often has some knowledge about the identity of the user in the form of *cookies*, that it might be willing to transfer together with the ad slot to the advertisers. Each advertiser can customize its ad if it has some (potentially partial) knowledge about the user, so it might be willing to pay more to purchase cookies. Each advertiser has a preference over outcomes (whether or not it receives the ad slot, the information transferred from the seller about the user, and payment), which could depend on its budget, preference over risk, and the value it assigns to cookies. The problem is no longer linear (budget constraints and/or attitude towards risk), nor single parameter (budget, risk, as well as values for different "configurations" of the

item: what cookies are transferred with the ad slot). Yet, the revenue maximization problem is well-defined by adopting the Bayesian model, which assumes the parameters are drawn from a distribution, and sets to maximize the expected revenue.

With *general preferences*, that is, beyond the linear single-parameter setting, even the single-agent problem (i.e., where there is no concern about the joint feasibility of outcomes for multiple agents) is either ill understood or is known to not admit well behaved solutions. These negative results and our lack of understanding of the single-agent problem have mostly prohibited studying the general multi-agent problem, resulting in a vacuum in our understanding of optimal auctions. In this thesis, in addition to studying the single-agent problem, we adopt a modular approach to the theory of mechanism design that separates the challenging task of solving the single-agent problem from the task of combining solutions to form a multi-agent solution, thus enabling us to make progress on the multi-agent front independently of the challenges of the single-agent problem.

There is a common drawback in attempting to solve for more and more general models of optimal auction design: naturally, as the model becomes more general, the optimal mechanisms could lose their nice properties, making them less applicable to the real world problems. Instead, one could attempt to *identify* key conditions of the model that enable strong predictions of the solution. The more natural those key conditions are, the more useful that theory would be. In fact, Myerson's result can be seen as identifying conditions such as linearity and risk neutrality enabling such strong predictions. In this thesis, equipped with the modular approach, we set to extend this theory:

Identify natural conditions common in real world problems implying practical properties of the solution, in order to extend design insights beyond what is provided by classical auction design theory.

1.2. Our Contributions

In this work, we formalize and study an abstract auction design problem which includes, but is more general than, the display advertising scenario discussed above. We adopt three approaches in Chapter 3 to Chapter 5. A key perspective common among these approaches is viewing the auction design task as one of understanding a simpler problem, designing optimal mechanisms for each agent, and then putting the pieces back together carefully to maintain feasibility while preserving optimality.

A modular analysis of the problem separates the challenges of understanding incentives (encapsulated inside the single-agent problem) from feasibility (combination of single-agent solutions to form a multi-agent solution).

The difference between the approaches is two-fold. The first difference is in focus. The first approach in Chapter 3 investigates the single-agent module, delving deep into the mathematics of incentive compatible design. The second and third approaches in Chapter 4 and Chapter 5 focus on the structural and computational task of combining modules together to form a solution by adopting *oracle models* that describe or compute the optimal multi-agent solution assuming existence of an oracle that solves the single-agent problem. Together, these approaches provide a comprehensive analysis of the problem ranging from micro-level studying of the incentives, to the more high-level task of understanding the feasibility of the auction.

The second difference between the approaches is in how restrictive the assumptions and the conclusions of the results are. The first result demands the most from the problem, but provides the strongest conclusions (a closed form for the optimal auction). The second result assumes that the solution to a single-agent problem is given and satisfies some conditions (weaker than the conditions in the first part), and provides structural insights into how the single-agent solutions can be combined to form multi-agent solutions, using a basic principle in microeconomics,

marginal revenue maximization, as a guideline.¹ The third result assumes that a computationally efficient solution to a more general single-agent problem is given as a black box, requires no properties of the problem, and provides a computationally efficient algorithm to solve the multi-agent problem. In what follows we discuss each of the three results in more detail.

Optimality of Single-dimensional Projections (Chapter 3). The starting point of this thesis is a close examination of the main module of the problem, the single agent problem. This examination allows us to closely understand the incentives, which in turn enables us to characterize the optimal single- and multi-agent auctions in closed forms. This is basically the approach undertaken by Myerson, and our results completely include Myerson’s results as a special case. Consider the problem of selling two items, say the ad slot with cookie and the ad slot without cookie, to a single *unit demand* agent (i.e., the value for a bundle of items is equal to the agent’s value for her favorite item in the bundle, here the ad slot with cookie). Assuming that the agent always favors owning the cookie to not owning it, and motivated by Myerson’s result, it is natural to form a conjecture: it is optimal to simply post a price for the ad slot together with cookie. However, Thanassoulis (2004) showed that this is not true: the optimal auction randomizes the allocation, sometimes allocating the less favored ad slot without cookie.²

We show in Theorem 4 that under a condition called *max-ratio affiliation* (more permissive than Myerson’s setting), the optimal auction for a single unit-demand agent simply posts the same price (a uniform price) for each of the items.³ In the ad auctions example, this would state that a uniform price is posted for the ad slot with cookie and the ad slot without cookie; the agent, of course, picks the ad slot with cookie, if at all. For multiple agents, we show in

¹This principle, for example, states that the optimal allocation of an extra unit of a good to two markets, each with a current level of consumption, is to allocate the good to the market with the highest marginal return.

²Thanassoulis’ example was symmetric: any of the two items could be favored. But it can be easily adopted to our example in which one of the items is always favored.

³Roughly speaking, max-ratio affiliation states that the following two parameters should be positively correlated with each other (they can be independent): the value for the favorite item, and the ratio of values of least favorite to most favorite item. The general condition does not require that one of the items is always favorite.

Theorem 10 that the optimal solution is to only run a single auction to pick the “winner,” in which we compare the virtual values for favorite items of agents. The winner then gets its favorite item (in the ad auctions example, the optimal auction only allocates and compares the virtual values of the ad slot together with cookie). We prove this result by designing a general framework to extend Myerson’s construction of virtual values to general preferences. We demonstrate the power of the framework by applying it, in addition to the unit demand problem above, to a multi-item setting with additive preferences, providing sufficient conditions for optimality of the auction that simply posts a price for the bundle of items.

Marginal Revenue Maximization (Chapter 4). As hinted above, numerous negative results show that even the single-agent problem in general does not admit solutions that have simple form (in structure by Vincent and Manelli 2007 and Thanassoulis 2004, and in *menu complexity* by Hart and Nisan 2013). In spite of that, it is still possible to argue about the structure of optimal auctions by adopting the modular view that separates the design of modules (single-agent mechanisms) from the rules that govern the combination of modules,

What are the rules governing the composition of (nearly optimal) single-agent solutions into (nearly optimal) multi-agent solutions?

Marginal revenue is a central concept in economics and optimization. A firm providing a unit of a divisible good to two markets, each described by a revenue function mapping quantity to profit, maximizes its profit by dividing the good to equate marginal revenues across the two markets. Myerson (1981) and Bulow and Roberts (1989) show that the same principle holds for optimal auctions in the linear single-parameter setting, by showing that each player can be thought of as a “market” represented by a revenue curve, and that the optimal auction allocates the item to the player with maximum marginal revenue. We identify *revenue linearity* as a property of single-agent settings (stating that a generalized revenue function, as a function of parameters restricting allocation of the agent, is linear) in which marginal revenue maximization

is revenue optimal (Proposition 4.2). Revenue linear settings include the linear single-parameter settings as well as the max-ratio affiliated settings considered in Chapter 3. By showing that optimality of marginal revenue maximization extends to revenue-linear settings, we can automatically apply many of the results from the theory of the linear single-parameter settings. In addition, we define the notion of *approximate* revenue-linearity, which implies that the multi-agent is approximately optimal, and use the result to show that many practical mechanisms are approximately optimal.

By adopting the modular view to the problem, we can bypass the complexity of the design of the single-agent problem, and still gain valuable design insights independently of our ability to analytically solve for the single-agent problem. For example, a practitioner might be able to find reasonable solutions to a single-agent problem using numerical methods or simulations. The above result motivates adopting the marginal revenue principle as a general guideline to form a multi-agent solution.

Computational Tractability (Chapter 5). Motivated by the first two results, it is natural to form a conjecture: we can always compute the multi-agent problem efficiently by first computing a solution to the single-agent problem, and then putting the pieces back together. However, various negative results state that even the single-agent problem is computationally hard to solve (Daskalakis et al., 2014). This directly implies that the multi-agent problem, which is only more general, is also computationally hard. Still, it is useful to know *how much harder* the multi-agent problem is compared to the single-agent problem. The concept of a *reduction* from computational complexity allows us to formalize this question. Given a computationally efficient solution to the single-agent design problem in the form of a black box, can we solve the multi-agent problem efficiently? We show in Proposition 14 that the answer is generally positive.

*Given oracles that efficiently compute solutions to (a generalization of) the single-agent problem, multi-agent optimal solutions can be efficiently computed in natural feasibility settings.*⁴

Both the results in Chapter 4 and Chapter 5 reduce the multi-agent problem to the single-agent problem (though they use different definitions of the single-agent problem). The reduction in Chapter 5 is the computational reduction prevalent in computer science, whereas the reduction in Chapter 4 is a novel *structural reduction*. However, the purpose of both reductions is the same: understand the multi-agent problem by abstracting away the complexity of the single-agent problem.

1.3. Practicality and Optimality

As discussed above, the goal of this thesis is to design mechanisms selling goods or services to maximize revenue, in the equilibrium induced by the behavior of strategic agents. On a first impression, this task seems extremely challenging: the space of all mechanisms could be huge, with mechanisms using complicated message passing protocols to determine the outcome; on top of that, one must be able to solve for equilibrium as a prediction of the outcome. Fortunately, the powerful yet simple *revelation principle* significantly simplifies this task. The revelation principle (Gibbard (1973); Myerson (1986)) states that the outcome of any mechanism can be achieved by a *direct revelation truthful mechanism*, that is, a mechanism where the space of messages equal the space of private knowledge of agents, and in which truthfully reporting the private knowledge is an equilibrium.

To maximize an objective such as revenue, we can use the revelation principle and formulate the problem as one of optimizing over truthful direct mechanisms, which is a more well structured class. However, even though this approach will let us formulate (and often solve for)

⁴More formally, in matroid feasibility settings, which significantly generalizes the single- and multi-ad slot auction discussed above.

the mechanism that maximizes revenue, the theoretically optimal solution might not necessarily be practical. It could be too complicated to be understood by or even communicated to the agents, or it could use undesirable features such as randomization. For the theory to be useful in practice, we should seek mechanisms that optimize our objective while satisfying some notion of simplicity or practicality. Myerson’s result, for example, is extremely valuable for not only “solving for” the optimal solution, but also identifying conditions (the linear single-parameter setting) implying the solution is simple: posted pricing for a single agent, second price with reserve for multiple identical agents, and virtual value optimization in general.⁵ To quote the report by the Nobel prize committee for the 2007 award going to Hurwicz, Maskin, and Myerson for their work on mechanism design theory,

“Optimization over the set of all direct mechanisms for a given allocation problem is a well-defined mathematical task, and once an optimal direct mechanism has been found, the researcher can ‘translate back’ that mechanism to a more realistic mechanism.”

The focus of this thesis is to identify the conditions of the preference domain that allow for generalization of the linear single-parameter theory while preserving the nice properties of the solution. Based on what is considered to be a “nice property,” we provide different answers:

- (1) Max-ratio affiliation of single agent problems implies that the optimal auction has a simple closed form: item pricing for a single agent, the second price auction for multiple identical agents, and virtual values maximization of favorite items in general. This naturally extends the form of the optimal auctions from the linear single-parameter

⁵In fact, using Myerson’s *revenue-equivalence* theorem, for identical agents the Dutch and English auctions, as well as sealed bid first and second price auctions with reserve can all be shown to be optimal.

theory.⁶ Here, the desired property is simplicity of the closed form, and the sufficient condition is max-ratio affiliation of the single-agent problem.

- (2) (Approximate) Revenue linearity of single agent problems, which is more general than max-ratio affiliation, implies that the (approximately) optimal multi-agent auctions follow the marginal revenue principle. We can then plug in many results using the theory of the linear single-parameter settings to extend the (approximate) optimality of many practical mechanisms such as sequential posted mechanisms, and mechanisms requiring limited access to the distribution, to revenue-linear settings. Here, the desired property is simplicity of rules governing the composition of multi-agent solutions, and the sufficient condition is revenue linearity of the single-agent problem.
- (3) Efficient computability of single agent problems implies that the optimal multi-agent auction is also efficiently computable in natural feasibility settings. Here, borrowing the notion of computational tractability from computer science literature as a minimal yet mathematically rigorous requirement for practicality, the desired property is the efficient computability of the multi-agent problem, and the sufficient condition is efficient computability of the single-agent problem.

1.4. Multi-Parameter Models

This thesis expands the reach of optimal auction design theory by studying the extension of linear single-parameter settings to a large class of multi-parameter settings, termed the *service-constrained* environments. Exemplified by the ad auctions scenario in the introduction, the important property of this class of environments is that the feasibility of the joint outcome of agents (e.g., allocation of ads slots to advertisers) is determined by the identities of agents receiving the item or service (e.g., the number of agents receiving ad-slots, regardless of the

⁶Again, a form of the revenue equivalence theorem can be used to show that Dutch, English, and first price auctions are also optimal for identical agents.

cookies transferred, or the budget and risk degrees of the advertisers). This model is different from multi-item auctions, which is perhaps what one would initially imagine as a generalization of Myerson’s single-item setting. The important observation here is that in generalizations of Myerson’s theory, and in numerous applications of optimal auction design, dimensions of preferences can be independent from the dimensions of feasibility.

In this thesis we choose to focus on generalized domain of preferences, whilst restricting the dimensions of feasibility. This allows us to focus on the challenge of general preferences by separating it from the challenge of general feasibility. We can then identify the properties of the preference domain that allow for generalization of the theory while preserving the nice properties of the solution.

1.5. Organization

As outlined in Section 1.2, Chapter 3 to Chapter 5 study the theory of optimal auctions in service-constrained environments, in progressively more general settings. The chapters can be read independently, and require no background: each chapter defines its required notation in its preliminary section (building upon the previous notation consistently), with Chapter 5 containing the service-constrained design problem in full generality. However, I strongly recommend reading Chapter 2 for background and warmup before reading those chapters. Chapter 2 defines the linear single-parameter and service-constrained models (Subsection 2.1.1 and Subsection 2.1.2), applies the main techniques and ideas of this thesis to the linear single-parameter model, and informally discusses the challenges of applying those techniques to the service-constrained model. In fact, a reader only wishing to obtain a quick understanding of the main tools and techniques in this thesis can do so by only reading Chapter 2.

This main chapters of this thesis, Chapter 3 to Chapter 5, are based on the following three papers, respectively:

- Nima Haghpanah and Jason Hartline. Reverse mechanism design. In submission, 2014.
- Saeed Alaei and Hu Fu and Nima Haghpanah and Jason Hartline. The economics of approximately optimal auctions. In *Foundations of Computer Science*, 2013.
- Saeed Alaei and Hu Fu and Nima Haghpanah and Jason Hartline and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*, 2012.

CHAPTER 2

Models, Background, and Warmup

Overview and Organization. This chapter provides a background review of the linear single-parameter setting, and also showcases our techniques by applying them to that setting. The chapter is divided into three sections. The following three sections, sections 2.1 to 2.3, each apply the techniques in one of the following chapters of the thesis, chapters 3 to 5, to the linear single-parameter settings for the case of selling a single item:

- (1) Section 2.1 defines the linear single-parameter as well as the service-constrained model, and discusses the method of virtual values. Section 2.1 explains Myerson’s original derivation to the linear single-parameter using the terminology that we will use later in Chapter 3 to extend the method to service-constrained setting. In Subsection 2.1.3 briefly discusses the extension of the method of virtual values to service-constrained environments. Our method of virtual values in Chapter 3, projected into the linear single-parameter setting, becomes precisely Myerson’s original derivation.
- (2) Section 2.2 proves optimality of the marginal revenue mechanism in the linear single-parameter model. This derivation is different from Bulow and Roberts (1989), and is intended to highlight the role of revenue-linearity, a condition we identify to generalize the linear single-parameter settings. Subsection 2.2.1 briefly discusses generalizing the marginal revenue mechanism to service-constrained environments.
- (3) Section 2.3 sketches a computationally efficient solution to the problem in the linear single-parameter model. It relies on a characterization of the space of feasible *interim* allocation rules, which are useful in reducing the dimension of the optimization problem.

Subsection 2.3.1 briefly discusses the extension of the techniques to service-constrained environments.

The results of the second two sections, Section 2.2 and Section 2.3, can be directly proved using Myerson's characterization in Section 2.1. However, the purpose of those sections is to derive those results by relying minimally on the properties of the linear single-parameter settings. This allows us to single out the crucial conditions in the derivation of results, which will consequently simplify reading the following chapters which apply the techniques to the more general service-constrained environments.

2.1. Models and the Method of Virtual Values

This section defines the linear single-parameter model in Subsection 2.1.1 as well as the service constrained setting in Subsection 2.1.2 and applies the method of virtual values values, due to Myerson (1981), to the linear single-parameter setting. Our method in Chapter 3 is a complete generalization of Myerson's techniques to a class of service-constrained environments (i.e., when projected back to the linear single-parameter settings, it recovers Myerson's method and results). Therefore, this section serves as a warmup for the content of Chapter 3, in addition to defining the model and the background on the linear single-parameter model. The generalization to service-constrained environments in Chapter 3, however, is significantly more challenging and requires a novel set of tools, briefly discussed in Subsection 2.1.3.

2.1.1. The Linear Single-Parameter Model and Myerson's Solution

The Single-agent Problem. A seller owns an item or service that it wishes to sell to a buyer. The buyer is specified by a set of possible types $T = [0, \bar{v}]$, where each $v \in T$ is a value for the item being sold.¹ The type of the agent is drawn from a known distribution with density f .

¹The analysis easily generalizes to allows for a lower bound of types $\underline{v} > 0$, as well as unbounded set of types $\bar{v} = \infty$

The cumulative distribution function of the type is denoted F . The *allocation* $x \in [0, 1]$ is the probability of receiving the item. The utility of the agent with type v for allocation $x \in [0, 1]$ and payment $p \in \mathbb{R}$ is $v \cdot x - p$, which is the reason we call the problem *linear*. The problem is *single-parameter* since there is only one unknown parameter, v . The problem is *single-agent* since there is only one potential buyer (we extend this to multi-agent problem later).

A single-agent mechanism is a pair of functions, the allocation function $x : T \rightarrow [0, 1]$ and the payment function $p : T \rightarrow \mathbb{R}$. A mechanism is *individually rational* if the utility of every type of the agent is at least zero,

$$(IR) \quad v \cdot x(v) - p(v) \geq 0, \quad \forall v.$$

A mechanism is *incentive compatible* if no type of the agent increases his utility by misreporting,

$$(IC) \quad v \cdot x(v) - p(v) \geq v \cdot x(\hat{v}) - p(\hat{v}), \quad \forall v, \hat{v}.$$

The revelation principle (Gibbard (1973); Myerson (1986)) states that the outcome of any mechanism in equilibrium can be achieved by an incentive compatible direct revelation mechanism. By invoking the revelation principle, the problem of designing the revenue-optimal mechanism becomes to find the revenue-optimal direct revelation incentive compatible and individually rational mechanism. That is, an incentive compatible and individually rational pair (x, p) maximizing the expected revenue, $\mathbf{E}_{v \sim F}[p(v)]$.

The following lemma connects the payment function of an IC mechanism with its allocation function.

Lemma 1 (Myerson, 1981; Rochet, 1985). *Function x is the allocation function of an agent in an incentive compatible mechanism if and only if x is monotone non-decreasing. With monotone nondecreasing x , the agent's payment is $p(v) = p(0) + vx(v) - \int_{z \leq v} x(z) dz$, or alternatively, the agents utility $u(v) = v \cdot x(v) - p(v)$ is $u(v) = \int_{z \leq v} x(z) dz - p(0)$.*

Proof. Directly from the definition of incentive compatibility, for all v and \hat{v} we must have,

$$vx(v) - p(v) \geq vx(\hat{v}) - p(\hat{v}),$$

$$\hat{v}x(\hat{v}) - p(\hat{v}) \geq \hat{v}x(v) - p(v).$$

Summing up the above two inequalities, we must have $(v - \hat{v})(x(v) - x(\hat{v})) \geq 0$ for all v and \hat{v} . In particular, this must be true when $\hat{v} = v - \epsilon$. As ϵ goes to zero, we can write

$$\begin{aligned} (v - \hat{v})(x(v) - x(\hat{v})) &= \epsilon \cdot (x(v) - (x(v) - \epsilon x'(v))) \\ &= \epsilon^2 \cdot x'(v). \end{aligned}$$

Therefore, the inequality $(v - \hat{v})(x(v) - x(\hat{v})) \geq 0$ can be satisfied only if $x'(v) \geq 0$. This proves the necessity of monotonicity of x .

Now assume that x is monotone, and define $p(v) = p(0) + vx(v) - \int_{z \leq v} x(z) dz$. We will prove that (x, p) is incentive compatible. Consider any v and \hat{v} . Using the definition of p we have

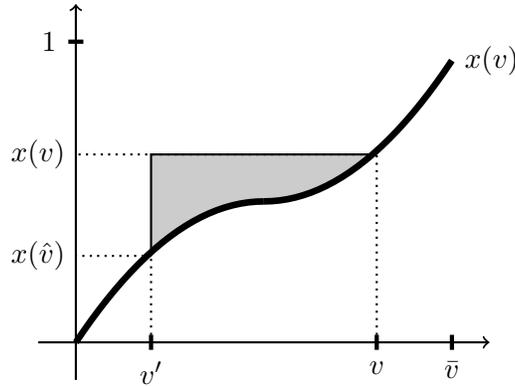


Figure 2.1. When $v \geq \hat{v}$, the shaded area is equal to $\int_{z=\hat{v}}^v [x(z) - x(\hat{v})] dz$, which is the difference between the utility of truth-telling, i.e. reporting v , and the utility of reporting \hat{v} . The case where $v \leq \hat{v}$ is similar.

$$\begin{aligned}
 vx(v) - p(v) - (vx(\hat{v}) - p(\hat{v})) &= x(\hat{v})(\hat{v} - v) + \int_{z=\hat{v}}^v x(z) dz \\
 &= \int_{z=\hat{v}}^v [x(z) - x(\hat{v})] dz \\
 &\geq 0,
 \end{aligned}$$

where the last inequality follows by monotonicity of x (see Figure 2.1). This implies that,

$$vx(v) - p(v) \geq vx(\hat{v}) - p(\hat{v}).$$

This completes the proof that the mechanism is incentive compatible. \square

Any optimal mechanism must satisfy $p(0) = 0$, since otherwise if $p(0) < 0$ we can increase all payments equally and improve revenue, and $p(0)$ can not be positive since otherwise the mechanism would violate individual rationality for the type $v = 0$. So in the rest of the section we assume $p(0) = 0$. The revenue maximization problem can then be written as the following

program.

$$(2.1) \quad \begin{aligned} \max_x \quad & \int_v \left[v \cdot x(v) - \int_{z \leq v} x(z) \, dz \right] f(v) \, dv \\ & \frac{d}{dv} x(v) \geq 0, \\ & x(v) \in [0, 1]. \end{aligned}$$

The first constraint is the monotonicity of the allocation rule, and the third is the feasibility of the outcome, since the probability of allocation can be at most one and at least zero (the feasibility will become more complicated with more agents). Notice that the above formulation contains both x and its integral. It will be more convenient to rewrite the objective function in terms of x only. Myerson used *integration by parts* to accomplish this task.

Lemma 2 (Single-dimensional Integration by Parts). *For any two differentiable functions $h, g : \mathbb{R} \rightarrow \mathbb{R}$,*

$$\int_{v=a}^b h(v)g'(v) \, dv = h(v)g(v) \Big|_{v=a}^b - \int_{v=a}^b h'(v)g(v) \, dv.$$

We can define $h(v) = \int_{z \leq v} x(z) \, dz$, $g(v) = -(1 - F(v))$ and invoke integration by parts to write

$$\begin{aligned} \int_v \int_{z \leq v} x(z) \, dz f(v) \, dv &= - \int_{z \leq v} x(z) \, dz (1 - F(v)) \Big|_{v=0}^{\bar{v}} + \int_v x(v)(1 - F(v)) \, dv \\ &= \int_v x(v)(1 - F(v)) \, dv. \end{aligned}$$

The second equality followed because $\int_{z \leq 0} x(z) \, dz = 0$ and $1 - F(\bar{v}) = 0$. By plugging in the above equation into the revenue function of Equation 2.1 we conclude the revenue of a mechanism with allocation x is equal to

$$\int_v x(v) \left[v - \frac{1 - F(v)}{f(v)} \right] f(v) \, dv$$

Lemma 3 (Strong amortization of revenue for the linear single-parameter setting). *The function $\phi(v) = v - \frac{1 - F(v)}{f(v)}$ is a strong amortization of revenue. That is, the revenue of a mechanism with allocation x can be written as*

$$\int_v x(v) \phi(v) f(v) \, dv.$$

Using the above lemma, we can therefore rewrite program 2.1,

$$(2.2) \quad \begin{aligned} \max_x \quad & \int_v x(v) \phi(v) f(v) \, dv \\ & \frac{d}{dv} x(v) \geq 0, \\ & x(v) \in [0, 1]. \end{aligned}$$

The revenue-maximization problem becomes to maximize $\mathbf{E}_v[x(v)\phi(v)]$ subject to monotonicity and feasibility of allocation. Consider a relaxation of the above problem in which the monotonicity constraint is removed.

$$(2.3) \quad \begin{aligned} \max_x \quad & \int_v x(v) \phi(v) f(v) \, dv \\ & x(v) \in [0, 1]. \end{aligned}$$

The relaxed problem is a *pointwise* optimization problem, that is, its solution is a function that maximizes $x(v)\phi(v)$ for each v . If the solution to the relaxed problem is a monotone

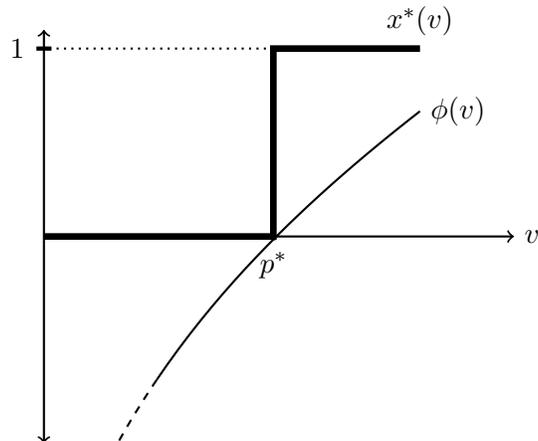


Figure 2.2. The function x^* optimizes $\mathbf{E}_v[x(v)\phi(v)]$ subject to the feasibility constraint $x \in [0, 1]$, for monotone ϕ . The allocation corresponds to the allocation of posting a price p^* for the item. In this figure, the distribution is $f(v) = 0.5 + v$ and $F(v) = 0.5v + 0.5v^2$, for $v \in [0, 1]$, the virtual value is $\phi(v) = \frac{1.5v^2 + v - 1}{v + 0.5}$, and $p^* = \frac{\sqrt{5}-1}{2}$. As a sanity check, you can verify that p^* indeed optimizes the revenue of pricing $p(1 - F(p))$.

allocation, then the same solution must be optimal for the more restricted problem (when monotonicity constraint is restored).

Definition 1 (Incentive compatibility and virtual value function). A function $\hat{\phi}$ is *incentive compatible* if a monotone allocation maximizes $x(v)\hat{\phi}(v)$ pointwise. The strong amortization function ϕ is a *virtual value function* if it is incentive compatible.

The above paragraph and definition directly imply the following:

Lemma 4. *If ϕ is a virtual value function, then the allocation function that pointwise optimizes ϕ is the allocation of the optimal mechanism.*

If ϕ is a monotone non-decreasing function, the solution to the pointwise optimization is a function x^* which is zero below a threshold p^* and one above that threshold, where p^* is such that $\phi(p^*) = 0$ (see Figure 2.2). Notice that x^* is monotone, and is the allocation function of

posting a price of p^* for the item: given a price p^* , the dominant strategy of the agent is to accept the price if $v \geq p^*$, and reject it otherwise.

Lemma 5. *If the strong amortization $\phi(v)$ is monotone non-decreasing, then it is a virtual value function. In that case, the allocation of posting a price p^* such that $\phi(p^*) = 0$ pointwise optimizes ϕ , and therefore pricing p^* is the optimal mechanism.*

We call a distribution *regular* if the strong amortization $\phi(v)$ is monotone non-decreasing. We call a price p^* such that $\phi(p^*) = 0$ the *monopoly reserve price* of the distribution. The above lemma can be restated to say that for regular distributions, posting the monopoly reserve price for the item is the optimal mechanism. We will next discuss the extension of the problem definition and the method of virtual values to the case of multiple buyers, termed the *multi-agent* problem.

The Multi-agent Single-item Problem. There are n agents labeled 1 to n . Agent i is assigned a distribution f_i . A profile of values is $\mathbf{v} = (v_1, \dots, v_n)$, with the value of each agent drawn independently of the others' from the corresponding distribution.² A multi-agent mechanism is a pair of functions $(\bar{\mathbf{x}}, \bar{\mathbf{p}})$ where the allocation function is $\bar{\mathbf{x}}(\mathbf{v}) = (\bar{x}_1(\mathbf{v}), \dots, \bar{x}_n(\mathbf{v}))$ in which $\bar{x}_i(\mathbf{v})$ is the allocation of agent i given the reported vector of types \mathbf{v} , and $\bar{\mathbf{p}}(\mathbf{v}) = (\bar{p}_1(\mathbf{v}), \dots, \bar{p}_n(\mathbf{v}))$ is the vector of payments.³ The constraint that there is a single item to allocate implies that $\sum_i \bar{x}_i(\mathbf{v}) \leq 1$.⁴

Given $\bar{\mathbf{x}}$, we can define the *interim allocation rule* x_i for agent i to be the probability of allocation when other agents' values are drawn at random, that is,

²Throughout this thesis, I will denote vectors using bold symbols, and keep their elements unbolded.

³Throughout this thesis, I will denote multi-agent functions, i.e., functions of profiles of types, with an *upper bar*, and single-agent functions without the bar.

⁴We will later generalize the feasibility constraint.

$$(2.4) \quad x_i(v_i) = \mathbf{E}_{\mathbf{v}_{-i} \sim F_{-i}}[\bar{x}_i(v_i, \mathbf{v}_{-i})],$$

where \mathbf{v}_{-i} is a vector from which the i 'th element has been removed. We use notation (v_i, \mathbf{v}_{-i}) to insert v_i as the i 'th element back into the vector. Similarly we can define *interim payment* p_i for agent i ,

$$(2.5) \quad p_i(v_i) = \mathbf{E}_{\mathbf{v}_{-i} \sim F_{-i}}[\bar{p}_i(v_i, \mathbf{v}_{-i})].$$

The interim mechanism (x_i, p_i) is a single-agent mechanism *induced* for agent i from the multi-agent mechanism $(\bar{\mathbf{x}}, \bar{\mathbf{p}})$. A multi-agent mechanism is *Bayesian incentive compatible* (BIC) if the induced interim mechanism of each agent is incentive compatible. A multi-agent mechanism is *interim individually rational* (IIR) if the induced interim mechanism of each agent is individually rational. Again, by applying the revelation principle, the problem is to design a BIC and IIR mechanism maximizing revenue,

$$\mathbf{E}_{\mathbf{v} \sim F_1 \times \dots \times F_n} \left[\sum_i \bar{p}_i(\mathbf{v}) \right].$$

We can now directly apply Lemma 1 to conclude that a mechanism is Bayesian incentive compatible if x_i is monotone non-decreasing for all i . In addition, the interim payment and allocation rules of Bayesian incentive compatible mechanisms relate as follows: $p_i(v_i) = v_i x_i(v_i) - \int_{z \leq v_i} x_i(z) dz$ (again, without loss of generality $p_i(0) = 0$ for the optimal solution).

We can now use linearity of expectation to rewrite the expected revenue of a mechanism

$$\begin{aligned}
\mathbf{E}_{\mathbf{v}} \left[\sum_i \bar{p}_i(\mathbf{v}) \right] &= \sum_i \mathbf{E}_{v_i} [p_i(v_i)] \\
&= \sum_i \mathbf{E}_{v_i} \left[v_i x_i(v_i) - \int_{z \leq v_i} x_i(z) dz \right] \\
&= \sum_i \mathbf{E}_{v_i} \left[x_i(v_i) \left[v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \right] \right] \\
&= \sum_i \mathbf{E}_{v_i} [x_i(v_i) \phi_i(v_i)],
\end{aligned}$$

where the third equality followed using integration by parts (similar to the single-agent derivation), and the last equality by defining $\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$.

As a result, the problem of optimizing revenue subject to Bayesian incentive compatibility and individual rationality becomes,

$$\begin{aligned}
(2.6) \quad & \max_{\bar{\mathbf{x}}} \quad \sum_i \mathbf{E}_{v_i} [x_i(v_i) \phi_i(v_i)] \\
& \frac{d}{dv_i} x_i(v_i) \geq 0, \\
& \bar{\mathbf{x}}(\mathbf{v}) \text{ is feasible.}
\end{aligned}$$

We can use linearity of expectation one more time and write

$$\sum_i \mathbf{E}_{v_i} [x_i(v_i) \phi_i(v_i)] = \mathbf{E}_{\mathbf{v}} \left[\sum_i \bar{x}_i(\mathbf{v}) \phi_i(v_i) \right],$$

and rewrite the optimization program

$$(2.7) \quad \max_{\bar{\mathbf{x}}} \quad \mathbf{E}_{\mathbf{v}} \left[\sum_i \bar{x}_i(\mathbf{v}) \phi_i(v_i) \right]$$

$$\frac{d}{dv_i} x_i(v_i) \geq 0,$$

$$\sum_i \bar{x}_i(\mathbf{v}) \leq 1.$$

Once again, we consider the problem without the the monotonicity constraint, and verify that the solution is indeed monotone. The problem becomes to optimize $\sum_i \bar{x}_i(\mathbf{v}) \phi_i(v_i)$ *pointwise*, that is, for each profile of types \mathbf{v} , subject to the feasibility constraint $\sum_i \bar{x}_i(\mathbf{v}) \leq 1$. We will naturally extend the definition of incentive compatibility of the strong amortization functions to require that the pointwise optimization of $\sum_i \bar{x}_i(\mathbf{v}) \phi_i(v_i)$ subject to the feasibility condition is an incentive compatible allocation. In that case, the allocation that pointwise optimizes $\sum_i \bar{x}_i(\mathbf{v}) \phi_i(v_i)$ is the allocation of the optimal mechanism.

Lemma 6. *If the distributions of all agents are regular, then the amortization functions ϕ_i are virtual value functions. As a result, pointwise optimization of virtual values is optimal.*

Proof. A solution is to set the allocation of the agent with the highest positive virtual value equal to one, with ties are broken at random, and the rest to zero. If the virtual value function is monotone, an agent can only increase its allocation by increasing its value. Therefore, the allocation functions are monotone. \square

Notice with monotone virtual values, the optimal mechanism satisfies stronger notions of incentive compatibility and individual rationality than Bayesian incentive compatibility and interim individual rationality, namely, *dominant strategy incentive compatibility*, and *ex-post individual rationality*. That is, fixing any profile of other agents (and not just in expectation),

truthfulness maximizes any agents utility (since the allocation is monotone), and also the utility of being truthful is non-negative.

Example 1 (Identical regular distributions). Assume that agents are identical, i.e., each agent's value is drawn independently from a distribution F . As a result, the virtual value functions are identical too. Assume further that F is regular, which means that the agent with highest value will have the highest virtual value. Since the virtual value of the winner can not be negative, the winner will be the agent with the highest value, conditioned on the value being above the monopoly reserve price. This is exactly the allocation of the second price auction with monopoly reserve. We conclude that for identical regular distributions, the second price auction with monopoly reserve price is optimal.

Ironing. Myerson (1981) used ironing to design virtual values for the case of non-regular distributions. The rest of this subsection explains the construction. Our method of virtual values in Chapter 3 uses similar techniques to handle non-regular distributions in service-constrained settings. We will start by considering a generic single agent and drop the label of the agent.

The geometry of single-dimensional auction theory is more readily apparent when we index an agent's private type by its strength relative to the distribution. The *quantile* of an agent is the measure of higher-valued types, i.e., an agent with type v has quantile $q = 1 - F(v)$. Notice that with this definition, higher values are mapped to lower quantiles. Perform a change of variables $q = 1 - F(v)$ and rewrite the expected virtual value,

$$\begin{aligned} \int_v x(v) \left[v - \frac{1 - F(v)}{f(v)} \right] f(v) \, dv &= - \int_v x(v) \frac{d}{dv} [v(1 - F(v))] \, dv \\ &= \int_{q=0}^1 x(F^{-1}(1 - q)) \frac{d}{dq} [F^{-1}(1 - q)q] \, dq. \end{aligned}$$

For brevity of notation let us overload x such that $x(q)$ is the allocation of the type that is mapped to the quantile q , that is $x(q) = x(F^{-1}(1 - q))$. Define the *revenue curve* $P : [0, 1] \rightarrow \mathbb{R}$, where $P(q)$ is the revenue of posting a price that is accepted with probability q , that is, $P(q) = qF^{-1}(1 - q)$. The next lemma follows directly.

Lemma 7 (Strong amortization in quantile space). *The marginal revenue function $\frac{d}{dq}P(q)$ is a strong amortization of revenue in quantile space. That is, the revenue of a mechanism can be written as its expected marginal revenue,*

$$\int_{q=0}^1 x(q) \frac{d}{dq}[P(q)] dq.$$

It is a simple exercise to observe that the regularity of the distribution is equivalent to concavity of the revenue curve. This suggests an alternative interpretation of the optimal mechanism, as *the marginal revenue mechanism* for concave revenue curves: map types to quantiles according to the mapping $q = 1 - F(v)$ for each agent, and maximize expected marginal revenue by allocating the item to the quantile with the highest non-negative marginal revenue. We will next define the concept of the *ironed revenue curve* to extend the analysis to irregular distributions.

Definition 2. The *ironed revenue curve* R is the concave hull of the revenue curve P .

The marginal ironed revenue $\frac{d}{dq}R(q)$ is monotone since R is concave, and therefore the optimization of $\frac{d}{dq}R(q)$ pointwise results in incentive compatible allocations. However, the ironed marginal revenue $\frac{d}{dq}R(q)$ is no longer an amortization of revenue. We will next show that a more careful analysis can be used to prove that optimizing $\frac{d}{dq}R(q)$ pointwise is optimal. The

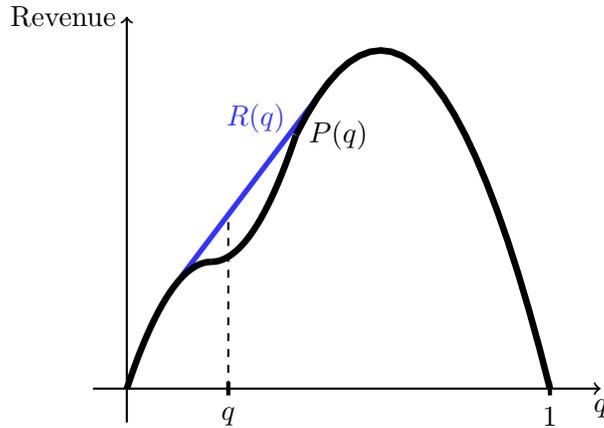


Figure 2.3. The curve R is the concave hull of the revenue curve P . By definition, it is concave, and satisfies $R \geq P$. Importantly, whenever $R(q) > P(q)$, the derivative of the concave hull $R'(q)$ is constant.

expression of revenue in terms of marginal revenue suggests an approach to characterize optimal mechanisms for non-regular distributions. Using integration by parts,

$$\begin{aligned}
 \int_q x(q)(R'(q) - P'(q)) &= x(q)(R(q) - P(q))\Big|_{q=0}^1 - \int_q x'(q)(R(q) - P(q)) \\
 (2.8) \qquad \qquad \qquad &= - \int_q x'(q)(R(q) - P(q)) \geq 0,
 \end{aligned}$$

where the inequality followed because $x' \leq 0$ for all incentive compatible allocations (notice that high values are mapped to low quantiles, therefore incentive compatible x is a non-increasing function of q), and $R \geq P$ since R is the concave hull of P . We can therefore use the following upper bound on revenue,

$$\int_q x(q)R'(q) \, dq.$$

We are now ready to study multi-agent mechanisms with agents 1 to n . Consider the allocation \bar{x} of the multi-agent mechanism that optimizes *ironed* marginal revenue, that is, after mapping types to quantiles, it allocates the item to the agent i with highest ironed marginal revenue $R'_i(q_i)$. We will argue that \bar{x} optimizes revenue. Notice that the interim allocation rules of this mechanism satisfy $\int_q x_i(q)R'_i(q) dq = \int_q x_i(q)P'_i(q) dq$. The reason is that $R_i(q_i) > P_i(q_i)$ implies that $R'_i(q_i)$ is a constant (see Figure 2.3), and therefore $x'_i(q_i) = 0$. This would imply that inequality 2.8 holds with equality, and therefore $\int_q x_i(q)R'_i(q) dq = \int_q x_i(q)P'_i(q) dq$. As a result, the expected ironed marginal revenue of \bar{x} , $\sum_i \int_q x_i(q)R'_i(q) dq$, is equal to its revenue. Since \bar{x} optimizes ironed marginal revenue pointwise and therefore in expectation, its ironed marginal revenue, and therefore its expected revenue, must be more than the ironed marginal revenue of any other mechanism. Since the expected marginal revenue of any mechanism upper bounds its revenue (Equation 2.8), \bar{x} optimizes revenue. We will call $\frac{d}{dq}R(q)$ satisfying the above properties a *weak amortization of revenue*.

Definition 3 (Weak amortization of revenue). The function $\frac{d}{dq}R(q)$ is a weak amortization of revenue. That is,

- The revenue of any mechanism can be upper bounded by its expected marginal ironed revenue $\mathbf{E}_q[x(q)R'(q)]$.
- The revenue of the mechanism that optimizes $\frac{d}{dq}R(q)$ pointwise is equal to its expected marginal ironed revenue.

To summarize, in this section we provided two methods of proving optimality of mechanism, based on strong or weak amortizations. We showed that if amortizations are in addition incentive compatible, then their pointwise optimization is the optimal mechanism.

2.1.2. The Service-Constrained Model

Single-agent Mechanisms. For the special case of a single agent, the single-service model is a fully general single-agent mechanism design problem, which we describe next. There is a set of outcomes W , with $w \in W$ being a potential outcome the agent could receive. The outcome in general encodes anything the agent could care about, including the items and services it acquires, as well as the payment the agent is required to make, denoted $\text{Payment}(w) \in \mathbb{R}$. In a randomized environment (e.g., randomness from a randomized mechanism or Bayesian environment) the outcome an agent receives is a random variable from a distribution over W . The space of all such distributions is denoted $\Delta(W)$.

The agent has a type t from a type space T . This type is drawn from distribution $f \in \Delta(T)$ and we equivalently denote by f the probability mass function. I.e., for every $t \in T$, $f(t)$ is the probability that the type is t . The utility function $u : T \times W \rightarrow \mathbb{R}$ maps the agent's type and the outcome to a real valued utility. The agent is a von Neumann–Morgenstern expected utility maximizer and we extend u to $\Delta(W)$ linearly, i.e., for $w \in \Delta(W)$, $u(t, w)$ is the expectation of u where the outcome is drawn according to w . We require no extra structure from u , and in particular, we do not require the usual assumption of quasi-linearity.

A single-agent mechanism, without loss of generality by the revelation principle, is just an *outcome rule*, a mapping from the agent's type to a distribution over outcomes. We denote an *outcome rule* by $w : T \rightarrow \Delta(W)$. We say that an outcome rule w is *incentive compatible* (IC) and *individually rational* (IR) if for all $t, t' \in T$, respectively,

$$(IC) \quad u(t, w(t)) \geq u(t, w(t')),$$

$$(IR) \quad u(t, w(t)) \geq 0.$$

The *expected* payment rule is $p : T \rightarrow \mathbb{R}$ is $p(t) = \mathbf{E}[\text{Payment}(w(t))]$. By invoking the revelation principle, the single-agent revenue maximization problem is to design an incentive compatible and individually rational mechanism maximizing expected revenue,

$$\mathbf{E}_{t \sim f, \hat{w} \sim w(t)}[\text{Payment}(\hat{w})] = \mathbf{E}_{t \sim f}[p(t)].$$

Notice that there is no assumption aside from the assumption that the agent is von Neumann-Morgenstern expected utility maximizer, and in particular, no feasibility structure. As a result, the service-constrained mechanism design problem for the case of a single agent is a fully general single-agent mechanism design problem.

We give three examples to illustrate the abstract model described above.

Example 2 (Linear Single-Parameter Model). Here the agent's type space is $T \subset \mathbb{R}_+$ where $t \in T$ represents the agent's valuation for the item. An outcome is $(\text{Assignment}(w), \text{Payment}(w)) \in \{0, 1\} \times \mathbb{R}_+ = W$, where $\text{Assignment}(w)$ indicates whether or not the item is assigned to the agent, and $\text{Payment}(w)$ indicates the payment. The agent's linear utility function is $u(t, w) = t \cdot \text{Assignment}(w) - \text{Payment}(w)$. We can define the assignment rule $x(t) = \mathbf{E}[\text{Assignment}(w(t))]$ denoting the probability of assignment of a type t . By risk neutrality, the agent's utility in a randomized mechanism is $t \cdot x(t) - p(t)$.

Example 3 (Multi-service, Unit Demand, Quasi-linear, and Risk-neutral Preferences). Here the type space is $T \subset \mathbb{R}_+^m$, where m is the number of services, and a type is an m -dimensional vector $\mathbf{t} \in T$ indicating the agent's valuation for each service. An outcome is $w = (\mathbf{Assignment}(w), \text{Payment}(w)) \in \{0, 1\}^m \times \mathbb{R}_+ = W$, where $\mathbf{Assignment}(w)$ is a vector of indicators of services assigned to the agent, and the outcome's payment is $\text{Payment}(w)$. We assume the agent is unit demand, meaning its value for a set of services is equal to its value

for the favorite service in the set. Therefore without loss of generality we can assume that the outcome assigns at most one service to the agent, that is $\sum_i \text{Assignment}_i(w) \leq 1$. The agent's quasi-linear utility function is $u(\mathbf{t}, w) = \mathbf{t} \cdot \mathbf{Assignment}(w) - \text{Payment}(w)$. We can define the assignment rule of a type $\pi(\mathbf{t}) = \mathbf{E}_w[\mathbf{Assignment}(w(\mathbf{t}))]$ to be the vector of assignment probabilities of services (we distinguish the case $m = 1$, the linear single-parameter setting, by using function $x(\cdot)$ instead of $\pi(\cdot)$, as in Example 2). By risk neutrality, the agent's utility in a randomized mechanism is $\mathbf{t} \cdot \pi(\mathbf{t}) - p(\mathbf{t})$. For example, in the ad auction scenario discussed in the introduction, each service is the ad slot accompanied by a set of cookies (possibly the empty set). Another example could be a seller that sells an item in one of the potential m "days", and the agent has different values for acquiring the item in each day. The linear single-parameter model is the special case when $m = 1$. In Chapter 3, we consider the problem with general m .

Example 4 (Non-linear Preferences). Beyond these two examples, our framework can easily incorporate more general agent preferences exhibiting, e.g., risk aversion or a budget limit. To do this, we can augment the set of types to include parameters of more general forms of utility function. For example, the set of types could be $T = \mathbb{R}^+ \times \mathbb{R}^+$, where $(v, B) \in T$ denotes the value that the agent has for the item, and its budget. The utility for an outcome $w \in W = \{0, 1\} \times \mathbb{R}_+$ is $v \cdot \text{Assignment}(w) - \text{Payment}(w)$ if $\text{Payment}(w) \leq B$, and negative infinity otherwise.

Multi-agent Single-Service Mechanisms. There are n independent agents. Agents need not be identical, i.e., agent i 's type space is T_i , the probability mass function for her type is f_i , her outcome space is W_i , and her utility function is u_i . The profile of agent types is denoted by $\mathbf{t} = (t_1, \dots, t_n) \in T_1 \times \dots \times T_n = \mathbf{T}$, the joint distribution on types is $\mathbf{f} \in \Delta(T_1) \times \dots \times \Delta(T_n)$, a vector of outcomes is $(w_1, \dots, w_n) \in \mathbf{W}$, where \mathbf{W} is the set of *jointly feasible* profile of outcomes, which is supposed to capture constraints on possible outcomes dictated by scarcity of resources. In single-service mechanisms, the set of jointly possible outcomes has

a specific structure. Importantly, each agent is assigned an *allocation function* $\text{Alloc}_i(w_i) \in \{0, 1\}$ denoting whether or not the agent is *served* in that outcome, and the mechanism has an inter-agent feasibility constraint that permits serving at most a single agent. That is, outcome (w_1, \dots, w_n) is feasible if and only if $\sum_i \text{Alloc}_i(w_i) \leq 1$.⁵ A mechanism that obeys this constraint is *feasible*. Notice that we require no correlation between the utility functions of agents and the allocation functions. For the special case of a single agent, the feasibility condition is automatically satisfied, and therefore the single-agent single-service problem is a general single-agent problem.

A mechanism maps type profiles to a (distribution over) outcome vectors via an *ex post outcome rule*, denoted $\bar{\mathbf{w}} : \mathbf{T} \rightarrow \Delta(\mathbf{W})$ where $\bar{w}_i(\mathbf{t})$ is the outcome obtained by agent i . We will similarly define $\bar{\mathbf{x}} : \mathbf{T} \rightarrow [0, 1]^n$ as $\bar{\mathbf{x}}(\mathbf{t}) = \mathbf{E}[\mathbf{Alloc}(\bar{\mathbf{w}}(\mathbf{t}))]$, as the *ex post allocation rule* (where $[0, 1] \equiv \Delta(\{0, 1\})$). The ex post allocation rule $\bar{\mathbf{x}}$ and the probability mass function \mathbf{f} on types induce *interim* outcome and allocation rules. For agent i with type t_i and $\mathbf{t} \sim \mathbf{Dist}_{\mathbf{t}}[\mathbf{t} | t_i]$ the interim outcome and allocation rules are $w_i(t_i) = \mathbf{Dist}_{\mathbf{t}}[\bar{w}_i(\mathbf{t}) | t_i]$ and $x_i(t_i) = \mathbf{Dist}_{\mathbf{t}}[\bar{x}_i(\mathbf{t}) | t_i] \equiv \mathbf{E}_{\mathbf{t}}[\bar{x}_i(\mathbf{t}) | t_i]$.⁶ A profile of interim allocation rules is feasible if it is derived from an ex post allocation rule as described above; the set of all profiles of feasible interim allocation rules is denoted by \mathbb{X} . A mechanism is Bayesian incentive compatible and interim individually rational if equations (IC) and (IR), respectively, hold for all i and all t_i .

Consider again the examples described previously of quasi-linear single-dimensional and unit-demand preferences. For the single-dimensional example, the multi-agent mechanism design problem is the standard single-item auction problem. The allocation and assignment indicators are the same, that is $x(t) = \pi(t)$, and denote whether the item is allocated or not. In this chapter we use notation $x(\cdot)$ as the allocation/assignment function for linear single-parameter settings.

⁵We will later generalize the model to allow for multiple agents being served.

⁶We use notation $\mathbf{Dist}[X | E]$ to denote the distribution of random variable X conditioned on the event E .

For the unit-demand example, the multi-agent mechanism design problem is an *attribute auction*. In this problem there is a single unit available and it can be configured in one of m ways. Importantly, the designer’s feasibility constraint places no restrictions on how the unit can be configured. Our two unit-demand examples fit naturally in this model: We define allocation rule $\text{Alloc}_i(w_i) = 1$ if the agent is assigned a service. The assignment rule, on the other hand, is a vector which identifies which service (if any) is assigned to the agent. Even though the assignment rule contains more information, feasibility can be verified using only the allocation rule. For the ad auctions example, there is a single ad to be sold, but there is no feasibility constraint on what cookies can accompany the ad. For the seller with multiple days to sell the item, feasibility of the allocation can be determined without specifying the day on which the item is to be sold. In addition, extra structure on the utility function, such as budget or risk, impose no extra structure on the feasibility of a mechanism.

2.1.3. Service-constrained Extension: The Method of Virtual Values

In Chapter 3, we extend the method of virtual values to a class of unit-demand service-constrained settings. For regular distributions, Myerson’s virtual values satisfy two key properties:⁷

- (1) Strong amortization: the revenue of any allocation can be written as its expected virtual values.
- (2) Incentive compatibility: An incentive compatible allocation optimizes virtual values pointwise.

The main tools used by Myerson in his analysis are expressing the payment function in terms of the allocation function and its integral using Lemma 1, and using integration by parts to rewrite revenue in terms of the allocation function only. Fortunately, Rochet’s original lemma can

⁷The more general discussion including the design of weak amortizations is deferred to Chapter 3.

be used directly to connect payment function and allocation function with general preferences. Also, the integration by parts lemma can be defined for high dimensional spaces. In fact, in higher dimensions, integration by parts can be applied with *degrees of freedom*, that is, in a parametric form. There is a simple reason for this. Integration by parts basically expresses the value of a function as the sum of marginal increments in the value of the function on a path between two points in space. In higher dimensional spaces, there are more paths between any two points (whereas in one dimension, there is only one path connecting any pair of points), and as a result there are more ways to integrate by parts. Therefore, in a sense, integrating by parts becomes *easier* in higher dimensions. As a result, the claim that designing virtual values in high dimensions is difficult might come as a surprise. The key to understand this difficulty is in considering the second property that virtual values must satisfy: pointwise optimization of virtual values must be incentive compatible. In high dimensional spaces, this heavily constrains admissible virtual values, since they must satisfy constraints in more dimensions. Our main contribution is a methodology to design virtual values, i.e., the right way to perform integration by parts, to satisfy the incentive compatibility requirement.

2.2. The Marginal Revenue Mechanism

Marginal revenue plays a fundamental role in microeconomic theory. For example, a monopolist providing a commodity to two markets each with its own concave revenue (as a function of the supply provided to that market) optimizes her profit by dividing her total supply to equate the marginal revenues across the two markets. Moreover this central economic principle also governs classical auction theory. Bulow and Roberts (1989) reinterpret Myerson's virtual value as the marginal revenue of a certain concave revenue curve. In Chapter 4, we extend the definition of the marginal revenue mechanism to service-constrained environments, and provide conditions proving its optimality.

In this section we warm up by giving a new proof that the marginal revenue mechanism is revenue optimal for agents with single-dimensional linear preferences. We have already seen a proof using Myerson’s characterization in Section 2.1, but a crucial part of that proof was the ability to tie payment and allocation functions closely to each other (Lemma 1), which is not possible in general. Therefore, in this section we present a new proof relying minimally on the details of the linear single-parameter setting. In this proof we will introduce many concepts that make our generalization possible (which were not present in previous proofs). The basic approach is as follows. We formulate an important class of *lottery pricing problems*, the solution to which define a revenue curve. We show that single-dimensional linear agents are *revenue linear* in the sense that it is optimal to decompose the allocation to any agent as a convex combination of the solutions to these lottery pricing problems. Finally, we observe that this decomposition implies that the optimal revenue can be expressed in terms of the *surplus of marginal revenue*: the sum of derivatives of the revenue curves of agents served evaluated at points corresponding to the agents’ types. The marginal revenue mechanism optimizes this latter term pointwise and, therefore, also in expectation. In the interest of brevity we will keep the discussion informal; many of the proofs in this section are subsumed by generalizations in Section 4.2 which are given formally. Except for the proof of revenue-linearity of the linear single-parameter setting, the rest of the composition does not rely on the details of the setting and is easily generalizable.

Similar to Section 2.1, we will start by studying single-agent problems. Such an analysis will be useful since most of the analysis for a multi-agent mechanism can be done via studying the induced interim mechanisms. Until multi-agent mechanisms are discussed later in the section, we discuss a generic single agent and drop the label of the agent.

Let $V(q) = F^{-1}(1 - q)$ be the *inverse demand curve*, i.e., $V(\hat{q})$ is the posted price that would be accepted by the \hat{q} measure of highest-valued agents (and rejected by all others). Importantly, for v drawn at random from the distribution F , $q = V^{-1}(v)$ is uniform on $[0, 1]$ (therefore,

expectations of functions of q are given by integrals with probability density one). From the perspective of an agent in a single-agent mechanism and as a function of the agent's report, the agent is served with some probability and makes some expected payment. Similar to Section 2.1, we may index the allocation and payment of each type by the quantile of that type, i.e., the agent with quantile q is assigned outcome $(x(q), p(q))$. We can write *incentive compatibility* and *individual rationality* constraints,

$$(IC) \quad V(q)x(q) - p(q) \geq V(q)x(q') - p(q'), \quad \forall q, q' \in [0, 1].$$

$$(IR) \quad V(q)x(q) - p(q) \geq 0, \quad \forall q \in [0, 1].$$

We call a mechanism satisfying the above conditions a *lottery pricing*.

Constrained Lottery Pricings. Given a lottery pricing and a distribution over the agent's value, an ex ante expected payment $\mathbf{E}_q[p(q)]$ and ex ante probability of allocation $\mathbf{E}_q[x(q)]$ are induced. The single-agent lottery pricing problem that forms the basis for the marginal revenue mechanism is the following. Given an ex ante constraint \hat{q} , find the lottery pricing that serves the agent with probability \hat{q} and maximizes revenue.

Definition 4. The *ironed revenue curve* $R(\hat{q})$ is defined for all $\hat{q} \in [0, 1]$ as the revenue of optimal lottery pricing x with ex ante allocation probability \hat{q} , $\mathbf{E}_q[p(q)] = \hat{q}$.

Recall the alternative definition of the ironed revenue curve as the concave hull of the revenue curve $P(\hat{q}) = \hat{q}V(\hat{q})$ in Definition 2. The two definitions are in fact equivalent. The proof of the equivalence is not required for the material of this section. We will here only provide some intuition by arguing that the ironed revenue curve of Definition 4 upper bounds the concave hull of the revenue curve, and for the rest of the section use the new definition in Definition 4. It is clear that R of Definition 4 satisfies $R \geq P$, since by definition $P(\hat{q})$ is the revenue of a price that is accepted with probability \hat{q} . In addition, given any \hat{q} , we can post a random price that is

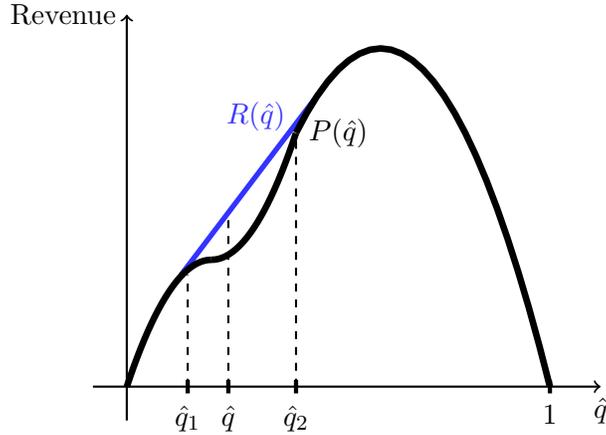


Figure 2.4. A mechanism with ex ante probability of service \hat{q} can be defined by randomizing prices $V(\hat{q}_1)$ and $V(\hat{q}_2)$ such that the expected probability of service is \hat{q} .

accepted in expectation with probability q , which implies that R must at least be as large as the concave hull of P (see Figure 2.4). The fact that R is equal to the concave hull of P (which we will not prove, and is not required for the rest of the material) states that a convex combination of at most two prices is guaranteed to obtain the highest revenue.

The ex ante lottery pricing problem gives an (equality) constraint on the total probability that the agent is served in expectation over all quantiles she may have. To get more fine-grained control over the lottery pricing we additionally allow upper bounds to be specified on the total probability of allocation to subsets of quantiles. Consider the following lottery pricing problem: Given a monotone concave function $\hat{X}(\hat{q})$, find the optimal lottery pricing where the ex ante probability of allocating to any \hat{q} measure of quantiles is at most $\hat{X}(\hat{q})$ for all $\hat{q} \in [0, 1)$ and exactly equal to $\hat{X}(1)$ at $\hat{q} = 1$. In particular, the problem is more restrictive than requiring an ex ante probability of service $\hat{X}(1)$.

Since an incentive compatible allocation rule is monotone, meaning stronger (i.e., lower) quantiles receive no lower probability of allocation than weaker quantiles, the only set of measure \hat{q} for which the constraint $\hat{X}(\hat{q})$ on allocation probability may be tight is the strongest \hat{q} measure

of quantiles, i.e., $[0, \hat{q}]$. For allocation rule $x(\cdot)$ the probability of allocation to the strongest \hat{q} measure of agents is exactly $X(\hat{q}) = \int_0^{\hat{q}} x(q) dq$. We refer to $X(\cdot)$ as the *cumulative allocation rule*. Thus, the allocation constraint is exactly, $X(\hat{q}) \leq \hat{X}(\hat{q})$ for all $\hat{q} \in [0, 1]$, with equality for $\hat{q} = 1$.

Notice that allocation rule $\hat{x} := \frac{d}{dq} \hat{X}$ satisfies the constraints given by \hat{X} with equality everywhere. Among allocation rules that satisfy \hat{X} as a constraint, \hat{x} has the highest probability on stronger quantiles. The allocation constraint \hat{X} is met by any allocation rule x that relatively has allocation probability shifted from stronger quantiles to weaker quantiles. In this case, we say that the allocation $\hat{x} = \frac{d}{dq} \hat{X}$ *majorizes* $x = \frac{d}{dq} X$.

Definition 5. We say an allocation rule x is *weaker* than another allocation rule \hat{x} if \hat{x} majorizes x . $\text{Rev}[\hat{x}]$ is defined for all allocation constraints \hat{x} as the revenue of the *optimal lottery pricing* with allocation rule weaker than \hat{x} .⁸

We will next explore the connection between the two versions of constrained lottery pricing problems of $R(\hat{q})$, the ex ante lottery pricing with constraint \hat{q} (Definition 4), and $\text{Rev}[\hat{x}]$, lottery pricing with allocation rules weaker than \hat{x} (Definition 5). Consider the allocation of posting a price $V(\hat{q})$. The allocation rule, denoted $\hat{x}^{\hat{q}}$, is the reverse step function that is one on quantiles $[0, \hat{q}]$ and then zero on $(\hat{q}, 1]$. We will argue that $\text{Rev}[\hat{x}^{\hat{q}}] = R(\hat{q})$. We do this by arguing that any allocation rule that is feasible for one problem is also feasible for the other. Any allocation x that is weaker than $\hat{x}^{\hat{q}}$ must by definition satisfy $X(1) = \hat{X}^{\hat{q}}(1) = \hat{q}$, and therefore has ex ante allocation probability \hat{q} . Conversely, any allocation x with ex ante probability of allocation \hat{q} satisfies $X(q) \leq \hat{X}^{\hat{q}}(q)$ for all q , and therefore is weaker than $\hat{x}^{\hat{q}}$ (see Figure 2.5). We will next

⁸From the agent's perspective in a multi-agent mechanism, the allocation constraint \hat{x} is applied at the interim stage of the mechanism, i.e., when the agent knows her own type but considers the types of other agents to be drawn from their respective distributions.

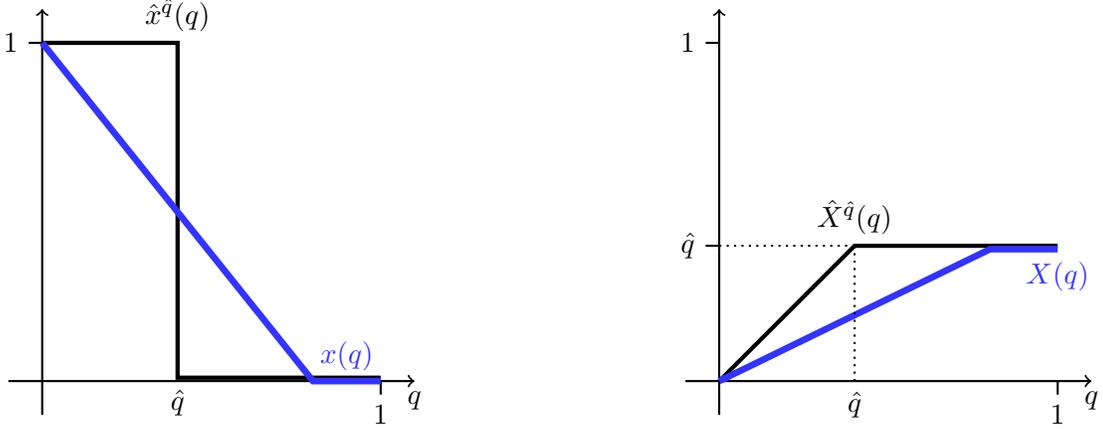


Figure 2.5. The equivalence of two constrained problems $\text{Rev}[\hat{x}^{\hat{q}}]$ and $R(\hat{q})$. Since any allocation x weaker than $\hat{x}^{\hat{q}}$ must satisfy $X(1) = \hat{X}^{\hat{q}}(1) = \hat{q}$, it will be feasible under ex-ante constraint \hat{q} . Conversely, since any cumulative allocation rule X is monotone non-decreasing and satisfies $\frac{d}{dq}X(q) \leq 1$, $X(1) = \hat{X}^{\hat{q}}(1)$ implies that $X(q) \leq \hat{X}^{\hat{q}}(q)$ everywhere and therefore any allocation that satisfies ex-ante constraint \hat{q} is also weaker than $\hat{x}^{\hat{q}}$.

extend the connection between $\text{Rev}[\cdot]$ and $R(\cdot)$ to general allocations beyond step functions $\hat{x}^{\hat{q}}$.

We start by defining the key concept of *revenue-linearity*.

Definition 6. An agent is *revenue linear* if $\text{Rev}[\cdot]$ is a linear functional, i.e., if the optimal revenue for allocation constraints $\hat{x} = \hat{x}^A + \hat{x}^B$ is $\text{Rev}[\hat{x}] = \text{Rev}[\hat{x}^A] + \text{Rev}[\hat{x}^B]$.

The following lemma extends the definition naturally to the case of decomposing \hat{x} into more than two functions. Informally, it states that under revenue linearity, the $\text{Rev}[\cdot]$ of a convex combination of a family of functions is equal to the convex combination of $\text{Rev}[\cdot]$ of those functions.

Lemma 8. Let \hat{x} be the convex combination of functions \hat{x}^α , where α is drawn from a distribution G , that is $\hat{x} = \mathbf{E}_{\alpha \sim G}[\hat{x}^\alpha]$. Revenue linearity implies that $\text{Rev}[\hat{x}] = \mathbf{E}_{\alpha \sim G}[\text{Rev}[\hat{x}^\alpha]]$.

Consider any \hat{x} that is a monotone non-increasing function. We can use reverse step functions to provide a basis for such a function, that is, view \hat{x} as a convex combination of reverse step

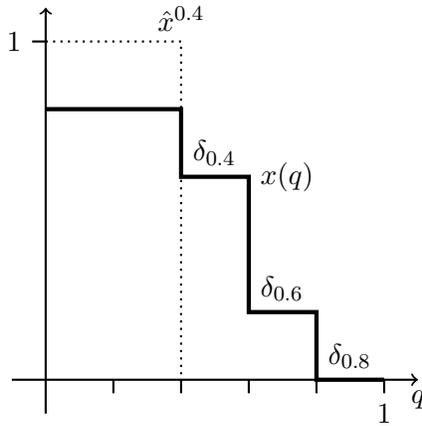


Figure 2.6. Function \hat{x} can be decomposed as $\delta_{0.4}\hat{x}^{0.4} + \delta_{0.6}\hat{x}^{0.6} + \delta_{0.8}\hat{x}^{0.8}$. In the limit, any function \hat{x} can be written as the convex combination of step functions where $\hat{x}^{\hat{q}}$ is sampled with probability $-\frac{d}{dq}\hat{x}(q)$.

functions. This convex combination can be sampled from by drawing \hat{q} at random from the distribution $G^{\hat{x}}$ with density $-\hat{x}'(q) = -\frac{d}{dq}\hat{x}(q)$ (see Figure 2.6). That is,

$$\hat{x} = \mathbf{E}_{\hat{q} \sim G^{\hat{x}}} \left[\hat{x}^{\hat{q}}(\cdot) \right].$$

We can now invoke Lemma 8 to conclude that under revenue linearity,

$$\begin{aligned} \text{Rev}[\hat{x}] &= \mathbf{E}_{\hat{q} \sim G^{\hat{x}}} \left[\text{Rev}[\hat{x}^{\hat{q}}] \right] \\ &= -\mathbf{E}_{\hat{q}} \left[\hat{x}'(\hat{q}) \text{Rev}[\hat{x}^{\hat{q}}] \right]. \end{aligned}$$

Recall that by definition, $R(\hat{q}) = \text{Rev}[\hat{x}^{\hat{q}}]$. Therefore, revenue-linearity implies that,

$$\text{Rev}[\hat{x}] = -\mathbf{E}_{\hat{q}}[\hat{x}'(\hat{q})R(\hat{q})].$$

Now using integration by parts,

$$\begin{aligned} \text{Rev}[\hat{x}] &= -\mathbf{E}_{\hat{q}}[\hat{x}'(\hat{q})R(\hat{q})] \\ &= [-\hat{x}(\hat{q})R(\hat{q})]_0^1 + \mathbf{E}_{\hat{q}}[R'(\hat{q})\hat{x}(\hat{q})] \\ &= \mathbf{E}_{\hat{q}}[R'(\hat{q})\hat{x}(\hat{q})]. \end{aligned}$$

The second equality follows from integration by parts and the third equality from $R(1) = R(0) = 0$. This construction motivates the following definition.

Definition 7. The *marginal revenue* for an agent with quantile q is $R'(q) = \frac{d}{dq}R(q)$; the marginal revenue for an allocation constraint \hat{x} is $\text{MR}[\hat{x}] = \mathbf{E}_q[R'(q)\hat{x}(q)]$.

The definition of revenue linearity and the above discussion immediately imply the following theorem.

Theorem 1. *For a revenue-linear agent, the optimal revenue for an allocation constraint is equal to its marginal revenue, i.e., for all \hat{x} , $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

Notice that $\text{Rev}[\hat{x}^{\hat{q}}] = R(\hat{q})$, which we previously showed to hold in general, is a special case of the above theorem. We will next show that the linear single-parameter settings are revenue linear. The only property of the linear single-parameter settings used in the proof is the existence of a strong amortization of revenue (see Alaei et al. (2013) for an alternative proof

of linearity of the linear single-parameter setting). We start by the next lemma which follows directly from the definition of majorization.

Lemma 9. *Consider allocation rules x^A and x^B that are weaker than \hat{x}^A and \hat{x}^B , respectively. Then $x^A + x^B$ is weaker than $\hat{x}^A + \hat{x}^B$. Conversely, for any x weaker than $\hat{x}^A + \hat{x}^B$, there exist x^A and x^B such that $x = x^A + x^B$, and that x^A and x^B are weaker than \hat{x}^A and \hat{x}^B , respectively.*

Proof. The first part follows directly from definitions, as for all \hat{q} , $X^A(q) + X^B(q) \leq \hat{X}^A(q) + \hat{X}^B(q)$.

For the second part, consider x weaker than \hat{x} . By definition, the graph of X is always *below* the graph of \hat{X} , and $X(0) = \hat{X}(0)$ and $X(1) = \hat{X}(1)$. Equivalently, the graph of X is *to the right* of graph \hat{X} (see Figure 2.7). More formally, we can define the *quantile shift operator* $\mathcal{Q} : [0, 1] \rightarrow [0, 1]$, such that $\mathcal{Q}(q) \leq q$, and $\hat{X}(\mathcal{Q}(q)) = X(q)$ (see Figure 2.7). Given \mathcal{Q} we can define x^A and x^B by shifting \hat{x}^A and \hat{x}^B , respectively. That is, $X^A(q) = \hat{X}^A(\mathcal{Q}(q))$ and $X^B(q) = \hat{X}^B(\mathcal{Q}(q))$. Now notice that x^A and x^B are weaker than \hat{x}^A and \hat{x}^B since $\mathcal{Q}(q) \leq q$ and \hat{X}^A and \hat{X}^B are increasing, and also that $X^A(q) + X^B(q) = (\hat{X}^A + \hat{X}^B)(\mathcal{Q}(q)) = \hat{X}(\mathcal{Q}(q)) = X(q)$, which implies that $x^A + x^B = x$. \square

Recall the strong amortization Lemma 7 which stated that in linear single-parameter settings, the revenue of any mechanism with allocation x can be written as $\mathbf{E}_q[x(q)P'(q)]$. The revenue linearity of linear single-parameter agents is a simple consequence of the strong amortization in Lemma 7 and the above lemma. In fact, the functional form of the revenue function $P(q)$ does not matter in the proof. All that matters is that a strong amortization of revenue exists.

Theorem 2. *An agent with single-dimensional linear utility is revenue linear.*

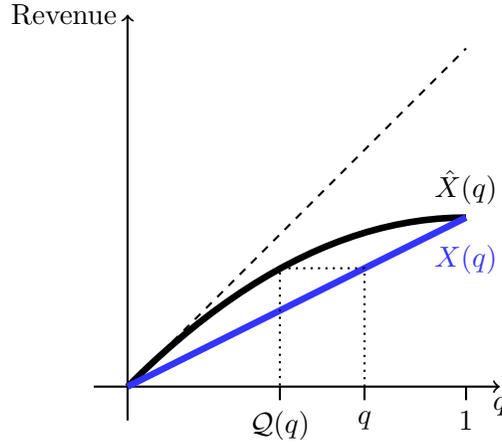


Figure 2.7. The quantile shift operator.

Proof. Consider $\hat{x} = \hat{x}^A + \hat{x}^B$. We first show that $\text{Rev}[\hat{x}] \geq \text{Rev}[\hat{x}^A] + \text{Rev}[\hat{x}^B]$. Consider optimal allocations x^A and x^B subject to \hat{x}^A and \hat{x}^B , respectively, and let p^A and p^B be their price functions. Define mechanism $(x, p) = (x^A + x^B, p^A + p^B)$. The mechanism is incentive compatible and individually rational, its allocation function is feasible subject to \hat{x} by Lemma 9, and its revenue is equal to the sum of the revenues of x^A and x^B .

Conversely, consider the optimal allocation x subject to \hat{x} . By Lemma 9, there exist $x^A \preceq \hat{x}^A$ and $x^B \preceq \hat{x}^B$ such that $x = x^A + x^B$. Again, using Lemma 9 and Lemma 7 we have

$$\begin{aligned} \text{Rev}[\hat{x}] &= \mathbf{E}_q[x(q)P'(q)] \\ &= \mathbf{E}_q[x^A(q)P'(q)] + \mathbf{E}_q[x^B(q)P'(q)] \\ &\leq \text{Rev}[\hat{x}^A] + \text{Rev}[\hat{x}^B]. \end{aligned}$$

□

Corollary 1. *For a single-dimensional linear agent, the optimal revenue for an allocation constraint is equal to its marginal revenue, i.e., for all \hat{x} , $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

Multi-agent Mechanisms. The conclusion of the preceding discussion is that the optimal revenue for any allocation constraint is equal to its marginal revenue. We will now discuss multi-agent mechanisms for agents labeled 1 to n .

Definition 8. Any mechanism and distribution over types induces a profile $\mathbf{x} = (x_1, \dots, x_n)$ of interim allocation rules. The *surplus of marginal revenue* is the sum of the marginal revenues of interim allocation rules of each agent $\sum_i \text{MR}[x_i]$.

Multi-agent mechanism design problems reduce to single-agent lottery pricing problems by the following standard argument. For an agent in the optimal mechanism, her contribution to the revenue is equal to the marginal revenue of her allocation rule (Corollary 1). We thus look for the mechanism that optimizes the surplus of marginal revenue. Consider relaxing the incentive constraints (namely: monotonicity of the allocation rule) and optimizing marginal revenue pointwise. Specifically, when the agent quantiles are $\mathbf{q} = (q_1, \dots, q_n)$ select the allocation $\mathbf{x} = (x_1, \dots, x_n)$ to maximize the surplus of marginal revenue $\sum_i R'_i(q_i) x_i$ subject to feasibility of \mathbf{x} (e.g., for a single-item auction, serve the agent with the highest positive marginal revenue, or none if the marginal revenues are all negative). We call this mechanism the *marginal revenue mechanism*. Now check that the previously relaxed incentive constraints are not violated. Notice that since revenue curves are concave, the marginal revenues are monotone non-increasing in quantile, for any agent a stronger (lower) quantile corresponds to a weakly higher marginal revenue, and so the induced allocation rule is monotone. Furthermore, as these allocations optimize marginal revenue pointwise for all profiles of agent quantiles, they certainly also maximize marginal revenue in expectation over the agent quantiles.

Theorem 3. *The marginal revenue mechanism is revenue optimal for single-dimensional linear agents.*

Proof. The optimal mechanism induces some profile \mathbf{x} of interim allocation rules. By revenue linearity, the expected revenue of this profile of interim allocation rules is equal to its surplus of marginal revenue. The marginal revenue mechanism selects its outcome to optimize surplus of marginal revenue pointwise for the feasibility constraint. Its expected surplus of marginal revenue is, thus, at least that of the optimal mechanism. \square

2.2.1. Service-constrained Extension: The Marginal Revenue Mechanism

To study the marginal revenue mechanism in service-constrained environments, the first challenge is, naturally, finding the right definition for the marginal revenue mechanism. A key step in the definition of the marginal revenue mechanism is finding the right mapping between types to quantiles. In the linear single-parameter setting, the mapping is clear because types are naturally ordered based on their *power*, their willingness to pay. For general types, such a ranking no longer clearly exists. For example, for the case of selling two items, how should we compare two types, one type valuing each item by 1, and the other one valuing item one by 0 and item two by 2? The answer to this question lies in the solution to the single-agent interim-constrained revenue maximization problem Definition 5. We show that for revenue-linear service-constrained settings, the solution to interim-constrained revenue maximization problem ranks types and allocates higher probability of service to types ranked higher. In general, for non-revenue-linear problems, a randomized ranking, and therefore a randomized mapping of types to quantiles, can be *read* from the way the single-agent interim-constrained solution allocates probabilities of allocation to types.

Clearly, once types are matched to quantile (sometimes at random), the type with the highest marginal revenue *wins the service*. One more question should be answered though: what should the attributes of the service be? For example, for the ad-auctions problem, what cookies should be transferred together with the ad slot? For the inter-temporal mechanism design problem,

what day should we allocate the item to the winner? Again, this question can be answered by solving the single-agent interim-constrained revenue maximization problem. The winner will be offered a menu of options, that are determined based on its *critical quantile*, the maximum winning quantile that it could have had.

The proof of optimality of the marginal revenue mechanism follows easily, basically from the outline of the proof in this chapter, once the right definitions for the generalization of the mechanism to service-constrained environments is figured out. We show that, in addition to describing the optimal multi-agent mechanism, our generalization can be used to apply many results from the linear single-parameter theory directly to service-constrained settings.

2.3. Interim Feasibility and Computation

This section discusses the outline of Chapter 5 applied to the linear single-parameter setting. The main technical insight from that chapter is reducing the size of the optimization problem by an interim representation of the mechanism, since the interim mechanism of each agent is a function of that agent's type only, whereas to express the mechanism ex-post, we need to describe outcomes for every possible profile of types. As we have seen already, both the expected revenue of the mechanism, as well as interim incentive compatibility of the mechanism, can be expressed and verified using interim mechanisms. The main challenge is to make sure that the mechanism represented by a profile of interim allocation rules is feasible, that is, it is indeed a profile of interim allocation rules induced from a feasible multi-agent mechanism (see definitions of induced interim mechanisms in equations 2.4 and 2.5). Border (1991) first characterized the space of feasible interim allocation rules. We extend that characterization to k -unit as well as more general feasibility settings. More importantly, the characterization provides us with insights that enable us to solve the interim optimization problem in polynomial time, and also map the solution back to an ex-post expression of the optimal mechanism in a computationally

efficient manner. The rest of this section highlights this approach for the case of the single-item linear single-parameter setting.

A profile of interim allocation rules $\mathbf{x} = (x_1, \dots, x_n)$ is *feasible* if there exists a feasible multi-agent mechanism $\bar{\mathbf{x}}$ that induces interim mechanism x_i for each agent i . Let \mathbb{X} denote the space of all feasible interim allocation rules $\mathbf{x} = (x_1, \dots, x_n)$. The following program upper bounds the optimal revenue

$$(2.9) \quad \max_{\mathbf{x} \in \mathbb{X}} : \sum_i \text{Rev}_i[x_i].$$

Consider the allocation rule of the optimal mechanism $\mathbf{x} \in \mathbb{X}$. By definition, the revenue of the mechanism is at most $\sum_i \text{Rev}_i[x_i]$, implying that the above program upper bounds optimal revenue. In fact, it can be shown that the solution to the program is exactly equal to optimal revenue. To see this, we define a multi-agent mechanism $(\bar{\mathbf{x}}, \bar{\mathbf{p}})$ and show that the revenue of the mechanism $(\bar{\mathbf{x}}, \bar{\mathbf{p}})$ is equal to the solution of the program. Let \mathbf{x} be the optimizer of the program, and let \mathbf{p} be its payment rule (remember that the payment rule is uniquely defined given \mathbf{x}). By definition of \mathbb{X} , there must exist an ex post allocation rule $\bar{\mathbf{x}}$ from which the interim allocation rule \mathbf{x} is derived. Define the multi-agent payment rule $\bar{p}_i(\mathbf{v}) = p_i(v_i)/x_i(v_i)$. The interim allocation and payment rule induced by this multi-agent mechanism, $(\bar{\mathbf{x}}, \bar{\mathbf{p}})$, is exactly (x_i, p_i) for each agent i .

The above argument suggests an approach to compute the optimal mechanism:

- (1) Compute optimal interim allocation and payments as a solution to program 2.9.
- (2) Compute a multi-agent allocation inducing the interim allocations.
- (3) Define multi-agent payments given the above formula.

If the first two steps can be performed in polynomial time, the approach would result in polynomial-time computation of the optimal mechanism. In what follows, we discuss the conditions implying polynomial time computability of the first two steps.

- Even though the number of constraints required to represent \mathbb{X} is generally exponential in the size of the type space, the theory of combinatorial optimization provides us with conditions implying the existence of computationally efficient solutions to program 2.9. In particular, if the objective function of a program is concave and it can be computed in polynomial time, and the feasibility polytope is a *polymatroid* (technical definition deferred), then the solution can be computed efficiently. We show that the objective function is concave in general. In fact, as shown above, the objective function is linear in the linear single-parameter settings. In addition, we provide a characterization of space \mathbb{X} which implies that \mathbb{X} is a polymatroid.
- The characterization of \mathbb{X} as a polymatroid proves useful in the problem of implementing the computed interim allocation rules with a multi-agent mechanism. In particular, there is a one-to-one correspondence between a simple class of ex post allocation rules, and the *vertices* of the polymatroid representing implementable interim allocation rules. The ex post allocation rule is parameterized by a complete order over a subset of type space of all agents, and given a profile of types of agents allocates to the item to the type that is highest in the order (the mechanism does not always allocate the item since not all types need to be included in the ranking). We show that we can express any feasible interim allocation rule as a convex combination of vertices of the polymatroid in a computationally efficient way. As a result, the optimal mechanism can be computed efficiently as a sample from the distribution over mechanisms that correspond to the vertices of the polymatroid.

2.3.1. Service-constrained Extension: Interim Feasibility and Efficient Computation

The approach in Section 2.3 minimally depended on the details of the linear single-parameter setting. We used the properties of the linear single-parameter setting to argue that the revenue function is concave. Proving concavity for general service-constrained problem is not much more complicated. Other than that, to ensure a computationally efficient solution using the above approach, the only property required from the single-agent problem is that the interim-constrained single-agent problem is solvable in a computationally efficient manner. Our general characterization of the feasible interim allocation rules is completely independent of the space of preferences, and therefore does not depend on the details of the linear single-parameter setting.

CHAPTER 3

The Single Agent Problem and Incentives

Overview and Organization. This chapter designs a framework to prove optimality of auctions by focusing on the single-agent problem. We give a framework for *reverse mechanism design*. Instead of solving for the optimal mechanism in general, we assume a (natural) specific form of the mechanism. As an example of the framework, for agents with unit-demand preferences, we restrict attention to mechanisms that sell each agent her favorite item or nothing. From this restricted form, we will derive multi-dimensional virtual values. These virtual values prove this form of mechanism is optimal for a large class of item-symmetric distributions over types. As another example of our framework, for bidders with additive preferences, we derive conditions for the optimality of posting a single price for the grand bundle. Our framework generalizes Myerson's design of virtual values for linear single-parameter settings, and can be used to prove optimality in multi-agent settings.

This chapter can be read independently of the background, Myerson's derivation of virtual values in Section 2.1, even though reading the background can simplify reading the material of this chapter. A limited statement of main results and an overview of related work is given in Section 3.1. The single-agent problem formulation appears in Section 3.2. The general framework to design virtual values is discussed in Section 3.3. We use the framework to provide conditions sufficient for optimality of uniform pricing for unit demand preferences in Section 3.4, and for optimality of bundle pricing for additive preferences in Section 3.5. Finally, in Section 3.6 we show how the analysis can be extended to multi-agent problems.

3.1. Introduction

Optimal mechanisms for agents with multi-dimensional preferences are generally complex. This complexity makes them challenging to solve for and impractical to run. In a typical mechanism design approach, a model is posited and then the optimal mechanism is designed for the model. Successful mechanism design gives mechanisms that one could at least imagine running. By this measure, multi-dimensional mechanism design has had only limited success. In this chapter we take the opposite approach, which we term *reverse mechanism design*. We start by imagining a restriction on mechanisms that would make them simple and reasonable to run, then we solve for sufficient conditions for a restricted mechanism to be optimal (among all mechanisms). Our approach is successful if the conditions under which restricted mechanisms are optimal are broad and representative of relevant settings.

This chapter has two main contributions. The first is in codifying the *method of virtual values* from single-dimensional auction theory and extending it to agents with multi-dimensional preferences. The second is in applying this method to two paradigmatic classes of multi-dimensional preferences. The first class is unit-demand preferences (e.g., a homebuyer who wishes to buy at most one house); for this class we give sufficient conditions under which the posting a uniform price for all items is optimal. This result generalizes one of Chapter 4 for a consumer with values uniform on interval $[0, 1]$, and contrasts with an example of Thanassoulis (2004) for a consumer with values uniform on interval $[5, 6]$ where uniform pricing is not optimal. The second class is additive preferences, for this class we give sufficient conditions under which the posting a price for the grand bundle is optimal. This result generalizes a recent result of Hart and Nisan (2013) and relates to work of Armstrong (1999). Similarly to an approach of Chapter 4, these results for single-agent pricing problems described above can be generalized naturally to multi-agent auction problems.

Myerson's (1981) characterization of revenue optimal auctions for single-dimensional agents is the cornerstone of modern auction theory and mechanism design. This characterization is successful in describing simple and practical mechanisms in simple environments where the agents preferences are independent and identically distributed according to a well-behaved distribution. In this case, the optimal auction is reserve price based. Myerson's characterization is also successful in describing the complex optimal mechanism for agents with preferences that are non-identically distributed or distributed according to an ill-behaved distribution. However, due to this complexity, the resulting mechanism has limited application. The consequence of our work is similar in that we characterize simple optimal mechanisms for well-behaved preferences; but distinct in that it does not characterize optimal mechanisms beyond the class of well-behaved preferences.

Recall from Section 2.1 that Myerson's approach is based on mapping agent values to appropriately defined virtual values and then optimizing the *virtual surplus*, i.e., the sum of the virtual values of agents served. Importantly, this approach replaces the global objective optimizing revenue in expectation over the distribution of agent values to the pointwise objective of optimizing virtual surplus on each profile of agent values. Furthermore, virtual surplus maximization leads to a simple and practical optimal mechanism in many environments. The simplicity of analysis by virtual values and of mechanisms resulting from optimizing virtual values has led to a rich theory single-dimensional auction theory. Our multi-dimensional virtual values similarly give a pointwise objective and their optimization results in simple optimal mechanisms.

Myerson solved for revenue optimal auctions for agents with single-dimensional preferences. He gives a definition of virtual values, claims that optimization of virtual surplus gives the optimal auction, and proves this claim by utilizing two properties of the virtual values. The following two properties collectively imply that the mechanism that optimizes virtual values is

indeed the optimal mechanism.¹ A *virtual value function* maps values pointwise to virtual values satisfying two properties (Definition 1):

- The pointwise optimization of virtual values gives an assignment rule that is *incentive compatible*. That is, there exist payments for this assignment rule that induce an agent to truthfully report his value.
- The virtual values are an *amortization* of the revenue. That is, the expected sum of the virtual values of the winners of any incentive compatible mechanism is equal to the expected revenue of that mechanism.

This definition of virtual value functions gives a roadmap to identifying the optimal mechanism: find a virtual value function (that satisfies the two conditions) and run the mechanism that maximizes virtual value pointwise. The identification of a virtual value function reduces the problem of optimization of the expected revenue (a global quantity) to the optimization of virtual surplus (a pointwise quantity).

The Challenge of Multi-dimensional Preferences. As described above, if a virtual value function that satisfies properties of incentive compatibility and amortization conditions can be identified then the optimal mechanism design problem is solved. For agents with single-dimensional preferences the function that satisfies the amortization property is unique and can be derived by a simple exercise. Omitting the details: uniqueness follows because there is only one path on the line between any value and the origin (and can be found via integration by parts). It then remains to check the incentive compatibility property, i.e., that pointwise optimization of virtual value is incentive compatible; by standard characterizations of incentive compatible mechanisms, this is a simple task as well. Multi-dimensional preferences are challenging because the amortization property does not uniquely pin down the virtual value function. Omitting

¹Subsequently, in Section 3.3, we will define a more permissive version of the amortization property which corresponds to the *ironed virtual values* of Myerson (1981). The simpler definition here will, nonetheless, be sufficient for our introductory discussion.

the details: non-uniqueness follows because from any point in a multi-dimensional space there are many paths between the point and the origin. This difficulty has prevented the design of mechanisms for multi-dimensional agents that follows the virtual-value-based approach.

Multi-dimensional Virtual Values. To resolve the non-uniqueness of functions that satisfy the amortization property we consider additional constraints that the optimality of a restricted mechanism would place on virtual values.

We walk through this approach for the example of a unit-demand agent and the restriction to mechanisms that post a uniform price for each item. Each item is assigned a virtual value function that maps the type of the agent (the values for all items) to a real number. On one hand, under uniform pricing, the agent will always choose to buy his favorite item, or no item if all values are below the price. On the other hand, a mechanism that optimizes a multi-dimensional virtual value (for a unit-demand agent) would serve the agent the item he has the highest positive virtual value for, or no item if all virtual values are negative. Synthesizing these constraints, the following conditions are sufficient for virtual surplus maximization to imply optimality of uniform pricing.

- The virtual value function is a *single-dimensional projection* if the virtual value for the favorite item corresponds to the single-dimensional virtual value for the distribution of the value for the favorite item (the distribution of the maximum value).²
- The virtual value function is *consistent with uniform pricing* if there is a price such that (a) when the value for the favorite item exceeds the price the virtual value for the favorite item is non-negative and at least the virtual value of any other items and (b) when the value for the favorite item is below the price both virtual values are non-positive.

²Rationale: The mechanism has effectively projected the agent's multi-dimensional preference onto a single dimension. In this single dimension the unique function that satisfies the amortization property is the one given by the single-dimensional virtual values of Myerson (1981)

Any virtual value function that satisfies the consistency-with-uniform-pricing conditions can be easily seen to satisfy the incentive compatibility requirement. Thus, the identification of a virtual value function becomes one of simultaneously resolving the three conditions of amortization, single-dimensional projection, and consistency with uniform pricing (which implies incentive compatibility). Conditions on the distributions over values that guarantee the existence of such a virtual value function are sufficient for the optimality of uniform pricing.

For two-dimensional preferences (i.e., the number of items is equal to two) the amortization and single-dimensional-projection restrictions pin down a two-dimensional virtual value function uniquely. Specifically, as the virtual value for the favorite item is fixed by the restriction, only the virtual value for the other item must be determined. Essentially, we are left with a single-dimensional problem and in a single dimension the function that satisfies the amortization condition are unique. Our task is then to give sufficient conditions under which this unique virtual value function is consistent with uniform pricing. For higher-dimensional preferences, virtual value functions for the non-favorite items are not pinned down. Instead of deriving a formula for them, we generalize the sufficient condition from the two-dimensional case and prove that there exists such a virtual value function that satisfies all three constraints.

The restriction of selling a unit-demand agent his favorite item provides our main example of the framework of reverse mechanism design. We will also apply the framework to agents with additive preferences and give sufficient conditions for pricing the grand bundle to be optimal. The techniques we develop can be similarly applied to other environments and appropriate restrictions, and these extensions are an important topic for future work.

Sufficient Conditions for Optimality of Simple Mechanisms. The opening example of this introduction of selling one of many houses to a homebuyer with value distributed independently, identically, and uniformly from $[0, 1]$ satisfies the conditions for optimality of the restriction to selling him only his favorite item. More generally, our framework identifies

sufficient conditions for the distribution over agent values. Importantly these conditions allow for positive correlation between the agent's value for distinct items, a.k.a, *affiliation*. Such correlation is natural when values correlate, e.g., with initial wealth. Below we state the sufficient conditions for optimality of uniform pricing for unit-demand agents and grand-bundle pricing for additive agents for the special case of two items; the general multi-item statements are deferred to later in the chapter.

A single-agent problem with unit-demand preferences is defined with a density function f on a bounded set of types normalized to $T = [0, 1]^2$, where (t_1, t_2) is the pair of values for the two items. We consider only symmetric distributions and therefore only define f on the subset of T satisfying $t_1 \geq t_2$, that is, where the first item is the favorite item. A uniform posted price mechanism posts a uniform price for each item, and lets the agent choose one of the two items or nothing based on his type. It is easiest to state the result in the following *max-ratio* representation of the distribution. The max-ratio representation of a density function f is the function f^{MR} that is defined by

$$f(t_1, t_2) = f^{MR}(t_1, t_2/t_1).$$

The following theorem states that uniform pricing is optimal if f^{MR} satisfies a supermodularity condition. Intuitively, this supermodularity implies that a higher value of t_1 signals a higher value of t_2/t_1 (see Figure 3.1).

Theorem 4. *Posting a uniform price for each item is optimal if the max-ratio representation of the density function is log-supermodular, that is,*

$$f^{MR}(t_1, \theta) \times f^{MR}(t_1, \theta') \leq f^{MR}(t_1, \theta') \times f^{MR}(t'_1, \theta), \quad t_1 \leq t'_1, \theta \geq \theta'.$$

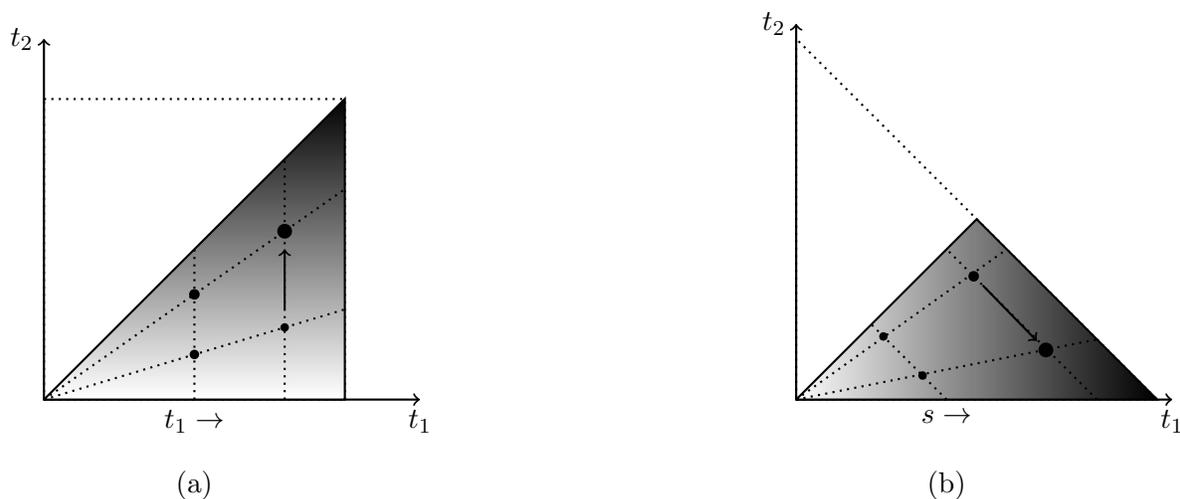


Figure 3.1. a) A distribution satisfying the supermodularity condition of Theorem 4. Dark color indicates relatively high mass. Roughly speaking, it states that the probability mass shifts upwards as the value for the favorite item, t_1 , increases. b) A distribution satisfying the submodularity condition of Theorem 5. It implies that probability mass shifts downwards as the sum of the values increases.

Example 5 (Inter-temporal Pricing with Privately Known Discount Rate). Consider the problem of inter-temporal pricing when the seller has commitment power: the seller owns an item that can be sold today or tomorrow. The buyer's value v for receiving the item today, and $v \times \delta$, where $\delta \leq 1$ for receiving the item tomorrow. The seller can commit to a selling mechanism in advance, that could involve lotteries over the allocation of the item. Since the value for the item tomorrow is less than the value for the item today, a natural guess for the optimal mechanism is that it only allocates the item today. In other words, it posts a uniform price for both days, and the buyer will naturally select the item today (if at all). When δ is known to the seller, this intuition is proved correct in several related models such as Stokey (1979) and Acquisti and Varian (2005). For our model, this can be proved by simply applying Myerson's theorem, and noting that the virtual value for selling the item today multiplied by δ is equal to the virtual value for selling the item tomorrow.

When δ is only privately known to the buyer, however, uniform pricing is no longer generally optimal (Thanassoulis (2004) can be easily adopted to show this for our asymmetric problem). A direct application of our main theorem states that uniform pricing is optimal when δ and v are affiliated, which includes but is more general than the case where δ and v are independent. For instance, when v is a uniform draw on the interval $[0, 1]$ and is independent of δ the two dimensional the virtual value function that satisfies the amortization and single-dimensional-projection constraints is pinned down as $2v - 1$ for the item today and $\delta(2v - 1)$ for the item tomorrow. For $\delta \in [0, 1]$ as defined, these virtual values satisfy are consistent with uniform pricing. Therefore, the existence of this virtual value function proves that uniform pricing is revenue optimal among all mechanisms.

For other i.i.d. distributions, the virtual values that amortization and single-dimensional-projection pin down do not generally satisfy consistency with uniform pricing (e.g., when values are i.i.d. $U[5, 6]$ the optimal mechanism does not sell the agent his favorite item); however, we show that consistency with uniform pricing is satisfied by for multi-dimensional values that are i.i.d. from any distribution that satisfies a geometric convexity property.

As the second main example of the application of our framework, we provide sufficient conditions for optimality of grand bundle pricing for agents with additive preferences. The *sum-ratio* representation of a density function f is the function f^{SR} that is defined by

$$f(t_1, t_2) = f^{SR}(t_1 + t_2, t_2/t_1).$$

Theorem 5. *Posting a price for the bundle of items is optimal if the sum-ratio representation of the density function is log-submodular, that is,*

$$f^{SR}(s, \theta) \times f^{SR}(s', \theta') \geq f^{SR}(s, \theta') \times f^{SR}(s', \theta), \quad s \leq s', \theta \geq \theta'.$$

Related Work. The starting point of work in multi-dimensional optimal mechanism design is the observation that an agent's utility must be a convex function of his private type (e.g., Rochet, 1985, cf. the envelope theorem). The second step is in writing revenue as the difference between the surplus of the mechanism and the agent's utility (e.g., McAfee and McMillan, 1988; Armstrong, 1996). The surplus can be expressed in terms of the gradient of the utility. The third step is in rewriting the objective in terms of either the utility (e.g., McAfee and McMillan, 1988; Manelli and Vincent, 2006; Hart and Nisan, 2013; Daskalakis et al., 2013; Wang and Tang, 2014; Giannakopoulos and Koutsoupias, 2014) or in terms of the gradient of the utility (e.g., Armstrong, 1996; Chapter 4; and this chapter). This manipulation follows from an integration by parts. The first category of papers (rewriting objective in terms of utility) performs the integration by parts independently in each dimension, and the second category (rewriting objective in terms of gradient of utility, except for ours) does the integration along rays from the origin (see below). In our approach, in contrast, the integration by parts is performed in general and is dependent on the distribution and the form of the mechanism we wish to show is optimal.

Closest to this chapter is Armstrong (1996) which uses integration by parts along paths that connect types with straight lines to the zero type (which has value zero for any assignment) to define virtual values. Armstrong (1996) finds properties on valuation functions (beyond linear ones considered in our work) and distributions, that when jointly satisfied, imply that the point-wise optimization of virtual values results in an incentive compatible mechanism. Armstrong gives some examples of mechanisms that result from this approach but does not generally interpret the form of the resulting mechanisms. Armstrong suggests generalizing his approach from rays from the origin to other kinds of paths; our approach, in contrast, proves the existence of

appropriate paths over which to integrate without requiring the form of the path to be specified in advance. When Armstrong's condition on the distribution is satisfied (which we refer to as independence in max-ratio coordinates), our solution is also equivalent to an integration along rays.

There has been work looking at properties of single-agent mechanism design problems that are sufficient for optimal mechanisms to make only limited use of randomization. For context, the optimal single-item mechanism is always deterministic (e.g., Myerson, 1981; Riley and Zeckhauser, 1983), while the optimal multi-item mechanism is sometimes randomized (e.g., Thanassoulis, 2004; Pycia, 2006). For agents with additive preferences across multiple items, McAfee and McMillan (1988), Manelli and Vincent (2006), and Giannakopoulos and Koutsoupas (2014) find sufficient conditions under which deterministic mechanisms, i.e., bundle pricings, are optimal. Pavlov (2011) considers more general preferences and a more general condition; for unit-demand preferences, this condition implies that in the optimal mechanism an agent deterministically receives an item or not, though the item received may be randomized.

A number of papers consider the question of finding closed forms for the optimal mechanism for an agent with additive preferences and independent values across the items. One such closed form is grand-bundle pricing. Our work for additive preferences contrasts in that we restrict to grand-bundle pricing and a particular family of correlated distributions. For the two item case, Hart and Nisan (2013) give sufficient conditions for the optimality of grand-bundle pricing; these conditions are further generalized by Wang and Tang (2014). Their results are not directly comparable to ours as our results apply to correlated distributions. Daskalakis et al. (2013) and Giannakopoulos and Koutsoupas (2014) give frameworks, similar to ours, for proving optimality of multi-dimensional mechanisms. Daskalakis et al. (2013) use their framework to find the close form for the optimal mechanism for several special cases; Giannakopoulos and Koutsoupas

(2014) give a closed form for the optimal mechanism when values are i.i.d. from the uniform distribution (with up to six items).

3.2. Preliminaries

3.2.1. The Setting

An agent is specified by a bounded set of possible types T normalized to be $T = [0, 1]^m$, where each $\mathbf{t} = (t_1, \dots, t_m) \in T$ is an m -dimensional vector of values for m items.³ The type of the agent is drawn from a known distribution with density f . For the special case that the type space is single-dimensional (i.e., $m = 1$), the cumulative distribution function of the type is denoted by F . We do not require that the values for items be drawn independently. We consider unit-demand and additive agents. For unit-demand agents the assignment $\boldsymbol{\pi} \in [0, 1]^m$ must satisfy $\sum_i \pi_i \leq 1$; for additive agents $\boldsymbol{\pi}$ must satisfy $\pi_i \leq 1$ for all i . The utility of the agent with type \mathbf{t} for assignment $\boldsymbol{\pi} \in [0, 1]^m$ and payment $p \in \mathbb{R}$ is $\mathbf{t} \cdot \boldsymbol{\pi} - p$.

A single-agent mechanism is a pair of functions, the assignment function $\boldsymbol{\pi} : T \rightarrow [0, 1]^m$ and the payment function $p : T \rightarrow \mathbb{R}$. A mechanism is *individually rational* if the utility of every type of the agent is at least zero,

$$\mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - p(\mathbf{t}) \geq 0, \quad \forall \mathbf{t}.$$

A mechanism is *incentive compatible* if no type of the agent increases his utility by misreporting,

³Recall that we maintain the convention of denoting a vector \mathbf{v} by a bold symbol and each of its components v_i by a non-bold symbol.

$$\mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - p(\mathbf{t}) \geq \mathbf{t} \cdot \boldsymbol{\pi}(\hat{\mathbf{t}}) - p(\hat{\mathbf{t}}), \quad \forall \mathbf{t}, \hat{\mathbf{t}}.$$

3.2.2. Multivariable Calculus Notation

For a function $h : \mathbb{R}^k \rightarrow \mathbb{R}$, we use $\partial_j h : \mathbb{R}^k \rightarrow \mathbb{R}$ to denote the *partial derivative* of function h with respect to its j 'th variable. The *gradient* of h is a vector field, denoted by $\nabla h : \mathbb{R}^k \rightarrow \mathbb{R}^k$, defined to be $\nabla h = (\partial_1 h, \dots, \partial_k h)$. The *divergence* of a vector field $\boldsymbol{\alpha} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is denoted by $\nabla \cdot \boldsymbol{\alpha} : \mathbb{R}^k \rightarrow \mathbb{R}$ and is defined to be

$$\nabla \cdot \boldsymbol{\alpha} = \partial_1 \alpha_1 + \dots + \partial_k \alpha_k.$$

We denote the integral of function $h : \mathbb{R}^k \rightarrow \mathbb{R}$ over a subset T of \mathbb{R}^k

$$\int_{\mathbf{t} \in T} h(\mathbf{t}) \, d\mathbf{t}.$$

Let ∂T be the boundary of set T and $\boldsymbol{\eta}(\mathbf{t})$ be the outward-pointing unit normal vector of T at point \mathbf{t} on ∂T . The multi-variable *integration by parts* for functions $h : \mathbb{R}^k \rightarrow \mathbb{R}$ and $\boldsymbol{\alpha} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is as follows

$$(3.1) \quad \int_{\mathbf{t} \in T} (\nabla h \cdot \boldsymbol{\alpha})(\mathbf{t}) \, d\mathbf{t} = \int_{\mathbf{t} \in \partial T} h(\mathbf{t})(\boldsymbol{\alpha} \cdot \boldsymbol{\eta})(\mathbf{t}) \, d\mathbf{t} - \int_{\mathbf{t} \in T} h(\mathbf{t})(\nabla \cdot \boldsymbol{\alpha})(\mathbf{t}) \, d\mathbf{t},$$

Setting h to be the constant function equal to 1 everywhere gives us the *divergence theorem*

$$\int_{\mathbf{t} \in T} (\nabla \cdot \boldsymbol{\alpha})(\mathbf{t}) \, d\mathbf{t} = \int_{\mathbf{t} \in \partial T} (\boldsymbol{\alpha} \cdot \boldsymbol{\eta})(\mathbf{t}) \, d\mathbf{t}.$$

When the dimension $k = 1$, the integration by parts has the familiar form of (Lemma 2)

$$\int_{x=a}^b h'(x)g(x) dx = h(x)g(x)|_{x=a}^b - \int_{x=a}^b h(x)g'(x) dx,$$

and the divergence theorem is the *fundamental theorem of calculus*.

$$\int_{x=a}^b h'(x) dx = h(b) - h(a).$$

3.2.3. Problem Formulation

A single agent mechanism $(\boldsymbol{\pi}, p)$ defines a utility function $u(\mathbf{t}) = \mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - p(\mathbf{t})$. The following lemma connects the utility function of an IC mechanism with its assignment function.

Lemma 10 (Rochet, 1985). *Function u is the utility function of an agent in an incentive compatible mechanism if and only if u is convex, and in that case, the agent's assignment is $\boldsymbol{\pi}(\mathbf{t}) = \nabla u(\mathbf{t})$.*

Notice that the payment function can be defined using the utility function and the assignment function as $p(\mathbf{t}) = \mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - u(\mathbf{t})$. Applying the above lemma, we can write payment to be $p(\mathbf{t}) = \mathbf{t} \cdot \nabla u(\mathbf{t}) - u(\mathbf{t})$. The revenue maximization problem can then be written as the following mathematical program, which is the starting point for the analysis of this chapter.

$$(3.2) \quad \max_{\boldsymbol{\pi}, u} \int_{\mathbf{t}} [\mathbf{t} \cdot \nabla u(\mathbf{t}) - u(\mathbf{t})] f(\mathbf{t}) d\mathbf{t}$$

$$\boldsymbol{\pi} = \nabla u; u \text{ is convex,}$$

$$\forall \mathbf{t}, \boldsymbol{\pi}(\mathbf{t}) \text{ is feasible assignment.}$$

Notice that when the dimension of the type space is $m = 1$, the above program is equivalent to the following familiar form from Myerson (1981) (Lemma 1),

$$\begin{aligned} \max_{\pi} \quad & \int_v [v\pi(v) - \int_{z \leq v} \pi(z) dz] f(v) dv \\ & \pi \text{ is monotone non-decreasing,} \\ & \forall v, \pi(v) \leq 1. \end{aligned}$$

In Section 3.6, we extend the above formulation and our results to multi-agent settings.

3.3. Amortization of Revenue

This section formalizes our codification of multi-dimensional virtual values for incentive compatible mechanism design and describes the working pieces of our framework. The main construct is the definition of multi-dimensional virtual value functions and the accompanying proposition, below.

Definition 9. A *vector field* $\hat{\phi} : [0, 1]^m \rightarrow \mathbb{R}^m$ that maps an an m -dimensional type to an m -dimensional vector is

- *incentive compatible* if the *virtual surplus maximizer* given by selecting the assignment π for type \mathbf{t} that optimizes virtual surplus $\pi \cdot \hat{\phi}(\mathbf{t})$ is incentive compatible;
- a *weak amortization of revenue* if, in expectation over types drawn from the distribution, the virtual surplus of any incentive compatible mechanism upper bounds is revenue, i.e., $\mathbf{E}[\hat{\phi}(\mathbf{t}) \cdot \pi(\mathbf{t})] \geq \mathbf{E}[p(\mathbf{t})]$, and with equality for the virtual surplus maximizer;
- a *strong amortization of revenue* if the inequality (of weak amortization) holds with equality for all incentive compatible mechanisms; and
- a *virtual value function* if it is incentive compatible and a weak or strong amortization.

Proposition 1. *In any environment for which a virtual value function exists, the virtual surplus maximizer is incentive compatible and revenue optimal.*

Proof. The expected revenue of the virtual surplus maximizer is equal to its expected virtual surplus (by weak amortization). This expected virtual surplus is at least the virtual surplus of any alternate mechanism (by definition of virtual surplus maximization). The expected virtual surplus is an upper bound on the expected revenue of the alternative mechanism (by weak amortization). Thus, the expected revenue of the virtual surplus maximizer is at least that of the alternative mechanism. Incentive compatibility follows directly from the definition of a virtual value function. \square

For a single-dimensional agent (i.e., $m = 1$), Myerson (1981) showed that the function $v - \frac{1-F(v)}{f(v)}$ is a strong amortization, when it is monotone it is incentive compatible, when it is non-monotone an ironing procedure can be applied to obtain from it a weak amortization function that is monotone and thus incentive compatible. Our approach will analogously enable the derivation of multi-dimensional virtual value functions (i.e., satisfying incentive compatibility and weak amortization) via the construction of a strong amortization function that is not necessarily incentive compatible.

We now give sufficient conditions for a vector field to be a strong amortization of revenue. At a high level, we derive first a strong amortization of utility and then, using the fact that revenue is value minus utility, derive a strong amortization of revenue. These strong amortizations will be building blocks for the derivation of virtual values for unit-demand and additive agents in the subsequent sections. The following lemma follows from integration by parts as per equation (3.1), the definition of strong amortization of utility (Definition 9, generalized to utility), and the fact that the gradient of utility is the assignment rule of the mechanism (Lemma 10).

Lemma 11. For type space T and distribution f , vector field α/f is a strong amortization of utility, i.e., $\mathbf{E}[u(\mathbf{t})] = \mathbf{E}[\alpha(\mathbf{t})/f(\mathbf{t}) \cdot \pi(\mathbf{t})]$ for all incentive compatible assignment rules π , if it satisfies

- divergence density equality, i.e., that $\nabla \cdot \alpha(\mathbf{t}) = -f(\mathbf{t})$ at any point $\mathbf{t} \in T$, and
- boundary orthogonality, i.e., that $\alpha(\mathbf{t}) \cdot \eta(\mathbf{t}) = 0$ for all $\mathbf{t} \neq \mathbf{0}$ on the boundary of type space ∂T with normal vector $\eta(\mathbf{t})$.

Proof. Write the expectation of $\mathbf{E}[\alpha(\mathbf{t})/f(\mathbf{t}) \cdot \pi(\mathbf{t})]$ as the integral $\int_{\mathbf{t} \in T} \alpha(\mathbf{t}) \cdot \pi(\mathbf{t}) \, d\mathbf{t}$. From Lemma 3, substitute ∇u for assignment π and apply integration by parts.

$$(3.3) \quad \int_{\mathbf{t} \in T} \alpha(\mathbf{t}) \cdot \nabla u(\mathbf{t}) \, d\mathbf{t} = \int_{\mathbf{t} \in \partial T} u(\mathbf{t})(\alpha \cdot \eta)(\mathbf{t}) \, d\mathbf{t} - \int_{\mathbf{t} \in T} u(\mathbf{t}) (\nabla \cdot \alpha(\mathbf{t})) \, d\mathbf{t}$$

$$(3.4) \quad = u(\mathbf{0}) + \int_{\mathbf{t} \in T} u(\mathbf{t}) f(\mathbf{t}) \, d\mathbf{t}$$

The second equality is derived from the first equality by employing the assumptions of the lemma on vector field α as follows.

- By divergence density equality, the second term simplifies by substituting $\nabla \cdot \alpha(\mathbf{t}) = -f(\mathbf{t})$.
- Recall that the divergence theorem is equivalent setting $u(\mathbf{t}) = 1$ in the formula (3.3), this gives

$$(3.5) \quad \int_{\mathbf{t} \in \partial T} (\alpha \cdot \eta)(\mathbf{t}) \, d\mathbf{t} = \int_{\mathbf{t} \in T} f(\mathbf{t}) \, d\mathbf{t} = 1,$$

the total probability of any type. Boundary orthogonality implies that the integrand in the boundary integral of equation (3.5) is identically zero everywhere except $\mathbf{t} = \mathbf{0}$. To integrate to one on the boundary, the function must be the Dirac delta function at $\mathbf{t} = \mathbf{0}$; thus, the integral of the first term in equation (3.3) is $u(\mathbf{0})$.

Without loss of generality for revenue optimal mechanisms $u(\mathbf{0}) = 0$. We can interpret the left- and right-hand sides of equation (3.4) as expectations, which gives $\mathbf{E}[\boldsymbol{\alpha}(\mathbf{t})/f(\mathbf{t}) \cdot \boldsymbol{\pi}(\mathbf{t})] = \mathbf{E}[u(\mathbf{t})]$, the definition of strong amortization of utility for $\boldsymbol{\alpha}/f$. \square

For a single-dimensional agent with value v in type space $T = [0, 1]$, the only function that satisfies the conditions of Lemma 11 and gives a strong amortization of utility is $\alpha(v)/f(v) = \frac{1-F(v)}{f(v)}$. For this formula, notice that the divergence of $\alpha(v) = 1 - F(v)$ is simply its derivative $-f(v)$. The boundary $\partial T \setminus \{0\}$ is the point $v = 1$, the upper bound of the distribution, and thus trivially satisfies orthogonality as $\alpha(1) = 0$. In classical auction theory the amortization of utility $\frac{1-F(v)}{f(v)}$ is often referred to as the agent's *information rent*.

The following lemma is immediate from the fact that revenue is the agent's surplus minus the agent's utility. For a single-dimensional agent it implies that $\phi(v) = v - \frac{(1-F(v))}{f(v)}$ is the strong amortization of revenue.

Lemma 12. *For type space T and distribution f , vector field ϕ is a strong amortization of revenue, i.e., $\mathbf{E}[p(\mathbf{t})] = \mathbf{E}[\phi(\mathbf{t}) \cdot \boldsymbol{\pi}(\mathbf{t})]$ for all incentive compatible assignment rules $\boldsymbol{\pi}$, if and only if $\phi(\mathbf{t}) = \mathbf{t} - \boldsymbol{\alpha}(\mathbf{t})/f(\mathbf{t})$ for all \mathbf{t} but a measure zero subset of T and $\boldsymbol{\alpha}/f$ is a strong amortization of utility.*

Unlike the case of a single-dimensional agent, for multi-dimensional agents there are many strong amortizations of utility and, consequentially, many strong amortizations of revenue. As an example, suppose we wish to show the optimality of a restricted form of mechanism via a strongly amortized virtual value function for an $m = 2$ dimensional agent. This virtual value function has two degrees of freedom. We can pin down one degree of freedom by equating virtual surplus to expected revenue for mechanisms with this restricted form. The divergence density equality for strong amortizations that Lemma 12 inherits from Lemma 11 gives a differential equation that then pins down the other degree of freedom. It remains to find sufficient conditions

on the distribution under which virtual surplus maximization identically gives mechanisms of the restricted form.

The approach above can be generalized to show optimality of mechanisms via a weakly amortized virtual value function. For example, such a generalization can be used to give proofs of optimality under more permissive distributional assumptions. To substitute weak amortization for strong amortization we need a way to relate the differential equations (from divergence density equality) that govern strong amortizations to any given weak amortization. Such a relationship follows directly from the definitions of both weak and strong amortization in terms of the expected revenue of any incentive compatible mechanism (Definition 9) and is summarized below as Lemma 13.

Lemma 13. *For type space T and distribution f , vector field $\hat{\phi}$ is a weak amortization of revenue if and only if there exists a strong amortization of revenue ϕ such that $\mathbf{E}[\hat{\phi}(\mathbf{t}) \boldsymbol{\pi}(\mathbf{t})] \geq \mathbf{E}[\phi(\mathbf{t}) \boldsymbol{\pi}(\mathbf{t})]$, for all incentive compatible mechanisms $\boldsymbol{\pi}$, with equality for $\boldsymbol{\pi}$ that pointwise maximizes $\hat{\phi}$.*

3.4. Optimality of Uniform Pricing for Unit Demand Preferences

In this section we study the existence of virtual values in order to prove optimality of uniform pricing for a single unit-demand agent. To simplify the exposition we focus on the case of two items and on the type space where item one is the favorite item, i.e., the lower right half of the unit square, $T = \{\mathbf{t} = (t_1, t_2) : 0 \leq t_2 \leq t_1 \leq 1\}$. Our conclusions extend easily to the $[0, 1]^2$ type space with symmetric distributions; other extensions are given at the end of the section. The general case of $m \geq 2$ items is considered in Section A.1.

The *single-dimensional projection for the favorite item* is given by distribution and density function for the agent's favorite item, $F_{\max}(v)$ and $f_{\max}(v)$. The distribution function is the integral of f over \mathbf{t} with $t_1 \geq v$. The density function is the integral of f of \mathbf{t} with $t_1 = v$, i.e.,

$f_{\max}(v) = \int_0^v f(v, z) dz$. As described in Section 3.3, the unique strong amortization of revenue for a single-dimensional agent (and thus for the single-dimensional projection) is $\phi_{\max}(v) = v - \frac{1-F_{\max}(v)}{f_{\max}(v)}$. The strong amortization of utility α_{\max}/f_{\max} requires $\alpha_{\max}(v) = 1 - F_{\max}(v)$.

Definition 10. The two-dimensional extension ϕ of the favorite-item projection ϕ_{\max} (satisfying $\phi_{\max}(\mathbf{t}) = t_1 - \frac{1-F_{\max}(t_1)}{f_{\max}(t_1)}$) is constructed as follows:

- (a) Set $\phi_1(\mathbf{t}) = \phi_{\max}(t_1)$ for all $\mathbf{t} \in T$.
- (b) Let $\alpha_1(\mathbf{t}) = (t_1 - \phi_1(\mathbf{t})) f(\mathbf{t}) = \frac{1-F_{\max}(t_1)}{f_{\max}(t_1)} f(\mathbf{t})$.
- (c) Let $\alpha_2(\mathbf{t}) = - \int_{y=0}^{t_2} (f(t_1, y) + \partial_1 \alpha_1(t_1, y)) dy$.
- (d) Set $\phi_2(\mathbf{t}) = t_1 - \alpha_2(\mathbf{t})/f(\mathbf{t})$.

In the remainder of this section we will show that for single-agent mechanism design to optimize revenue minus a fixed non-negative cost for selling either item, that this two-item extension of the favorite-item projection is a strong amortization of revenue that proves the optimality of uniform pricing.⁴

An informal justification of the steps of the construction is as follows:

- (a) First, for fixed t_1 and as a function of t_2 , $\phi_1(\mathbf{t})$ must be constant (i.e., on a vertical line in T); otherwise, there is a cost c for which virtual surplus maximization with respect to ϕ serves one such type and not other is not the other which is not a uniform pricing. Second, the revenue of any mechanism that only ever sells the favorite item or nothing has revenue given by the favorite-item projection and must satisfy $\phi_1(\mathbf{t}) = \phi_{\max}(t_1)$ (given the first point).
- (b) We obtain α_1 from ϕ_1 by Lemma 12. Orthogonality of the right boundary ($t_1 = 1$) requires that $\alpha \cdot (1, 0) = 0$, and therefore $\alpha_1(1, t_2) = 0$. By definition, $\alpha_1(1, t_2) = \frac{\alpha_{\max}(1)}{f_{\max}(1)} f(1, t_2) = 0$

⁴The extra constraint imposed by a non-negative cost of service will enable this method to be extended to multi-agent settings, see Section 3.6. This strong amortization is unique on the portion of type space for which $\phi_1(\mathbf{t}) > 0$.

and follows directly from boundary orthogonality of the favorite-item projection at $t_1 = 1$ which required $\alpha_{\max}(1) = 0$ and was satisfied.

(c) The derivatives of α_1 (with respect to t_1) and α_2 (with respect to t_2) are related by the divergence density equality; integrating and employing boundary orthogonality on the bottom boundary ($t_2 = 0$) of the type space, which requires that $\alpha_2(t_2, 0) = 0$, gives the formula; these constraints are required by Lemma 11.

(d) We obtain ϕ_2 from α_2 by Lemma 12.

For ϕ to prove optimality of uniform pricing, it must be that virtual surplus maximization would never assign the agent the non-favorite item, i.e., item two. This requirement is simply $\phi_1(\mathbf{t}) \geq \phi_2(\mathbf{t})$ for any type $\mathbf{t} \in T$ for which either $\phi_1(\mathbf{t})$ or $\phi_2(\mathbf{t})$ is positive. A little algebra shows that this condition is implied by the angle of $\boldsymbol{\alpha}(\mathbf{t})$ being at most the angle of \mathbf{t} (with respect to the horizontal t_2 axis; see Lemma 15, below). Importantly, in relation to the prior work of Armstrong (1996), the direction of $\boldsymbol{\alpha}$ corresponds to the paths on which incentive compatibility constraints are considered. The approach we are taking does not fix the direction, it allows any direction that satisfies the above constraint on angles. The condition on angles is equivalent to the dot product between $\boldsymbol{\alpha}$ and the upward orthogonal vector to \mathbf{t} being non-positive. The following lemma is proved by the divergence theorem.

Lemma 14. *Vector field $\boldsymbol{\alpha}/f$ in the definition of the two-dimensional extension of the favorite-item projection is a strong amortization of utility and satisfies*

$$\theta \alpha_1(t_1, t_1\theta) - \alpha_2(t_1, t_1\theta) = (1 - F_{\max}(t_1)) \frac{d}{dt_1} \left[\frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{f_{\max}(t_1)} \right]$$

for all $t_1, \theta \in [0, 1]$ (and thus $(t_1, t_1\theta) \in T$).

Proof. The informal justifications of Steps (b) and (c) show that $\boldsymbol{\alpha}$ satisfies the divergence density equality and bottom and right boundary orthogonality. This proof starts with these

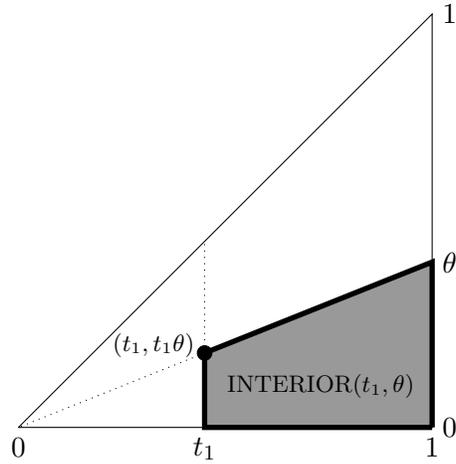


Figure 3.2. The trapezoidal set parameterized by t_1 and θ , and the four curves that define its boundary, $\{\text{TOP, RIGHT, BOTTOM, LEFT}\}(t_1, \theta)$

assumption and derives the identity of the lemma. Notice that for $\theta = 1$ the numerator and the denominator in the derivative of the identity are equal for all t_1 , the right-hand side is zero, and therefore $\alpha_1(t_1, t_1) = \alpha_2(t_1, t_1)$, and boundary orthogonality holds for the diagonal boundary. Thus, α/f is a strong amortization of utility.

The strategy for the proof of the identity is as follows. We fix t_1 and θ and apply the divergence theorem to α on the trapezoidal subspace of type space defined by types \mathbf{t}' with $t'_1 \geq t_1$, $t'_2/t'_1 \leq \theta$, $t'_2 \geq 0$, and $t'_1 \leq 1$ (Figure 3.2). The divergence theorem equates the the integral of the vector field α on the boundary of the subspace to the integral of its divergence within the subspace. As the upper boundary of this trapezoidal subspace has slope t_2/t_1 , one term in this equality is the integral of $\alpha(\mathbf{t}')$ with the upward orthogonal vector to \mathbf{t} . Differentiating this integral and evaluating at $\mathbf{t}' = (t_1, t_1\theta)$ gives the desired quantity.

Applying the divergence theorem to $\boldsymbol{\alpha}$ on the trapezoid and expressing the top boundary as the interior divergence minus the other three boundaries gives:

$$\begin{aligned} \int_{\mathbf{t}' \in \text{TOP}(t_1, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' &= \int_{\mathbf{t}' \in \text{INTERIOR}(t_1, \theta)} \nabla \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' \\ &\quad - \int_{\mathbf{t}' \in \{\text{RIGHT}, \text{BOTTOM}, \text{LEFT}\}(t_1, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}'. \end{aligned}$$

Since $\boldsymbol{\alpha}/f$ is a strong amortization of utility, the divergence density equality and boundary orthogonality of Lemma 11 imply that the integral over the interior simplifies and the integrals over the right and bottom boundary are zero, respectively. We have,

$$\int_{\mathbf{t}' \in \text{TOP}(t_1, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' = - \int_{\mathbf{t}' \in \text{INTERIOR}(t_1, \theta)} f(\mathbf{t}') \, d\mathbf{t}' - \int_{\mathbf{t}' \in \text{LEFT}(t_1, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}'.$$

For the trapezoid at \mathbf{t} these integrals are,

$$\begin{aligned} \int_{t'_1=t_1}^1 (-\theta \alpha_1(t'_1, t'_1\theta) + \alpha_2(t'_1, t'_1\theta)) \, dt'_1 \\ = - \int_{t'_1=t_1}^1 \int_{t'_2=0}^{t'_1\theta} f(\mathbf{t}') \, dt'_2 \, dt'_1 + \int_{t'_2=0}^{t_1\theta} \alpha_1(t_1, t'_2) \, dt'_2. \end{aligned}$$

Differentiating with respect to t_1 gives,

$$\theta \alpha_1(t_1, t_1\theta) - \alpha_2(t_1, t_1\theta) = \int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) \, dt'_2 + \frac{d}{dt_1} \int_{t'_2=0}^{t_1\theta} \alpha_1(t_1, t'_2) \, dt'_2.$$

On the right-hand side, multiply first term by $\frac{f_{\max}(t_1)}{f_{\max}(t_1)} = 1$ and plug in the strong amortization of utility for the two-dimensional extension as $\alpha(\mathbf{t}) = \frac{1-F_{\max}(t_1)}{f_{\max}(t_1)} f(\mathbf{t})$ to the second term. Notice that the integral of the second term is only on t'_2 therefore we can bring the terms related to t_1 outside the integral. These two terms then simplify by the product rule for differentiation to give the identity of the lemma.

$$\begin{aligned} \theta \alpha_1(t_1, t_1\theta) - \alpha_2(t_1, t_1\theta) &= f_{\max}(t_1) \frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{f_{\max}(t_1)} \\ &\quad + \frac{d}{dt_1} \left[(1 - F_{\max}(t_1)) \frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{f_{\max}(t_1)} \right] \\ &= (1 - F_{\max}(t_1)) \frac{d}{dt_1} \left[\frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{f_{\max}(t_1)} \right]. \quad \square \end{aligned}$$

Lemma 15. *For ϕ and α defined by the two-dimensional extension of the favorite-item projection, if $\frac{t_2}{t_1} \alpha_1(t_1, t_2) - \alpha_2(t_1, t_2)$ is non-positive and $\phi_1(\mathbf{t})$ is monotone non-decreasing in \mathbf{t} , then virtual surplus maximization with respect to ϕ and any non-negative service cost c gives a uniform pricing.*

Proof. From the assumption $\frac{t_2}{t_1} \alpha_1(t_1, t_2) - \alpha_2(t_1, t_2) \leq 0$ and Definition 10 we have

$$\frac{t_2}{t_1} \phi_1(\mathbf{t}) = \frac{t_2}{t_1} \left(t_1 - \frac{\alpha_1(\mathbf{t})}{f(\mathbf{t})} \right) = \frac{t_2}{t_1} \left(t_1 - \frac{\alpha_1(\mathbf{t})}{f(\mathbf{t})} \right) \geq t_2 - \frac{\alpha_2(\mathbf{t})}{f(\mathbf{t})} = \phi_2(\mathbf{t}).$$

Thus, for \mathbf{t} with $\phi_1(\mathbf{t}) \geq c$, $\phi_1(\mathbf{t}) \geq \phi_2(\mathbf{t})$ and virtual surplus maximization serves the agent item one. Since $\phi_1(\mathbf{t})$ is a function only of t_1 (Definition 10), its monotonicity implies that there is a smallest t_1 such that all greater types are served. Also, if $\phi_1(\mathbf{t}) \leq c$, again the above calculation implies that $\phi_2(\mathbf{t}) \leq c$ and therefore the type is not served. This outcome is one of a uniform pricing. \square

We are now ready to state the main theorems of this section. In the next section we will give an interpretation of the main technical condition as a supermodularity condition on the density function.

Theorem 6. *Uniform pricing is revenue optimal for any service cost c and any distribution for which the favorite-item projection has monotone non-decreasing strong amortization $\bar{\phi}_{\max}(t_1) = t_1 - \frac{1-F_{\max}(t_1)}{f_{\max}(t_1)}$ and $\frac{d}{dt_1} \left[\frac{1}{f_{\max}(t_1)} \int_0^{t_2} f(t_1, y) dy \right]$ is non-positive.*

The conditions of Corollary 2 can be further relaxed by constructing a weak amortization $\hat{\phi}$ from the strong amortization ϕ , above. The following is a special case of the more general Lemma 25 given in Subsection 3.6.3.

Theorem 7. *Uniform pricing is revenue optimal for any service cost c and any distribution for which $\frac{d}{dt_1} \left[\frac{1}{f_{\max}(t_1)} \int_0^{t_2} f(t_1, y) dy \right]$ is non-positive.*

In Section A.1 we employ using the natural generalization of the condition to extend the above results to $m \geq 2$ items.

3.4.1. Max-ratio log-supermodular distributions

Theorem 7 and Theorem 6 in the preceding section require the non-positivity of the expression $\frac{d}{dt_1} \left[\frac{1}{f_{\max}(t_1)} \int_0^{t_2} f(t_1, y) dy \right]$ which is not easy to directly check; in this section we give a simpler sufficient condition. To specify the condition, it is useful to view the distribution in the *max-ratio* (MR) coordinates. In particular the max-ratio representation f^{MR} of a density function f is defined to be

$$f^{MR}(t_1, \theta) = f(t_1, \theta t_1), \quad \forall t_1, \theta, 0 \leq t_1, \theta \leq 1.$$

Equivalantly,

$$f(t_1, t_2) = f^{MR}(t_1, t_2/t_1) \quad \forall t_1, t_2, 0 \leq t_1, t_2 \leq 1.$$

We call a distribution MR-log-supermodular, if its max-ratio representation is log-supermodular,

$$f^{MR}(t_1, \theta) \times f^{MR}(t'_1, \theta') \geq f^{MR}(t_1, \theta') \times f^{MR}(t'_1, \theta), \quad \forall t_1 \leq t'_1, \theta \leq \theta'.$$

Notice that, for example, *MR-independent* distributions are MR-log-supermodular. A MR-independent distribution f^{MR} is such that $f^{MR}(t_1, \theta_1) = f_1(t_1) \times f_\theta(\theta)$ (for arbitrary f_1 and f_θ), and is MR-log-supermodular because

$$\begin{aligned} f^{MR}(t_1, \theta) \times f^{MR}(t'_1, \theta') &= f_1(t_1) f_\theta(\theta) f_1(t'_1) f_\theta(\theta') \\ &= f_1(t_1) f_\theta(\theta') f_1(t'_1) f_\theta(\theta) \\ &= f^{MR}(t_1, \theta') \times f^{MR}(t'_1, \theta). \end{aligned}$$

We now prove that $\frac{d}{dt_1} \left[\frac{1}{f_{\max}(t_1)} \int_0^{t_2} f(t_1, y) dy \right]$ is non-positive for MR-log-supermodular distributions. Recall that the density function of the favorite-item projection is defined from the two-item density function as $f_{\max}(t_1) = \int_0^{t_1} f(t_1, y) dy$.

Lemma 16. *If f is a MR-log-supermodular function, then for any t_1 and θ ,*

$$\frac{d}{dt_1} \left[\frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{\int_{t'_2=0}^{t_1} f(t_1, t'_2) dt'_2} \right] \leq 0,$$

with equality if distribution is MR-independent.

Proof. We prove that for any θ, t_1 , and t'_1 such that $t_1 < t'_1$,

$$\frac{\int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2}{\int_{t'_2=0}^{t_1} f(t_1, t'_2) dt'_2} \leq \frac{\int_{t'_2=0}^{t'_1\theta} f(t'_1, t'_2) dt'_2}{\int_{t'_2=0}^{t'_1} f(t'_1, t'_2) dt'_2}.$$

The proof first converts the above form into max-ratio coordinates, then applies MR-log-supermodularity, and then transforms back to the standard form. Before applying MR-log-supermodularity, we break down the integral set into two set, and apply MR-log-supermodularity to only one of the integrals. More particularly, notice that

$$\begin{aligned} & \int_{t'_2=0}^{t_1\theta} f(t_1, t'_2) dt'_2 \times \int_{t'_2=0}^{t'_1} f(t'_1, t'_2) dt'_2 \\ &= \int_{\theta'=0}^{\theta} f^{MR}(t_1, \theta') t_1 d\theta' \times \int_{\theta''=0}^1 f^{MR}(t'_1, \theta'') t'_1 d\theta'' \quad (\text{change of variables}) \\ &= \int_{\theta'=0}^{\theta} \int_{\theta''=0}^{\theta} f^{MR}(t_1, \theta') t_1 f^{MR}(t'_1, \theta'') t'_1 d\theta'' d\theta' \\ & \quad + \int_{\theta'=0}^{\theta} \int_{\theta''=\theta}^1 f^{MR}(t_1, \theta') t_1 f^{MR}(t'_1, \theta'') t'_1 d\theta'' d\theta' \quad (\text{separate double integral into two sets}) \\ &\leq \int_{\theta'=0}^{\theta} \int_{\theta''=0}^{\theta} f^{MR}(t_1, \theta'') t_1 f^{MR}(t'_1, \theta') t'_1 d\theta' d\theta'' \quad (\text{rename variables } \theta' \text{ and } \theta'') \\ & \quad + \int_{\theta'=0}^{\theta} \int_{\theta''=\theta}^1 f^{MR}(t_1, \theta'') t_1 f^{MR}(t'_1, \theta') t'_1 d\theta'' d\theta' \quad (\text{apply MR-log-supermodularity}) \\ &= \int_{\theta''=0}^1 f^{MR}(t_1, \theta'') t_1 d\theta'' \int_{\theta'=0}^{\theta} f^{MR}(t'_1, \theta') t'_1 d\theta' \quad (\text{merge integrals}) \\ &= \int_{t'_2=0}^{t_1} f(t_1, t'_2) dt'_2 \times \int_{t'_2=0}^{t'_1\theta} f(t'_1, t'_2) dt'_2. \end{aligned}$$

□

By combining the above lemma and Corollary 2 we get the following corollary.

Corollary 2. *Uniform pricing is revenue optimal for any service cost c and any max-ratio log-supermodular distribution.*

To understand MR-log-supermodular distributions better, we next show properties under which product distributions are MR-log-supermodular. We call a function g geometric-geometric (GG) convex if

$$g(z_1^\lambda z_2^{1-\lambda}) \geq g(z_1)^\lambda g(z_2)^{1-\lambda}, \quad \forall \lambda \in [0, 1], z_1, z_2.$$

For example, the function $g(x) = x^k$ is GG-convex. More generally, for any convex function h and constant c , the function $g(x) = c \cdot e^{h(\log(x))}$ is GG-convex. The proof of the following lemma can be found in Section A.2.

Lemma 17. *The product distribution on i.i.d. draws from a distribution with geometric-geometric convex density is max-ratio log-supermodular.*

3.4.2. Bundle Pricing for General Distributions

We can use the generality of unit-demand settings and the above result to find conditions under which bundle pricing is optimal for general distributions. More precisely, notice that any problem with finite outcome space can be converted to a unit-demand problem with dot-product utility function by letting \mathbf{t} and $\boldsymbol{\pi}$ be vectors of size equal to the size of the outcome space, with t_i being the value for outcome i . For multi-item settings with m items, the size of the outcome space is equal to 2^m . The input to the problem is a density function $f(t_1, \dots, t_k)$ over vectors of size $k = 2^m$. We assume that the density function is *monotone*, meaning that the density of a type is non-zero only if the valuation for the grand bundle is at least as much as the value for any other subset of items. The max-ratio representation of the density function is $f^{MR}(t_1, \theta_2, \dots, \theta_k)$, where t_1 , normalized to be at most 1, is the value for the grand bundle, and $\theta_i \leq 1$ is the

ratio of the value for the i -th subset, over the value of the grand bundle. Now a bundle pricing mechanism corresponds to a mechanism in the unit demand setting which only sells the favorite outcome, which is the grand bundle. We therefore have the following theorem.

Theorem 8. *Bundle pricing is optimal if the max-ratio representation of the density function is MR-log-supermodular.*

3.5. Optimality of Bundle Pricing for Additive Preferences

In this section we provide sufficient conditions for optimality of grand bundle pricing for agents with additive utilities. Similar to Section 3.4, we only focus on constructing a proof assuming that item one is the favorite item, as the proof generalizes to symmetric distributions easily by mirroring the construction for the other half of set of types. It is also easiest to express the results of this section when the sum of the values for the items are normalized to be at most one. We thus define the set of types to be $T = \{(t_1, t_2) | t_1, t_2 \geq 0, t_1 \geq t_2, t_1 + t_2 \leq 1\}$.

The *single-dimensional projection for the sum of values* is given by distribution and density function for the agent's favorite item, $F_{sum}(v)$ and $f_{sum}(v)$. The distribution function is the integral of f over \mathbf{t} with $t_1 + t_2 \leq v$. The density function $f_{sum}(v)$ is the derivative of $F_{sum}(v)$ with respect to v . As described in Section 3.3, the unique strong amortization of revenue for a single-dimensional agent (and thus for the single-dimensional projection) is $\phi_{sum}(v) = v - \frac{1 - F_{sum}(v)}{f_{sum}(v)}$. The strong amortization of utility α_{sum}/f_{sum} requires $\alpha_{sum}(v) = 1 - F_{sum}(v)$.

In this section, we prove optimality of grand bundle pricing by constructing a weak amortization of revenue $\hat{\phi}$ (as opposed to Section 3.4 in which we directly constructed a strong amortization), which we call the *two-dimensional extension of the sum-of-values projection*. We will first define $\hat{\phi}$ such that the virtual surplus of bundle pricing with respect to $\hat{\phi}$ is equal to its revenue (Lemma 18), for which grand bundle pricing maximizes virtual surplus (Lemma 19). The rest of the section proves that $\hat{\phi}$ is indeed a weak amortization of revenue by using Lemma 13

to show existence of a strong amortization which lower bounds the virtual surplus, with respects to $\hat{\phi}$, of any incentive compatible mechanism.

Define $\hat{\phi}_1$ and $\hat{\phi}_2$ as follows

$$\begin{aligned}\hat{\phi}_1(\mathbf{t}) &= \frac{t_1}{t_1 + t_2} \phi_{sum}(t_1 + t_2) = t_1 - \frac{1 - F_{sum}(t_1 + t_2)}{f_{sum}(t_1 + t_2)}, \\ \hat{\phi}_2(\mathbf{t}) &= \frac{t_2}{t_1 + t_2} \phi_{sum}(t_1 + t_2) = t_2 - \frac{1 - F_{sum}(t_1 + t_2)}{f_{sum}(t_1 + t_2)}.\end{aligned}$$

This definition ensures that $\hat{\phi}$ satisfies two important properties. First, the virtual surplus of bundle pricing is equal to the virtual surplus of the single-dimensional projection for sum of values, that is, $\hat{\phi}_1(\mathbf{t}) + \hat{\phi}_2(\mathbf{t}) = \phi_{sum}(t_1 + t_2)$. As a result of this property, the virtual surplus of grand bundle pricing can be shown to be equal to its revenue (see the next lemma). Second, the two components of the virtual value function have the same sign, which together with the monotonicity of ϕ_{sum} would imply that bundle pricing maximizes virtual surplus (Lemma 19).

Lemma 18 (single-dimensional projection). *The expected revenue of a bundle pricing is equal to its expected virtual surplus with respect to the two-dimensional extension of the sum-of-values projection $\hat{\phi}$ satisfying $\hat{\phi}_1(\mathbf{t}) + \hat{\phi}_2(\mathbf{t}) = \phi_{sum}(t_1 + t_2)$ where $\phi_{sum}(v) = v - \frac{1 - F_{sum}(v)}{f_{sum}(v)}$ is the strong amortization for agent's sum-of-values projection.*

Functions $\hat{\phi}_1$ and $\hat{\phi}_2$ have the same sign as ϕ_{sum} , and therefore as long as ϕ_{sum} is a non-decreasing function of p , there exists a price p^* such that the assignment corresponding to price p^* ex-post optimizes $\hat{\phi}$, according to the following lemma.

Lemma 19 (consistency with virtual surplus maximization; incentive compatibility). *Virtual surplus maximization according to vector field $\hat{\phi}$ gives a bundle pricing p (and is incentive compatible) if and only if*

- $\hat{\phi}_1(\mathbf{t}), \hat{\phi}_2(\mathbf{t}) \geq 0$ when $t_1 + t_2 \geq p$ and $\hat{\phi}_1(\mathbf{t}), \hat{\phi}_2(\mathbf{t}) \leq 0$ otherwise.

Proof. We need to show that for the uniform price p , the assignment function π of posting a price p for the bundle optimizes ϕ pointwise. Pointwise optimization of $\pi \cdot \hat{\phi}$ will result in $\pi = (1, 1)$ whenever $\hat{\phi}_1, \hat{\phi}_2 \geq 0$, and will result in $\pi = (0, 0)$ whenever $\phi_1, \phi_2 \leq 0$. \square

In order to complete the proof we should ensure that $\hat{\phi}$ is an amortization of revenue. We already constructed $\hat{\phi}$ such that the expected virtual surplus of the virtual surplus maximizer, bundle pricing, is equal to its revenue. So we only need to show that the virtual surplus of any incentive compatible assignment is at least its revenue. We do this by constructing a strong amortization of revenue ϕ and invoking Lemma 13. That is, we will show that there exists a strong amortization ϕ such that for any incentive compatible assignment π ,

$$(3.6) \quad \int_t \pi(\mathbf{t}) \cdot [\hat{\phi}(\mathbf{t}) - \phi(\mathbf{t})] f(\mathbf{t}) \, d\mathbf{t} \geq 0.$$

To do this, it would be sufficient to ensure that $\hat{\phi}_1 \geq \phi_1$ and $\hat{\phi}_2 \geq \phi_2$ for any type. Instead, in Lemma 20 we suggest a more permissive sufficient condition using integration by parts, and the properties of incentive compatible assignments. Lemma 22 later defines conditions that would imply that a strong amortization satisfying the conditions of Lemma 20 can be constructed, using an approach very similar to the proof of Lemma 14.

For a function h on type space T , define h^{SR} to be its transformation to sum-ratio coordinates, that is

$$h(t_1, t_2) = h^{SR}\left(t_1 + t_2, \frac{t_2}{t_1}\right)$$

Sum-ratio coordinates are convenient to work with since it allows us to refer to a type using its value for the bundle, and the ratio of values of item one and item two. The proof of the following lemma first transforms the expression in inequality (3.6) to sum-ratio coordinates, and

invokes single-dimensional integration by parts along the lines with constant sum of values to get an expression in terms of derivative of π times a (weighted) integral of $\hat{\phi} - \phi$ (see expression (3.7)). We then use a property of incentive compatible assignments to verify the sign of the given expression.

Lemma 20. *The two-dimensional extension of the sum-of-values projection $\hat{\phi}$ is a weak amortization of revenue if there exists a strong amortization ϕ such that*

$$(1) \phi_1(\mathbf{t}) + \phi_2(\mathbf{t}) = \phi_{sum}(t_1 + t_2).$$

$$(2) \phi_1(\mathbf{t}) \frac{t_2}{t_1} \leq \phi_2(\mathbf{t}).$$

Proof. We invoke Lemma 13 and construct a strong amortization ϕ that shows that the expected virtual surplus, with respect to $\hat{\phi}$ of any incentive compatible reverse is at least its revenue. Thus we need to show that the following expression for all incentive compatible π ,

$$\int_{\mathbf{t}} \pi(\mathbf{t}) \cdot [\hat{\phi}(\mathbf{t}) - \phi(\mathbf{t})] f(\mathbf{t}) \, d\mathbf{t} \geq 0.$$

By symmetry of the distribution, the optimal mechanism is also symmetric. Therefore, we prove the lemma only for symmetric incentive compatible assignments (in particular, we assume that $x_1(t_1, t_1) = x_2(t_1, t_1)$ for all t_1). Perform a change of variables $s = t_1 + t_2$, and $\theta = t_2/t_1$ on the left hand side of the above expression (3.6)

$$\begin{aligned} & \int_{s=0}^1 \int_{\theta=0}^1 \mathbf{x}^{SR}(s, \theta) \cdot [\hat{\phi} - \phi]^{SR}(s, \theta) f^{SR}(s, \theta) \frac{s}{1 + \theta} \, d\theta \, ds \\ &= \int_{s=0}^1 \int_{\theta=0}^1 \mathbf{x}^{SR}(s, \theta) \cdot \frac{d}{d\theta} \int_{\theta'=0}^{\theta} [\hat{\phi} - \phi]^{SR}(s, \theta') f^{SR}(s, \theta') \frac{s}{1 + \theta'} \, d\theta' \, d\theta \, ds. \end{aligned}$$

To simplify exposition, let

$$\gamma(s, \theta) = \int_{\theta'=0}^{\theta} [\hat{\phi} - \phi]^{SR}(s, \theta') f^{SR}(s, \theta') d\theta',$$

to rewrite the left hand side of inequality (3.6)

$$\begin{aligned} &= \int_{s=0}^1 \int_{\theta=0}^1 \mathbf{x}^{SR}(s, \theta) \cdot \frac{d}{d\theta} \gamma(s, \theta) d\theta ds \\ &= \int_{s=0}^1 \mathbf{x}^{SR}(s, \theta) \cdot \gamma(s, \theta) \Big|_{\theta=0}^1 - \int_{\theta=0}^1 \frac{d}{d\theta} \mathbf{x}^{SR}(s, \theta) \cdot \gamma(s, \theta) d\theta ds, \end{aligned}$$

using integration by parts. The first term of the above expression is zero. For $\theta = 0$, we have $\gamma = 0$ since the interval of integration is empty. For $\theta = 1$, $x_1(s, 1) = x_2(s, 1)$ by symmetry, and $\hat{\phi}_1 - \phi_1 + \hat{\phi}_2 - \phi_2 = 0$ by the assumption of the lemma. Therefore, the left hand side of inequality (3.6) simplifies to

$$(3.7) \quad = - \int_{s=0}^1 \int_{\theta=0}^1 \frac{d}{d\theta} \mathbf{x}^{SR}(s, \theta) \cdot \gamma(s, \theta) d\theta ds.$$

Lemma 21 shows that $-\frac{d}{d\theta} \mathbf{x}^{SR}(s, \theta) \cdot (1, -1) \geq 0$, and therefore we can complete the proof by showing that $\gamma_1 = -\gamma_2$ and $\gamma_1 \geq 0$ everywhere. The first property holds by the first assumption of the lemma, since $\phi_1 + \phi_2 = \phi_{sum}(t_1 + t_2) = \hat{\phi}_1 + \hat{\phi}_2$ implies that $\hat{\phi}_1 - \phi_1 = \phi_2 - \hat{\phi}_2$. Also, $\phi_1 + \phi_2 = \hat{\phi}_1 + \hat{\phi}_2$ together with $\hat{\phi}_1 \frac{t_2}{t_1} = \hat{\phi}_2$ (by definition) and $\phi_1 \frac{t_2}{t_1} \leq \phi_2$ (by assumption of the lemma) implies that $\hat{\phi}_1 \geq \phi_1$. \square

Lemma 21. *The assignment of any incentive compatible mechanism satisfies*

$$\frac{d}{d\theta} \mathbf{x}^{SR}(s, \theta) \cdot (-1, 1) \geq 0$$

Proof. Incentive compatibility implies that for any s, θ , and ϵ ,

$$(\mathbf{x}^{SR}(s, \theta + \epsilon) - \mathbf{x}^{SR}(s, \theta)) \cdot (-1, 1) \frac{\theta \epsilon}{s} \geq 0.$$

Letting ϵ approach zero implies the claim. \square

So in order to complete the proof that $\hat{\phi}$ is a weak amortization, we should construct strong amortization ϕ satisfying conditions of Lemma 20. Recall that $\alpha_1 \frac{t_2}{t_1} \geq \alpha_2$ implies that $\phi_1 \frac{t_2}{t_1} \leq \phi_2$ because

$$\frac{t_2}{t_1} \phi_1(\mathbf{t}) = \frac{t_2}{t_1} \left(t_1 - \frac{\alpha_1(\mathbf{t})}{f(\mathbf{t})} \right) = \frac{t_2}{t_1} \left(t_1 - \frac{\alpha_1(\mathbf{t})}{f(\mathbf{t})} \right) \leq t_2 - \frac{\alpha_2(\mathbf{t})}{f(\mathbf{t})} = \phi_2(\mathbf{t}).$$

So the goal is to construct ϕ such that $\alpha_1 \frac{t_2}{t_1} \geq \alpha_2$. Extend the definition of the projection for the sum of values as follows. Define $F_{sum}(v, \theta)$ to be the probability of the set of types \mathbf{t} such that $t_1 + t_2 \geq v$, and $t_2/t_1 \leq \theta$. This implies that $F_{sum}(v) = F_{sum}(v, 1)$. Define $f_{sum}(v, \theta)$ to be the density function of the distribution, that is $f_{sum}(v, \theta) = \frac{d}{dv} F_{sum}(v, \theta)$.

Lemma 22. *There exists unique strong amortization $\phi = tf - \alpha$ satisfying $\phi_1 + \phi_2 = \phi_{sum}(t_1 + t_2)$. This amortization also satisfies*

$$\theta \alpha_1^{SR}(s, \theta) - \alpha_2^{SR}(s, \theta) = -(1 + \theta)(1 - F_{sum}(s)) \frac{d}{ds} \left[\frac{f_{sum}(s, \theta)}{f_{sum}(s)} \right].$$

Proof. We will here sketch the proof of the above lemma, and the complete proof appears in Section A.2.

The proof is similar to the proof of Lemma 14, and takes the derivative of the divergence theorem. We assume ϕ exists, use the divergence theorem and properties of ϕ to derive a closed

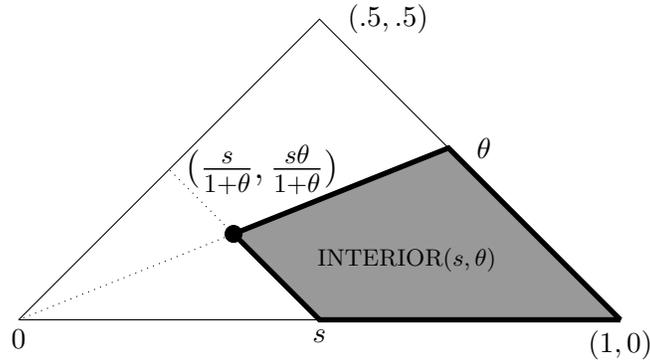


Figure 3.3. The trapezoidal set parameterized by s and θ , and the four curves that define its boundary, $\{\text{TOP, RIGHT, BOTTOM, LEFT}\}(s, \theta)$

form for it, and then verify that ϕ indeed satisfies all the required properties. We apply the divergence theorem to α on the trapezoidal subspace of type space defined by types \mathbf{t}' with $t'_1 + t'_2 \geq s$, $t'_2/t'_1 \leq \theta$, $t'_2 \geq 0$, and $t'_1 + t'_2 \leq 1$ (Figure 3.3). The divergence theorem equates the integral of the vector field α on the boundary of the subspace to the integral of its divergence within the subspace. The integral of the vector field over right and bottom boundaries equates zero by boundary orthogonality, and its value over the left boundary will be specified given the equation $\phi_1 + \phi_2 = \phi_{sum}(t_1 + t_2)$ (the outward pointing vector is $(-1, -1)$). As the upper boundary of this trapezoidal subspace has slope t_2/t_1 , one term in this equality is the integral of $\alpha(\mathbf{t}')$ with the upward orthogonal vector to \mathbf{t} . Differentiating this integral and evaluating at $\mathbf{t}' = (\frac{s}{1+\theta}, \frac{s\theta}{1+\theta})$ gives a closed form expression for $\frac{t_2}{t_1}\alpha_1 - \alpha_2$. We then verify that the specified ϕ satisfies all the properties. \square

A distribution f is SR-log-submodular if it is log-submodular in sum-ratio coordinates, that is,

$$f^{SR}(s, \theta) \times f^{SR}(s', \theta') \geq f^{SR}(s', \theta) \times f^{SR}(s, \theta'), \quad \forall s \leq s', \theta \leq \theta',$$

and is SR-independent if the above holds with equality everywhere. Similar to Section 3.4 we prove the following lemma (notice that the sign is the opposite of the sign in Section 3.4, which is the reason that supermodularity is replaced by sub modularity).

Lemma 23. *If f is a SR-log-submodular function, then for any s and θ ,*

$$\frac{d}{ds} \left[\frac{f_{sum}(s, \theta)}{f_{sum}(s)} \right] \geq 0,$$

with equality if distribution is SR-independent.

Theorem 9. *Bundle pricing is revenue optimal for any SR-log-submodular distribution for which the sum-of-values projection has monotone strong amortization ϕ_{sum} .*

3.6. Multi Agent Extensions

3.6.1. The Setting

A multi-agent problem is defined by n agents, each agent i associated with a distribution f^i , and a feasibility setting $\mathcal{S} \subseteq \{0, 1\}^{n \times m}$. Types of agents are drawn independent of each other from corresponding distributions. A multi-agent mechanism is a sequence of assignment functions $(\bar{\pi}^1(\mathbf{t}^1, \dots, \mathbf{t}^n), \dots, \bar{\pi}^n(\mathbf{t}^1, \dots, \mathbf{t}^n)) \in \mathcal{S}$ in which $\bar{\pi}^i(\mathbf{t}^1, \dots, \mathbf{t}^n) \in \mathbb{R}^m$ is the assignment of agent i , together with a sequence of payment function $(\bar{p}^1(\mathbf{t}^1, \dots, \mathbf{t}^n), \dots, \bar{p}^n(\mathbf{t}^1, \dots, \mathbf{t}^n)) \in \mathbb{R}^n$ in which $\bar{p}^i(\mathbf{t}^1, \dots, \mathbf{t}^n) \in \mathbb{R}$ is the payment of agent i . A multi-agent mechanism is *ex-post individually rational* (EPIR) if the utility of every type of every agent is always positive,

$$\mathbf{t}^i \cdot \bar{\pi}^i(\mathbf{t}^i, \mathbf{t}^{-i}) - \bar{p}^i(\mathbf{t}^i, \mathbf{t}^{-i}) \geq 0, \quad \forall i, \mathbf{t}^i, \mathbf{t}^{-i},$$

where \mathbf{t}^{-i} is the vector of other agents types. A multi-agent mechanism is *dominant strategy incentive compatible* (DSIC) if no type of any agent increases his utility by misreporting,

$$\mathbf{t}^i \cdot \bar{\pi}^i(\mathbf{t}^i, \mathbf{t}^{-i}) - \bar{p}^i(\mathbf{t}^i, \mathbf{t}^{-i}) \geq \mathbf{t}^i \cdot \bar{\pi}^i(\hat{\mathbf{t}}^i, \mathbf{t}^{-i}) - \bar{p}^i(\hat{\mathbf{t}}^i, \mathbf{t}^{-i}), \quad \forall i, \mathbf{t}^i, \hat{\mathbf{t}}^i, \mathbf{t}^{-i}.$$

The *interim mechanism* for agent i is a single-agent mechanism $(\boldsymbol{\pi}^i, p^i)$ defined to be

$$\begin{aligned} \boldsymbol{\pi}^i(\mathbf{t}^i) &= E_{\mathbf{t}^{-i} \sim f^{-i}}[\bar{\boldsymbol{\pi}}^i(\mathbf{t}^i, \mathbf{t}^{-i})], \\ p^i(\mathbf{t}^i) &= E_{\mathbf{t}^{-i} \sim f^{-i}}[\bar{p}^i(\mathbf{t}^i, \mathbf{t}^{-i})], \end{aligned}$$

where \mathbf{t}^{-i} is the vector of other agents types, drawn at random from $f^{-i} := f^1 \times \dots \times f^{i-1} \times f^{i+1} \times \dots \times f^n$. A mechanism is *interim individually rational* (IIR) if the interim mechanism of every agent is individually rational. A mechanism is *Bayesian incentive compatible* (BIC) if the interim mechanism of every agent is incentive compatible. Notice that EPIR implies IIR, and DSIC implies BIC. The goal is to design feasible IIC and IIR mechanisms that maximize expected payments. We however prove stronger results that the optimal mechanisms are in fact EPIR and DSIC.

3.6.2. Problem Formulation

A single agent mechanism $(\boldsymbol{\pi}, p)$ defines a utility function $u(\mathbf{t}) = \mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - p(\mathbf{t})$. Recall the following lemma from Subsection 3.2.3 which connects the utility function of an IC mechanism with its assignment function.

Lemma 3 (Rochet (1985)). *Function u is the utility function of an agent in an incentive compatible mechanism if and only if u is convex, and in that case, the agent's assignment is $\boldsymbol{\pi}(\mathbf{t}) = \nabla u(\mathbf{t})$.*

The payment function can be defined using the utility function and the assignment function as $p(\mathbf{t}) = \mathbf{t} \cdot \boldsymbol{\pi}(\mathbf{t}) - u(\mathbf{t})$. Using the above lemma, we can write payment to be $p(\mathbf{t}) = \mathbf{t} \cdot \nabla u(\mathbf{t}) - u(\mathbf{t})$. This allows us to write the following optimization problem as our revenue maximization problem, which is the starting point of the analysis of this section.

$$(3.8) \quad \max_{\bar{\boldsymbol{\pi}}^1, \dots, \bar{\boldsymbol{\pi}}^n, u^1, \dots, u^n} \quad \sum_i \int_{\mathbf{t}} [\mathbf{t} \cdot \nabla u^i(\mathbf{t}) - u^i(\mathbf{t})] f^i(\mathbf{t}) \, d\mathbf{t}$$

$$\forall i, \boldsymbol{\pi}^i = \nabla u^i; u^i \text{ is convex,}$$

$$\forall \mathbf{t}, (\bar{\boldsymbol{\pi}}^1(\mathbf{t}^1, \dots, \mathbf{t}^n), \dots, \bar{\boldsymbol{\pi}}^n(\mathbf{t}^1, \dots, \mathbf{t}^n)) \text{ is feasible assignment.}$$

Notice that when $m = 1$, the above program is equivalent to the following familiar form of

$$\max_{\bar{\boldsymbol{\pi}}^1, \dots, \bar{\boldsymbol{\pi}}^n} \quad \sum_i \int_v [v \boldsymbol{\pi}^i(v) - \int_{z \leq v} \boldsymbol{\pi}^i(z) dz] f(v) \, dv$$

$$\forall i, \boldsymbol{\pi}^i \text{ is monotone non-decreasing,}$$

$$\forall \mathbf{t}, (\bar{\boldsymbol{\pi}}^1(\mathbf{t}^1, \dots, \mathbf{t}^n), \dots, \bar{\boldsymbol{\pi}}^n(\mathbf{t}^1, \dots, \mathbf{t}^n)) \text{ is feasible assignment.}$$

Now assume that for each agent i , we have constructed $\boldsymbol{\alpha}^i(\mathbf{t})$ according to Lemma 11. The above analysis and Program 3.8 tells us that we can write revenue as

$$\begin{aligned} \sum_i \int_{\mathbf{t}^i} \boldsymbol{\pi}^i(\mathbf{t}^i) \cdot \boldsymbol{\phi}^i(\mathbf{t}^i) f^i(\mathbf{t}^i) \, d\mathbf{t}^i &= \sum_i \int_{\mathbf{t}^i} \left[\int_{\mathbf{t}^{-i}} \bar{\boldsymbol{\pi}}^i(\mathbf{t}^i, \mathbf{t}^{-i}) f^{-i}(\mathbf{t}^{-i}) \, d\mathbf{t}^{-i} \right] \cdot \boldsymbol{\phi}^i(\mathbf{t}^i) f^i(\mathbf{t}^i) \, d\mathbf{t}^i \\ &= \sum_i \int_{\mathbf{t}} \bar{\boldsymbol{\pi}}^i(\mathbf{t}^i, \mathbf{t}^{-i}) \cdot \boldsymbol{\phi}^i(\mathbf{t}^i) f(\mathbf{t}) \, d\mathbf{t} \\ &= \int_{\mathbf{t}} \left[\sum_i \bar{\boldsymbol{\pi}}^i(\mathbf{t}^i, \mathbf{t}^{-i}) \cdot \boldsymbol{\phi}^i(\mathbf{t}^i) \right] f(\mathbf{t}) \, d\mathbf{t}. \end{aligned}$$

We say that a mechanism $\bar{\pi}$ optimizes an objective ϕ *ex-post*, if for any input $\mathbf{t}^1, \dots, \mathbf{t}^m$ it selects an outcome that maximizes $\sum_i \bar{\pi}^i(\mathbf{t}^1, \dots, \mathbf{t}^m) \cdot \phi^i(\mathbf{t}^i)$.

Notice that in a multi-agent problem, the existence of α^i depends only on the distribution f^i . As a result, in the rest of this section we study construction of α^i for single agents, and drop the index i . By showing the existence of proofs for each agent separately, we can apply the meta theorem to get the structure of the multi-agent optimal solution.

3.6.3. Ex-post Optimization and Ironing

Having proved that ϕ is a strong amortization of revenue by proving it satisfies conditions of Lemma 11, we now study its incentive compatibility in multi-agent settings. First we define *service-constrained* settings. A service-constrained environment is characterized by a set system $\mathcal{S} \subseteq 2^n$. An assignment is feasible if $\{i \mid \sum_{j=m} x_{ij} = 1\} \in \mathcal{S}$. A service-constrained environment is downward-closed if for any $S, S' \subseteq S$ such that $S \in \mathcal{S}$, we also have $S' \in \mathcal{S}$.

Notice that $\phi_2 \leq \frac{t_2}{t_1} \phi_1$ implies that when $\phi_1 \geq 0$, we have $\phi_2 \leq \phi_1$, and when $\phi_1 \leq 0$, we also have $\phi_2 \leq 0$. This implies that in service-constrained settings, an ex-post optimizer of ϕ will assign positive probability only to the favorite item. In addition, if $\phi(t)$ is a non-decreasing function of t_1 , then so will be the assignment probability for the favorite item. This implies that the assignment rule is incentive compatible because

$$\begin{aligned} (\boldsymbol{\pi}(\mathbf{t}) - \boldsymbol{\pi}(\mathbf{t}')) \cdot (\mathbf{t} - \mathbf{t}') &= (\boldsymbol{\pi}_1(\mathbf{t}) - \boldsymbol{\pi}_1(\mathbf{t}'))(t_1 - t'_1) \\ &\geq 0, \end{aligned}$$

which shows the convexity of the utility function, which in turn proves incentive compatibility (see Lemma 3). We have therefore proved the following lemma.

Lemma 24. *If $\phi(\mathbf{t})$ is a non-decreasing function of t_1 , and $\phi_1 \geq \phi_2$ whenever $\phi_2 \geq 0$, then ϕ is incentive compatible for any multi-agent downward-closed service-constrained environment.*

Unfortunately, as it happens in single-dimensional settings, the derivative of revenue is not necessarily monotone. However, we can fix this problem by *ironing* the virtual value function in a similar manner to the single-dimensional problem. An ironing of multi-dimensional virtual values should be performed carefully in order to keep the properties of virtual values. More precisely, we would like to maintain the property that $\phi_2 \leq \phi_1$ when $\phi_2 \geq 0$. In particular, we can prove the following lemma (proof is in Section A.2).

Lemma 25. *There exists an ironed virtual value function ϕ that is a weak amortization of revenue and is incentive compatible (by satisfying the conditions of Lemma 24) for MR-log-supermodular distributions.*

The above lemma directly implies our main theorem.

Theorem 10. *In multi-agent downward-closed service-constrained environments with MR-log-supermodular distributions, the optimal auction selects a feasible set maximizing the sum of virtual values for favorite items.*

CHAPTER 4

Single-Agent to Multi-Agent Solutions: The Structure

Overview and Organization. This chapter studies the problem of composing single-agent solutions to form a multi-agent solution. In particular, we investigate the role of marginal revenue maximization as a rule to combine single-agent solutions. Our contributions are three fold: we characterize the settings for which marginal revenue maximization is optimal (by identifying an important condition that we call *revenue linearity*), we give simple procedures for implementing marginal revenue maximization in general, and we show that marginal revenue maximization is approximately optimal. Our approximation factor smoothly degrades in a term that quantifies how far the environment is from an ideal one (i.e., where marginal revenue maximization is optimal). Because the marginal revenue mechanism is optimal for well-studied single-dimensional agents, our generalization immediately approximately extends many results for single-dimensional agents to more general preferences.

In Section 2.2 we reviewed the Myerson-Bulow-Roberts single-dimensional linear agent model and gave a new proof that the marginal revenue mechanism is revenue optimal. The proof followed from an argument that for single-dimensional linear agents a class of single-agent lottery pricing problems satisfies a natural revenue-linearity property. This chapter is the generalization of the result to service constrained setting. In Section 4.1 we discuss the setting, the marginal revenue mechanism, and an overview of main results. In Section 4.2 we formalize the service constrained model for general preferences and generalize the marginal revenue derivation to general preferences that satisfy the previously identified revenue-linearity property. In Section 4.3 we give general methods for implementing the marginal revenue mechanism for general preferences

regardless of revenue linearity, and in Section 4.4 we show that approximate revenue linearity, properly defined, implies approximate optimality. In Section 4.5 we suggest numerous extensions of results in the single-dimensional mechanism design literature to general preferences that are direct consequences of the marginal revenue mechanism framework.

4.1. Introduction

Marginal revenue plays a fundamental role in microeconomic theory. For example, a monopolist providing a commodity to two markets each with its own concave revenue (as a function of the supply provided to that market) optimizes her profit by dividing her total supply to equate the marginal revenues across the two markets. Moreover this central economic principle also governs classical auction theory. Myerson (1981) characterizes profit maximizing single-item auction as formulaically optimizing the *virtual value* of the winner; Bulow and Roberts (1989) reinterpret Myerson's virtual value as the marginal revenue of a certain concave revenue curve.

Because it is simple and intuitive, the Myerson-Bulow-Roberts approach provides the basis for most of Bayesian auction theory. Unfortunately though, this theory has been limited to settings where agents have linear single-parameter preferences, i.e., where an agent's utility is given by her value for service less her payment. Consequently, Bayesian auction theory is often similarly limited. With more general forms of agent preferences; especially multi-dimensionality, e.g., for multi-item auctions, or non-linearity, e.g., risk aversion or budgets; auction theory is complex, less versatile, and often not well understood.

Our main result is to show that hidden under the complexity of optimal mechanism design problems for agents with multi-dimensional and non-linear (henceforth: general) preferences is marginal revenue maximization. The approach of marginal revenue maximization decomposes a multi-agent mechanism design problem as a composition of simple single-agent mechanism design problems, specifically, from the construction of the appropriate notion of revenue curves.

This new approach for general preferences uncovers a condition we refer to as *revenue linearity* that is satisfied by all linear single-dimensional preferences and governs the performance of the marginal revenue mechanism more generally. When the single-agent problems are revenue linear, marginal revenue maximization is optimal and the Myerson-Bulow-Roberts mechanism generalizes exactly. When the single-agent problems are approximately revenue linear, marginal revenue maximization is approximately optimal (though the composition of the single-agent mechanisms to implement marginal revenue maximization requires new techniques). Finally, because the marginal revenue approach is structurally similar to the classical approach, many results from classical auction theory approximately and automatically extend to general preferences.

A central result in classical auction theory is derived from an interpretation of the Myerson-Bulow-Roberts mechanism (i.e., for maximizing marginal revenue) in the special case of symmetric agents. Our generalization admits a similar interpretation. In the classical setting there is a single item for sale and agents with i.i.d. values for it; in our setting there is a single item for sale which the seller can configure in one of several ways and agents have i.i.d. values for each configuration, e.g., a car that can be painted red or blue (importantly, the seller sets the configuration and the buyer cannot change it).¹

Selling a car.: Classical auction theory says that (a) the optimal way to sell an object (henceforth: a car) to a single agent with value drawn from a uniform distribution on $[0, 1]$ is to post a take-it-or-leave-it price of $1/2$, (b) the optimal way to sell a car to one of multiple agents with uniformly distributed values is to run a second-price auction with reserve price $1/2$, and (c) more generally the optimal way to sell the car

¹The red-or-blue car example is slightly unnatural as a forward auction (i.e., when the auctioneer is selling); however, the analogous reverse auction (i.e., the auctioneer is buying) is an important problem in procurement. For instance the government may wish to hire a contractor to build a bridge. Contractors can build different kinds of bridges. From bids of the contractors over the different bridges the auctioneer selects a kind of bridge to procure, which contractor to procure it from, and how much is to be paid. Our results for reverse auctions are analogous to those for forward auctions; interested readers can find the details in Section B.5.

to multiple agents with i.i.d. values is to run the second price auction with the same reserve price that would be offered as a take-it-or-leave-it price to one agent (assuming the distribution satisfies a mild assumption).

Selling a red-or-blue car.: Consider selling a car that, on sale, can be painted one of two colors, red or blue.² Our theory says that (a) the optimal way to sell a red-or-blue car to a single agent with values for the different colors each drawn independently and uniformly from $[0, 1]$ is to post a take-it-or-leave-it price of $\sqrt{1/3}$ for either color, (b) the optimal way to sell a red-or-blue car to one of multiple agents each with i.i.d. uniform values for each color is to run the second-price auction with reserve $\sqrt{1/3}$ and allow the winning agent to choose her favorite color on sale, and (c) more generally to sell a red-or-blue car to one of multiple agents each with values drawn i.i.d. (from a distribution that satisfies the same mild assumption as above) for each color, the second price auction with the reserve price equal to the same price that would be offered to a single agent is (at worst) a 4-approximation to the optimal auction, that is, the second price auction with reserve achieves at least a quarter of the revenue of the optimal auction.

It should be noted that reducing a multi-dimensional preference to a single-dimensional preference by always selling the winning agent her favorite color is very natural and practical; however, it is not generally optimal. For example, when the agent's values for each color is distributed uniformly on $[5, 6]$, the analysis of Thanassoulis (2004) shows that the optimal auction does not sell the agent her favorite color subject to a reserve (in fact, it is not even deterministic). However, many relevant distributions, including the uniform distribution on $[5, 6]$, satisfy the mild assumption sufficient for the four approximation, above.

²In this example we give the reserve price for $m = 2$ colors; however, with the appropriate reserve price these results hold for any number of colors.

Approach. We focus on service constrained environments where, in any outcome the mechanism produces, each agent is either considered served or unserved. The designer has a feasibility constraint that governs which subset of agents can be simultaneously served, but the other aspects of the outcome, e.g., payments, are unconstrained. This model allows additional unconstrained attributes of the service (e.g., the color of the car in the previous red-or-blue car example, or the grade or quality of a service). We assume that the space of mechanisms is closed under convex combination, which allows for randomized mechanisms.

The agents in the mechanism have independently but not necessarily identically distributed preferences (a.k.a., types). We do not place any assumption on the agent preferences other than they are expected utility maximizers. This includes the most challenging preference models in Bayesian mechanism design such as multi-dimensionality, public or private budgets, and risk-aversion (e.g., as given by a concave utility function).

Revenue curves result from the following single-agent mechanism design problem. Consider a single agent with private type drawn from a known distribution. Via the taxation principle (see e.g., Wilson, 1997) the outcomes of a mechanism, for all possible reports the agent might make in the mechanism, can be viewed as a menu where the agent selects her favorite outcome by making the appropriate report. This menu may contain outcomes that are randomized and for this reason we refer to it as a *lottery pricing*. Ex ante, i.e., in expectation over the distribution of the agent's type, a lottery pricing induces a probability with which the agent receives an outcome that corresponds to service, and an expected payment, i.e., revenue.

As every lottery pricing induces an ex ante service probability and expected revenue, we can ask the optimization question of identifying the lottery pricing with a given ex ante service probability that has the highest expected revenue. As a function of the ex ante service probability this optimal revenue induces a *revenue curve*. Important in the construction of revenue curves are the lottery pricings, i.e., single-agent mechanisms, that give the optimal revenue for each

ex ante service probability. As the space of (mechanisms and hence) lottery pricings is closed under convex combination, the revenue curves are always concave. The marginal revenue curve is given by the derivative of the revenue curve with respect to ex ante service probability.

As discussed in the opening paragraph, the standard economic intuition suggests that a monopolist splitting the sale of a commodity between two markets should do so to equate marginal revenue. There is an intuitive algorithmic reinterpretation of this fact. If we break the allocation to each market into tiny pieces ordered by willingness to pay and attribute to each piece the change in revenue from adding that piece (i.e., the marginal revenue), then the total revenue of an allocation is the sum of the marginal revenues of each piece. A simple algorithm for optimizing this *surplus of marginal revenue* is to repeatedly allocate a tiny amount to the market that has the highest marginal revenue at its current allocation (until the good is totally allocated or marginal revenues are non-positive). Clearly this results in a final allocation where the markets marginal revenues are roughly equal as in the microeconomic interpretation. This allocation is optimal.

The main contribution of this chapter is a methodology for constructing multi-agent mechanisms from the simple single-agent lottery pricings that define the revenue curve. The main task of such a construction is to specify a method for combining the single-agent mechanisms into a multi-agent mechanism that is both feasible with respect to the service constraint and obtains good revenue.

Definition 11. The family of *marginal revenue mechanisms* take the following form:

- (1) Map each agent's private type (which may lie in an arbitrary type space) to a *quantile* in $[0, 1]$.
- (2) Calculate the marginal revenue of each agent as the derivative of the revenue curve at her quantile.

- (3) Select for service the set of agents that maximize the *surplus of marginal revenue*, i.e., the total marginal revenue of agents served, subject to feasibility.
- (4) Calculate for each agent the appropriate non-service aspects of the outcome, e.g., payments.

Thus far in the discussion only Steps 2 and 3 should be clear. The remaining steps are non-trivial in general and a main issue that we will be resolving. For the special case of the selling-a-car example; where the agents' values are independently, identically, and uniformly drawn from the $[0, 1]$ interval; the marginal revenue mechanism is instantiated as follows.

For an agent with value drawn uniformly from the $[0, 1]$ interval, the optimal lottery pricing for ex ante service probability \hat{q} is to post a take-it-or-leave-it price of $\hat{v} = 1 - \hat{q}$. The revenue from such pricing is the price times the probability that it is accepted. Therefore, the revenue curve is $R(\hat{q}) = (1 - \hat{q}) \times \hat{q}$ and the marginal revenue curve is its derivative $R'(\hat{q}) = 1 - 2\hat{q}$.

The optimal lottery for ex ante probability \hat{q} serves the agent if her value v is on interval $[1 - \hat{q}, 1]$. This is the strongest \hat{q} measure of the values from the distribution. This motivates, in Step 1, mapping value v to quantile $q = 1 - v$. Composing this mapping from value to quantile with the above mapping from quantile to marginal revenue gives a mapping from value v to marginal revenue as $2v - 1$.

For a single-item auction, in Step 3 the surplus of marginal revenue is maximized by serving nobody if all have negative marginal revenues and, otherwise, by serving the agent with the highest marginal revenue. As the agents are symmetric and marginal revenue is monotone in value, equivalently, the highest-valued agent wins as long as her value is at least $1/2$ (the value for which marginal revenue is zero, i.e., solving $2v - 1 = 0$).

The appropriate calculation of payments for Step 4 is the following. All losers have payments equal to zero. The payment of a winner is the minimum value she could declare and still win in Step 3, i.e., it is the maximum of the the second highest agent value and $1/2$.

This auction, as claimed in the earlier discussion of the selling-a-car example, is the second-price auction with reserve $1/2$. Moreover, the mapping from value to marginal revenue is identical to the virtual values in the derivation of Myerson (1981).

Results. This chapter generalizes the marginal-revenue approach for agents with single-dimensional linear preferences (Bulow and Roberts, 1989) to general preferences. Our main algorithmic contribution is to generalize Steps 1 and 4 thereby allowing the construction of service constrained multi-agent mechanisms from single-agent ex ante lottery pricings. There are a number of challenges in this endeavor. First, revenue equivalence does not hold for general preferences (which is used in the proof of optimality for single-dimensional preferences).³ Second, there is not a natural ordering on types for general preferences (making it difficult to map types to quantiles). Third, the set of agents served by the marginal revenue mechanism may be randomized. None of these issues are present for single-dimensional linear preferences.

Orthogonal to the question of implementing the marginal revenue mechanism for general preferences are questions of quantifying its performance. Via the Myerson-Bulow-Roberts analysis it is known that for single-dimensional linear preferences, the marginal revenue mechanism is optimal. As a first step in understanding the performance of the mechanism more generally we give a new derivation of the optimality for single-dimensional agents. Our derivation exposes a previously unobserved property of single-dimensional linear preferences which we refer to as *revenue linearity*. Generally, i.e., beyond single-dimensional linear preferences, the optimality of the marginal revenue mechanism is implied by revenue linearity. Moreover, if the single-agent

³For single-dimensional linear preference agents, the revenue equivalence theorem states that any two auctions with the same allocation in expectation have the same expected revenue.

problems are α -approximately revenue linear, then the marginal revenue mechanism is an α approximation to the optimal mechanism.

Revisiting our red-or-blue car example above, (a) is a description of the optimal unconstrained lottery pricing, (b) is a consequence of the revenue-linearity of unit-demand preferences that are uniformly distributed on a multi-dimensional hypercube, and (c) is a consequence of 4-approximate revenue linearity for agents with unit-demand preferences drawn from any product distribution.

One of the main benefits of considering the marginal revenue mechanism for approximately optimal mechanism design is that, as its structure is similar to optimal mechanisms for single-dimensional environments, many results from the extensive single-dimensional mechanism design literature can be easily generalized. The following are some of the most important consequences.

Algorithmic mechanism design.: When weighted optimization is hard we can replace an exact algorithm for weighted maximization (Step 3 of Definition 11) with any approximation algorithm using either of the single-dimensional black-box reductions of Hartline and Lucier (2010) and Hartline et al. (2011).

Sequential posted pricing.: Sequential posted pricing mechanisms of Chawla et al. (2010b) and Yan (2011) that are approximately optimal for single-dimensional agents are approximately optimal for general agents (in the same service constrained environment) and the same approximation factor is guaranteed. Moreover, these sequential posted pricing bounds give another bound on the approximation factor of the marginal revenue mechanism. The marginal revenue mechanism is in fact optimal within a class of mechanisms that contains the sequential posted pricing mechanisms; therefore, its approximation factor is no worse. As an example, for the single-item service constraint, a sequential posted pricing bound implies an $e/(e-1)$ -approximation regardless of approximate revenue linearity of the single-agent problems.

Simple versus optimal.: While our marginal revenue mechanism is already generally much simpler than the optimal mechanism, we can get even simpler approximation mechanisms by applying methods developed for single-dimensional preferences to prove that simple mechanisms approximate the marginal revenue mechanism. In particular, in single-dimensional environments maximizing marginal revenue is more complex than simple reserve-price-based mechanisms, i.e., mechanisms that maximize welfare subject to a reserve price. Nonetheless, Hartline and Roughgarden (2009) show that reserve-price-based mechanisms are often approximately optimal. When uniform pricing is approximately optimal, e.g., in generalizations of the red-or-blue car example, these mechanisms extend to general preferences.

Single-sample mechanisms.: Approaches above have been for Bayesian optimal mechanism design where the designer optimizes a mechanism given a distribution of preferences. Dhangwatnotai et al. (2010) relax the assumption that the distribution is known and show that a mechanism based on drawing a single sample from the distribution gives a good approximation to the Bayesian optimal mechanism. Again, the single-sample framework extends to general preferences for which uniform pricing is approximately optimal.

It is important to contrast the simplicity of the marginal revenue approach with recent algorithmic results in Bayesian mechanism design for general agent preferences. Chapter 5 and Cai et al. (2012b,c, 2013) give polynomial-time mechanisms for large classes of Bayesian mechanism design problems. Chapter 5 focuses on service-constrained settings, whereas Cai et al. (2012b,c, 2013) consider additive settings. The two main conclusions of Chapter 5 and these works is that (a) optimal mechanisms continue to have weighted maximization at their core, and (b) the appropriate weights (i.e., virtual values) are stochastic and can be solved for as a convex optimization problem, e.g., via the ellipsoid method, that takes into account the feasibility

constraint and the distribution over types of all agents. (This latter result is simply because the space of mechanisms is convex, any point in the interior of a convex set can be implemented by a convex combination of vertices, and vertices correspond to linear, a.k.a., weighted, optimization.) There are a number of important distinctions between our work and these algorithmic results. First, the weights in our derivation have a natural economic interpretation as marginal revenues. Second, the weights in our derivation can be found easily from solutions to the single-agent lottery pricing problems and are not derived from the solution to an additional multi-agent optimization problem. Third, in most cases, the weights in our derivation depend only on the single-agent problem and not on the multi-agent feasibility constraint or presence of other agents. Therefore, our approach affords significant structural simplification and interpretation that enables the consequences previously enumerated. Finally, one of the biggest open questions in the above algorithmic work is in developing approaches that are not brute-force in each agent’s type space. As an example that breaks this barrier, our approach gives approximately optimal mechanisms for multi-dimensional unit-demand agents with values from product distributions; these mechanisms are easy to compute with a computational complexity that scales linearly with the dimensionality of the type space (i.e., logarithmically in the size of the type space).

4.2. Multi-dimensional and Nonlinear Preferences⁴

Bayesian mechanism design. An agent has a private type t from type space T drawn from distribution F with density function f . The agent may be assigned outcome w from outcome space W . This outcome encodes what kind of service the agent receives and any payments she must make for the service. In particular the payment specified by an outcome w is denoted by $\text{Payment}(w)$. The agent has a von Neumann–Morgenstern utility function: for type t and deterministic outcome w her utility is $u(t, w)$, and when w is drawn from a distribution her

⁴The model and notation is consistent with Subsection 2.1.2, and is included here to ease reading.

utility is $\mathbf{E}_w[u(t, w)]$.⁵ We will extend the definition of the utility function to distributions over outcomes $\Delta(W)$ linearly. For a random outcome w from a distribution, $\text{Payment}(w)$ will denote the expected payment.

Example 6 (A unit-demand quasi-linear-utility agent). There are m alternatives and the agent's type is given by a vector (v^1, \dots, v^m) representing her value for each alternative. An outcome is of the form (p, π^1, \dots, π^m) , where p denotes the payment, and each $\pi^j \in \{0, 1\}$ indicates whether the agent gets the alternative j , with $\sum_j \pi^j \leq 1$. The agent's utility at such an outcome is then given by the linear form $\sum_j v^j \pi^j - p$. When randomizing over such outcomes, we relax the π^j 's to be in $[0, 1]$, still with $\sum_j \pi^j \leq 1$. Such a distribution with a price p is called a *lottery*.

Example 7 (A single-dimensional public-budget agent). The agent has a publicly known budget B , and her type is given by her private value v for an item being auctioned. An outcome $w = (x, p)$ indicates by $x \in \{0, 1\}$ whether the agent gets the item, and by p the amount of payment she makes. In contrast to the single-dimensional linear-utility agents of Section 2.2, this agent's utility is $v \cdot x - p$ only if $p \leq B$, and negative infinity otherwise.

There are n agents indexed $\{1, \dots, n\}$ and each agent i may have her own distinct type space T_i , utility function u_i , etc. The agents types are indepently distributed. A *direct revelation* mechanism takes as its input a profile of types $\mathbf{t} = (t_1, \dots, t_n) \in T_1 \times \dots \times T_n$ and outputs ex post outcome $\bar{\mathbf{w}}(\mathbf{t}) \in \Delta(W_1 \times \dots \times W_n)$. Agent i 's *ex post outcome rule* is denoted by $\bar{w}_i(\mathbf{t})$ and, with the other agents' types drawn from the distribution, her *interim outcome rule* $w_i(t_i)$ is distributed as $\bar{w}_i(t_i, \mathbf{t}_{-i})$ with $t_j \sim F_j$ for each $j \neq i$. We say that a mechanism is *Bayesian*

⁵This form of utility function allows for encoding of budgets and risk aversion; we do not require quasi-linearity.

incentive compatible if

$$(BIC) \quad u_i(t_i, w_i(t_i)) \geq u_i(t_i, w_i(t'_i)), \quad \forall i, \forall t_i, t'_i \in T_i.$$

A mechanism is *interim individually rational* if

$$(IIR) \quad u_i(t_i, w_i(t_i)) \geq 0, \quad \forall i, \forall t_i \in T_i.$$

The mechanism designer seeks to optimize an objective subject to BIC, IIR, and ex post feasibility. We consider the objective of expected revenue, i.e., $\mathbf{E}_t[\sum_i \text{Payment}(w_i(t_i))]$; however, any objective that separates linearly across the agents can be considered. Below we discuss the mechanism's feasibility constraint.

Service constrained environments. In a *service constrained environment* the outcome w provided to an agent is distinguished as being either a *service* or a *non-service* outcome, respectively, with $\text{Alloc}(w) = 1$ or $\text{Alloc}(w) = 0$. There is a feasibility constraint restricting the set of agents that may be simultaneously served; there is no feasibility constraint on how an agent is served. With respect to the feasibility constraint any outcome $w \in W$ with $\text{Alloc}(w) = 1$ is the same. For example, payments are part of the outcome but are not constrained by the environment. An agent may have multi-dimensional and non-linear preferences over distinct service and non-service outcomes.

From least rich to most rich, standard service constrained environments are *single-unit environments* where at most one agent can be served, *multi-unit environments* where at most a fixed number of agents can be served, *matroid environments* where the set of agents served must be an independent set of a given matroid, *downward-closed environments* where the set of agents served can be specified by an arbitrary set systems for which all subsets of a feasible set

are feasible, and *general environments* where the feasible subsets of agents can be given by an arbitrary set system that may not even be downward closed.

Ex Ante Lottery Pricings and Revenue Curves. The only aspect of the marginal revenue approach that translates identically from single-dimensional preferences to general preferences is the definition of the ex ante optimal pricing for allocation probabilities $\hat{q} \in [0, 1]$. This is the lottery pricing (i.e., collection of outcomes where the agent is permitted to choose her type-dependent favorite) denoted $\tilde{w}^{\hat{q}}(\cdot)$ that optimizes revenue subject to the constraint that $\mathbf{E}_t[\text{Alloc}(\tilde{w}^{\hat{q}}(t))] = \hat{q}$. The revenue curve for the agent is then $R(\hat{q}) = \mathbf{E}_t[\text{Payment}(\tilde{w}^{\hat{q}}(t))]$ as per Definition 4.

Allocation rules. The first challenge in generalizing the marginal revenue approach to general preferences is determining the mapping from types to quantiles. This challenge arises as there is no explicit ordering of an agent's type space T by strength. E.g., if the type is multi-dimensional then it is unclear which is stronger, a higher value in one dimension and lower in another or vice versa. In fact, which is stronger often depends on the context, e.g., the competition from other agents.

Our approach is based on two observations. First, relative to a mechanism and for a particular agent, the relevant part of the mechanism is the (interim) outcome rule $w(\cdot)$. For a given outcome rule $w(\cdot)$ an ordering on types by strength can be defined. Simply, a type that is more likely to be served is stronger than a type that is less likely to be served. I.e., t is stronger than t' relative to $w(\cdot)$ if $\text{Alloc}(w(t)) \geq \text{Alloc}(w(t'))$. This definition induces a mapping from the type space to quantile space; moreover, the distribution of quantiles induced by this mapping and the distribution on types is uniform.⁶ Second, (by the above mapping) any outcome rule $w(\cdot)$ induces an allocation rule $x(\cdot)$ that maps quantile to service probability. This allocation

⁶Quantiles are uniformly distributed when ties in allocation probability are measure zero; when there is a measurable probability of ties, quantiles can be defined by drawing uniformly from the interval containing the tie.

rule has a simple intuition in discrete type spaces: For each type $t \in T$ make a rectangle of width equal to the probability of the type $f(t)$ and height equal to the service probability of the type $\text{Alloc}(w(t))$. Sort the types in decreasing order of heights; the resulting monotone non-increasing piecewise constant function from $[0, 1]$ to $[0, 1]$ is the allocation rule. This is generalized for continuous distributions as follows.

Definition 12. For an agent with $t \in T$ drawn from distribution F and outcome rule $w(\cdot)$, the *allocation rule* mapping quantiles to service probabilities is given by $x(\hat{q}) = \sup\{y : \Pr_{t \sim F}[\text{Alloc}(w(t)) \geq y] \leq \hat{q}\}$.

Optimal Lottery Pricing. With the definition of allocation rules for any lottery pricing above, allocation constrained lottery pricings generalize naturally. Even though the order on types may change from one lottery pricing to another, we can still ask for the lottery pricing with the optimal revenue subject to a constraint on its allocation rule. The optimal lottery pricing for allocation constraint \hat{x} with cumulative allocation constraint \hat{X} is given by the outcome rule $w(\cdot)$ that optimizes expected revenue subject to its corresponding allocation rule x with cumulative allocation rule X satisfying $X(\hat{q}) \leq \hat{X}(\hat{q})$ for $\hat{q} \in [0, 1]$ with equality at $\hat{q} = 1$. As per Definition 5 the optimal revenue for allocation constraint \hat{x} is denoted $\text{Rev}[\hat{x}]$.

We will generally denote by x the optimal allocation rule for constraint \hat{x} . The ex ante constraint on total service probability by \hat{q} is given by the reverse step function at \hat{q} denoted $\hat{x}^{\hat{q}}$; the corresponding allocation rule of the \hat{q} *ex ante optimal pricing* is denoted $x^{\hat{q}}$.

Revenue Linearity and Marginal Revenue. Revenue linearity and marginal revenue have the same definitions (Definition 6 and Definition 7) as for single-dimensional preferences. The marginal revenue of an allocation constraint is $\text{MR}[\hat{x}] = \mathbf{E}_q[R'(q) \hat{x}(q)]$. By its construction as the revenue of the appropriate convex combination of ex ante optimal pricings it is a lower bound on the optimal revenue, i.e., $\text{Rev}[\hat{x}] \geq \text{MR}[\hat{x}]$. Again by its construction, revenue linearity

would imply that its revenue is equal to the optimal revenue (Theorem 1). We will describe the marginal-revenue approach for non-revenue-linear agents by analogy to the single-dimensional case.

Definition 13. The *single-dimensional analog* of a service constrained environment for general agents is the environment with single-dimensional linear agents with the same revenue curves. The *optimal marginal revenue* for a service constrained environment for general agents is the optimal revenue of the single-dimensional analog (which is equal to its surplus of marginal revenue).

Our approach to multi-agent mechanism design via the single-dimensional analog is to look at the profile of interim allocation rules induced by maximization of surplus of marginal revenue and then to construct a mechanism for general agents that looks to each agent like the convex combination of the ex ante optimal pricings for her allocation rule. For revenue curves R_1, \dots, R_n , draw quantiles $\mathbf{q} = (q_1, \dots, q_n)$ uniformly from $[0, 1]^n$, serve to maximize surplus of marginal revenue pointwise as $\sum_i R'_i(q_i) x_i$ for feasible $\mathbf{x} = (x_i[1], \dots, x_i[n])$. We can interpret the allocation rules induced by this process as allocation constraints for the general environment and denote them by $\hat{\mathbf{x}}^{MR} = (\hat{x}_1^{MR}, \dots, \hat{x}_n^{MR})$. As for single-dimensional linear agents (see Section 2.2), one way to serve an agent subject to allocation constraint \hat{x} is to draw a quantile \hat{q} from the distribution $G^{\hat{x}}$ with density $-\frac{d}{dq} \hat{x}(q)$ and run the ex ante optimal pricing for ex ante constraint \hat{q} . This approach suggests attempting to implement the general mechanism with outcome rules that correspond to allocation rules of the single-dimensional analog. Denoting the outcome rule for the \hat{q} ex ante optimal pricing for agent i by $w_i^{\hat{q}}(t_i)$. The agent's outcome rule corresponding to constraint \hat{x}_i^{MR} is $w_i^{MR}(t_i) = \int_0^1 w_i^{\hat{q}}(t_i) (-d\hat{x}_i^{MR}(q))$. There may be multiple ways to implement this profile of outcome rules ex post; however, the direct approach employed for single-dimensional linear agents in Section 2.2 does not always generalize.

Definition 14. The *marginal revenue outcome rule* of an allocation rule x is $w(t) = \int_0^1 w^{\hat{q}}(t) (-dx(q))$. A *marginal revenue mechanism* is one with interim outcome rules equal to the marginal revenue outcome rules corresponding to the optimal marginal revenue.

Implementation with Revenue Linearity. We show now that the marginal revenue mechanism generalizes exactly for general preferences that satisfy revenue linearity. Moreover, we show that in this case the marginal revenue mechanism inherits all of the nice properties of the marginal revenue mechanism for single-dimensional preferences. Namely, it deterministically selects the set of agents to serve, it is dominant strategy incentive compatible (truthful reporting is a best response for any actions of the other agents), and the mapping from types to quantiles to marginal revenues is deterministic and *context free*⁷ in that it does not depend on the feasibility constraint or other agents in the mechanism. The mechanism, however, is optimal among the larger class of randomized and Bayesian incentive compatible mechanisms. As motivation for this result, we will show subsequently that there are multi-dimensional preferences that are revenue linear, e.g., when multi-dimensional values are uniformly distributed on a hypercube.

The main challenge of implementing the marginal revenue mechanism is in specifying Step 1, i.e., the mapping from types to quantiles, and Step 4, i.e., selecting the appropriate outcomes for the set of agents that are served. If, however, each agent’s types are orderable by the following definition, then both steps are essentially identical to the single-dimensional case.

Definition 15. A single-agent problem is *orderable* if there is an equivalence relation on the types, and there is an ordering on the equivalence classes, such that for any allocation

⁷Note that this contrasts with recent algorithmic work in multi-dimensional optimal mechanism design where the optimal mechanism is characterized by mapping types stochastically to “virtual values” and this mapping is solved for from the feasibility constraint and the distributions of all agents types. See Chapter 5 and Cai et al. (2012b,c).

constraint \hat{x} , the optimal outcome rule w induces an allocation rule that is greedy by this ordering with ties between types in a same equivalence class broken uniformly at random.⁸

Orderability may look like a stringent and unlikely condition to hold generally. We note that it holds for single-dimensional agents and we show now, more generally, that it is a consequence of revenue linearity.

Theorem 11. *For any single-agent problem, revenue linearity implies orderability.*

The theorem is proved by the following two lemmas which characterize the structure of optimal lottery pricings.

Lemma 26. *For a revenue-linear single-agent problem, let x be the optimal allocation rule subject to some constraint \hat{x} . Then, for any \hat{q} such that $R''(\hat{q}) \neq 0$ we have $X(\hat{q}) = \hat{X}(\hat{q})$.*

Proof. Since x is the optimal allocation rule subject to \hat{x} , we have $\text{Rev}[x] = \text{Rev}[\hat{x}]$. Linearity implies that

$$\text{MR}[\hat{x}] = \int_0^1 x(q)R'(q) \, dq = \int_0^1 \hat{x}(q)R'(q) \, dq = \text{MR}[x].$$

Integrating by parts, we have

$$(4.1) \quad \left[X(q)R'(q) \right]_0^1 - \int_0^1 X(q)R''(q) \, dq = \left[\hat{X}(q)R'(q) \right]_0^1 - \int_0^1 \hat{X}(q)R''(q) \, dq.$$

Note that \hat{x} and x have the same ex ante probability of allocation $\hat{X}(1) = X(1)$; also by definition $X(0) = \hat{X}(0) = 0$. Combining these observations with (4.1) we have

$$\int_0^1 X(q)R''(q) \, dq = \int_0^1 \hat{X}(q)R''(q) \, dq,$$

⁸By greedy by the given ordering, we mean process each equivalence class in order and serve the corresponding types with as much probability as possible subject to the allocation constraint. If all equivalence classes are measure zero, then the resulting allocation rule is equal to the allocation constraint.

and therefore,

$$(4.2) \quad \int_0^1 [X(q) - \hat{X}(q)]R''(q) dq = 0.$$

Notice that for any q , $X(q) - \hat{X}(q)$ and $R''(q)$ are non-positive (by domination and concavity, respectively) so their product is non-negative. Therefore, (4.2) can be satisfied only if $[X(q) - \hat{X}(q)]R''(q) = 0$ for all q . This implies that if $R''(q) < 0$, then we must have $X(q) = \hat{X}(q)$, which completes the proof. \square

Lemma 26 in particular implies that for \hat{q} with $R''(\hat{q}) \neq 0$ the \hat{q} ex ante optimal pricing (i.e., with allocation constraint given by the reverse step function $\hat{x}^{\hat{q}}$) has allocation rule $x^{\hat{q}} = \hat{x}^{\hat{q}}$. I.e., the \hat{q} ex ante optimal pricing has only full lotteries (all types are served with either probability one or zero).

For any such \hat{q} , define $T_{\hat{q}}$ to be the set of types allocated (with full lotteries) in the optimal allocation subject to $\hat{x}^{\hat{q}}$. The following lemma shows that these sets are nested.

Lemma 27. *For a revenue-linear single-agent problem, for any $\hat{q}_1 > \hat{q}_2$ and $R''(\hat{q}_1), R''(\hat{q}_2) \neq 0$, we must have $T_{\hat{q}_1} \supseteq T_{\hat{q}_2}$.*

Proof. Assume for contradiction that $T_{\hat{q}_2} \setminus T_{\hat{q}_1} \neq \emptyset$. Let $\alpha = F(T_{\hat{q}_2} \setminus T_{\hat{q}_1}) > 0$. Consider the following allocation constraint

$$\hat{x}(q) = \begin{cases} 1 & q \leq \hat{q}_2 \\ 1/2 & \hat{q}_2 < q \leq \hat{q}_1 \\ 0 & \hat{q}_1 < q. \end{cases}$$

By revenue linearity, the revenue of the optimal auction subject to \hat{x} is $[R(\hat{q}_1) + R(\hat{q}_2)]/2$. Notice that the mechanism that runs $R(\hat{q}_1)$ and $R(\hat{q}_2)$ each with probability 1/2 achieves this revenue.

The allocation rule x of this mechanism is

$$x(q) = \begin{cases} 1 & q \leq q_2 - \alpha \\ 1/2 & q_2 - \alpha \leq q \leq q_1 + \alpha \\ 0 & q_1 + \alpha \leq q. \end{cases}$$

Notice that this allocation rule is dominated by \hat{x} , and achieves the optimal revenue. Yet, we have

$$\hat{X}(\hat{q}_1) = \int_{q=0}^{\hat{q}_1} \hat{x}(q) \, dq > \int_{q=0}^{\hat{q}_1} x(q) \, dq = X(\hat{q}_1).$$

This contradicts Lemma 26. □

PROOF OF THEOREM 11. By Lemma 27, all \hat{q} ex ante optimal pricings order the types by the same equivalence classes. By revenue linearity the optimal lottery pricing for an allocation constraint \hat{x} is a convex combination of the \hat{q} ex ante optimal pricings. Therefore, it allocates greedily to types by the same equivalence classes. □

Given orderability and the fact that (by Lemma 26) the optimal \hat{q} ex ante optimal pricings are full lotteries for \hat{q} for which $R(\hat{q})$ is locally non-linear, the marginal revenue mechanism is easy to define.

Definition 16. The *marginal revenue mechanism for orderable agents* works as follows.

- (1) Map reported types $\mathbf{t} = (t_1, \dots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \dots, q_n)$ via the implied ordering.⁹
- (2) Calculate the marginal revenue of each agent i as $R'_i(q_i)$.

⁹This ordering can be found by calculating the optimal single-agent mechanism for allocation constraint $\hat{x}(q) = 1 - q$.

- (3) For each agent i , calculate the maximum quantile \hat{q}_i that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).
- (4) Offer each agent i the \hat{q}_i ex ante optimal pricing.

Proposition 2. *The marginal revenue mechanism deterministically selects a feasible set of agents to serve and is dominant strategy incentive compatible.*

Proof. Because ties are broken consistently, critical values cannot fall in intervals where the revenue curve is locally linear (and the marginal revenue curve is locally constant). Therefore, the lottery pricings offered to each agent are full lotteries; each type is deterministically served or not served. Feasibility follows as the set of agents that select service outcomes is exactly the marginal revenue maximizing set subject to feasibility. To verify the dominant strategy incentive compatibility consider any agent i 's perspective. The parameter \hat{q}_i is a function only of the other agents' reports; the agent's outcome is determined by the \hat{q}_i ex ante optimal pricing which is incentive compatible for any \hat{q}_i . \square

Proposition 3. *In service constrained environment with revenue-linear agents, the marginal revenue mechanism obtains the optimal marginal revenue (which equals the optimal revenue).*

Testing Revenue Linearity. Revenue linearity is computationally easy to test. From the concavity of $\text{Rev}[\cdot]$ and equality of revenue and marginal revenue for allocation constraints $\hat{x}^{\hat{q}}$ which are a basis for general allocation constraints, it suffices to check the equality of revenue and marginal revenue, i.e., $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$, for any allocation constraint \hat{x} with positive derivative (\hat{x} as a convex combination of $\hat{x}^{\hat{q}}$ has positive density on each \hat{q}). For example, $\hat{x}(q) = 1 - q$ is such an allocation constraint. Since the theorem facilitates testing the property, we discretize the quantile space to $Q_N = \{0, \frac{1}{N}, \dots, \frac{N-1}{N}, 1\}$ for an arbitrary integer $N > 0$.

Theorem 12. *Let $\hat{x} : Q_N \rightarrow [0, 1]$ be any strictly decreasing function, (e.g., $\hat{x}(\hat{q}) = 1 - \hat{q}$). Then a set of single-agent pricings are revenue linear (or, more precisely, $\text{Rev}[\cdot]$ is a linear functional for non-increasing functions mapping Q_N to $[0, 1]$), if $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$.*

Proof. Consider the $N + 2$ reverse step functions that “steps down” from 1 to 0 at a point in Q_N . Any non-increasing function mapping Q_N to $[0, 1]$ is a convex combination of these base functions, and a strictly decreasing function can be written uniquely as such a convex combination. Therefore $\text{Rev}[\hat{x}] = \text{MR}[\hat{x}]$ amounts to saying that $\text{Rev}[\cdot]$ is linear on one interior point in a simplex, and the theorem states that $\text{Rev}[\cdot]$ is linear on the whole simplex. If we shift $\text{Rev}[\cdot]$ by a linear functional such that it is zero on all the base functions, then this theorem follows from the simple fact that, if a concave function g is 0 on all vertices of a simplex *and* one interior point A , then g is uniformly 0 on the simplex. To see this, suppose on point B in the simplex, $g(B) \neq 0$. By concavity, $g(B) > 0$. A can be written as a convex combination of B and vertices of the simplex with a strictly positive coefficient on B . (E.g., connect B and A with a straight line and extend it to intersect at one facet of the simplex formed by $N - 1$ vertices, then A can be written as a convex combination of B and these $N - 1$ vertices, where the coefficient on B in the decomposition is strictly positive.) But the concavity of g implies $g(A) > 0$, a contradiction. \square

Example 8 (A multi-dimensional revenue-linear example). The example of the seller who can paint her car red or blue as she sells it to agents with independent and uniform values for each color is revenue linear (proof given in Section B.4). Therefore, the marginal revenue mechanism is optimal and its simple form can be derived from Definition 16 as follows. For a unit-demand agent with values for m variants of a service (i.e., possible colors of the car) distributed uniformly on $[0, 1]^m$, we show that the ex ante optimal mechanism for constraint \hat{q} is to post a price of $\sqrt[m]{1 - \hat{q}}$ for any service. Notice that such a price will be accepted with

probability \hat{q} , and therefore the revenue function is $R(\hat{q}) = \hat{q} \sqrt[m]{1 - \hat{q}}$, and the marginal revenue function is $R'(\hat{q}) = (1 - \hat{q})^{1/m-1}(1 - \hat{q} - \hat{q}/m)$. The quantile of each type is $t = (t^1, \dots, t^m)$ to be $q = 1 - (\max_i t^i)^m$. Notice that both the mapping and the marginal revenue function are monotone. Therefore serving the agent with the highest marginal revenue (Definition 16) means serving the player with the highest value for any kind of service and charging her the minimum she needs to bid to exceed the second-highest value (subject to the reserve of $\sqrt[m]{\frac{1}{m+1}}$ which is where the marginal revenue becomes zero). Revenue-linearity implies that this mechanism is optimal.

4.3. Implementation

The marginal revenue mechanism for agents with orderable types (Definition 16) does not extend to general agents. In this section we give two approaches for defining the marginal revenue mechanism more generally. The first approach assumes that the parameterized family of \hat{q} ex ante optimal pricings satisfies a natural monotonicity requirement: that the probability that an agent with a given type is served is monotone in the ex ante constraint \hat{q} . Key to this construction is a *randomized* mapping from an agent's types to quantiles that is determined by the agent's type space and distribution alone, and is therefore context free, i.e., unaffected by the presence of other agents and the feasibility constraints. Consequently, (a) the resulting mechanism is dominant strategy incentive compatible but, (b) the set of winners is generally a randomized function of the profile of types. The second approach is brute-force but easily computable and completely general. It results in a Bayesian incentive compatible mechanism. Both these mechanisms will differ from the marginal revenue mechanism for orderable types only in the first (mapping types to quantiles) and last (serving each agent if her quantile is at most her critical quantile) steps; these changes can be mix-and-matched for different agents in the same mechanism.

We conclude this section by describing a relevant class of agents for which the ex ante optimal pricings satisfy the monotonicity property required by the first approach. The example considers single-dimensional agents with a public budget that constrains their maximum payment.

4.3.1. Monotone ex ante optimal pricings

We consider agents whose ex ante optimal pricings satisfy the following natural monotonicity property.

Definition 17. An agent has *monotone ex ante optimal pricings* if, given her type, the probability she wins in the \hat{q} ex ante optimal pricing is monotone non-decreasing in \hat{q} .

Suppose that the \hat{q} ex ante optimal pricing for an agent each consists of a menu of full lotteries. I.e., for any type of the agent she will choose a lottery that either serves her with probability one or zero. In this case the monotone ex ante optimal pricings assumption would require that the sets of types served for each \hat{q} are nested. There is a simple deterministic mapping from types to quantiles in this case: set the quantile of a type to be the minimum \hat{q} such that the \hat{q} ex ante optimal pricing serves the type. Below, we generalize this selection procedure to the case of partial lotteries (where types may be probabilistically served).

Recall that the \hat{q} ex ante optimal pricing, as a function of the agent's type, has an allocation and outcome rule $x^{\hat{q}}$ and $w^{\hat{q}}$, respectively. Fix the type of the agent as t and consider the function $G_t(\hat{q}) = x^{\hat{q}}(t)$ which, by the monotonicity condition above, can be interpreted as a cumulative distribution function. Recall that \hat{q} ex ante optimal pricing has probability of service $\mathbf{E}_t[x^{\hat{q}}(t)] = \hat{q}$. Therefore, if t is drawn from the type distribution and then q drawn from G_t then the distribution of q is uniform on $[0, 1]$.

Lemma 28. *If $t \sim F$ and $q \sim G_t$ then q is $U[0, 1]$.*

Definition 18. The *marginal revenue mechanism for agents with monotone ex ante optimal pricings* works as follows.

- (1) Map reported types $\mathbf{t} = (t_1, \dots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \dots, q_n)$ by sampling q_i from the distribution with cumulative distribution function $G_{t_i}(q) = x_i^{\hat{q}}(t_i)$.
- (2) Calculate the marginal revenue of each agent i as $R'_i(q_i)$.
- (3) For each agent i , calculate the maximum quantile \hat{q}_i that she could possess and be in the marginal revenue maximizing feasible set (breaking ties consistently).
- (4) For each agent i , offer the \hat{q}_i ex ante optimal pricing conditioned so that i is served if $q_i \leq \hat{q}_i$ and not served otherwise.

The last step of the marginal revenue mechanism warrants an explanation. In the \hat{q}_i ex ante optimal pricing, the outcome that i would obtain with type t_i may be a partial lottery, i.e., it may probabilistically serve i or not. The probability that i is served is $x_i^{\hat{q}_i}(t_i) = \mathbf{Pr}_{q_i}[q_i \leq \hat{q}_i] = G_{t_i}(\hat{q}_i)$ by our choice of q_i . When we offer agent i the \hat{q}_i ex ante optimal pricing we must draw an outcome from the distribution given by $w_i^{\hat{q}_i}(t_i)$. Some of these outcomes are service outcomes, some of these are non-service outcomes. If $q_i \leq \hat{q}_i$ then we draw an outcome from the distribution $w_i^{\hat{q}_i}(t_i)$ conditioned on service; if $q_i > \hat{q}_i$ then we draw an outcome conditioned on no-service. Notice that, while it may not be feasible to serve all agents who receive non-trivial partial lottery, this method coordinates across the partial lotteries which agents to serve to maintain the right distribution on agent outcomes and ensure feasibility.

Proposition 4. *The marginal revenue mechanism for agents with monotone step mechanisms is feasible and dominant strategy incentive compatible.*

Proof. Feasibility follows as the set of agents that select service outcomes is exactly the marginal revenue maximizing set subject to feasibility. To verify the dominant strategy incentive compatibility consider any agent i 's perspective. The parameter \hat{q}_i is a (randomized) function

only of the other agents' reports; the agent's outcome is determined by the \hat{q}_i ex ante optimal pricing which is incentive compatible for any \hat{q}_i . \square

Theorem 13. *The marginal revenue mechanism for agents with monotone ex ante optimal pricings implements marginal revenue maximization (Definition 14).*

Proof. From each agent i 's perspective, the other agents' quantiles are distributed independently and uniformly on $[0, 1]$ (Lemma 28). Therefore, this agent faces a distribution over ex ante optimal pricings that is identical to the distribution of "critical quantiles" in the maximization of marginal revenue, i.e., with density $\frac{d}{d\hat{q}}G_i^{MR}(\hat{q})$. \square

4.3.2. General ex ante optimal pricings

For general agents for whom the ex ante optimal pricings do not satisfy the monotonicity condition (Definition 17), we give in Section B.1 an simple procedure to implement the marginal revenue mechanism (recall Definition 14). This mechanism is given by Definition 39 in Section B.1. The key to the proof of Theorem 14 is a variation of the technique of vector majorization (Hardy et al., 1929).

Theorem 14. *For service constrained environments, there is a simple Bayesian incentive compatible implementation of the marginal revenue mechanism.*

4.3.3. Example: single-dimensional agents with public budgets

In this section we exhibit a class of single-agent problems with non-linear utilities that has monotone ex ante optimal pricings (Definition 17). Consider an agent with a single-dimensional value for receiving a good but has a public budget that limits the payment she could make. Her utility is her value for receiving the good minus her payment as long as her payment is at most

her budget. We show that under standard conditions on the agent's valuation distribution, this single agent problem has monotone ex ante optimal pricings.

The following proposition is a consequence of techniques developed by Laffont and Robert (1996) and Pai and Vohra (2008); for completeness we provide a proof in Section B.2 whose steps largely resemble the ones in these two references.

Proposition 5. *For regular distribution F with non-decreasing density, budget B , and $\hat{q} \leq 1 - F(B)$, the \hat{q} ex ante optimal pricing offers a single take-it-or-leave-it lottery for price B that serves with probability π , where π is the solution to the equation $\hat{q} = \pi[1 - F(B/\pi)]$. This lottery is bought by the agent when her value is at least B/π which happens with probability $1 - F(B/\pi)$.*

For $\hat{q} > 1 - F(B)$, it is easy to see that the budget does not bind and the \hat{q} ex ante optimal pricing is the same as when there is no budget.

Notice that the allocation rule of the mechanism satisfying Proposition 5 is a function that steps from 0 to π at value B/π . The required payment B can be viewed as “the area above the allocation curve” which is given by a rectangle with width B/π and height π . If π increases, B/π decreases and more types are served and with a higher probability; thus, the ex ante probability of service is increased. Analogously, if we increase the ex ante probability of service, we enlarge the set of types served and their probability of service. We conclude with the following consequence.

Theorem 15. *An agent with value drawn from a regular distribution with non-increasing density has monotone ex ante optimal pricings.*

Proof. The only case not argued by the text above is when $\hat{q} \geq 1 - F(B)$. In this case, the budget is not binding and the ex ante optimal pricing posts price p that satisfies $\hat{q} = 1 - F(p)$ and serves agents willing to pay this price with probability one. The ex ante optimal pricings are monotone over these quantiles as well. \square

Example 9 (Implementation with public budgets). The following procedure implements the marginal revenue mechanism in a single item auction for bidders each with a publicly known budget B and value drawn uniformly from $[0, 1]$. The auction is easier to describe separately for the two cases when $B < 1/2$ and $B \geq 1/2$.

For $B < 1/2$, if no bidder bids above B , the item is not sold; if only one bidder bids above B , she wins the item and makes a payment of B ; if at least two bidders bid above B , the winner of the item will be decided among these bidders by a random procedure described shortly. The winner always makes a payment of B .

For $B \geq 1/2$, if no bidder bids above B , a second price auction is run with a reserve price of $1/2$; if one bidder bids above B , she wins the item and makes the same payment as in a second price auction with reserve $1/2$; if at least two bidders bid above B , one of them is decided to be the winner by a random procedure, but all bidders that bid above a randomly chosen threshold also makes a payment of B .

Now we describe the random procedure used to determine the winner in both cases. Note again that only bidders who bid at least B will enter this procedure. Each such bidder i draws a random number r_i uniformly from $[0, 1]$, and her quantile q_i will be $\max\{r_i, B/v_i\} - B$. Whichever bidder i^* having the smallest quantile is declared the winner. The threshold above which other bidders make the payment is $B/(B + q_{i^*})$.

Section B.2 gives the derivation showing that this is the instantiation of Definition 18.

Note the role played by the random mapping in this example. When multiple bidders bid above B , the highest bidder is not guaranteed to win the item. Her higher value helps her obtain a lower quantile by posing a smaller B/v_i , but with positive probability she may lose to a lower bidder.

4.4. Approximation

In previous sections, we have shown that for any service constrained environment the marginal revenue mechanism can be implemented. In Section 4.2 we have also shown that for revenue linear agents, it obtains the optimal revenue. In this section, we show that, quite generally, the optimal marginal revenue is a good approximation to the optimal revenue.

We will give two approaches for approximation bounds. The first kind of bound is based on the single-agent problem, i.e., the distribution and type space of each agent: if for all allocation constraints \hat{x} , the marginal revenue $\text{MR}[\hat{x}]$ is a good approximation to the optimal revenue $\text{Rev}[\hat{x}]$, then the marginal revenue mechanism is a good approximation to the optimal mechanism. The second approach will derive approximation bounds from the feasibility constraint. With no feasibility constraint, marginal revenue maximization is optimal; for matroid environments, it remains a $1 - 1/e$ approximation; and for general downward-closed environments with n quasi-linear-preference agents, it gives an $O(\log n)$ approximation.

Of course, if we are in an environment where agent-based arguments imply an α approximation and feasibility-based arguments imply a β approximation, the marginal revenue mechanism is in fact a $\min(\alpha, \beta)$ approximation. For revenue linear agents, $\alpha = 1$ (and the optimal marginal revenue gives the optimal revenue); the approximation smoothly degrades in α as the environment becomes less revenue linear until it reaches the approximation bound β given by the feasibility constraint.

4.4.1. Agent-based Approximation

If, for all allocation constraints, the marginal revenue is close to the optimal revenue, then marginal revenue maximization is approximately optimal. One approach to deriving such a bound is to give a linear upper bound on the optimal revenue and a lower bound through a class of, what we refer to as, ex ante pseudo pricings. An ex ante pseudo pricing respects an ex

ante constraint but may not be optimal. If for every ex ante service probability \hat{q} the \hat{q} ex ante pseudo pricing approximates the linear upper bound, then for all allocation constraints \hat{x} , the marginal revenue $\text{MR}[\hat{x}]$ approximates the optimal revenue $\text{Rev}[\hat{x}]$. Furthermore, these ex ante pseudo pricings can be directly optimized over and the same approximation factor is obtained. Such an approach might be desirable if the ex ante pseudo pricings are better behaved than the (optimal) ex ante optimal pricings, e.g., if they are easy to compute, respect an ordering on types (à la Definition 15), or are monotone (à la Definition 17). This approach is formalized by the following sequence of definitions and propositions.

Proposition 6. *If for any agent i and allocation constraint \hat{x}_i , the marginal revenue $\text{MR}[\hat{x}_i]$ is at least an α approximation to the optimal revenue $\text{Rev}[\hat{x}_i]$, then the marginal revenue mechanism in the multi-agent setting is an α approximation to the optimal mechanism.*

Definition 19. An *linear revenue bound*, UB , is a function mapping an allocation constraint to a revenue, which is

- (1) linear in the allocation constraint, i.e., for all allocation constraints $\hat{x} = \hat{x}^A + \hat{x}^B$, $\text{UB}(\hat{x}) = \text{UB}(\hat{x}^A) + \text{UB}(\hat{x}^B)$; and
- (2) an upper bound on revenue for all allocation constraints, i.e., $\forall \hat{x}$, $\text{UB}(\hat{x}) \geq \text{Rev}[\hat{x}]$, and

Definition 20. A *ex ante pseudo pricing* is one that respects an ex ante service probability constraint but is not necessarily revenue optimal for such a constraint. The revenue of a \hat{q} ex ante pseudo pricing is denoted $\tilde{R}(\hat{q})$; and the *pseudo marginal revenue* for allocation constraint \hat{x} is $\text{PMR}[\hat{x}] = \mathbf{E}[\tilde{R}'(\hat{q})\hat{x}(\hat{q})]$.

We can assume without loss of generality that the pseudo marginal revenue \tilde{R} is concave. If it is not we could always redefine the class by taking its closure with respect to convex combination and letting the \hat{q} ex ante pseudo pricing be the revenue-optimal lottery pricing in the closure

that serves with ex ante probability \hat{q} . This construction is analogous to the ironing method of Myerson (1981).

Proposition 7. *For a given linear revenue bound UB , if for all $\hat{q} \in [0, 1]$ the \hat{q} ex ante pseudo pricing α approximates the bound on the \hat{q} ex ante constrained revenue $UB(\hat{x}^{\hat{q}})$, then the pseudo marginal revenue α approximates the optimal revenue for all allocation constraints.*

Proof. This proposition follows from linearity of both the revenue bound and pseudo marginal revenue. □

Definition 21. The *pseudo marginal revenue mechanism* is the one that maximizes pseudo marginal revenue via any of the approaches of Definition 16, Definition 18, or Definition 39 that applies.

Pseudo ex ante optimal pricing for downward-closed unit-demand agents. We illustrate the methodology proposed above for the example of downward-closed service-constrained environments and unit-demand agents. Recall for unit-demand agents a service outcome is one of m alternatives. An agent's type is described by the vector (v^1, \dots, v^m) , her valuations for each of the m alternatives, and her utility for obtaining alternative j with payment p is simply $v^j - p$. The agent's type is drawn from a product distribution over the distinct alternatives.

Chawla et al. (2010b,a) show, for a single-unit demand agent with values for distinct alternatives from a product distribution and no feasibility constraint, that individually pricing alternatives is a four approximation to optimal lottery pricing. Our approach in this section will be to extend this result to settings with ex ante and interim allocation constraints. Our generalization preserves the approximation bound of four and exposes the approximate linearity condition required by Proposition 6.

Consider the syntactically-related problem of selling a single item to one of m single-dimensional agents with values drawn from a product distribution, i.e., the value v_i of agent

i is drawn independently from F_i . As described earlier (Section 2.2), the optimal auction for this single-dimensional problem is well understood. Agent values are mapped to virtual values (equivalent to each agent’s marginal revenue), and the agent with the highest positive virtual value is selected as the winner of the auction. We refer to this auction environment as the single-dimensional *representative environment*, the revenue obtained by the optimal auction as the *optimal representative revenue*, and the agents participating in the auction as *representatives*.

Notice that if these representatives were all colluding together the problem would be identical to our original single-agent unit-demand problem where the alternatives correspond to the identity of the winning representative. We refer to this original environment as the *unit-demand environment* and the revenue of the optimal lottery pricing as the *optimal unit-demand revenue*. Chawla et al. (2010b,a) considered quantifying the performance of optimal unit-demand lottery pricings relative to the optimal representative revenue. The approach of Chawla et al. (2010b) is to set individual prices for each alternative in the unit-demand environment so as to mimic the outcome of the optimal auction for the representative environment. As the optimal auction in the representative environment orders representatives by virtual values, a natural approach to pricing the alternatives in the unit-demand environment is to set a uniform virtual price, i.e., the price for each alternative has the same virtual value (with respect to the distribution from which the agent’s value for that alternative is drawn).¹⁰ The prices in value space are generally distinct when the agent’s value distributions for the alternatives are non-identical. Chawla et al. (2010b) show that the unit-demand revenue of such a pricing is a 2-approximation to the optimal representative revenue; Chawla et al. (2010a) show that the optimal unit-demand revenue (e.g., from lottery pricings) is at most twice the optimal representative revenue. Combining these two results, *uniform virtual pricing* is a 4-approximation to the optimal unit-demand revenue.

¹⁰As mentioned above, a representative’s virtual value is equal to their marginal revenue. For clarity of discussion and to disambiguate the marginal revenue of the unit demand agent versus that of his representatives we will refer to a representative’s marginal revenue as his virtual value.

We generalize the approach above to the single-agent problem of serving an agent with independent values for m alternatives subject to an allocation constraint \hat{x} . In particular, twice the optimal representative revenue is a linear revenue bound (Definition 19), and for any allocation constraint it upper bounds the optimal (unit-demand) revenue. We define a class of ex ante pseudo pricings where the \hat{q} ex ante pseudo pricing is given by a uniform virtual pricing that sells with probability \hat{q} . Since the virtual values are weakly increasing in the representative agents' values, the sets of types served by these ex ante pseudo pricings respect an ordering on types (Definition 15). Therefore, the pseudo marginal revenue mechanism can be implemented via the marginal revenue mechanism for orderable agents (Definition 16). Finally, we show that for all \hat{q} the \hat{q} ex ante pseudo pricing is a four approximation to the linear upper bound given by twice the optimal representative revenue. This result, with Proposition 7, implies that the pseudo marginal revenue mechanism is a four approximation to the optimal revenue for any allocation constraint. The proof of Theorem 16, below, is a non-trivial but straightforward extension of Chawla et al. (2010b,a) and we include it in Section B.3.

Definition 22. The \hat{q} ex ante pseudo pricing for a unit-demand agent with values for alternatives drawn independently from F^1, \dots, F^m is given by the pricing that sets a uniform virtual price for the alternatives such that the probability that the agent buys any alternative is equal to \hat{q} . (If this class does not have a concave pseudo revenue curve we take its closure with respect to convex combination to make it concave; if this class does not have a monotone non-decreasing pseudo revenue curve $\tilde{R}(\cdot)$ we invoke downward closure to make it monotone.)¹¹

¹¹For showing approximate linearity for downward-closed environments it is expedient to incorporate the downward closure into the outcome space by duplicating each non-service outcome and relabeling the duplicate outcome as a service outcome. This transformation is allowed for downward closed environments because we are always allowed to withhold service to an agent who would otherwise be served and this withholding will not violate the feasibility constraint. Of course, with such a transformation the revenue curves are non-decreasing.

Theorem 16. *In downward-closed (service constrained) environments with unit-demand agents, both the pseudo marginal revenue mechanism and the marginal revenue mechanism give 4 -approximations to the optimal revenue.*

4.4.2. Feasibility-based Approximation

We now show that feasibility constraints imply approximation bounds. As a first trivial observation, if there is no feasibility constraint (e.g., for digital good environments) then marginal revenue maximization is optimal. With no feasibility constraint, each agent can be considered separately. For any agent i , suppose the revenue optimal mechanism serves with probability \hat{q}_i , by definition the revenue it obtains is equal to that of the \hat{q}_i ex ante optimal pricing. The optimal revenue $\sum_i \text{Rev}[\hat{x}_i^{\hat{q}}]$ is equal to the marginal revenue $\sum_i \text{MR}[\hat{x}_i^{\hat{q}}] = \sum_i R_i(\hat{q}_i)$. This observation approximately generalizes as follows. The marginal revenue mechanism is an $e/(e-1)$ approximation in service-constrained matroid environments and an $O(\log n)$ bound for downward-closed environments on n quasi-linear-utility agents.

Matroid environments, by single-dimensional-agent reduction. Marginal revenue maximization is an $e/(e-1)$ approximation for service-constrained matroid environments, i.e., when the feasibility constraint is induced by independent sets of a matroid set system. Multi-unit environments, where at most a fixed number k of the agents can be simultaneously served, are a special case of matroid environments (corresponding to the k -uniform matroid). For the $k=1$ unit environment, which corresponds to a single-item auction, the bound remains $e/(e-1)$; for general k the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k}-1)$, as simplified by Stirling's approximation. These results follow by reduction to the correlation gap theorem of Yan (2011).

Our approach is to reduce the question of approximation of the optimal mechanism by the marginal revenue mechanism to a question of approximation in the single-dimensional analog environment (recall Definition 13). In particular, we consider relaxing the feasibility constraint

to hold *ex ante* instead of *ex post*. Such a relaxation potentially enables a higher revenue to be obtained. The single-dimensional-agent approximation question is to quantify the extent to which the optimal mechanism for the *ex post* feasibility constraint approximates the optimal mechanism for the *ex ante* feasibility constraint.

Definition 23. A profile of *ex ante* service probabilities $\hat{\mathbf{q}} = (\hat{q}_1, \dots, \hat{q}_n)$ is *ex ante feasible* if there is a distribution over feasible subsets of agents such that for each i , \hat{q}_i is the (*ex ante*) probability agent i is in the subset. The *ex ante optimal mechanism* is the one that maximizes $\sum_i R_i(\hat{q}_i)$ subject to *ex ante* feasibility of $\hat{\mathbf{q}}$.

Proposition 8. *The ex ante optimal revenues for a general service constrained environment and its single-dimensional analog are equal and an upper bound on the (ex post feasible) optimal revenues (which may not be equal). If the optimal mechanism is a β -approximation to the ex ante optimal revenue in the single-dimensional analog environment, then the marginal revenue mechanism is a β -approximation to the optimal revenue in the original environment.*

Proof. The *ex ante* optimal revenue is defined only in terms of revenue curves and feasibility for the service constrained environment; therefore, a general environment and its single-dimensional analog have the same *ex ante* optimal revenue. By Definition 13 the (*ex post* feasible) optimal revenue in the single-dimensional analog is equal to the (*ex post* feasible) optimal marginal revenue of the original environment. To show the reduction, then, it suffices to observe that the *ex ante* optimal revenue is an upper bound on the optimal revenue in the original environment. As every *ex post* feasible mechanism is *ex ante* feasible (i.e., the latter is a relaxation of the former), the observation holds. \square

The following single-dimensional agent theorem is an immediate consequence of results of Yan (2011); his results, in fact, gave a specific (*ex post* feasible but non-optimal) mechanism

that satisfies the claimed bound. Of course, then, the (ex post feasible) optimal mechanism satisfies the bound too. We obtain our desired result for general agents as a corollary of this theorem and Proposition 8.

Theorem 17. *For matroid environments with single-dimensional linear agents, the (ex post feasible) optimal mechanism is an $e/(e - 1)$ approximation to the ex ante optimal mechanism; in any k -unit environment the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k} - 1)$.*

Corollary 3. *In any service constrained matroid environment, the marginal revenue mechanism is an $e/(e - 1)$ approximation to the optimal mechanism; in any service constrained k -unit environment the bound improves to $\sqrt{2\pi k}/(\sqrt{2\pi k} - 1)$.*

Downward-closed environments. In this section we show that in downward-closed environments and for a large class of agent preferences, the optimal marginal revenue is a logarithmic approximation, in the number of agents, to the optimal revenue. For example, this class includes quasi-linear preferences. In contrast to Subsection 4.4.1 where we gave a four approximation for unit-demand preferences with a product distribution (over alternatives), the results here apply, for example, to agents with correlated value distributions over alternatives and to quasi-linear preferences beyond unit demand.

To show this result we will incorporate the downward closure of the environment in the single-agent lottery pricing problems. Specifically, it is without loss of generality for downward-closed environments to duplicate every non-service outcome and label the duplicate a service outcome. This transformation implies that revenue is monotone in the allocation constraint, i.e., weaker constraints give no lower revenue.

A summary of the construction in the proof is as follows. If we consider allocation constraints with a minimum probability of 2^{-K} for allocating to any type, then the allocation constraint can be partitioned into K pieces such that the highest and lowest probabilities of allocation in

each piece are within a factor of two of each other. If the single-agent lottery pricing problems satisfy a natural scalability property then the revenue of each piece can be approximated by a \hat{q} ex ante optimal pricing scaled appropriately so that it is dominated by the original allocation constraint. The optimal revenue, then, is at most an $O(K)$ multiple of the revenue of the best such scaled ex ante optimal pricing. By downward closure, the optimal marginal revenue exceeds this revenue and is thus an $O(K)$ approximation. We obtain a logarithmic approximation by observing that attention can be restricted to allocation constraints for which $K \approx \log n$.

Recall that the revenue operator $\text{Rev}[\cdot]$ is concave in its argument and therefore, for any $\gamma \in [0, 1]$, $\text{Rev}[\gamma\hat{x}] \geq \gamma \text{Rev}[\hat{x}] + (1 - \gamma) \text{Rev}[0]$, where $\text{Rev}[0] = R(0)$ is the optimal revenue when the agent is never served. We assume for simplicity of exposition that $\text{Rev}[0] = R(0) = 0$, i.e., that an agent who is not served generates no revenue. The revenue scalability property we need is the opposite of this inequality, which, if the property holds, must therefore be an equality. In fact, revenue scalability can be viewed as a very permissive relaxation of revenue linearity.

Definition 24. An agent is *revenue scalable* if for any $\gamma \in [0, 1]$ and any allocation constraint \hat{x} , the optimal revenue for $\gamma\hat{x}$ is equal to γ times the optimal revenue for \hat{x} . I.e.,

$$\text{Rev}[\gamma\hat{x}] = \gamma \text{Rev}[\hat{x}].$$

For example, as we will show, quasi-linear agents satisfy revenue scalability, but are not generally revenue linear. Moreover, if individual rationality is assumed, which usually implies that the utility and payment of an agent for non-service outcomes is zero, then even non-quasi-linear agents are revenue scalable. These observations are formalized in the following lemma.

Lemma 29. *Both (a) quasi-linear agents with no value for non-service outcomes and (b) agents with no utility and payment for any non-service outcome are revenue scalable.*

Proof. A key property of agents that are quasi-linear or have no utility and payment for non-service outcomes is that their utility and payment for any non-service outcome can be arbitrarily scaled upward. If an agent's utility and payment for a non-service outcome is zero, then scaling it upwards is trivial; if an agent is quasi-linear then his value for a non-service outcome is (minus) his payment and quasi-linearity requires that scaled payments translate to scaled utility. Thus, it suffices to show that agents with scalable utility and payment for non-service outcomes are revenue scalable.

Consider any allocation constraint \hat{x} and the optimal lottery pricing for the scaled constraint $\gamma\hat{x}$. Denote by L the set of priced lotteries. As $\gamma\hat{x}(\hat{q}) \leq \gamma$ for all \hat{q} , the probability of a service outcome in any of the lotteries of L is at most γ . The theorem holds if we can define an alternative set of priced lotteries L' that meets the constraint \hat{x} where the utility and payment of any type for any lottery is scaled upward a $1/\gamma \geq 1$ multiple. This is achieved by scaling the probability of any service outcome in any lottery upwards by a $1/\gamma$ multiple (without changing its payment), scaling the remaining probability of non-service outcomes down (so that the total probability is one), and scaling the utility and payment for non-service outcomes so that it is $1/\gamma$ multiple of that for the original lottery (which is possible by the assumption on scalability for non-service outcomes). Let $\gamma x(\cdot)$ denote the optimal allocation rule for constraint $\gamma\hat{x}$; the allocation rule from this construction is x and it is feasible for constraint \hat{x} . \square

We now show that the marginal revenue approximates the optimal revenue for revenue-scalable agents in downward-closed service-constrained environments.

Lemma 30. *For a revenue-scalable agent, any allocation constraint with minimum allocation probability $\hat{x}(1) \geq 2^{-K}$ has revenue $\text{Rev}[\hat{x}]$ at most $2K \text{MR}[\hat{x}]$.*

Proof. Let $R^* = \text{Rev} \hat{x}$ be the optimal revenue under allocation constraint \hat{x} . Let $x \preceq \hat{x}$ (where notation $x \preceq \hat{x}$ denotes allocation rules whose cumulative allocation rules satisfy $X(\hat{q}) \leq$

$\hat{X}(\hat{q})$ for $\hat{q} \in [0, 1]$; importantly, $X(1) = \hat{X}(1)$ is not required) be the allocation of optimal mechanism subject to \hat{x} . Therefore, $\text{Rev } x = \text{Rev } \hat{x}$. If we prove the claim for x , the proof for \hat{x} follows because $\text{Rev } \hat{x} = \text{Rev } x \leq 2K \text{MR}[x] \leq 2K \text{MR}[\hat{x}]$, where the last inequality follows by definition of dominance and concavity of the revenue function. Therefore, in the rest of the proof we can assume without loss of generality that the optimal allocation subject to \hat{x} is \hat{x} itself.

Define a sequence of quantiles $0 = q_0 \leq q_1 \leq \dots \leq q_K = 1$ such that $\hat{x}(q_{j-1}) \leq 2\hat{x}(q_j)$, for $j \in \{1, \dots, K\}$. Define R_j^* to be the expected revenue from types that are mapped to a quantile in $[q_{j-1}, q_j]$, where the quantile of a type t is the probability that a type drawn at random has a higher probability of service than that of t (as per Definition 12 in Section 4.2). Therefore, the revenue of the mechanism is $R^* = \sum_{j=1}^K R_j^*$. Then there must exist j^* such that $R^* \leq KR_{j^*}^*$. In what follows, we define allocation rules $z_j(\cdot)$ for all j , such that $z_j \preceq \hat{x}$ (where notation $x \preceq \hat{x}$ denotes allocation rules whose cumulative allocation rules satisfy $X(\hat{q}) \leq \hat{X}(\hat{q})$ for $\hat{q} \in [0, 1]$; importantly, $X(1) = \hat{X}(1)$ is not required), and also $R_j^* \leq 2 \text{MR}[z_j]$. In particular, for j^* we will have $z_{j^*} \preceq \hat{x}$, and $2 \text{MR}[z_{j^*}] \geq R_{j^*}^* \geq R^*/K$, which will imply that

$$2 \max_{z \preceq \hat{x}} \text{MR}[z] \geq 2 \text{MR}[z_{j^*}] \geq R^*/K.$$

Define function $z_j(\cdot)$ to be $z_j(q) = \hat{x}(q_{j+1})$ if $q \leq q_{j+1} - q_j$, and 0 otherwise. Notice that for any q , we have $z_j(q) \leq \hat{x}(q)$, and therefore $z_j \preceq \hat{x}$, by the definition of dominance in downward-closed environments.

The main technical component of the proof is to show that, for z_j defined above, $R_j^* \leq 2 \text{MR}[z_j]$. By construction of z_j , and recalling that $\hat{x}(q_j) \leq 2\hat{x}(q_{j+1})$,

$$\begin{aligned} 2 \text{MR}[z_j] &= 2 \int_0^1 z_j(q) R'(q) \, dq \\ &= 2\hat{x}(q_{j+1})R(q_{j+1} - q_j) \\ &\geq \hat{x}(q_j)R(q_{j+1} - q_j) \end{aligned}$$

It is therefore sufficient to show that $\hat{x}(q_j)R(q_{j+1} - q_j) \geq R_j^*$. Recall that R_j^* is the revenue from types that are mapped to quantiles in $[q_j, q_{j+1}]$. Any type in $[q_j, q_{j+1}]$ is allocated in \hat{x} with probability at most $\hat{x}(q_j)$. Now define L to be the set of lotteries chosen by types in $[q_j, q_{j+1}]$, and offer only these lotteries to the agent.¹² Notice that types with quantiles in $[q_j, q_{j+1}]$ choose the same lottery in L as they did in \hat{x} (whereas other types that used to choose a lottery either switch to some lottery in L or no longer choose one if none in L give them non-negative utility). As a result, the measure of the types that choose some lottery in L is at least $q_{j+1} - q_j$. Now remove lotteries from L , from the one with lowest price, until the measure of types that choose some lottery is exactly $q_{j+1} - q_j$.¹³ Call this new set of lotteries L' . Notice that the revenue from L' is at least R_j^* . Now recall that all the lotteries in L , and therefore L' , allocate with probability at most $\hat{x}(q_j)$. Revenue scalability implies that the revenue of L' is at most $\hat{x}(q_j)R(q_{j+1} - q_j)$.

To complete the proof, recall that for downward-closed environments revenue curves are monotone non-decreasing and so marginal revenues are non-negative. Therefore, by the definition of marginal revenue and dominance, $\text{MR}[\hat{x}] \geq \text{MR}[z_j]$ for all j . \square

¹²Recall that, by the taxation principle, any incentive compatible mechanism consists of a set of lotteries, from which the agent chooses the one maximizing her utility.

¹³This requires continuity of the type space. We assume continuity for simplicity, but the proof can be easily generalized to handle discrete types.

Theorem 18. *In downward-closed revenue-scalable environments with n agents, the optimal marginal revenue is a $4 \log n$ approximation to the optimal revenue.*

Proof. Consider an alternative mechanism that runs the optimal mechanism with probability $1/2$, and otherwise picks an agent at random and outputs an arbitrary outcome that serves that agent, regardless of his type and without charging him. This alternative mechanism is obviously incentive compatible, and its revenue is half of the optimal. Let x_1, \dots, x_n be the allocation rules for the alternative mechanism. Notice also that by construction, for each i and any $q \in [0, 1]$ we have $x_i(q) \geq 1/2n$. Therefore we can invoke Lemma 30 with $K = \log 2n$ to conclude that the revenue of the alternative mechanism is at most

$$2 \log n \sum_i \text{MR}_i[x_i]. \quad \square$$

4.5. Single Dimensional Extension Theorems

The marginal revenue approach allows natural generalizations of techniques developed for single-dimensional linear agent environments. We will focus here on results for the approximation of optimal mechanisms by simple mechanisms. In such a study we are not free to arbitrarily design the simple mechanism. Instead, we show that performance guarantees for simple mechanisms are often obtainable by relating them to marginal revenue mechanisms.

Consider a variant of the red-or-blue car example from the introduction. There are n agents, k cars, and m possible colors. The social surplus maximizing mechanism (a.k.a. VCG; see Vickrey, 1961; Clarke, 1971; and Groves, 1973) selects the k agents whose values for their favorite color are the highest, serves these agents, and paints each car as the agent prefers. We consider this mechanism simple and practical, and we compare its revenue against the optimal revenue (cf. Hartline and Roughgarden, 2009). Shortly we will give conditions under which this mechanism is approximately optimal.

One feature of the VCG mechanism in service constrained environments is that, ex post, i.e., after agents make reports to the mechanism, each agent faces a *uniform* price over the alternatives. This price is equal to the favorite-color value among the other agents. Thus, in the interim stage each agent faces a distribution over uniform prices. We show that the VCG mechanism has near-optimal revenue in two steps. First, we show that the VCG revenue is close to the revenue of the optimal mechanism that only offers agents uniform prices. Second, we show that the latter revenue is close to the optimal revenue by any mechanism. An important observation is that the intermediate revenue in between these two steps is the optimal pseudo marginal revenue with uniform ex ante pseudo pricings (cf. Definition 20).

For the first step of the argument, the gap between the VCG revenue and the optimal pseudo marginal revenue is governed by the single-dimensional theory. Both mechanisms operate on the type space given by projection of each unit-demand agent's multi-dimensional type into the single-dimensional space given by his value for his favorite alternative. In particular, the VCG revenue for the unit-demand agents is equal to its revenue under this single-dimensional projection, and the optimal pseudo marginal revenue for the unit-demand agents is equal to the optimal revenue for the projection. For the second step in our argument, by the theory of agent-based approximation we developed in Section 4.4 (e.g., Proposition 7), we need only analyze how good the uniform pricings are, as ex ante pseudo pricings.

We consider a concrete simple case before developing the general theory. In the above car-selling example, let k be 1, and each agent's value for each color be i.i.d. (i.i.d. among agents and i.i.d. across the alternatives). Since each agent's values for the alternatives are i.i.d., a uniform price for an agent is also a uniform virtual price (see Definition 22). Thus, Theorem 16 implies that the optimal pseudo marginal revenue (with uniform ex ante pseudop pricing) is a four approximation to the optimal revenue. This constitutes the second step of our planned argument. For the first step, the standard single-dimensional theory. If the distribution of the

i.i.d. unit-demand agent's favorite-alternative value satisfies the regularity condition of Myerson (1981), then the theorem of Bulow and Klemperer (1996) implies that, for its single-dimensional analog, the second-price auction is an $\frac{n}{n-1}$ approximation to the single-dimensional optimal revenue, which is in turn equal to the optimal pseudo marginal revenue for the unit-demand agents. Combining the two steps, we have shown that the VCG mechanism for unit-demand agents is a $\frac{4n}{n-1}$ approximation. This discussion is formalized and generalized below.

Definition 25. A unit-demand agent is β *uniformly priceable* if, for any allocation constraint \hat{x} , a distribution over uniform pricings gives a β approximation to the optimal lottery pricing. The *uniform ex ante pseudo pricings* are the optimal of these pricings for ex ante constraints.

Definition 26. The *favorite-alternative single-dimensional analog* of a unit-demand service constrained environment is given by projecting the values of each unit-demand agent to the value of his favorite alternative. The *favorite-alternative extension* of a single-dimensional mechanism is a mechanism for unit-demand agents that simulates the given single-dimensional mechanism on reported values for favorite alternatives (ignoring the other values). It serves the winners of the simulation their favorite alternatives at the prices of the simulation.

Proposition 9. *For any service constrained environment, unit-demand β -uniformly-priceable agents, and α -approximation mechanism \mathcal{M} for the favorite-alternative single-dimensional analog environment; the favorite-alternative extension of \mathcal{M} is an $\alpha\beta$ approximation for the original environment.*

Proof. By construction, the revenue of the favorite-alternative extension of \mathcal{M} in the original environment is equal to the revenue of \mathcal{M} for the favorite-alternative single-dimensional analog environment. By assumption of the proposition, this revenue is an α -approximation to the optimal revenue for the favorite-alternative single-dimensional analog. This single-dimensional

optimal revenue is equal to the optimal pseudo marginal revenue (with uniform ex ante pseudo pricings) in the unit-demand environment. By the assumption that the unit-demand agents are β uniformly priceable, Proposition 7 implies that this optimal pseudo marginal revenue is a β approximation to the unit-demand optimal revenue. These bounds combine to imply that the revenue of the favorite-alternative extension is an $\alpha\beta$ approximation to the optimal revenue. \square

To draw single-dimensional extension theorems as consequences to Proposition 9, we first claim that unit-demand agents with values for each alternative independently drawn from (not necessarily identical) regular distributions are eight uniformly priceable. After this, we apply tools from the single-dimensional theory to provide approximation mechanisms for the favorite-alternative single-dimensional analog, and draw immediate corollaries.

Uniform Priceability. As described above, the i.i.d. special case of Theorem 16 implies that any unit-demand agent with i.i.d. values for distinct alternatives is four uniformly priceable. This result approximately generalizes to non-i.i.d. distributions that satisfy the regularity condition of Myerson (1981) as follows.

Definition 27. A distribution specified by distribution function F and density function f is *regular* if $v - \frac{1-F(v)}{f(v)}$ is monotone non-decreasing in v . A single-dimensional linear agent is regular if his value is drawn from a regular distribution.¹⁴

Lemma 31. *A unit-demand agent with values for alternatives drawn independently from (not necessarily identical) regular distributions is eight uniformly priceable.*

PROOF SKETCH. The proof will follow the template given by Proposition 7 with the following main ingredients.

¹⁴Single-dimensional regularity is equivalent (a) to $P(\hat{q}) = R(\hat{q})$ for all \hat{q} (see the proof of Theorem 2), and (b) to the revenue-optimal allocation rule for \hat{x} being \hat{x} itself (see Lemma 26). These properties of regular distributions enable approximation of optimal mechanisms by simple mechanisms in single-dimensional environments.

- Twice the optimal revenue of the representative environment (where the unit-demand agent is replaced by single-dimensional representatives for each alternative) is a linear upper bound on the optimal revenue for any allocation constraint (by Lemma 38).
- Uniform pricing in the representative environment with regular distributions gives a four approximation to the optimal representative revenue; the argument is as follows. Hartline and Roughgarden (2009) show that the second-price auction with a uniform (a.k.a., anonymous) reserve price is a four approximation to the optimal revenue. In fact, this result can be strengthened (a) using a prophet-inequality-like proof to give the same bound for uniform pricing and (b) to allow an ex ante constraint on the probability that any representative is served. These extensions follow from a relatively straightforward modification of Lemma 39 which we omit.
- Uniform pricing in the original environment has the same revenue as uniform pricing in the representative environment. \square

Below we will make use of the following slight strengthening of the regularity condition of Definition 27.

Definition 28. A unit-demand agent is *favorite-alternative regular* if the distribution of the agent's value for favorite alternative is regular; a unit-demand agent is *individual-alternative regular* if the agent's value for each alternative is regular; a unit-demand agent is *regular* if he is both favorite- and individual-alternative regular.¹⁵

Note that Lemma 31 requires only individual-alternative regularity.

Monopoly and Anonymous Reserve Pricing. For a single-dimensional single-agent problem, the *monopoly price* is the price that optimizes revenue. For single-dimensional, i.i.d.,

¹⁵Neither favorite-alternative nor individual-alternative regularity imply the other.

regular, matroid environments the surplus maximizing mechanism (a.k.a. VCG) with the monopoly reserve price (for the distribution) is revenue optimal. Hartline and Roughgarden (2009) approximately extend this result to non-identical distributions. They show that with regular single-dimensional agents, the revenue of the surplus maximizing mechanism with monopoly reserves is a two approximation to the optimal revenue. The following is a corollary of the above development and their theorems.

Corollary 4. *For independent unit-demand favorite-alternative-regular β -linearly-priceable agents and matroid service-constrained environments, the surplus maximizing mechanism with monopoly reserves (for distributions of favorite alternatives) is a 2β approximation to the optimal revenue. For regular unit-demand agents, $\beta = 8$.*

For single-item environments a similar approximation bound holds for an anonymous reserve price, i.e., one that is the same across the distinct agents. Hartline and Roughgarden (2009) show that with regular single-dimensional linear agents in single-item environments, the revenue of the second-price auction with an appropriate anonymous reserve is a four approximation to the optimal revenue. From this result we obtain the following corollary.

Corollary 5. *For independent unit-demand favorite-alternative-regular β -linearly-priceable agents and single-item service-constrained environments, the surplus maximizing mechanism with a suitably chosen anonymous reserve price is a 4β approximation to the optimal revenue. For regular unit-demand agents, $\beta = 8$.*

Market Expansion. Bulow and Klemperer (1996) show that the revenue of the single-item second-price auction for n i.i.d. regular agents is at least the revenue of the optimal auction for $n-1$ agents. An interpretation of this result is that the revenue loss of running the (surplus-optimal) second-price auction instead of the revenue-optimal auction can be made up by recruiting one

more agent to the auction. This result generalizes to matroid environments, see e.g., Dughmi et al. (2009), where the revenue of the surplus maximizing mechanism is at least the revenue of the optimal auction after removing a *base* of the matroid.¹⁶ The corollary, below, extends the ($k = 1$) multi-unit result described informally in the beginning of this section.

Corollary 6. *For i.i.d. unit-demand favorite-alternative-regular β -linearly-priceable agents and matroid service-constrained environments, the surplus maximizing mechanism is a β approximation to the optimal revenue with any base of the matroid removed. For regular unit-demand agents, $\beta = 8$.*

Prior-Independent Mechanisms. Dhangwatnotai et al. (2010) show that, in regular single-dimensional matroid environments, the surplus maximizing auction where each agent faces a reserve price randomly drawn from his value distribution is a four approximation to the optimal auction. If the agents' values are identically distributed then the approximation factor improves to two. Moreover, as long as there are at least two agents with values drawn from the same distribution, this approximation result can be obtained by a prior-independent mechanism, i.e., one that is not parameterized by the prior-distribution. We summarize the consequences of the i.i.d. result in general service-constrained matroid environments as follows.

Corollary 7. *For i.i.d. unit-demand favorite-alternative-regular β -linearly-priceable agents and matroid service-constrained environments, there is a prior-independent mechanism that is a 2β approximation to the optimal revenue. For regular unit-demand agents (whose values for alternatives are drawn independently from not necessarily identical distributions), $\beta = 8$.*

These results are meant as examples of single-dimensional results with automatic extensions to unit-demand service constrained environments. Many other single-dimensional results also can be extended.

¹⁶A base of a matroid is a feasible set with maximum cardinality.

CHAPTER 5

Single-Agent to Multi-Agent Solutions: Computation

Overview and Organization. This chapter the computational complexity of optimal multi-agent auction design. In particular, the focus is on the computational complexity of composition of single-agent solutions to form a multi-agent solution.

An optimal multi-agent mechanism can be computed by a linear/convex program on interim allocation rules by simultaneously optimizing several single-agent mechanisms subject to joint feasibility of the allocation rules. For single-unit auctions, Border (1991) showed that the space of all jointly feasible interim allocation rules for n agents is a D -dimensional convex polytope which can be specified by 2^D linear constraints, where D is the total number of all agents' types. Consequently, efficiently solving the mechanism design problem requires a separation oracle for the feasibility conditions and also an algorithm for ex-post implementation of the interim allocation rules. We show that the polytope of jointly feasible interim allocation rules is the projection of a higher dimensional polytope which can be specified by only $O(D^2)$ linear constraints. Furthermore, our proof shows that finding a preimage of the interim allocation rules in the higher dimensional polytope immediately gives an ex-post implementation. We generalize Border's result to the case of k -unit and matroid auctions.

A brief outline of the approach was discussed in Section 2.3. In Section 5.2 we describe single- and multi-agent mechanism design problems. In Section 5.3 we give algorithms for solving two kinds of single-agent problems: multi-item unit-demand preferences and private-value private-budget preferences. In Section 5.4, we give a high-level description of the multi- to single-agent reduction which allows for efficiently compute optimal mechanisms for many service constrained

environments. The key step therein, an efficient algorithm that implements any jointly feasible set of interim allocation rules, is presented in Section 5.5. This section is divided into three parts which address single-unit, multi-unit, and matroid feasibility constraints, respectively. Conclusions and extensions are discussed in Section 5.6.

5.1. Introduction

Classical economics and game theory give fundamental characterizations of the structure of competitive behavior. For instance, Nash's (1951) theorem shows that mixed equilibrium gives a complete description of strategic behavior, and the Arrow-Debreu (1954) theorem shows the existence of market clearing prices in multi-party exchanges. In these environments computational complexity has offered further perspective. In particular, mixed equilibrium in general games can be computationally hard to find (Chen and Deng, 2006; Daskalakis et al., 2009), whereas market clearing prices are often easy to find (Devanur et al., 2008; Jain, 2004). In this chapter we investigate an analogous condition for auction theory due to Kim Border (1991), give a computationally constructive generalization that further illuminates the structure of auctions, and thereby show that the theory of optimal auctions is tractable.

Consider an abstract optimal auction problem. A seller faces a set of agents. Each agent desires service and there may be multiple ways to serve each agent (e.g., when renting a car, you can get a GPS or not, you can get various insurance packages, and you will pay a total price). Each agent has preferences over the different possible ways she can be served and we refer to this preference as her type. The seller is restricted by the *feasibility* constraint that at most one agent can be served (e.g., only one car in the rental shop). When the agents' types are drawn independently from a known prior distribution, the seller would like to design an auction to optimize her objective, e.g., revenue, in expectation over this distribution, subject to

feasibility. Importantly, in this abstract problem we have not made any of the following standard assumptions on the agents' preferences: quasi-linearity, risk-neutrality, or single-dimensionality.

We assume that agents behave strategically and we will analyze an auction's performance in *Bayes-Nash equilibrium*, i.e., where each agent's strategy is a best response to the other agents' strategies and the distribution over their preferences. Without loss of generality the revelation principle (Myerson, 1981) allows for the restriction of attention to *Bayesian incentive compatible* (BIC) mechanisms, i.e., ones where the truth-telling strategy is a Bayes-Nash equilibrium.

Any auction the seller proposes can be decomposed across the agents as follows. From an agent's perspective, the other agents are random draws from the known distribution. Therefore, the composition of these random draws (as inputs), the bid of the agent, and the mechanism induce an *interim allocation rule* which specifies the probability the agent is served as a function of her bid. BIC implies that the agent is at least as happy to bid her type as any other bid.

Applying the same argument to each agent induces a profile of interim allocation rules. These interim allocation rules are jointly feasible in the sense that there exists an auction that, for the prior distribution, induces them. As an example, suppose an agent's type is high or low with probability $1/2$ each. Consider two interim allocation rules: rule (a) serves the agent with probability one when her type is high and with probability zero otherwise, and rule (b) serves the agent with probability $1/2$ regardless of her type. It is feasible for both agents to have rule (b) or for one agent to have rule (a) and the other to have rule (b); on the other hand, it is infeasible for both agents to have rule (a). This last combination is infeasible because with probability one quarter both agents have high types but we cannot simultaneously serve both of them. An important question in the general theory of auctions is to decide when a profile of interim allocation rules is feasible, and furthermore, when it is feasible, to find an auction that implements it.

A structural characterization of the necessary and sufficient conditions for the aforementioned *interim feasibility* is important for the construction of optimal auctions as it effectively allows the auction problem to be decomposed across agents. If we can optimally serve a single agent for a given interim allocation rule and we can check feasibility of a profile of interim allocation rules, then we can optimize over auctions. Effectively, we can reduce the multi-agent auction problem to a collection of single-agent auction problems.

We now informally describe Border's (1991) characterization of interim feasibility. A profile of interim allocation rules is implementable if for any subspace of the agent types the expected number of items served to agents in this subspace is at most the probability that there is an agent with type in this subspace. Returning to our infeasible example above, the probability that there is an agent with a high type is $3/4$ while the expected number of items served to agents with high types is one; Border's condition is violated.

The straightforward formulation of interim feasibility via Border's characterization has exponentially many constraints. Nonetheless, it can be simplified to a polynomial number of constraints in single-item auctions with symmetric agents (where the agents' type space and distribution are identical). This simplification of the characterization has led to an analytically tractable theory of auctions when agents have budgets (Laffont and Robert, 1996) or are risk averse (Matthews, 1984; Maskin and Riley, 1984).

Results. Our main theorem is to show computationally tractable (i.e., in polynomial time in the total number of agents' types) methods for each of the following problems. First, a given profile of interim allocation rules can be checked for interim feasibility. Second, for any feasible profile of interim allocation rules, an auction (i.e., ex post allocation rule) that induces these interim allocation rules can be constructed. In particular, for problems where the seller can serve at most one agent, we show that the exponentially-faceted polytope specified by the interim feasibility constraints is a projection of a quadratically-faceted polytope in a higher dimension,

and an ex post allocation rule implementing a feasible profile of interim allocation rules is given immediately by the latter's preimage in the higher dimensional polytope. In particular, this implies that optimal interim allocation rules can be computed by solving a quadratically sized linear/convex program. These results combine to give a (computationally tractable) reduction from the multi-agent auction problem to a collection of single-agent problems. Furthermore, our algorithmic procedure characterizes every single service auction as implementable by a simple token passing game.

We also generalize the interim feasibility characterization and use it to design optimal auctions for the cases where the seller faces a k -unit feasibility constraint, i.e., at most k agents can be served, and more generally to matroid feasibility constraints. Our generalization of the feasibility characterization is based on a simpler network-flow-based approach. Although the number of constraints in this characterization is exponential in the sizes of type spaces, we observe that the constraints define a polymatroid, whose vertices correspond to particularly simple auctions which can be implemented by simple determinist allocation rules based on ranking. Algorithms for submodular function minimization give rise to fast separation oracles which, given a set of interim allocation rules, detect a violated feasibility constraint whenever there is one; expressing any point in the polymatroid as a convex combination of the vertices allows us to implement any feasible interim allocation rule as a distribution over the simple auctions. These enable us again to reduce the multi-agent problem to single-agent problems.

Auction theory is very poorly understood outside the standard single-dimensional quasi-linear revenue maximization environment of Myerson (1981). The main consequence of this chapter is that even without analytical understanding, optimal auctions can be computationally solved for in environments that include non-quasi-linear utility (e.g., budgets or risk aversion) and multi-dimensional preferences (assuming that the corresponding single-agent problem can

be solved). Furthermore, unlike most work in auction theory with budgets or risk-aversion, our framework permits the budgets or risk parameters to be private to the agents.

Related Work. Myerson (1981) characterized Bayesian optimal auctions in environments with quasi-linear risk-neutral single-dimensional agent preferences. Bulow and Roberts (1989) reinterpreted Myerson’s approach as reducing the multi-agent auction problem to a related single-agent problem. Our work generalize this reduction-based approach to single-item multi-unit auction problems with general preferences.

An important aspect of our approach is that it can be applied to general multi-dimensional agent preferences. Multi-dimensional preferences can arise as distinct values different configurations of the good or service being auctioned, in specifying a private budget and a private value, or in specifying preferences over risk. We briefly review related work for agent preferences with multiple values, budgets, or risk parameters.

Multi-dimensional valuations are well known to be difficult. For example, Rochet and Chone (1998), showed that, because *bunching*¹ can not be ruled out easily, the optimal auctions for multi-dimensional valuations are dramatically different from those for single dimensional valuations. Because of this, most results are for cases with special structure (e.g., Armstrong, 1996; Wilson, 1994; McAfee and McMillan, 1988) and often, by using such structures, reduce the problems to single-dimensional ones (e.g., Spence, 1980; Roberts, 1979; Mirman and Sibley, 1980). Our framework does not need any such structure.

A number of papers consider optimal auctions for agents with budgets (see, e.g., Pai and Vohra, 2008; Che and Gale, 1995; Maskin, 2000). These papers rely on budgets being public or the agents being symmetric; our technique allows for a non-identical prior distribution and private budgets. Mechanism design with risk averse agents was studied by Maskin and Riley

¹Bunching refers to the situation in which a group of distinct types are treated the same way in by the mechanism.

(1984) and Matthews (1983). Both works assume i.i.d. prior distributions and have additional assumptions on risk attitudes; our reduction does not require these assumptions.

Our work is also related to a line of work on approximating the Bayesian optimal mechanism. These works tend to look for simple mechanisms that give constant (e.g., two) approximations to the optimal mechanism. Chawla et al. (2007), Briest et al. (2010), and Cai and Daskalakis (2011) consider item pricing and lottery pricing for a single agent; the first two give constant approximations the last gives a $(1 + \epsilon)$ -approximation for any ϵ . These problems are related to the single-agent problems we consider. Chawla et al. (2010b) and Bhattacharya et al. (2010) extend these approaches to multi-agent auction problems. The point of view of reduction from multi- to single-agent presented in this chapter bears close relationship to recent work by Alaei (2011) who gives a reduction from multi- to single-agent mechanism design that loses at most a constant factor of the objective. Our reductions, employing entirely different techniques, give rise to optimal mechanisms instead of approximations thereof.

Characterization of interim feasibility plays a vital role in this chapter. For single-item single-unit auctions, necessary and sufficient conditions for interim feasibility were developed through a series of works (Maskin and Riley, 1984; Matthews, 1984; Border, 1991, 2007; Mierendorff, 2011); this characterization has proved useful for deriving properties of mechanisms, Manelli and Vincent (2010) being a recent example. Border (1991) characterized symmetric interim feasible auctions for single-item auctions with identically distributed agent preferences. His characterization is based on the definition of “hierarchical auctions.” He observes that the space of interim feasible mechanisms is given by a polytope, where vertices of this polytope corresponding to hierarchical auctions, and interior points corresponding to convex combinations of vertices. Mierendorff (2011) generalize Border’s approach and characterization to asymmetric single-item auctions. The characterization via hierarchical auctions differs from our characterization via

ordered subset auctions in that hierarchical auctions allow for some types to be relatively unordered with the semantics that these unordered types will be considered in a random order; it is important to allow for this when solving for symmetric auctions. Of course convex combinations over hierarchical auctions and ordered subset auctions provide the same generality. Our work generalizes the characterization from asymmetric single-unit auctions to asymmetric multi-unit and matroid auctions.

Our main result provides computational foundations to the interim feasibility characterizations discussed above. We show that interim feasibility can be checked, that interim feasible allocation rules can be optimized over, and that corresponding ex post implementations can be found. Independently and contemporaneously Cai et al. (2012b) provided similar computational foundations for the single-unit auction problem. Their approach to the single-unit auction problem is most comparable to our approach for the multi-unit and matroid auction problems where the optimization problem is written as a convex program which can be solved by the ellipsoid method; while these methods result in strongly polynomial time algorithms they are not considered practical. In contrast, our single-unit approach, when the single-agent problems can be solved by a linear program, gives a single linear program which can be practically solved.

While our work gives computationally tractable interim feasibility characterizations in “service based” environments like multi-unit auctions and matroid auctions; Cai et al. (2012b) generalize the approach to multi-item auctions with agents with additive preferences. The problem of designing an optimal auction for agents with multi-dimensional additive preferences is considered one of the main challenges for auction theory and their result, from a computational perspective, solves this problem.

5.2. Preliminaries²

Single-agent Mechanisms. We consider the provisioning of an abstract service. This service may be parameterized by an *attribute*, e.g., quality of service, and may be accompanied by a required payment. We denote the outcome obtained by an agent as $w \in W$. We view this outcome as giving an indicator for whether or not an agent is served and as describing attributes of the service such as quality of service and monetary payments. Let $\text{Alloc}(w) \in \{0, 1\}$ be an indicator for whether the agent is served or not; let $\text{Payment}(w) \in \mathbb{R}$ denote any payment the agent is required to make. In a randomized environment (e.g., randomness from a randomized mechanism or Bayesian environment) the outcome an agent receives is a random variable from a distribution over W . The space of all such distributions is denoted $\Delta(W)$.

The agent has a type t from a finite type space T . This type is drawn from distribution $f \in \Delta(T)$ and we equivalently denote by f the probability mass function. I.e., for every $t \in T$, $f(t)$ is the probability that the type is t . The utility function $u : T \times W \rightarrow \mathbb{R}$ maps the agent's type and the outcome to real valued utility. The agent is a von Neumann–Morgenstern expected utility maximizer and we extend u to $\Delta(W)$ linearly, i.e., for $w \in \Delta(W)$, $u(t, w)$ is the expectation of u where the outcome is drawn according to w . We do not require the usual assumption of quasi-linearity.

A single-agent mechanism, without loss of generality by the revelation principle, is just an *outcome rule*, a mapping from the agent's type to a distribution over outcomes. We denote an *outcome rule* by $w : T \rightarrow \Delta(W)$. We say that an outcome rule w is *incentive compatible* (IC)

²The model and notation is consistent with Subsection 2.1.2, and is included here to ease reading.

and *individually rational* (IR) if for all $t, t' \in T$, respectively,

$$(IC) \quad u(t, w(t)) \geq u(t, w(t')),$$

$$(IR) \quad u(t, w(t)) \geq 0.$$

We refer to restriction of the outcome rule to the indicator for service as the *allocation rule*. As the allocation to each agent is a binary random variable, distributions over allocations are fully described by their expected value. Therefore the allocation rule $x : T \rightarrow [0, 1]$ for a given outcome rule w is $x(t) = \mathbf{E}[\text{Alloc}(w(t))]$.

We give two examples to illustrate the abstract model described above. The first example is the standard quasi-linear risk-neutral preference which is prevalent in auction theory. Here the agent's type space is $T \subset \mathbb{R}_+$ where $t \in T$ represents the agent's valuation for the item. The outcome space is $W = \{0, 1\} \times \mathbb{R}_+$ where an outcome w in this space indicates whether or not the item is sold to the agent, by $\text{Alloc}(w)$, and at what price, by $\text{Payment}(w)$. The agent's quasi-linear utility function is $u(t, w) = t \cdot \text{Alloc}(w) - \text{Payment}(w)$. The second example is that of an m -item unit-demand (also quasi-linear and risk-neutral) preference. Here the type space is $T \subset \mathbb{R}_+^m$ and a type $t \in T$ indicates the agent's valuation for each of the items when the agent's value for no service is normalized to zero. An outcome space is $W = \{0, \dots, m\} \times \mathbb{R}_+$. The first coordinate of w specifies which item the agent receives or none and $\text{Alloc}(w) = 1$ if it is non-zero; the second coordinate of w specifies the required payment $\text{Payment}(w)$. The agent's utility for w is the value the agent attains for the item received less her payment. Beyond these two examples, our framework can easily incorporate more general agent preferences exhibiting, e.g., risk aversion or a budget limit.

Consider the following single-agent mechanism design problem. A feasibility constraint is given by an upper bound $x(t)$ on the probability that the agent is served as a function of her

type t ; the distribution on types in T is given by f . The single-agent problem is to find the outcome rule w^* that satisfies the allocation constraint of x and maximizes the performance, e.g., revenue. This problem is described by the following program:

$$\begin{aligned}
 \text{(SP)} \quad & \max_w : \mathbf{E}_{t \sim f, w(t)}[\text{Payment}(w(t))] \\
 & \text{s.t.} \quad \mathbf{E}_{w(t)}[\text{Alloc}(w(t))] \leq x(t), \quad \forall t \in T \\
 & w \text{ is IC and IR.}
 \end{aligned}$$

We denote the outcome rule w^* that optimizes this program by $\text{Outcome}(x)$ and its revenue by $\text{Rev}[x] = \mathbf{E}_{t \sim f, w^*(t)}[\text{Payment}(w^*(t))]$. We note that, although this chapter focuses on revenue maximization, the same techniques presented can be applied to maximize (or minimize) general separable objectives such as social welfare.

Multi-agent Mechanisms. There are n independent agents. Agents need not be identical, i.e., agent i 's type space is T_i , the probability mass function for her type is f_i , her outcome space is W_i , and her utility function is u_i . The profile of agent types is denoted by $\mathbf{t} = (t_1, \dots, t_n) \in T_1 \times \dots \times T_n = \mathbf{T}$, the joint distribution on types is $\mathbf{f} \in \Delta(T_1) \times \dots \times \Delta(T_n)$, a vector of outcomes is $(w_1, \dots, w_n) \in \mathbf{W}$, and an allocation is $(x_1, \dots, x_n) \in \{0, 1\}^n$. The mechanism has an inter-agent feasibility constraint that permits serving at most k agents, i.e., $\sum_i x_i \leq k$.³ A mechanism that obeys this constraint is *feasible*. The mechanism has no inter-agent constraint on attributes or payments.

A mechanism maps type profiles to a (distribution over) outcome vectors via an *ex post outcome rule*, denoted $\bar{\mathbf{w}} : \mathbf{T} \rightarrow \Delta(\mathbf{W})$ where $\bar{w}_i(\mathbf{t})$ is the outcome obtained by agent i . We will similarly define $\bar{\mathbf{x}} : \mathbf{T} \rightarrow [0, 1]^n$ as the *ex post allocation rule* (where $[0, 1] \equiv \Delta(\{0, 1\})$). The ex post allocation rule $\bar{\mathbf{x}}$ and the probability mass function \mathbf{f} on types induce *interim* outcome

³Furthermore, in Section 5.5.3, we review the theory of *matroids* and extend our basic results environments with feasibility constraint derived from a matroid set system.

and allocation rules. For agent i with type t_i and $\mathbf{t} \sim \mathbf{Dist}_{\mathbf{t}}[\mathbf{t} | t_i]$ the interim outcome and allocation rules are $w_i(t_i) = \mathbf{Dist}_{\mathbf{t}}[\bar{w}_i(\mathbf{t}) | t_i]$ and $x_i(t_i) = \mathbf{Dist}_{\mathbf{t}}[\bar{x}_i(\mathbf{t}) | t_i] \equiv \mathbf{E}_{\mathbf{t}}[\bar{x}_i(\mathbf{t}) | t_i]$.⁴ A profile of interim allocation rules is feasible if it is derived from an ex post allocation rule as described above; the set of all feasible interim allocation rules is denoted by \mathbb{X} . A mechanism is Bayesian incentive compatible and interim individually rational if equations (IC) and (IR), respectively, hold for all i and all t_i .

Consider again the examples described previously of quasi-linear single-dimensional and unit-demand preferences. For the single-dimensional example, the multi-agent mechanism design problem is the standard single-item k -unit auction problem. For the unit-demand example, the multi-agent mechanism design problem is an *attribute auction*. In this problem there are k -units available and each unit can be configured in one of m ways. Importantly, the designer's feasibility constraint restricts the number of units sold to be k but places no restrictions on how the units can be configured. E.g., a restaurant has k tables but each diner can order any of the m entrees on the menu.

A reduction from multi-agent mechanism design to single-agent mechanism design as we have described above would assume that for any typespace T_i , any probability mass function f_i , and interim allocation rule x_i , the optimal outcome rule $\text{Outcome}(x_i)$ and its performance $\text{Rev}[x_i]$ can be found efficiently (see Section 5.3 for examples). The goal then is to construct an optimal multi-agent auction from these single-agent mechanisms. Our approach to such a reduction is as follows.

- (1) Optimize, over all feasible profiles of interim allocation rules $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{X}$, the sum of performances of the allocation rules $\sum_i \text{Rev}[x_i]$.
- (2) Implement the profile of interim outcome rules \mathbf{w} given by $w_i = \text{Outcome}(x_i)$ with a feasible ex post outcome rule $\bar{\mathbf{w}}$.

⁴We use notation $\mathbf{Dist}[X | E]$ to denote the distribution of random variable X conditioned on the event E .

Two issues should be noted. First, Step 2 requires an argument that the existence of a feasible ex post outcome rule for a given profile of interim allocation rules implies the existence of one that combines the optimal interim outcome rules from $\text{Outcome}(\cdot)$. We address this issue in Section 5.4. Second, Step 1 requires that we optimize over jointly feasible interim allocation rules, and after solving for \mathbf{x} , its implementation by an ex post allocation rule is needed to guide Step 2. For single-unit (i.e., $k = 1$) auctions a characterization of the necessary and sufficient condition for interim feasibility was provided by Kim Border.

Theorem 19 (Border, 1991). *In a single-item auction environment, interim allocation rules \mathbf{x} are feasible (i.e., $\mathbf{x} \in \mathbb{X}$) if and only if the following holds:*

(MRMB)

$$\forall S_1 \subseteq T_1, \dots, \forall S_n \subseteq T_n : \quad \sum_{i=1}^n \mathbf{E}[x_i(t_i) \mid t_i \in S_i] \cdot \Pr[t_i \in S_i] \leq \Pr_{\mathbf{t} \sim \mathbf{f}}[\exists i \in [n] : t_i \in S_i]$$

5.3. The Single-agent Problem

Given an allocation rule $x(\cdot)$ as a constraint the single-agent problem is to find the (possibly randomized) outcome rule $w(\cdot)$ that allocates no more frequently than $x(\cdot)$, i.e., $\forall t \in T$, $\mathbf{E}_{w(t)}[\text{Alloc}(w(t))] \leq x(t)$, with the maximum expected performance. Recall that the optimal such outcome rule is denoted $\text{Outcome}(x)$ and its performance (e.g., revenue) is denoted $\text{Rev}[x]$. We first observe that $\text{Rev}[\cdot]$ is concave.

Proposition 10. *$\text{Rev}[\cdot]$ is a concave function in x .*

Proof. Consider any two allocation rules x and x' , and any $\alpha \in [0, 1]$. Define x'' to be $\alpha x + (1 - \alpha)x'$. We will show that $\alpha \text{Rev}[x] + (1 - \alpha) \text{Rev}[x'] \leq \text{Rev}[x'']$, which proves the claim. To see this, let w and w' be $\text{Outcome}(x)$ and $\text{Outcome}(x')$, respectively. Define w'' to be the outcome rule that runs w with probability α , and w' with probability $1 - \alpha$. The incentive

compatibility of outcome rules w and w' imply the incentive compatibility of w'' , since for any $t, t' \in T$,

$$\begin{aligned} \mathbf{E}[u(t, w''(t))] &= \alpha \mathbf{E}[u(t, w(t))] + (1 - \alpha) \mathbf{E}[u(t, w'(t))] \\ &\geq \alpha \mathbf{E}[u(t, w(t'))] + (1 - \alpha) \mathbf{E}[u(t, w'(t'))] \\ &= \mathbf{E}[u(t, w''(t'))]. \end{aligned}$$

Also, w'' is feasible as $\mathbf{E}[\text{Alloc}(w''(t))] = \alpha \mathbf{E}[\text{Alloc}(w(t))] + (1 - \alpha) \mathbf{E}[\text{Alloc}(w'(t))] \leq x''(t)$ for all $t \in T$. As a result, $\text{Rev}[x'']$ is at least the revenue of w'' , which is in turn equal to $\alpha \text{Rev}[x] + (1 - \alpha) \text{Rev}[x']$. \square

We now give two examples for which the single-agent problem is computationally tractable. Both of these examples are multi-dimensional. The first example is that of a standard multi-item unit-demand preferences. The second example is that of a single-item with a private budget. For both of these problems the single-agent problem can be expressed as a linear program with size polynomial in the cardinality of the agent's type space.

5.3.1. Quasi-linear Unit-demand Preferences

There are m items available. There is a finite type space $T \subset \mathbb{R}_+^m$; the outcome space W is the direct product between an assignment to the agent of one of the m items, or none, and a required payment. $\Delta(W)$ is the cross product of a probability distribution over which item the agent receives and a probability distribution over payments. Without loss of generality for a quasi-linear agent such a randomized outcome can be represented as $w = (w_1, \dots, w_m, w_p)$ where for $j \in [m]$, w_j is the probability that the agent receives item j and w_p is the agent's required payment.

A single-agent mechanism assigns to each type an outcome as described above. An outcome rule specifies an outcome for any type t of the agent as $w(t) = (w_1(t), \dots, w_m(t), w_p(t))$. This gives $m+1$ non-negative real valued variables for each of $|T|$ types. The following linear program, which is a simple adaptation of one from Briest et al. (2010) to include the feasibility constraint given by x , solves for the optimal single-agent mechanism:

$$\begin{aligned}
\max : & \quad \sum_{t \in T} f(t)w_p(t) \\
\text{s.t.} & \quad \sum_j w_j(t) \leq x(t) && \forall t \in T \\
& \quad \sum_j t_j w_j(t) - w_p(t) \geq \sum_j t_j w_j(t') - w_p(t') && \forall t, t' \in T \\
& \quad \sum_j t_j w_j(t) - w_p(t) \geq 0 && \forall t \in T.
\end{aligned}$$

The optimal outcome rule from this program is $w^* = \text{Outcome}(x)$ and its performance is $\text{Rev}[x] = \mathbf{E}_{t \sim f} [w_p^*(t)]$.

Proposition 11. *The single-agent m -item unit-demand problem can be solved in polynomial time in m and $|T|$.*

5.3.2. Private budget preferences.

There is a single item available. The agent has a private value for this item and a private budget, i.e., $T \subset \mathbb{R}_+^2$; we will denote by t_v and t_b this value and budget respectively. The outcome space is $W = \{0, 1\} \times \mathbb{R}$ where for $w \in W$ the first coordinate w_x denotes whether the agent receives the item or not and the second coordinate w_p denotes her payment. The agent's utility is

$$u(t, w) = \begin{cases} t_v w_x - w_p & \text{if } w_p \leq t_b, \text{ and} \\ -\infty & \text{otherwise.} \end{cases}$$

Claim 1 below implies that when optimizing over distributions on outcomes we can restrict attention to $[0, 1] \times [0, 1] \times \mathbb{R}_+ \subset \Delta(W)$ where the first coordinate denotes the probability that the agent receives the item, the second coordinate denotes the probability that the agent makes a non-zero payment, and the third coordinate denotes the non-zero payment made.

Claim 1. *Any incentive compatible and individually rational outcome rule can be converted into an outcome rule above with the same expected revenue.*

As a sketch of the argument to show this claim, note that if an agent with type t receives randomized outcome w she is just as happy to receive the item with the same probability and pay her budget with probability equal to her previous expected payment divided by her budget. Such a payment is budget feasible and has the same expectation as before. Furthermore, this transformation only increases the maximum payment that any agent makes which means that the relevant incentive compatibility constraints are only fewer. Importantly, the only incentive constraints necessary are ones that prevent types with higher budgets from reporting types with lower budgets.

A single-agent mechanism assigns to each type an outcome as described above. We denote the distribution over outcomes for t by $w(t) = (w_x(t), w_\rho(t), t_b)$ where only the first two coordinates are free variables. This gives two non-negative real valued variables for each of $|T|$ types. The

following linear program solves for the optimal single-agent mechanism:

$$\begin{aligned}
\max : & \quad \sum_{t \in T} f(t) t_b w_\rho(t) \\
\text{s.t.} & \quad w_x(t) \leq x(t) && \forall t \in T \\
& \quad t_v w_x(t) - t_b w_\rho(t) \geq t_v w_x(t') - t'_b w_\rho(t') && \forall t, t' \in T \text{ with } t'_b \leq t_b \\
& \quad t_v w_x(t) - t_b w_\rho(t) \geq 0 && \forall t \in T \\
& \quad w_\rho(t) \leq 1 && \forall t \in T.
\end{aligned}$$

The optimal outcome rule from this program is $w^* = \text{Outcome}(x)$ and its performance is $\text{Rev}[x] = \mathbf{E}_{t \sim f} [t_b w_\rho^*(t)]$.

Proposition 12. *The single-agent private budget problem can be solved in polynomial time in $|T|$.*

5.4. Multi- to Single-agent Reductions

An ex post allocation rule $\bar{\mathbf{x}}$ takes as its input a profile of types $\mathbf{t} = (t_1, \dots, t_n)$ of the agents, and indicates by $\bar{x}_i(\mathbf{t})$ a set of at most k winners. Agent i 's type $t_i \in T_i$ is drawn independently at random from distribution $f_i \in \Delta(T_i)$. An ex post allocation rule implements an interim allocation rule $x_i : T_i \rightarrow [0, 1]$, for agent i , if the probability of winning for agent i conditioned on her type $t_i \in T_i$ is exactly $x_i(t_i)$, where the probability is taken over the random types other agents and the random choices of the allocation rule. A profile of interim allocation rules $\mathbf{x} = (x_1, \dots, x_n)$ is feasible if and only if it can be implemented by some ex post allocation rule. \mathbb{X} denotes the space of all feasible profiles of interim allocation rules.

The optimal performance (e.g., revenue) of the single-agent problem with allocation constraint given by x is denoted $\text{Rev}[x]$. The outcome rule corresponding to this optimal revenue is $\text{Outcome}(x)$. Given any feasible interim allocation rule $\mathbf{x} \in \mathbb{X}$ we would like to construct an

auction with revenue $\sum_i \text{Rev}[x_i]$. We need to be careful because $\text{Outcome}(x_i)$, by definition, is only required to have allocation rules *upper bounded* by x_i (see (SP) in Section 5.2), while the ex post allocation rule \bar{x}_i implements x_i exactly, and hence we may need to scale down \bar{x}_i accordingly. This is defined formally as follows.

Definition 29. An optimal auction $\bar{\mathbf{w}}^*$ for feasible interim allocation rule \mathbf{x} (with corresponding ex post allocation rule $\bar{\mathbf{x}}$) is defined as follows on \mathbf{t} . For agent i :

- (1) Let $w_i^* = \text{Outcome}(x_i)$ be the optimal outcome rule for allocation constraint x_i .
- (2) Let $x_i^* = \mathbf{E}[\text{Alloc}(w_i^*)]$ be the allocation rule corresponding to outcome rule w_i^* .
- (3) If $\bar{x}_i(\mathbf{t}) = 1$, output

$$\bar{w}_i^*(\mathbf{t}) \sim \begin{cases} \mathbf{Dist}[w_i^*(t_i) \mid \text{Alloc}(w_i^*(t_i)) = 1] & \text{w.p. } x_i^*(t_i)/x_i(t_i), \text{ and} \\ \mathbf{Dist}[w_i^*(t_i) \mid \text{Alloc}(w_i^*(t_i)) = 0] & \text{otherwise.} \end{cases}$$

- (4) Otherwise (when $\bar{x}_i(\mathbf{t}) = 0$), output $\bar{w}_i^*(\mathbf{t}) \sim \mathbf{Dist}[w_i^*(t_i) \mid \text{Alloc}(w_i^*(t_i)) = 0]$.

Proposition 13. For any feasible interim allocation rule $\mathbf{x} \in \mathbb{X}$, the optimal auction for this rule has expected revenue $\sum_i \text{Rev}[x_i]$.

Proof. The ex post outcome rule $\bar{\mathbf{w}}^*$ of the auction, by construction, induces interim outcome rule \mathbf{w}^* for which the revenue is as desired. \square

The optimal multi-agent auction is the solution to optimizing the cumulative revenue of individual single-agent problems subject to the joint interim feasibility constraint given by $\mathbf{x} \in \mathbb{X}$.

Proposition 14. The optimal revenue is given by the convex program

$$(CP) \quad \max_{\mathbf{x} \in \mathbb{X}} : \sum_i \text{Rev}_i[x_i].$$

Proof. This is a convex program as $\text{Rev}[\cdot]$ is concave and \mathbb{X} is convex (convex combinations of feasible interim allocation rules are feasible). By Proposition 13 this revenue is attainable; therefore, it is optimal. \square

5.5. Optimization and Implementation of Interim Allocation Rules

In this section we address the computational issues pertaining to (i) solving optimization problems over the space of feasible interim allocation rules, and (ii) ex post implementation of such a feasible interim allocation rule. We present computationally tractable methods for both problems.

Normalized interim allocation rules. It will be useful to “flatten” the interim allocation rule \mathbf{x} for which $x_i(t_i)$ denotes the probability that i with type t_i is served (randomizing over the mechanism and the draws of other agent types); we do so as follows. Without loss of generality, we assume that the type spaces of different agents are disjoint.⁵ Denoting the set of all types by $T_N = \bigcup_i T_i$, the interim allocation rule can be flattened as a vector in $[0, 1]^{T_N}$.

Definition 30. The *normalized interim allocation rule* $\tilde{x} \in [0, 1]^{T_N}$ corresponding to interim allocation rule \mathbf{x} under distribution \mathbf{f} is defined as

$$\tilde{x}(t_i) = x_i(t_i)f_i(t_i) \quad \forall t_i \in T_N$$

For the rest of this section, we refer to interim allocation rules via \tilde{x} instead of \mathbf{x} . Note that there is a one-to-one correspondence between \tilde{x} and \mathbf{x} as specified by the above linear equation; so any linear or convex optimization problem involving \mathbf{x} can be written in terms of \tilde{x} without affecting its linearity or convexity. As \mathbb{X} denotes the space of feasible interim allocation rules \mathbf{x} , we will use $\tilde{\mathbb{X}}$ to denote the space of feasible normalized interim allocation rules.

⁵This can be achieved by labeling all types of each agent with the name of that agent, i.e., for each $i \in [n]$ we can replace T_i with $T'_i = \{(i, t) | t \in T_i\}$ so that T'_1, \dots, T'_n are disjoint.

In the remainder of this section we characterize interim feasibility and show that normalized interim allocation rules can be optimized over and implemented in polynomial time.

5.5.1. Single Unit Feasibility Constraints

In this section, we consider environments where at most one agent can be allocated to. For such environments, we characterize interim feasibility as implementability via a particular, simple *stochastic sequential allocation* mechanism. Importantly, the parameters of this mechanism are easy to optimize efficiently.

A stochastic sequential allocation mechanism is parameterized by a stochastic transition table. Such a table specifies the probability by which an agent with a given type can steal a token from a preceding agent with a given type. For simplicity in describing the process we will assume the token starts under the possession of a “dummy agent” indexed by 0; the agents are then considered in the arbitrary order from 1 to n ; and the agent with the token at the end of the process is the one that is allocated (or none are allocated if the dummy agent retains the token).

Definition 31 (stochastic sequential allocation mechanism). Parameterized by a stochastic transition table π , the *stochastic sequential allocation mechanism (SSA)* computes the allocations for a type profile $\mathbf{t} \in \mathbf{T}$ as follows:

- (1) Give the token to the dummy agent 0 with dummy type t_0 .
- (2) For each agent i : (in order of 1 to n)

If agent i' has the token, transfer the token to agent i with probability $\pi(t_{i'}, t_i)$.

- (3) Allocate to the agent who has the token (or none if the dummy agent has it).

First, we present a dynamic program, in the form of a collection of linear equations, for calculating the interim allocation rule implemented by SSA for a given π . Let $y(t_{i'}, i)$ denote

the ex-ante probability of the event that agent i' has type $t_{i'}$ and is holding the token at the end of iteration i . Let $z(t_{i'}, t_i)$ denote the ex-ante probability in iteration i of SSA that agent i has type t_i and takes the token from agent i' who has type $t_{i'}$.

The following additional notation will be useful in this section. For any subset of agents $N' \subseteq N = \{1, \dots, n\}$, we define $T_{N'} = \bigcup_{i \in N'} T_i$ (Recall that without loss of generality agent type spaces are assumed to be disjoint.). The shorthand notation $t_i \in S$ for $S \subseteq T_N$ will be used to quantify over all types in S and their corresponding agents (i.e., $\forall t_i \in S$ is equivalent to $\forall i \in N, \forall t_i \in S \cap T_i$).

The normalized interim allocation rule \tilde{x} resulting from the SSA is exactly given by the dynamic program specified by the following linear equations.

$$(S.1) \quad y(t_0, 0) = 1,$$

$$(S.2) \quad y(t_i, i) = \sum_{t_{i'} \in T_{\{0, \dots, i-1\}}} z(t_{i'}, t_i), \quad \forall t_i \in T_{\{1, \dots, n\}}$$

$$(S.3) \quad y(t_{i'}, i) = y(t_{i'}, i-1) - \sum_{t_i \in T_i} z(t_{i'}, t_i), \quad \forall i \in \{1, \dots, n\}, \forall t_{i'} \in T_{\{0, \dots, i-1\}}$$

$$(\pi) \quad z(t_{i'}, t_i) = y(t_{i'}, i-1) \pi(t_{i'}, t_i) f_i(t_i), \quad \forall t_i \in T_{\{1, \dots, n\}}, \forall t_{i'} \in T_{\{0, \dots, i-1\}}$$

$$\tilde{x}(t_i) = y(t_i, n), \quad \forall t_i \in T_{\{1, \dots, n\}}$$

Note that π is the only adjustable parameter in the SSA algorithm, so by relaxing the equation (π) and replacing it with the following inequality we can specify all possible dynamics of the SSA algorithm.

$$(S.4) \quad 0 \leq z(t_{i'}, t_i) \leq y(t_{i'}, i-1) f_i(t_i), \quad \forall t_i \in T_{\{1, \dots, n\}}, \forall t_{i'} \in T_{\{0, \dots, i-1\}}$$

Let \mathbb{S} denote the convex polytope captured by the 4 sets of linear constraints (S.1) through (S.4) above, i.e., $(y, z) \in \mathbb{S}$ iff y and z satisfy the aforementioned constraints. Note that every $(y, z) \in \mathbb{S}$ corresponds to some stochastic transition table π by solving equation (π) for $\pi(t_i, t_{i'})$. We show that \mathbb{S} captures all feasible normalized interim allocation rules, i.e., the projection of \mathbb{S} on $\tilde{x}(\cdot) = y(\cdot, n)$ is exactly $\tilde{\mathbb{X}}$, as formally stated by the following theorem.

Theorem 20. *A normalized interim allocation rule \tilde{x} is feasible if and only if it can be implemented by the SSA algorithm for some choice of stochastic transition table π . In other words, $\tilde{x} \in \tilde{\mathbb{X}}$ iff there exists $(y, z) \in \mathbb{S}$ such that $\tilde{x}(t_i) = y(t_i, n)$ for all $t_i \in T_N$.*

Corollary 8. *Given a blackbox for each agent i that solves for the optimal expected revenue $\text{Rev}_i[x_i]$ for any feasible interim allocation rule \mathbf{x} , the optimal interim allocation rule can be computed by the following convex program which is of quadratic size in the total number of types.*

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^n \text{Rev}_i[x_i] \\ & \text{subject to} && y(t_i, n) = \tilde{x}(t_i) = x_i(t_i) f_i(t_i), \quad \forall t_i \in T_N \\ & && (y, z) \in \mathbb{S}. \end{aligned}$$

Furthermore, given an optimal assignment for this program, the computed interim allocation rule can be implemented by SSA using the the stochastic transition table defined by:⁶

$$\pi(t_{i'}, t_i) = \frac{z(t_{i'}, t_i)}{y(t_{i'}, i-1) f_i(t_i)}, \quad \forall t_i \in T_{\{1, \dots, n\}}, \forall t_{i'} \in T_{\{0, \dots, i-1\}}.$$

Next, we present a few definitions and lemmas that are used in the proof of Theorem 20. Two transition tables π and π' are considered *equivalent* if their induced normalized interim

⁶If the denominator is zero, i.e., $y(t_i, i' - 1) = 0$, we can set $\pi(t_i, t_{i'})$ to an arbitrary value in $[0, 1]$.

allocation rules for SSA are equal. Type t_i is called *degenerate* for π if in the execution of SSA the token is sometimes passed to type t_i but it is always taken away from t_i later, i.e., if $y(t_i, i) > 0$ but $y(t_i, n) = 0$. The stochastic transition table π is degenerate if there is a degenerate type. For π , type t_i is *augmentable* if there exists a π' (with a corresponding y') which is *equivalent* to π for all types except t_i and has $y(t_i, n) > y'(t_i, n)$.⁷

Lemma 32. *For any stochastic transition table π there exists an equivalent π' that is non-degenerate.*

Lemma 33. *For any non-degenerate stochastic transition table π , any non-augmentable type t_i always wins against any augmentable type $t_{i'}$. I.e.,*

- *if $i' < i$ and $t_{i'}$ has non-zero probability of holding the token then $\pi(t_{i'}, t_i) = 1$, i.e., t_i always takes the token away from $t_{i'}$, and*
- *if $i < i'$ and t_i has non-zero probability of holding the token then $\pi(t_i, t_{i'}) = 0$, i.e., $t_{i'}$ never takes the token away from t_i .*

It is possible to view the token passing in stochastic sequential allocation as a network flow. From this perspective, the augmentable and non-augmentable types form a minimum-cut and Lemma 33 states that the token must eventually flow from the augmentable to non-augmentable types. We defer the proof of this lemma to Appendix C.1 where the main difficulty in its proof is that the edges in the relevant flow problem have dynamic(non-constant) capacities.

PROOF OF THEOREM 20. Any normalized interim allocation rule that can be implemented by the SSA algorithm is obviously feasible, so we only need to prove the opposite direction. The proof is by contradiction, i.e., given a normalized interim allocation rule \tilde{x} we show that if there is no $(y, z) \in \mathbb{S}$ such that $\tilde{x}(\cdot) = y(\cdot, n)$, then \tilde{x} must be infeasible. Consider the following linear

⁷We define t_0 to be augmentable unless the dummy agent never retains the token in which case all agents are non-augmentable (and for technical reasons we declare the dummy agent to be non-augmentable as well).

program for a given \tilde{x} (i.e., \tilde{x} is constant).

$$\begin{aligned} & \text{maximize} && \sum_{t_i \in T_{\{1, \dots, n\}}} y(t_i, n) \\ & \text{subject to} && y(t_i, n) \leq \tilde{x}(t_i), \quad \forall t_i \in T_{\{1, \dots, n\}} \\ & && (y, z) \in \mathbb{S}. \end{aligned}$$

Let (y, z) be an optimal assignment of this LP. If the first set of inequalities are all tight (i.e., $\tilde{x}(\cdot) = y(\cdot, n)$) then \tilde{x} can be implemented by the SSA, so by contradiction there must exist a type $\tau^* \in T_N$ for which the inequality is not tight. Note that τ^* cannot be augmentable — otherwise, by the definition of augmentability, the objective of the LP could be improved. Partition T_N to augmentable types T_N^+ and non-augmentable types T_N^- . Note that T_N^- is non-empty because $\tau^* \in T_N^-$. Without loss of generality, by Lemma 32 we may assume that (y, z) is non-degenerate.⁸

An agent wins if she holds the token at the end of the SSA algorithm. The ex ante probability that some agent with non-augmentable type wins is $\sum_{t_i \in T_N^-} y(t_i, n)$. On the other hand, Lemma 33 implies that the first (in the order agents are considered by SSA) agent with non-augmentable type will take the token from her predecessors and, while she may lose the token to another non-augmentable type, the token will not be relinquished to any augmentable type. Therefore, the probability that an agent with a non-augmentable type is the winner is exactly equal to the probability that at least one such agent exists, therefore

$$\Pr_{\mathbf{t} \sim \mathbf{f}} [\exists i : t_i \in T_N^-] = \sum_{t_i \in T_N^-} y(t_i, n) < \sum_{t_i \in T_N^-} \tilde{x}(t_i).$$

⁸By Lemma 32, there exists a non-degenerate assignment with the same objective value.

The second inequality follows from the assumption above that τ^* satisfies $y(\tau^*, n) < \tilde{x}(\tau^*)$. We conclude that \tilde{x} requires an agent with non-augmentable type to win more frequently than such an agent exists, which is a contradiction to interim feasibility of \tilde{x} . \square

The contradiction that we derived in the proof of Theorem 20 yields a necessary and sufficient condition, as formally stated in the following corollary, for feasibility of any given normalized interim allocation rule.

Corollary 9. *A normalized interim allocation rule \tilde{x} is feasible if and only if*

$$(MRMB) \quad \sum_{\tau \in S} \tilde{x}(\tau) \leq \mathbf{Pr}_{\mathbf{t} \sim f} [\exists i : t_i \in S], \quad \forall S \subseteq T_N$$

The necessity of condition (MRMB) is trivial and its sufficiency was previously proved by Border (1991). This condition implies that the space of all feasible normalized interim allocation rules, $\tilde{\mathbb{X}}$, can be specified by 2^D linear constraints on D -dimensional vectors \tilde{x} . An important consequence of Theorem 20 is that $\tilde{\mathbb{X}}$ can equivalently be formulated by only $O(D^2)$ variables and $O(D^2)$ linear constraints as a projection of \mathbb{S} , therefore any optimization problem over $\tilde{\mathbb{X}}$ can equivalently be solved over \mathbb{S} .

5.5.2. k -Unit Feasibility Constraints

In this section, we consider environments where at most k agents can be simultaneously allocated to. First, we generalize Border's characterization of interim feasibility to environments with k -unit feasibility constraint. Our generalization implies that the space of feasible normalized interim allocation rules is a polymatroid. Second, we observe that optimization problems can be efficiently solved over polymatroids; this allows us to optimize over feasible interim allocation rules. Third, we show that the normalized interim allocation rules corresponding to the vertices of this polymatroid are implemented by simple deterministic ordered-subset-based allocation

mechanisms. Furthermore, for any point in this polymatroid, the corresponding normalized interim allocation rule can be implemented by, (i) expressing it as a convex combination of the vertices of the polymatroid, (ii) sampling from this convex combination, and (iii) using the ordered subset mechanism corresponding to the sampled vertex. We present an efficient randomized rounding routine for rounding a point in a polymatroid to a vertex which combines the steps (i) and (ii). These approaches together yield efficient algorithms for optimizing and implementing interim allocation rules.

Polymatroid Preliminaries. Consider an arbitrary set function $\mathcal{F} : 2^U \rightarrow \mathbb{R}_+$ defined over an arbitrary finite set U ; let $P_{\mathcal{F}}$ denote the polytope associate with \mathcal{F} defined as

$$P_{\mathcal{F}} = \{y \in \mathbb{R}_+^U \mid \forall S \subseteq U : y(S) \leq \mathcal{F}(S)\}$$

where $y(S)$ denotes $\sum_{s \in S} y(s)$. The convex polytope $P_{\mathcal{F}}$ is called a *polymatroid* if \mathcal{F} is a submodular function. Even though a polymatroid is defined by an exponential number of linear inequalities, the separation problem for any given $y \in \mathbb{R}_+^U$ can be solved in polynomial time as follows: find $S^* = \arg \min_S \mathcal{F}(S) - y(S)$; if y is infeasible, the inequality $y(S^*) \leq \mathcal{F}(S^*)$ must be violated, and that yields a separating hyperplane for y . Note that $\mathcal{F}(S) - y(S)$ is itself submodular in S , so it can be minimized in strong polynomial time. Consequently, optimization problems can be solved over polymatroids in polynomial time. Next, we describe a characterization of the vertices of a polymatroid. This characterization plays an important role in our proofs and also in our ex post implementation of interim allocation rules.

Definition 32 (ordered subset). For an arbitrary finite set U , an *ordered subset* $\pi \subset U$ is given by an ordering on elements $\pi = (\pi_1, \dots, \pi_{|\pi|})$ where shorthand notation $\pi_r \in \pi$ denotes the r th element in π .

Proposition 15. *Let $\mathcal{F} : 2^U \rightarrow \mathbb{R}_+$ be an arbitrary non-decreasing submodular function with $\mathcal{F}(\emptyset) = 0$ and let $P_{\mathcal{F}}$ be the associated polymatroid with the set of vertices $\text{VERTEX}(P_{\mathcal{F}})$. Every ordered subset π of U (see Definition 32) corresponds to a vertex of $P_{\mathcal{F}}$, denoted by $\text{VERTEX}(P_{\mathcal{F}}, \pi)$, which is computed as follows.*

$$\forall s \in U : \quad y(s) = \begin{cases} \mathcal{F}(\{\pi_1, \dots, \pi_r\}) - \mathcal{F}(\{\pi_1, \dots, \pi_{r-1}\}) & \text{if } s = \pi_r \in \pi \\ 0 & \text{if } s \notin \pi \end{cases}$$

Furthermore, for every $y \in \text{VERTEX}(P_{\mathcal{F}})$ there exist a corresponding π .

It is easy to see that for any $y \in \text{VERTEX}(P_{\mathcal{F}})$, an associated π can be found efficiently by a greedy algorithm (see Schrijver (2003) for a comprehensive treatment of polymatroids).

Ordered subset allocation mechanisms. The following class of allocation mechanisms are of particular importance both in our characterization of interim feasibility and in ex post implementation of interim allocation rules.

Definition 33 (ordered subset allocation mechanism). Parameterized by a *ordered subset* π of T_N (see Definition 32), the *ordered subset mechanism*, on profile of types $\mathbf{t} \in \mathbf{T}$, orders the agents based on their types according to π , and allocates to the agents greedily (e.g., with k units available the k first ordered agents received a unit). If an agent i with type $t_i \notin \pi$ is never served.⁹

Characterization of interim feasibility. Border's characterization of interim feasibility for $k = 1$ unit auctions states that the probability of serving a type in a subspace of type space is no more than the probability that a type in that subspace shows up. This upper bound is

⁹The virtual valuation maximizing mechanisms from the classical literature on revenue maximizing auctions are ordered subset mechanisms, see, e.g., Myerson (1981), an observation made previously by Edith (2007). The difference between these ordered subset mechanisms and the classic virtual valuation maximization mechanisms is that our ordered subset will come from solving an optimization on the whole auction problem where as the Myerson's virtual values come directly from single-agent optimizations.

equivalent to the expected minimum of one and the number of types from the subspace that show up; furthermore, this equivalent phrasing of the upper bound extends to characterize interim feasibility in k -unit auctions.

Consider expressing an ex post allocation for type profile \mathbf{t} by $\bar{x}^{\mathbf{t}} \in \{0, 1\}^{T_N}$ as follows. For all $t'_i \in T_N$, $\bar{x}^{\mathbf{t}}(t'_i) = 1$ if player i is served and $t_i = t'_i$ and 0 otherwise. This definition of ex post allocations is convenient as the normalized interim allocation rule is calculated by taking its expectation, i.e., $\tilde{x}(t'_i) = \mathbf{E}_{\mathbf{t}}[\bar{x}^{\mathbf{t}}(t'_i)]$. Ex post feasibility requires that,

$$(5.1) \quad \bar{x}^{\mathbf{t}}(S) \leq \min(|\mathbf{t} \cap S|, k), \quad \forall \mathbf{t} \in \mathbf{T}, \forall S \subseteq T_N$$

In other words: for any profile of types \mathbf{t} , the number of types in S that are served by $\bar{x}^{\mathbf{t}}$ must be at most the number of types in S that showed up in \mathbf{t} and the upper bound k . Taking expectations of both sides of this equation with respect to \mathbf{t} motivates the following definition and theorem.

Definition 34. The *expected rank function* for distribution \mathbf{f} and subspace $S \subset T_N$ is

$$(g_k) \quad g_k(S) = \mathbf{E}_{\mathbf{t} \sim \mathbf{f}}[\min(|\mathbf{t} \cap S|, k)]$$

where $\mathbf{t} \cap S$ denotes $\{t_1, \dots, t_n\} \cap S$.

Theorem 21. For supply constraint k and distribution \mathbf{f} , the space of all feasible normalized interim allocation rules, $\tilde{\mathbf{X}}$, is the polymatroid associated with g_k , i.e., $\tilde{\mathbf{X}} = P_{g_k}$, i.e., for all $\tilde{x} \in \tilde{\mathbf{X}}$,

$$(5.2) \quad \tilde{x}(S) \leq \mathbf{E}_{\mathbf{t}}[\min(|\mathbf{t} \cap S|, k)] = g_k(S), \quad \forall S \subseteq T_N$$

The proof of this theorem will be deferred to the next section where we will derive a more general theorem. A key step in the proof will be relating the statement of the theorem to the polymatroid theory described already. To show that the constraint of the theorem is a polymatroid, we observe that the expected rank function is submodular.

Lemma 34. *The expected rank function g_k is submodular.*

Proof. Observe that for any fixed \mathbf{t} , $\min(\mathbf{t} \cap S, k)$ is obviously a submodular function in S , and therefore g_k is a convex combination¹⁰ of submodular functions, so g_k is submodular. \square

We now relate vertices of the polymatroid to ordered subset allocation mechanisms.

Theorem 22. *For supply constraint k , if $\tilde{x} \in \tilde{\mathbf{X}}$ is the vertex $\text{VERTEX}(P_{g_k}, \pi)$ of the polymatroid P_{g_k} the unique ex post implementation is the ordered subset mechanism induced by π (Definition 33).*

Proof. Let $\tilde{x} = \text{VERTEX}(P_{g_k}, \pi)$ be an arbitrary vertex of P_{g_k} with a corresponding ordered subset π ; by Proposition 15, such a π exists for every vertex of a polymatroid. For every integer $r \leq |\pi|$, define $S^r = \{\pi_1, \dots, \pi_r\}$ as the r -element prefix of the ordering. By Proposition 15, inequality (5.2) must be tight for every S^r which implies that inequality (5.1) must also be tight for every S^r and every $\mathbf{t} \in \mathbf{T}$. Observe that inequality (5.1) being tight for a subset S of types implies that any ex post allocation mechanism implementing \tilde{x} must allocate as much as possible to types in S . By definition, an ordered subset mechanism allocates to as many types as possible (up to k) from each S^r ; this is the unique outcome given that inequality (5.1) is tight for every S^r . \square

¹⁰Note that taking the expectation is the same as taking a convex combination.

Optimization over feasible interim allocation rules. The characterization of interim feasibility as a polymatroid constraint immediately enables efficient solving of optimization problems over the feasible interim allocation rules as long as we can compute g_k efficiently (see Schrijver (2003) for optimization over polymatroids). The following lemma states that g_k can be computed efficiently.

Lemma 35. *For independent agent (i.e., if \mathbf{f} is a product distribution), $g_k(S)$ can be exactly computed in time $O((n + |S|) \cdot k)$ for any $S \in T_N$ using dynamic programming.*

Ex post implementation of feasible interim allocation rules. We now address the task of finding an ex post implementation corresponding to any $\tilde{x} \in \tilde{\mathbf{X}}$. By Theorem 25, if \tilde{x} is a vertex of $\tilde{\mathbf{X}}$, it can be easily implemented by an ordered subset allocation mechanism (Definition 33). As any point in the polymatroid (or any convex polytope) can be specified as a convex combination of its vertices, to implement the corresponding interim allocation rule it is enough to show that this convex combination can be efficiently sampled from. An ex post implementation can then be obtained by sampling a vertex and using the ordered subset mechanism corresponding to that vertex. Instead of explicitly computing this convex combination, we present a general randomized rounding routine $\text{RANDROUND}(\cdot)$ which takes a point in a polymatroid and returns a vertex of the polymatroid such that the expected value of every coordinate of the returned vertex is the same as the original point. This approach is formally described next.

Definition 35 (randomized ordered subset allocation mechanism). Parameterized by a normalized interim allocation rule $\tilde{x} \in \tilde{\mathbf{X}}$, a *randomized ordered subset allocation mechanism (RRA)* computes the allocation for a profile of types $\mathbf{t} \in \mathbf{T}$ as follows.

- (1) Let $(\tilde{x}^*, \pi^*) \leftarrow \text{RANDROUND}(\tilde{x})$.
- (2) Run the ordered subset mechanism (Definition 33) with ordered subset π^* .

Theorem 23. *Any normalized interim allocation rule $\tilde{x} \in \tilde{\mathbb{X}}$ can be implemented by the randomized ordered subset allocation mechanism (Definition 35) as a distribution over deterministic ordered subset allocation mechanisms.*

Proof. The proof follows from linearity of expectation. \square

Randomized rounding for polymatroids. We describe $\text{RANDROUND}(\cdot)$ for general polymatroids. First, we present a few definitions and give an overview of the rounding operator. Consider an arbitrary finite set U and a polymatroid $P_{\mathcal{F}}$ associated with a non-decreasing submodular function $\mathcal{F} : 2^U \rightarrow \mathbb{R}_+$ with $\mathcal{F}(\emptyset) = 0$. A set $S \subseteq U$ is called *tight* with respect to a $y \in P_{\mathcal{F}}$, if and only if $y(S) = \mathcal{F}(S)$. A set $\mathbb{S} = \{S^0, \dots, S^m\}$ of subsets of U is called a *nested family of tight sets* with respect to $y \in P_{\mathcal{F}}$, if and only the elements of \mathbb{S} can be ordered and indexed such that $\emptyset = S^0 \subset \dots \subset S^m \subseteq U$, and such that S^r is tight with respect to y (for every $r \in [m]$).

$\text{RANDROUND}(y)$ takes an arbitrary $y \in P_{\mathcal{F}}$ and iteratively makes small changes to it until a vertex is reached. At each iteration ℓ , it computes $y^\ell \in P_{\mathcal{F}}$, and a nested family of tight sets \mathbb{S}^ℓ (with respect to y^ℓ) such that

- $\mathbf{E}[y^\ell | y^{\ell-1}] = y^{\ell-1}$, and
- y^ℓ is closer to a vertex (compared to $y^{\ell-1}$) in the sense that either the number of non-zero coordinates has decreased by one or the number of tight sets has increased by one.

Observe that the above process must stop after at most $2|U|$ iterations¹¹. At each iteration ℓ of the rounding process, a vector $\hat{y} \in \mathbb{R}^U$ and $\delta, \delta' \in \mathbb{R}_+$ are computed such that both $y^{\ell-1} + \delta \cdot \hat{y}$ and $y^{\ell-1} - \delta' \cdot \hat{y}$ are still in $P_{\mathcal{F}}$, but closer to a vertex. The algorithm then chooses a random $\delta'' \in \{\delta, -\delta'\}$ such that $\mathbf{E}[\delta''] = 0$, and sets $y^\ell \leftarrow y^{\ell-1} + \delta'' \cdot \hat{y}$.

¹¹In fact we will show that it stops after at most $|U|$ iterations.

Definition 36 ($\text{RANDROUND}(y)$). This operator takes as its input a point $y \in P_{\mathcal{F}}$ and returns as its output a pair (y^*, π^*) , where y^* is a random vertex of $P_{\mathcal{F}}$ and π^* is its associated ordered subset (see Proposition 15), and such that $\mathbf{E}[y^*] = y$.

The algorithm modifies y iteratively until a vertex is reached. It also maintains a nested family of tight sets \mathbb{S} with respect to y . As we modify \mathbb{S} , we always maintain an ordered labeling of its elements, i.e., if $\mathbb{S} = \{S^0, \dots, S^m\}$, we assume that $\emptyset = S^0 \subset \dots \subset S^m \subseteq U$; in particular, the indices are updated whenever a new tight set is added. For each $s \in U$, define $\mathbf{1}_s \in [0, 1]^U$ as a vector whose value is 1 at coordinate s and 0 everywhere else.

(1) Initialize $\mathbb{S} \leftarrow \{\emptyset\}$.

(2) Repeat each of the following steps until no longer applicable:

- If there exist distinct $s, s' \in S^r \setminus S^{r-1}$ for some $r \in [m]$:

(a) Set $\hat{y} \leftarrow \mathbf{1}_s - \mathbf{1}_{s'}$, and compute $\delta, \delta' \in \mathbb{R}_+$ such that $y + \delta \cdot \hat{y}$ has a new tight set S and $y - \delta' \cdot \hat{y}$ has a new tight set S' , i.e.,

$$- \text{ set } S \leftarrow \arg \min_{S^{r-1}+s \subseteq S \subseteq S^r-s'} \mathcal{F}(S) - y(S), \text{ and } \delta \leftarrow \mathcal{F}(S) - y(S);$$

$$- \text{ set } S' \leftarrow \arg \min_{S^{r-1}+s' \subseteq S' \subseteq S^r-s} \mathcal{F}(S') - y(S'), \text{ and } \delta' \leftarrow \mathcal{F}(S') - y(S').$$

(b) $\left\{ \begin{array}{l} \text{with prob. } \frac{\delta}{\delta+\delta'}: \quad \text{set } y \leftarrow y + \delta \cdot \hat{y}, \text{ and add } S \text{ to } \mathbb{S}. \\ \text{with prob. } \frac{\delta'}{\delta+\delta'}: \quad \text{set } y \leftarrow y - \delta' \cdot \hat{y}, \text{ and add } S' \text{ to } \mathbb{S}. \end{array} \right.$

- If there exists $s \in U \setminus S^m$ for which $y(s) > 0$:

(a) Set $\hat{y} \leftarrow \mathbf{1}_s$, and compute $\delta, \delta' \in \mathbb{R}_+$ such that $y + \delta \cdot \hat{y}$ has a new tight set S and $y - \delta' \cdot \hat{y}$ has a zero at coordinate s , i.e.,

$$- \text{ set } S \leftarrow \arg \min_{S \supseteq S^m+s} \mathcal{F}(S) - y(S), \text{ and } \delta \leftarrow \mathcal{F}(S) - y(S);$$

$$- \text{ set } \delta' \leftarrow y(s).$$

(b) $\left\{ \begin{array}{l} \text{with prob. } \frac{\delta}{\delta+\delta'}: \quad \text{set } y \leftarrow y + \delta \cdot \hat{y}, \text{ and add } S \text{ to } \mathbb{S}. \\ \text{with prob. } \frac{\delta'}{\delta+\delta'}: \quad \text{set } y \leftarrow y - \delta' \cdot \hat{y} \end{array} \right.$

- (3) Set $y^* \leftarrow y$ and define the ordered subset $\pi^* : S^m \rightarrow [m]$ according to $\$$, i.e., for each $r \in [m]$ and $s \in S^r \setminus S^{r-1}$, define $\pi^*(s) = r$.
- (4) Return (y^*, π^*) .

Theorem 24. *For any non-decreasing submodular function $\mathcal{F} : 2^U \rightarrow \mathbb{R}_+$ and any $y \in P_{\mathcal{F}}$, the operator $\text{RANDROUND}(y)$ returns a random (y^*, π^*) such that $y^* \in \text{VERTEX}(P_{\mathcal{F}})$, and π^* is the ordered subset corresponding to y^* (see Proposition 15), and such that $\mathbf{E}[y^*] = y$. Furthermore, the algorithm runs in strong polynomial time. In particular, it runs for $O(|U|)$ iterations where each iteration involves solving two submodular minimizations.*

5.5.3. Matroid Feasibility Constraints

In this section, we consider environments where the feasibility constraints are encoded by a matroid $\mathcal{M} = (T_N, \mathcal{I})$. For every type profile $\mathbf{t} \in \mathbf{T}$, a subset $S \subseteq \{t_1, \dots, t_n\}$ can be simultaneously allocated to if and only if $S \in \mathcal{I}$. We show that the results of Subsection 5.5.2 can be easily generalized to environments with matroid feasibility constraints.

Matroid Preliminaries. A matroid $\mathcal{M} = (U, \mathcal{I})$ consists of a ground set U and a family of independent sets $\mathcal{I} \subseteq 2^U$ with the following two properties.

- For every $I, I' \in \mathcal{I}$, if $I' \subset I$, then $I' \in \mathcal{I}$.
- For every $I, I' \in \mathcal{I}$, if $|I'| < |I|$, there exists $s \in I \setminus I'$ such that $I' \cup \{s\} \in \mathcal{I}$.

For every matroid \mathcal{M} , the rank function $r_{\mathcal{M}} : 2^U \rightarrow \mathbb{N} \cup \{0\}$ is defined as follows: for each $S \subseteq U$, $r_{\mathcal{M}}(S)$ is the size of the maximum independent subset of S . A matroid can be uniquely characterized by its rank function, i.e., a set $I \subseteq U$ is an independent set if and only if $r_{\mathcal{M}}(I) = |I|$. A matroid rank function has the following two properties:

- $r_{\mathcal{M}}(\cdot)$ is a non-negative non-decreasing integral submodular function.
- $r_{\mathcal{M}}(S) \leq |S|$, for all $S \subseteq U$.

Furthermore, every function with the above properties defines a matroid.

Any set $S \subseteq U$ can be equivalently represented by its incidence vector $\chi^S \in \{0, 1\}^U$ which has a 1 at every coordinate $s \in S$ and 0 everywhere else.

Proposition 16. *Consider an arbitrary finite matroid $\mathcal{M} = (U, \mathcal{I})$ with rank function $r_{\mathcal{M}}(\cdot)$. Let $P_{r_{\mathcal{M}}}$ denote the polymatroid associated with $r_{\mathcal{M}}(\cdot)$ (see Subsection 5.5.2); the vertices of $P_{r_{\mathcal{M}}}$ are exactly the incidence vectors of the independent sets of \mathcal{M} .*

See Schrijver (2003) for a comprehensive treatment of matroids.

Characterization of interim feasibility. We now generalize the characterization of interim feasibility as the polymatroid given by the expected rank of the matroid. From this generalization the computational results of the preceding section can be extended from k -unit environments to matroids.

Let b denote the random bits used by an ex post allocation rule, and let $\bar{x}^{\mathbf{t}, b} \in \{0, 1\}^{T_N}$ denote the ex post allocation rule (i.e., the incidence vector of the subset of types that get allocated to) for type profile $\mathbf{t} \in \mathbf{T}$ and random bits b . It is easy to see that $\bar{x}^{\mathbf{t}, b}$ is a feasible ex post allocation if and only if it satisfies the following class of inequalities.

$$(5.3) \quad \bar{x}^{\mathbf{t}, b}(S) \leq r_{\mathcal{M}}(\mathbf{t} \cap S), \quad \forall \mathbf{t} \in \mathbf{T}, \forall S \subseteq T_N$$

The above inequality states that the subset of types that get allocated to must be an independent set of the restriction of matroid \mathcal{M} to $\{t_1, \dots, t_n\}$. The expectation of the left-hand-side is exactly the normalized interim allocation rule, i.e., for any $t'_i \in T_N$, $\tilde{x}(t'_i) = \mathbf{E}_{\mathbf{t}, b}[\bar{x}^{\mathbf{t}, b}(t'_i)]$. Taking expectations of both sides of (5.3) then motivates the following definition and theorem that characterize interim feasibility.

Definition 37. The *expected rank* for distribution \mathbf{f} , subspace $S \subset T_N$, and matroid \mathcal{M} with rank function $r_{\mathcal{M}}$ is

$$(g_{\mathcal{M}}) \quad g_{\mathcal{M}}(S) = \mathbf{E}_{\mathbf{t} \sim \mathbf{f}} [r_{\mathcal{M}}(\mathbf{t} \cap S)]$$

where $\mathbf{t} \cap S$ denotes $\{t_1, \dots, t_n\} \cap S$.

Theorem 25. For matroid \mathcal{M} and distribution \mathbf{f} , the space of all feasible normalized interim allocation rules, $\tilde{\mathbf{X}}$, is the polymatroid associated with $g_{\mathcal{M}}$, i.e., $\tilde{\mathbf{X}} = P_{g_{\mathcal{M}}}$, i.e., for all $\tilde{x} \in \tilde{\mathbf{X}}$,

$$(5.4) \quad \tilde{x}(S) \leq \mathbf{E}_{\mathbf{t}} [r_{\mathcal{M}}(\mathbf{t} \cap S)] = g_{\mathcal{M}}(S), \quad \forall S \subseteq T_N$$

Theorem 26. For matroid \mathcal{M} , if $\tilde{x} \in \tilde{\mathbf{X}}$ is the vertex $\text{VERTEX}(P_{g_{\mathcal{M}}}, \pi)$ of the polymatroid $P_{g_{\mathcal{M}}}$ the unique ex post implementation is the ordered subset mechanism induced by π (Definition 33).

To prove the above theorems, we use the following decomposition lemma which applies to general polymatroids.

Lemma 36 (Polymatroidal Decomposition). Let U be an arbitrary finite set, $\mathcal{F}^1, \dots, \mathcal{F}^m : 2^U \rightarrow \mathbb{R}_+$ be arbitrary non-decreasing submodular functions, and $\mathcal{F}^* = \sum_{j=1}^m \lambda^j \mathcal{F}^j$ be an arbitrary convex combination of them. For every y^* the following holds: $y^* \in P_{\mathcal{F}^*}$ if and only if it can be decomposed as $y^* = \sum_{j=1}^m \lambda^j y^j$ such that $y^j \in P_{\mathcal{F}^j}$ (for each $j \in [m]$). Furthermore, if y^* is a vertex of $P_{\mathcal{F}^*}$, this decomposition is unique. More precisely, if $y^* = \text{VERTEX}(P_{\mathcal{F}^*}, \pi)$ for some ordered subset π , then $y^j = \text{VERTEX}(P_{\mathcal{F}^j}, \pi)$ (for each $j \in [m]$).

Proof. First, observe that the only-if part is obviously true, i.e., if $y^j \in P_{\mathcal{F}^j}$ (for each $j \in [m]$), we can write

$$(5.5) \quad y^j(S) \leq \mathcal{F}^j(S) \quad \forall S \subseteq U,$$

multiplying both sides by λ^j and summing over all $j \in [m]$ we obtain

$$(5.6) \quad y^*(S) = \sum_{i=1}^m \lambda^i y^i(S) \leq \sum_{j=1}^m \lambda^j \mathcal{F}^j(S) = \mathcal{F}^*(S) \quad \forall S \subseteq U,$$

which implies that $y^* \in P_{\mathcal{F}^*}$.

Next, we prove that for every $y^* \in P_{\mathcal{F}^*}$ such a decomposition exists. Note that a polymatroid is a convex polytope, so any $y^* \in P_{\mathcal{F}^*}$ can be written as a convex combination of vertices as $y^* = \sum_{\ell} \alpha^{\ell} y^{*\ell}$, where each $y^{*\ell}$ is a vertex of $P_{\mathcal{F}^*}$; consequently, if we prove the claim for the vertices of $P_{\mathcal{F}^*}$, i.e., that $y^{*\ell} = \sum_{j=1}^m \lambda^j y^{\ell,j}$ for some $y^{\ell,j} \in P_{\mathcal{F}^j}$, then a decomposition of $y^* = \sum_{j=1}^m \lambda^j y^j$ can be obtained by setting $y^j = \sum_{\ell} \alpha^{\ell} y^{\ell,j}$.

Next, we prove the second part of the theorem which also implies that a decomposition exists for every vertex of $P_{\mathcal{F}^*}$. Let $y^* = \text{VERTEX}(P_{\mathcal{F}^*}, \pi)$ be an arbitrary vertex of $P_{\mathcal{F}^*}$ with a corresponding ordered subset π ; by Proposition 15, such a π exists for every vertex of a polymatroid. For every integer $r \leq |\pi|$, define $S^r = \{\pi_1, \dots, \pi_r\}$ as the r -element prefix of the ordering. By Proposition 15, inequality (5.6) is tight for every S^r , which implies that inequality (5.5) must also be tight for each S^r and for every $j \in [m]$. Consequently, for each $r \in [|\pi|]$, and each $j \in [m]$, by taking the difference of the inequality (5.6) for S^r and S^{r-1} , given that they are tight, we obtain

$$y^j(S) = \mathcal{F}^j(S^r) - \mathcal{F}^j(S^{r-1})$$

Furthermore, for each $s \notin \pi$, $y^*(s) = 0$ which implies that $y^j(s) = 0$ for every $j \in [m]$. Observe that we have obtained a unique y^j for each $j \in [m]$ which is exactly the vertex of $P_{\mathcal{F}^j}$ corresponding to π as described in Proposition 15. It is easy to verify that indeed $y^* = \sum_{j=1}^m \lambda^j y^j$. \square

PROOF OF THEOREM 25. The inequality in equation (5.3) states that the subset of types that get allocated to must be an independent set of the restriction of matroid \mathcal{M} to $\{t_1, \dots, t_n\}$. Define $r_{\mathcal{M}}^{\mathbf{t}}(S) = r_{\mathcal{M}}(\mathbf{t} \cap S)$ (for all $(S \subseteq T_N)$). Notice that $r_{\mathcal{M}}^{\mathbf{t}}$ is a submodular function. The above inequality implies that $\bar{x}^{\mathbf{t},b} \in P_{r_{\mathcal{M}}^{\mathbf{t}}}$. Define $\bar{x}^{\mathbf{t}} = \mathbf{E}_b[\bar{x}^{\mathbf{t},b}]$. Observe that $\tilde{x} = \mathbf{E}_{\mathbf{t}}[\bar{x}^{\mathbf{t}}]$, so \tilde{x} is a feasible normalized interim allocation rule if and only if it can be decomposed as $\tilde{x} = \sum_{\mathbf{t} \in \mathbf{T}} f(\mathbf{t}) \bar{x}^{\mathbf{t}}$ where $\bar{x}^{\mathbf{t}} \in P_{r_{\mathcal{M}}^{\mathbf{t}}}$ for every $\mathbf{t} \in \mathbf{T}$; by Lemma 36, this is equivalent to $\tilde{x} \in P_{g_{\mathcal{M}}}$ where $g_{\mathcal{M}}(S) = \sum_{\mathbf{t} \in \mathbf{T}} f(\mathbf{t}) r_{\mathcal{M}}^{\mathbf{t}}(S) = \mathbf{E}_{\mathbf{t}}[r_{\mathcal{M}}^{\mathbf{t}}(S)]$ (for all $S \subseteq T_N$), as defined in Definition 37. That completes the proof of the first part of the theorem. \square

PROOF OF THEOREM 26. Suppose $\tilde{x} = \text{VERTEX}(P_{g_{\mathcal{M}}}, \pi)$ for some ordered subset π . By Lemma 36, the decomposition of \tilde{x} is unique and is given by $\bar{x}^{\mathbf{t}} = \text{VERTEX}(P_{r_{\mathcal{M}}^{\mathbf{t}}}, \pi)$. Notice that this is the same allocation obtained by the deterministic rank-based allocation mechanism which ranks according to π (see Definition 33). \square

Optimization over feasible interim allocation rules. As in Subsection 5.5.2, the characterization of interim feasibility as a polymatroid constraint immediately enables efficient solving of optimization problems over the feasible normalized interim allocation rules as long as we can compute $g_{\mathcal{M}}$ efficiently (see Schrijver (2003) for optimization over polymatroids). Depending on the specific matroid, it might be possible to exactly compute $g_{\mathcal{M}}$ in polynomial time (e.g., as in Lemma 35); otherwise, it can be computed approximately within a factor of $1 - \epsilon$ and with probability $1 - \delta$, by sampling, in time polynomial in $\frac{1}{\epsilon}$ and $\frac{1}{\delta}$.

Ex post implementation of feasible interim allocation rules. An ex post implementation for any $\tilde{x} \in \tilde{X}$ can be obtained exactly as in Subsection 5.5.2.

Corollary 10. *Any normalized interim allocation rule $\tilde{x} \in \tilde{X}$ can be implemented by the randomized rank-based allocation mechanism (Definition 35) as a distribution over deterministic rank-based allocation mechanisms (Definition 33).*

5.6. Conclusions and Extensions

In this chapter we have focused on binary allocation problems where an agent is either served or not served. For these binary allocation problems distributions over allocations are given by a single number, i.e., the probability that the agent is served. Our results can be extended to environments with multi-unit demand when the agents utility is linear in the expected number of units the agent receives.

In Section 5.5 we described algorithms for optimizing over feasible interim allocation rules and for (ex post) implementation of the resulting rules. Neither these algorithms nor the generalization of Border's condition require the types of the agents to be independently distributed. However, our formulation of incentive compatibility for interim allocation rules does require independence. For correlated distributions the interim allocation rule is a function of the actual type of the agent (which conditions the types of the other agents) and the reported type of the agent. Therefore, this generalization of our theorem to correlated environments has little relevance for mechanism design.

The algorithms in Section 5.5 do not require the feasibility constraint to be known in advance. A simple example where this generalization is interesting is a multi-unit auction where the supply k is stochastically drawn from a known distribution. Our result shows that the optimal auction in such an environment can be described by picking the random ordering on types and allocating

greedily by this ordering while supplies last. We do not know of many examples other than this where this generalization is interesting.

Our techniques can also be used in conjunction with the approach of Cai et al. (2012a) for solving multi-item auction problems for agents with additive values.

One important extension of our work is to scenarios where the type space and distribution are only available via oracle access (and can be very large or even infinite). Given a polynomial time approximation scheme (PTAS) for a variant of the single agent problem we can construct a polynomial time approximation scheme for the multi-agent problem. Such a model is important, for instance, when the agents type space is multi-dimensional but succinctly describable, e.g., for unit demand agents with independently distributed values for various items for sale. Such a type space would be exponentially large in the number of items but succinctly described in polynomial space in the number of items. While our reduction can be applied to this scenario, we do not know of any solution to the optimal single-agent problem with which to instantiate the reduction.

Finally, in symmetric environments, i.e., when the agents' type distribution and the designer's feasibility constraint are symmetric, e.g., for i.i.d. multi-unit auctions, the optimal interim allocation constraint imposed by feasibility is symmetric; furthermore, the constraint is given by a simple formula. Therefore, symmetric multi-agent problems reduce to solving the single-agent problem for a very specific constraint on the interim allocation rule. In comparison to the computational task of optimizing over feasible interim allocation rules and solving for ex post implementations, e.g., from Section 5.5, this multi-agent reduction for symmetric environments is computationally trivial.

Bibliography

- Alessandro Acquisti and Hal R Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, 2005.
- Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *FOCS*, pages 512–521, 2011.
- Saeed Alaei, Hu Fu, Nima Haghpanah, Jason Hartline, and Azarakhsh Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*, 2012.
- Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. The simple economics of approximately optimal auctions. In *FOCS*, 2013.
- Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, January 1996. URL <http://ideas.repec.org/a/ecm/emetrp/v64y1996i1p51-75.html>.
- Mark Armstrong. Price discrimination by a many-product firm. *The Review of Economic Studies*, 66(1):151–168, 1999.
- Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
- Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget constrained auctions with heterogeneous items. In *44th ACM Symposium on Theory of Computing*, pages 379–388, 2010.
- Kim Border. Reduced form auctions revisited. *Economic Theory*, 31(1):167–181, 2007.

- Kim C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59(4):1175–1187, 1991.
- Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing randomized allocations. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 585–597, 2010.
- Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996.
- Jeremy Bulow and John Roberts. The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–1090, October 1989.
- Y. Cai, C. Daskalakis, and M. Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Symposium on Theory of Computation*, 2012a.
- Y. Cai, C. Daskalakis, and M. Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Symposium on Theory of Computation*, 2012b.
- Y. Cai, C. Daskalakis, and M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *Foundations of Computer Science*, 2012c.
- Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531, 2011.
- Yang Cai, Constantinos Daskalakis, and Matthew Weinberg. Understanding incentives: Mechanism design becomes algorithm design. In *FOCS*, 2013.
- S. Chawla, D. Malec, and B. Sivan. The power of randomness in bayesian optimal mechanism design. In *EC*, pages 149–158, 2010a.
- Shuchi Chawla, Jason Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In Jeffrey K. MacKie-Mason, David C. Parkes, and Paul Resnick, editors, *ACM Conference on Electronic Commerce*, pages 243–251. ACM, 2007. ISBN 978-1-59593-653-0.

- Shuchi Chawla, Jason Hartline, David Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *42nd ACM Symposium on Theory of Computing*, pages 311–320, 2010b.
- Y.K. Che and I. Gale. The optimal mechanism for selling to budget-constrained consumers. Technical report, 1995.
- Xi Chen and Xiaotie Deng. Settling the complexity of two-player nash equilibrium. In *IEEE 47th Annual Symposium on Foundations of Computer Science*, pages 261–272, 2006.
- E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 269–286. ACM, 2013.
- Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *SODA*, pages 1302–1318, 2014.
- Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *J. ACM*, 55(5), 2008.
- Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138, 2010.
- Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 243–252. ACM, 2009.
- Edith Edith. Designing and learning optimal finite support auctions. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '07, pages 736–745, 2007.

- Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Electronic Commerce*, 2014.
- Allan Gibbard. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society*, pages 587–601, 1973.
- T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- G. Hardy, J. Littlewood, and G. Pólya. Some simple inequalities satisfied by convex functions. *Messenger of Math*, 58:145–152, 1929.
- Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. *arXiv preprint arXiv:1304.6116*, 2013.
- Jason Hartline and Brendan Lucier. Bayesian algorithmic mechanism design. In *STOC*, pages 301–310, 2010.
- Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234, 2009.
- Jason D. Hartline, Robert Kleinberg, and Azarakhsh Malekian. Bayesian incentive compatibility and matchings. In *SODA*, 2011.
- Kamal Jain. A polynomial time algorithm for computing the arrow-debreu market equilibrium for linear utilities. In *IEEE 45th Annual Symposium on Foundations of Computer Science*, pages 286–294, 2004.
- Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186, 1996.
- Alejandro M. Manelli and Daniel R. Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, March 2006. URL <http://ideas.repec.org/a/eee/jetheo/v127y2006i1p1-35.html>.
- Alejandro M. Manelli and Daniel R. Vincent. Bayesian and dominant-strategy implementation in the independent private-values model. *Econometrica*, 78(6):1905–1938, November 2010.

- Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica*, 52(6): 1473–1518, 1984.
- Eric S. Maskin. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44(4-6):667–681, 2000.
- Steven Matthews. Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30(2):370–400, August 1983.
- Steven A. Matthews. On the implementability of reduced form auctions. *Econometrica*, 52(6): pp. 1519–1522, 1984.
- R. Preston McAfee and John McMillan. Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory*, 46(2):335–354, 1988.
- Konrad Mierendorff. Asymmetric reduced form auctions. *Economics Letters*, 110(1):41–44, January 2011.
- Leonard J. Mirman and David Sibley. Optimal nonlinear prices for multiproduct monopolies. *Bell Journal of Economics*, 11(2):659–670, 1980.
- Roger Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73, 1981.
- Roger B Myerson. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pages 323–358, 1986.
- John Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
- Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 59–60. ACM, 2011.
- Mallesh Pai and Rakesh Vohra. Optimal auctions with financially constrained bidders. Discussion papers, Northwestern University, August 2008.

- Gregory Pavlov. A property of solutions to linear monopoly problems. *The BE Journal of Theoretical Economics*, 11(1), 2011.
- Marek Pycia. Stochastic vs deterministic mechanisms in multidimensional screening, 2006.
- John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289, 1983.
- Kevin W S Roberts. Welfare considerations of nonlinear pricing. *Economic Journal*, 89(353): 66–83, March 1979.
- J.C. Rochet. The taxation principle and multi-time hamilton-jacobi equations. *Journal of Mathematical Economics*, 14(2):113 – 128, 1985. ISSN 0304-4068.
- Jean-Charles Rochet and Philippe Chone. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, July 1998.
- Alexander Schrijver. *Combinatorial Optimization : Polyhedra and Efficiency (Algorithms and Combinatorics)*. Springer, July 2003. ISBN 3540204563. URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/3540204563>.
- A Michael Spence. Multi-product quantity-dependent prices and profitability constraints. *Review of Economic Studies*, 47(5):821–41, 1980.
- Nancy L Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, pages 355–371, 1979.
- John Thanassoulis. Hagglng over substitutes. *J. Economic Theory*, 117(2):217–245, 2004.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16: 8–37, 1961.
- Vincent and Manelli. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.
- Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *ACM Conference on Electronic Commerce*, 2014.

Robert Wilson. Nonlinear pricing. *Journal of Political Economy*, 102(6):1288–1291, 1994.

Robert Wilson. *Nonlinear Pricing*. Oxford University Press, 1997.

Qiqi Yan. Mechanism design via correlation gap. In *Symposium on Discrete Algorithms*, 2011.

APPENDIX A

The Single Agent Problem and Incentives

A.1. Extension for $m > 2$ Items

In this section, we extend the definition of α given in Section 3.4 from two items to $m > 2$. The approach to derive such general α is similar to that of Section 3.4, and therefore in this section we only define α and verify its properties. Assume that we are in a partition of the space in which $t_1 \geq t_i$ for all i . Define $T(t_1, \theta_2, \dots, \theta_m) = \{\mathbf{t}' | t'_1 \geq t_1, t'_i \leq \theta_i t_1, \forall i \geq 2\}$. Define $C(t_1, \theta_2, \dots, \theta_m) = \{\mathbf{t}' | t'_1 = t_1, t'_i \leq t'_i \theta_i, \forall i \geq 2\}$. Now define

$$\alpha_1(\mathbf{t}) = t_1 - \frac{\int_{\mathbf{t}' \in T(t_1, 1, \dots, 1)} f(\mathbf{t}') d\mathbf{t}'}{\int_{\mathbf{t}' \in C(t_1, 1, \dots, 1)} f(\mathbf{t}') d\mathbf{t}'}$$

We also define α_i for $i \geq 2$ as follows

$$\alpha_i(\mathbf{t})f(\mathbf{t}) = t_i f(\mathbf{t}) + \frac{1}{m-1} \int_{y=0}^{t_i} f(y, \mathbf{t}_{-i}) + \frac{d}{dt_1} [(t_1 - \alpha_1(y, \mathbf{t}_{-i}))f(y, \mathbf{t}_{-i})] dy$$

The above definition implies that

$$\frac{d}{dt_i} (t_i f(\mathbf{t}) - \alpha_i(\mathbf{t})f(\mathbf{t})) = -\frac{1}{m-1} (f(\mathbf{t}) + \frac{d}{dt_1} [(t_1 - \alpha_1(\mathbf{t}))f(\mathbf{t})]).$$

As a result,

$$\begin{aligned}
\nabla \cdot \phi &= \sum_i \frac{d}{dt_i} (t_i f(\mathbf{t}) - \alpha_i(\mathbf{t}) f(\mathbf{t})) \\
&= \frac{d}{dt_1} [(t_1 - \alpha_1(\mathbf{t})) f(\mathbf{t})] - (m-1) \times \frac{1}{m-1} (f(\mathbf{t}) + \frac{d}{dt_1} [(t_1 - \alpha_1(\mathbf{t})) f(\mathbf{t})]) \\
&= -f
\end{aligned}$$

We now verify that ϕ satisfies the boundary conditions. This holds because when $t_i = 0$, $\alpha_i(\mathbf{t}) = 0$, and also when $t_1 = 1$, $\alpha_1(\mathbf{t}) = f(\mathbf{t})$.

Finally, we verify that $\alpha_i(\mathbf{t}) \leq \frac{t_i}{t_1} \alpha_1(\mathbf{t})$. This is again done in a manner similar to Section 3.4. Fix values of $\theta_3, \dots, \theta_m$, let $T_{\theta_3, \dots, \theta_n}(t_1, \theta_2)$ be the projection of type space into set of types such that each type \mathbf{t}' satisfies $t'_1 \geq t_1$, $t'_1 \leq t'_2 \theta$, and $t'_i = t_1 \theta_i$. Now we can invoke the divergence theorem to conclude that $\alpha_i(\mathbf{t}) \leq \frac{t_i}{t_1} \alpha_1(\mathbf{t})$ if

$$\frac{d}{dt_1} \left(\frac{\int_{y=0}^{t_1 \theta_2} f(t_1, y, \theta_3, \dots, \theta_n) dy}{\int_{y=0}^{t_1} f(t_1, y, \theta_3, \dots, \theta_n) dy} \right) \leq 0.$$

Notice that fixing $\theta_3, \dots, \theta_n$, the above property is exactly what was required in two dimensions, which we showed follows from MR-log-supermodularity. As a result, if the function is MR-log-supermodular in every pair of variables t_1, θ_i for $i \geq 2$, then we $\alpha_i(\mathbf{t}) \leq \frac{t_i}{t_1} \alpha_1(\mathbf{t})$. Notice that this property is implied by MR-log-supermodularity of the distribution in all its variables, and therefore is a less demanding condition. We have therefore proved the following lemma.

Lemma 37. *If the distribution is MR-log-supermodular in every pair of variables t_1, θ_i for $i \geq 2$, then the revenue of an assignment can be upper bounded by*

$$\int_{\mathbf{t}} \boldsymbol{\pi} \cdot (1, \frac{t_2}{t_1}, \dots, \frac{t_n}{t_1}) \alpha(t_1) f(\mathbf{t}) d\mathbf{t},$$

where

$$\alpha_1(\mathbf{t}) = t_1 - \frac{\int_{\mathbf{t}' \in \mathcal{T}(t_1, 1, \dots, 1)} f(\mathbf{t}') \, d\mathbf{t}'}{\int_{\mathbf{t}' \in \mathcal{C}(t_1, 1, \dots, 1)} f(\mathbf{t}') \, d\mathbf{t}'}$$

A.2. Missing Proofs

A.2.1. Proof from Section 3.4

PROOF OF LEMMA 17. We need to prove

$$f^{MR}(t_1, \theta_1) \times f^{MR}(t'_1, \theta'_1) \leq f^{MR}(t_1, \theta'_1) \times f^{MR}(t'_1, \theta_1), \quad \forall t_1 \leq t'_1, \theta_1 \geq \theta'_1.$$

Recall that $f^{MR}(t_1, \theta_1) = f(t_1, t_1\theta_1)$. Since the distribution is a product one, this implies that $f^{MR}(t_1, \theta_1) = f_1(t_1)f_2(t_1\theta_1)$. Notice that pair of values $t\theta'$ and $t'\theta$ have the same geometric mean as the pair $t\theta, t'\theta'$. Also given the assumptions, $t\theta' \leq t\theta, t'\theta' \leq t'\theta$. GG-convexity implies that

$$f_2(t\theta) \times f_2(t'\theta') \leq f_2(t\theta') \times f_2(t'\theta).$$

Multiplying both sides by $f_1(t_1) \times f_1(t'_1)$ we get

$$f_1(t_1)f_2(t\theta) \times f_1(t'_1)f_2(t'\theta') \leq f_1(t_1)f_2(t\theta') \times f_1(t'_1)f_2(t'\theta),$$

which since the distribution is a product distribution implies that

$$f^{MR}(t_1, \theta_1) \times f^{MR}(t'_1, \theta'_1) \leq f^{MR}(t_1, \theta'_1) \times f^{MR}(t'_1, \theta_1).$$

□

A.2.2. Proofs from Section 3.5

PROOF OF LEMMA 18. Let π^p be the assignment corresponding to posting price p for the bundle, that is $\pi_1^p(\mathbf{t}) = \pi_2^p(\mathbf{t}) = 1$ if $t_1 + t_2 \geq p$, and $\pi_1^p(\mathbf{t}) = \pi_2^p(\mathbf{t}) = 0$ otherwise. We will show that the virtual surplus of π^p is equal to the revenue of posting price p , $R(p) = p(1 - F_{sum}(p))$. The virtual surplus is

$$\begin{aligned} \int_{\mathbf{t} \in T} (\pi^p \cdot \phi f)(\mathbf{t}) \, d\mathbf{t} &= \int_{\mathbf{t} \in T} \pi^p(t_1, t_2) \cdot \hat{\phi}(t_1, t_2) f(t_1, t_2) \, d\mathbf{t} \\ &= \int_{\mathbf{t} \in T, t_1 + t_2 \geq p} \phi_{sum}(t_1 + t_2) f(t_1, t_2) \, d\mathbf{t}. \end{aligned}$$

For a function h on T , define h^{SR} to be its transformation to sum-ratio coordinates, that is

$$h(t_1, t_2) = h^{SR}(t_1 + t_2, \frac{t_2}{t_1})$$

By performing a change of variables $s = t_1 + t_2$, and $\theta = t_2/t_1$, the virtual surplus of π^p can be written as (the Jacobian of the transformation is $\frac{s}{(1+\theta)^2}$)

$$\begin{aligned} \int_{s=p}^1 \int_{\theta=0}^1 \phi_{sum}(s) f^{SR}(s, \theta) \frac{s}{(1+\theta)^2} \, d\theta ds &= \int_{s=p}^1 \phi_{sum}(s) \int_{\theta=0}^1 f^{SR}(s, \theta) \frac{s}{(1+\theta)^2} \, d\theta ds \\ &= \int_{s=p}^1 \phi_{sum}(s) f_{sum}(s) \, ds \end{aligned}$$

Replacing ϕ_{sum} by its definition,

$$\begin{aligned}
&= \int_{s \geq p} t_1 f_{sum}(s) - (1 - F_{sum}(s)) \, ds \\
&= - \int_{s \geq p} \frac{d}{ds} (s(1 - F_{sum}(s))) \, ds \\
&= R(p) - R(1) = R(p).
\end{aligned}$$

□

PROOF OF LEMMA 22. We assume that ϕ satisfying the requirements of the lemma exists, derive the closed form suggested in the lemma, and then verify that the derived ϕ indeed satisfies all the required properties. We apply the divergence theorem to α on the trapezoidal subspace of type space defined by types \mathbf{t}' with $s \leq t'_1 + t'_2 \leq 1$, $t'_2/t'_1 \leq \theta$, and $0 \leq t'_1, t'_2 \leq 1$ (Figure 3.3). The divergence theorem equates the the integral of the vector field α on the boundary of the subspace to the integral of its divergence within the subspace. As the upper boundary of this trapezoidal subspace has slope t_2/t_1 , one term in this equality is the integral of $\alpha(\mathbf{t}')$ with the upward orthogonal vector to \mathbf{t} . Differentiating this integral gives the desired quantity.

Applying the divergence theorem to α on the trapezoid and expressing the top boundary as the interior divergence minus the other three boundaries gives:

$$\begin{aligned}
\int_{\mathbf{t}' \in \text{TOP}(s, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' &= \int_{\mathbf{t}' \in \text{INTERIOR}(s, \theta)} \nabla \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' \\
&\quad - \int_{\mathbf{t}' \in \{\text{RIGHT, BOTTOM, LEFT}\}(s, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}'.
\end{aligned}$$

Since α/f is a strong amortization of utility, the divergence density equality and boundary orthogonality imply that the integral over the interior simplifies and the integrals over the right and bottom boundary are zero, respectively. We have,

$$\int_{\mathbf{t}' \in \text{TOP}(s, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}' = - \int_{\mathbf{t}' \in \text{INTERIOR}(s, \theta)} f(\mathbf{t}') \, d\mathbf{t}' - \int_{\mathbf{t}' \in \text{LEFT}(s, \theta)} \boldsymbol{\eta}(\mathbf{t}') \cdot \boldsymbol{\alpha}(\mathbf{t}') \, d\mathbf{t}'.$$

For the trapezoid parameterized by (s, θ) these integrals are (recall that the Jacobian of the transformation from \mathbf{t} to (s, θ) is $\frac{s}{(1+\theta)^2}$),

$$\begin{aligned} \int_{s'=s}^1 \frac{\boldsymbol{\alpha}^{SR}(s', \theta) \cdot (-\theta, 1)}{1 + \theta} \, ds' &= - \int_{s'=s}^1 \int_{\theta'=0}^{\theta} \frac{f^{SR}(s', \theta') \cdot s}{(1 + \theta')^2} \, d\theta' \, ds' \\ &+ \int_{\theta'=0}^{\theta} \frac{\boldsymbol{\alpha}^{SR}(s, \theta') \cdot (-1, -1)s}{(1 + \theta')^2} \, d\theta'. \end{aligned}$$

Differentiating with respect to s gives,

$$\frac{\boldsymbol{\alpha}^{SR}(s, \theta) \cdot (-\theta, 1)}{1 + \theta} = \int_{\theta'=0}^{\theta} \frac{f^{SR}(s, \theta') \cdot s}{(1 + \theta')^2} \, d\theta' - \frac{d}{ds} \int_{\theta'=0}^{\theta} \frac{\boldsymbol{\alpha}^{SR}(s, \theta') \cdot (-1, -1)s}{(1 + \theta')^2} \, d\theta'.$$

On the right-hand side, multiply first term by $\frac{f_{sum}(s)}{f_{sum}(s)} = 1$. The assumption that $\phi_1 + \phi_2 = \phi_{sum}(t_1 + t_2)$ implies that $\alpha_1^{SR}(s, \theta) + \alpha_2^{SR}(s, \theta) = \frac{1 - F_{sum}(s)}{f_{sum}(s)} f^{SR}(s, \theta)$ for all s and θ . These two terms then simplify by the product rule for differentiation to give the identity of the lemma.

$$\begin{aligned} \frac{\boldsymbol{\alpha}^{SR}(s, \theta) \cdot (-\theta, 1)}{1 + \theta} &= f_{sum}(s) \frac{\int_{\theta'=0}^{\theta} \frac{f^{SR}(s, \theta') \cdot s}{(1 + \theta')^2} \, d\theta'}{f_{sum}(s)} \\ &+ \frac{d}{ds} \left[(1 - F_{max}(s)) \frac{\int_{\theta'=0}^{\theta} \frac{f^{SR}(s, \theta') \cdot s}{(1 + \theta')^2} \, d\theta'}{f_{max}(t_1)} \right]. \end{aligned}$$

By definition and change of variables,

$$F(s, \theta) = \int_{\mathbf{t}: t_1+t_2 \leq s, t_2/t_1 \leq \theta} f(\mathbf{t}) \, d\mathbf{t} = \int_{s' \leq s} \int_{\theta' \leq \theta} f^{SR}(s', \theta') \frac{s'}{(1+\theta')^2} \, d\theta' ds'.$$

Therefore, $f(s, \theta) = \int_{\theta' \leq \theta} f^{SR}(s, \theta') \frac{s}{(1+\theta')^2} \, d\theta'$. Plugging this definition into the above equation, we get

$$\begin{aligned} \frac{\alpha^{SR}(s, \theta) \cdot (-\theta, 1)}{1 + \theta} &= f_{sum}(s) \frac{f_{sum}(s, \theta)}{f_{sum}(s)} + \frac{d}{ds} \left[(1 - F_{sum}(s)) \frac{f_{sum}(s, \theta)}{f_{sum}(s)} \right]. \\ &= (1 - F_{sum}(s)) \frac{d}{ds} \left[\frac{f_{sum}(s, \theta)}{f_{sum}(s)} \right]. \end{aligned}$$

As a result,

$$\alpha^{SR}(s, \theta) \cdot (-\theta, 1) = (1 + \theta)(1 - F_{sum}(s)) \frac{d}{ds} \left[\frac{f_{sum}(s, \theta)}{f_{sum}(s)} \right].$$

We can now use the above equation, together with $\alpha_1^{SR}(s, \theta) + \alpha_2^{SR}(s, \theta) = \frac{1 - F_{sum}(s)}{f_{sum}(s)} f^{SR}(s, \theta)$ to solve for α_1 . □

PROOF OF LEMMA 23. We prove that for any θ, s , and s' such that $s < s'$,

$$\frac{f_{sum}(s, \theta)}{f_{sum}(s, 1)} \leq \frac{f_{sum}(s', \theta)}{f_{sum}(s', 1)}$$

The proof first converts the above form into max-ratio coordinates, then applies SR-log-submodularity, and then transforms it back to the standard form. Before applying SR-log-submodularity, we break down the integral set into two set, and apply SR-log-submodularity to only one of the integrals. More particularly, notice that

$$\begin{aligned}
& f_{sum}(s, \theta) \times f_{sum}(s', 1) \\
&= \int_{\theta'=0}^{\theta} f^{SR}(s, \theta') \frac{s}{(1+\theta')^2} d\theta' \times \int_{\theta''=0}^1 f^{SR}(s', \theta'') \frac{s'}{(1+\theta'')^2} d\theta'' \quad (\text{change of variables}) \\
&= \int_{\theta'=0}^{\theta} \int_{\theta''=0}^{\theta} f^{SR}(s, \theta') \frac{s}{(1+\theta')^2} f^{SR}(s', \theta'') \frac{s'}{(1+\theta'')^2} d\theta'' d\theta' \\
&\quad + \int_{\theta'=0}^{\theta} \int_{\theta''=\theta}^1 f^{SR}(s, \theta') \frac{s}{(1+\theta')^2} f^{SR}(s', \theta'') \frac{s'}{(1+\theta'')^2} d\theta'' d\theta' \quad (\text{separate double integral}) \\
&\geq \int_{\theta'=0}^{\theta} \int_{\theta''=0}^{\theta} f^{SR}(s, \theta'') \frac{s}{(1+\theta'')^2} f^{SR}(s', \theta') \frac{s'}{(1+\theta')^2} d\theta' d\theta'' \quad (\text{rename } \theta' \text{ and } \theta'') \\
&\quad + \int_{\theta'=0}^{\theta} \int_{\theta''=\theta}^1 f^{SR}(s, \theta'') \frac{s}{(1+\theta'')^2} f^{SR}(s', \theta') \frac{s'}{(1+\theta')^2} d\theta'' d\theta' \quad (\text{MR-log-supermodularity}) \\
&= \int_{\theta''=0}^1 f^{SR}(s, \theta'') \frac{s}{(1+\theta'')^2} d\theta'' \int_{\theta'=0}^{\theta} f^{SR}(s', \theta') \frac{s'}{(1+\theta')^2} d\theta' \quad (\text{merge integrals}) \\
&= f_{sum}(s, 1) \times f_{sum}(s', \theta).
\end{aligned}$$

□

A.2.3. Proofs from Subsection 3.6.3

PROOF OF LEMMA 25. Since $\phi_2(\mathbf{t}) \leq \frac{t_2}{t_1} \phi_1(\mathbf{t})$, the revenue of an assignment function can be upper bounded as follows,

$$\int \int \boldsymbol{\pi}(\mathbf{t}) \cdot (\phi_1(\mathbf{t}), \phi_2(\mathbf{t})) f(\mathbf{t}) d(\mathbf{t}) \leq \int \int \boldsymbol{\pi}(\mathbf{t}) \cdot (1, \frac{t_2}{t_1}) \phi_1(\mathbf{t}) f(\mathbf{t}) d\mathbf{t}$$

After a changing variables according to $\theta = \frac{t_2}{t_1}$, the upper bound becomes

$$\int_{t_1} \int_{\theta=0}^1 \mathbf{x}^{MR}(t_1, \theta) \cdot (1, \theta) \phi^{MR}(t_1, \theta) f^{MR}(t_1, \theta) t_1 d\theta dt_1,$$

where $\phi^{MR}(t_1, \theta) = \phi(t_1, t_1\theta)$. For a function h on T , define h^{QR} to be its transformation to quantile-ratio coordinates, that is

$$h(t_1, t_2) = h^{QR}\left(\int_{T=T(t_1,1)} f(\mathbf{t}) \, d\mathbf{t}, \frac{t_2}{t_1}\right)$$

We perform another change of variables according to $q = \int_{T=T(t_1,1)} f(\mathbf{t}) \, d\mathbf{t}$, and rewrite the upper bound as

$$\int_q \int_\theta \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) \phi^{QR}(q) \, d\theta \, dq.$$

Using integration by parts,

$$\int_q \int_\theta \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) \phi^{QR}(q) \, d\theta \, dq = \int_q \int_\theta \frac{d}{dq} \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) R(q) \, d\theta \, dq$$

where $R(q)$ is the revenue of posting a price that is accepted with probability q . Now let \bar{R} be the convex hull of function R . In particular, $\bar{R}'(q)$ is non-decreasing and $\bar{R}(q) \geq R(q)$. We next show that $\frac{d}{dq} \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) \geq 0$ for any incentive compatible assignment rule.

Recall that the incentive compatibility condition states that, for all \mathbf{t}, \mathbf{t}' ,

$$(\boldsymbol{\pi}(\mathbf{t}') - \boldsymbol{\pi}(\mathbf{t})) \cdot (\mathbf{t}' - \mathbf{t}) \geq 0.$$

Setting $\mathbf{t}' = \mathbf{t}(1 + \epsilon)$, this becomes

$$(\boldsymbol{\pi}(\mathbf{t}(1 + \epsilon)) - \boldsymbol{\pi}(\mathbf{t})) \cdot (\mathbf{t}\epsilon) \geq 0.$$

When ϵ goes to zero, this give us the following property

$$\frac{d}{dt_1} \boldsymbol{\pi}(\mathbf{t}) \cdot \mathbf{t} \geq 0.$$

Now notice that $q = \int_{T=T(t_1)} f(\mathbf{t}) \, d\mathbf{t}$ is a monotone function of t_1 , which implies the desired property

$$\frac{d}{dq} \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) \geq 0.$$

The facts that $\frac{d}{dq} \mathbf{x}^{QR}(q, \theta) \cdot (1, \theta) \geq 0$ and $\bar{R}(q) \geq R(q)$ imply that we can upper bound revenue as follows

$$\int_q \int_\theta \frac{d}{dq} \mathbf{x}^{QR}(q, \theta) (1, \theta) R(q) \, d\theta \, dq \leq \int_q \int_\theta \frac{d}{dq} \mathbf{x}^{QR}(q, \theta) (1, \theta) \bar{R}(q) \, d\theta \, dq$$

We now study the assignment rule that optimizes virtual value function $(1, \theta)R'(q)$. Notice that when $\bar{R}(q) > R(q)$, we have that $\bar{R}'(q)$ is a constant and therefore $\frac{d}{dq} \mathbf{x}^{QR} = 0$. This implies that $\hat{R}'(q)$ can be used as an upper bound on revenue, and for the expected virtual value of the optimal assignment exactly equals revenue. \square

APPENDIX B

Single-Agent to Multi-Agent Solutions: The Structure**B.1. Proofs from Subsection 4.3.2**

We now give a procedure for implementing the marginal revenue mechanism (Definition 14) with general agents. Recall that in the marginal revenue mechanism, each agent faces a distribution over ex ante optimal pricings, where the distribution is given by marginal revenue maximization over single-dimensional analog agents having the same revenue curves. This maximization over single-dimensional analogs gives rise to an allocation constraint \hat{x}^{MR} , and then the \hat{q} ex ante optimal pricing occurs with probability $-\frac{d}{d\hat{q}}\hat{x}^{MR}(\hat{q})$. We show that this mixture of ex ante optimal pricings is implementable within the marginal revenue mechanisms family (Definition 11), and the randomized mapping from type to quantile (Step 1) in this implementation is efficiently computable.

What properties are needed for such a mapping? First, for each agent, we need the quantile to be uniformly distributed over $[0, 1]$. This way, the distribution over marginal revenues faced by each agent is as if the competing agents are single-dimensional linear agents with the same revenue curves. This guarantees that, if we map an agent's type to a quantile q , the probability that she revenue wins a service in the marginal revenue mechanism is equal to $\hat{x}^{MR}(q)$. In other words, this property would designate an allocation probability $\hat{x}^{MR}(q)$ to each quantile q , and therefore in order to get the desired allocation rules for types, we need only to come up with appropriate mappings of types to quantiles. Secondly, we would like the allocation rules obtained by this procedure to match the allocation rule given by the previously described mixture over ex ante optimal pricings. To be specific, recall that each ex ante optimal pricing is derived from

optimizing revenue subject to a step function constraint. The resulting normalized allocation rule may not be a step function and is in general weaker. When these ex ante optimal pricings are composed into the mixture, the resulting allocation rule, which we denote by x^{MR} , is dominated by and not necessarily equal to \hat{x}^{MR} . It is the allocation rule x^{MR} , and not \hat{x}^{MR} , that we would like to produce.

Recall from the discussion of Definition 14 that our goal is to implement the outcome rule w^{MR} . If we order the types according to $\text{Alloc}(w^{MR}(\cdot))$, we get a natural mapping from types to quantiles: $\text{Quant}(t) \triangleq \Pr_{t' \sim \mathcal{D}}[\text{Alloc}(w^{MR}(t')) \geq \text{Alloc}(w^{MR}(t))]$. This mapping will have the first property¹, i.e., $\text{Quant}(t)$ will be uniform on $[0, 1]$, but it does not have the second property. This is because by definition the probability of type t winning in the marginal revenue mechanism with this mapping is $\hat{x}^{MR}(\text{Quant}(t))$. Overall, we get the allocation rule \hat{x}^{MR} and not the weaker x^{MR} . If we could keep the first property, then the problem reduces to the following: given two non-increasing functions $x^{MR}, \hat{x}^{MR} : [0, 1] \rightarrow [0, 1]$, such that x^{MR} is weaker than \hat{x}^{MR} (in the sense that $\hat{X}^{MR} \geq X^{MR}$ pointwise), is there a randomized function $g : [0, 1] \rightarrow [0, 1]$, such that $\mathbf{E}[\hat{x}^{MR}(g(q))] = x^{MR}(q)$ for every $q \in [0, 1]$, and $g(q)$ is uniform on $[0, 1]$ when q is uniform on $[0, 1]$? This is a problem addressed by the theory of *majorization* (see, e.g. Hardy et al., 1929), and has a general solution. In our context, we give a particularly simple interval resampling procedure that gives this mapping g , which is to be composed with $\text{Quant}(\cdot)$ for the eventual randomized mapping from types to quantiles.

Definition 38. For allocation constraint \hat{x} and dominated allocation rule x satisfying $\hat{X}(1) = X(1)$ on m discrete types, the *interval resampling sequence construction* starts with $x^{(0)} = \hat{x}$ and calculates $x^{(j+1)}$ from $x^{(j)}$ while $x^{(j)} \neq x$ as follows.

- (1) Find the highest quantile q where $x(q) \neq x^{(j)}(q)$.

¹As before, we break ties appropriately.

- (2) Let $q' > q$ be the quantile at which the line tangent to X at q with slope $x(q)$ crosses $X^{(j)}$.²
- (3) The j th resampling interval is $[q, q']$.
- (4) Let $x^{(j+1)}$ be $x^{(j)}$ averaged on $[q, q']$.

Proposition 17. *The interval sampling sequence construction gives a sequence of at most m intervals such that the composition of \hat{x} with the sequence of resamplings applied to $\text{Quant}(\cdot)$ is equal to x .*

Proof. The proof is by induction on j where the j th step assumes the first $j - 1$ types, in order of $\text{Quant}(\cdot)$, satisfy $x^{(j-1)}(\text{Quant}(t)) = x(\text{Quant}(t))$. Consider step j . The assumption that $\hat{X}(1) = X(1)$ ensures that the intersection of the tangent happens at a $q' \leq 1$. The line segment connecting interval $[q, q']$ of $X^{(j)}$ has slope equal to $x(q)$, by definition. Therefore, the j th step in the construction leaves $x^{(j)}(\text{Quant}(t)) = x(\text{Quant}(t))$ for the j th type. The procedure is linear time as both \hat{x} and x are, without loss of generality, piece-wise constant with m pieces, and in each step q and q' are increasing and at least one piece from \hat{x} or x is processed. \square

The final ingredient in the construction of the marginal revenue mechanism for agents with general types is in converting the allocation rule back into an outcome rule. This can be done exactly as in Alaei et al. (2012): if an agent with type t is served by the allocation rule, sample from service outcomes of $w^{MR}(t)$, otherwise sample from non-service outcomes of $w^{MR}(t)$.

Definition 39. The *marginal revenue mechanism* for general agents works as follows.

- (1) Map reported types $\mathbf{t} = (t_1, \dots, t_n)$ of agents to quantiles $\mathbf{q} = (q_1, \dots, q_n)$ by, for each agent, composing the interval resampling transformation with $\text{Quant}(\cdot)$.

²For discrete type, this intersection may happen at a quantile q' that does not correspond to the boundary between two types. When this happens split the type into two types each occurring with the same total probability and with the boundary between them at q' .

- (2) Calculate the marginal revenue of each agent i as $R'_i(q_i)$.
- (3) Calculate the set of agents to be served by marginal revenue maximization.
- (4) Calculate outcomes for each agent i as:
 - sample $w_i \sim w_i^{MR}(t_i)$ conditioned on $\text{Alloc}(w_i) = 1$ if i is to be served, or
 - sample $w_i \sim w_i^{MR}(t_i)$ conditioned on $\text{Alloc}(w_i) = 0$ if i is not to be served.

Note that instead of calculating outcome rules by mixing over step mechanisms we could, from the allocation constraint \hat{x}^{MR} for an agent, calculate the optimal mechanism subject to that constraint, i.e., with outcome rule $\text{Outcome}(\hat{x}^{MR})$ and revenue $\text{Rev}[\hat{x}^{MR}]$. The construction above can be invoked with this outcome rule in place of w^{MR} without modification; this change generally improves revenue.

B.2. Proofs from Subsection 4.3.3

The technique for the proof of Proposition 5 largely comes from Laffont and Robert (1996) and can be viewed as a consequence of that work. We remark that the condition of concavity of F (or, equivalently, the monotonicity of f), which was not used in the original paper of Laffont and Robert (1996), was in fact needed for their characterization, as correctly pointed out by Pai and Vohra (2008).

PROOF OF PROPOSITION 5. For this proof we will only use allocations for types (instead of quantiles), and to simplify notation we let $x(v)$ be the allocation probability for type v . Without loss of generality, we assume that the highest valuation in the support of F is 1. The standard incentive compatibility condition for single-dimensional linear preferences (monotonicity of the allocation rule and the payment identity) still holds. In particular, for $v > v'$, if $x(v) > x(v')$, then the payment of v is also strictly larger than that of v' . Therefore, if the budget constraint is binding (as we assumed), then there is a \bar{v} such that the allocation probability is a constant

for all types above \bar{v} , and the payment for all these types is B . The proposition then states that, in \hat{q} ex ante optimal pricing, the allocation for types smaller than \bar{v} is constantly 0.

Payment identity states that the payment of type v is $vx(v) - \int_0^v x(z) dz$. We therefore would like to maximize the objective function

$$(B.1) \quad \max \int_0^{\bar{v}} f(v)x(v)\phi(v) dv + [1 - \mathcal{D}(\bar{v})]\bar{v}x(\bar{v}),$$

where $\phi(v)$ is the standard virtual valuation function $v - \frac{1-\mathcal{D}(v)}{f(v)}$, subject to the constraints:

$$(B.2) \quad \bar{v}x(\bar{v}) - \int_0^{\bar{v}} x(v) dv = B,$$

$$(B.3) \quad \int_0^{\bar{v}} x(v) dv + [1 - \mathcal{D}(\bar{v})]x(\bar{v}) = \hat{q},$$

$$(B.4) \quad \forall v, \quad x(v) \geq 0,$$

$$(B.5) \quad x(\bar{v}) \leq 1.$$

We consider the first-order conditions for the above program. We use δ for the Lagrangian variable for the budget condition (B.2); λ for the ex ante selling probability constraint (B.3); Π_v for condition (B.4) for each v ($\Pi_v \leq 0$); η for condition (B.5) ($\eta \geq 0$). The first order condition gives

$$(B.6) \quad f(v) \left[\phi(v) + \lambda - \frac{\delta}{f(v)} \right] + \Pi_v = 0, \quad \forall v < \bar{v};$$

$$(B.7) \quad [1 - \mathcal{D}(\bar{v})] \left[\bar{v} + \lambda + \frac{\bar{v}\delta}{1 - \mathcal{D}(\bar{v})} \right] + \Pi_{\bar{v}} + \eta = 0.$$

By complementary slackness, for any v such that $x(v) > 0$, we have $\Pi_v = 0$. (In particular, $\Pi_{\bar{v}} = 0$.) We next argue that δ is negative. Assume there is a $v < \bar{v}$ such that $x(v) > 0$. Then

we have

$$\begin{aligned}\phi(v) + \lambda - \frac{\delta}{f(v)} &= 0; \\ \bar{v} + \lambda + \frac{\bar{v}\delta}{1 - \mathcal{D}(\bar{v})} + \frac{\eta}{1 - \mathcal{D}(\bar{v})} &= 0.\end{aligned}$$

We can therefore solve for δ :

$$(B.8) \quad \delta = \left[\phi(v) - \bar{v} - \frac{\eta}{1 - \mathcal{D}(\bar{v})} \right] / \left[\frac{\bar{v}}{1 - \mathcal{D}(\bar{v})} + \frac{1}{f(v)} \right] < 0.$$

Now, if for two different $v, v' < \bar{v}$ such that their allocation probabilities are both strictly positive, then $\Pi_v = \Pi_{v'} = 0$, and we will have

$$\phi(v) - \frac{\delta}{f(v)} = \phi(v') - \frac{\delta}{f(v')},$$

or

$$(B.9) \quad \phi(v) - \phi(v') = \delta \left(\frac{1}{f(v)} - \frac{1}{f(v')} \right).$$

Suppose $v < v'$, then $f(v) \geq f(v')$ by our assumption. Since the distribution is regular, we have $\phi(v) \leq \phi(v')$. Additionally, we know that $\delta < 0$, and so (B.9) can hold only if $f(v) = f(v')$, but then the equation says $f(v)(v - v') + \mathcal{D}(v) - \mathcal{D}(v') = 0$, which cannot be true since $\mathcal{D}(v) < \mathcal{D}(v')$. Therefore (B.9) cannot hold under our assumptions.

So far we have shown that in the optimal solution to the above linear program, there can be at most one value $v < \bar{v}$ such that $x(v) > 0$. But then lowering $x(v)$ to 0 affects neither the objective function nor the constraints, and so we obtain a monotone allocation rule.³ Therefore

³As a standard practice, we have relaxed the monotonicity condition in the formation of the linear program, and only observe that the optimal solution satisfies the monotonicity condition under the assumptions on the valuation distribution.

the solution to the program gives rise to an incentive compatible mechanism, which satisfies Proposition 5. \square

Derivation of Example 9. We first derive the \hat{q} ex ante optimal pricings for $\hat{q} < 1 - F(B) = 1 - B$. By Proposition 5, a lottery that costs B is offered, and, when bought, it sells the item with probability $\pi = B + \hat{q}$. A type with value at least $B/(B + \hat{q})$ will buy the lottery (and hence wins with probability $B + \hat{q}$). For $\hat{q} > 1 - B$, the budget does not bind and the item is sold at a price of $1 - \hat{q}$; all types with $v \geq 1 - \hat{q}$ wins the item with certainty. This immediately shows us the shape of $G_v(\hat{q})$, the probability of allocation as a function of \hat{q} for a fixed type v . From the perspective of a given type $v \geq B$, $G_v(\hat{q})$ jumps starts at $\hat{q} = \frac{B}{v} - B$, increases linearly with \hat{q} and saturates at $\hat{q} = 1 - B$. For $v < B$, the budget never binds, and the corresponding $G_v(\hat{q})$ is the familiar step function.

Calculating the revenue curve is straightforward. For $\hat{q} < 1 - B$, $R(\hat{q})$ is $B \cdot (1 - \frac{B}{B+\hat{q}})$. Its derivative, i.e., the marginal revenue, $\frac{B^2}{(B+\hat{q})^2}$, is strictly positive. (Note that, for $B < 1/2$, this is different from the linear preference case. There, the marginal revenue would be negative for $\hat{q} > 1/2$.) For $\hat{q} \geq 1 - B$, $R(\hat{q})$ is $\hat{q}(1 - \hat{q})$. Its derivative is positive for $\hat{q} < 1/2$ and negative for $\hat{q} > 1/2$.

By Step 1 of Definition 18, whenever $v' < B$, its quantile $1 - v'$ will be larger with probability 1 than the quantile of a type $v \geq B$, which is distributed between $\frac{B}{v} - B$ and $1 - B$. Therefore, such smaller v' 's are only considered when there is no bidder bidding above B ; when this happens, since the budget does not bind, the auction is the optimal one in the linear preference case, i.e., a second price auction with reserve price $\frac{1}{2}$. When there are bidders bidding above B , the way q_i is computed in Example 9 is simply sampling by $G_v(\hat{q})$ as stipulated in Definition 18. When i^* is the sole bidder bidding above B , she should pay the ex ante optimal pricing at her critical quantile. When $B < 1/2$, this critical quantile is $1 - B$, and she pays B ; when $B > 1/2$, the critical quantile is $\min\{1/2, 1 - \max_{i \neq i^*} v_i\}$, and she pays $\max_{i \neq i^*} \{1/2, v_i\}$. When there are

other bidders bidding above B , the critical quantile for i^* is always smaller than $1 - B$, and she will pay B in the corresponding ex ante optimal pricing. A losing bidder i bidding above B faces a critical quantile q_{i^*} . If q_{i^*} is larger than the quantile $\frac{B}{v_i} - B$ at which $G_{v_i}(\hat{q})$ jump starts, she will need to make a payment of B . This happens for $v_i \geq B/(q_{i^*} + B)$.

B.3. Proofs for unit-demand approximation

Theorem 16 is a consequence of the two lemmas below and Proposition 7.

Lemma 38. *Twice the optimal representative revenue is a linear upper bound on the optimal unit-demand revenue.*

Proof. Linearity follows simply from the revenue linearity of single-dimensional linear agents. We consider the collection of representatives as a whole (or, say, a single market), and we can ask what is the optimal revenue from this market given an ex ante selling probability \hat{q} or an allocation constraint \hat{x} . Both terms are easy to find. Consider the distribution of the maximum virtual value (or zero if the maximum virtual value is negative) in the representative environment. Index this distribution by quantile as $\psi_{\max}(\hat{q})$. The optimal revenue for any allocation constraint \hat{x} is $\mathbf{E}_{\hat{q}}[\psi_{\max}(\hat{q})\hat{x}(\hat{q})]$ which is linear in \hat{x} ; this follows from the proof that the optimal revenue in single-dimensional environments is the virtual surplus maximizer.

We now show that, under any allocation constraint, twice the optimal representative revenue upper bounds the optimal unit-demand revenue. To do this we will give two auctions for the representative environment with the allocation constraint \hat{x} and show that the sum of these auctions' revenue upper bounds the optimal unit-demand revenue for the same constraint. Of course, the optimal representative revenue in turn upper bounds each of these auctions' revenue.

A mechanism for the unit-demand problem is simply a lottery pricing, i.e., it is a set of lotteries L with a lottery for each type t taking the form of $(p(t), \pi^1(t), \dots, \pi^m(t))$ with $\sum_j \pi^j(t) \leq 1$.

The semantics of a lottery is that the agent pays the price $p(t)$ and then is allocated an alternative j at random with probability $\pi^j(t)$; the semantics of the collection of lotteries L is that the agent, upon drawing her type t from the distribution, chooses the lottery $(p(t), \pi^1(t), \dots, \pi^m(t))$ that corresponds to her type.

Given any collection of lotteries L that satisfies the allocation constraint \hat{x} we define two auctions for the representative environment that have combined revenue at least that of the collection of lotteries in the unit-demand environment.

The *L mimicking auction* considers the profile of values $\mathbf{v} = (v^1, \dots, v^m)$ of the representatives and the lottery that would have been selected by the unit-demand agent with these values. It serves the representative j with the highest value with probability $\pi^j(t)$ and charges her (no matter whether we serve her or not) $p(t) - \sum_{j' \neq j} \pi^{j'}(t)v^{j'} + \mu(\mathbf{v}^{(2)})$ where $\mu(\mathbf{v}^{(2)})$ is the expected utility of the unit-demand agent with valuation profile $\mathbf{v}^{(2)}$ which is \mathbf{v} with v_j replaced with $\max_{j' \neq j} v^{j'}$. Notice that the utility of the winning representative j in this auction is exactly the same as the unit-demand agent less an amount that is a function only of the values of the other representatives, \mathbf{v}^{-j} . As the utility of the unit-demand agent is monotone in her value for each alternative, the utility each representative has for winning is positive when she is the highest valued representative and negative if she is not (and were to misreport and pretend she were). Therefore, this auction is incentive compatible, has revenue at least $p(t) - \sum_{j' \neq j} \pi^{j'}(t)v^{j'}$ on valuation profile \mathbf{v} where j is the highest valued representative, and satisfies allocation constraint \hat{x} . For a given valuation profile, call the second term in the winning agent's payment, $\sum_{j' \neq j} \pi^{j'}(t)v^{j'}$, the *deficit* of the *L mimicking auction*.

The motivation for the next auction is that we want to obtain back the deficit lost by the *L mimicking auction*. Notice that the procedure that charges the highest valued representative the second highest value and serves with probability $\sum_j \pi^j(t)$ satisfies the allocation constraint \hat{x} and more than balances the deficit; however, it may not be incentive compatible.

The *allocation constrained second-price auction* sells to the highest valued representative at the second highest representative's value so as to maximize revenue subject to the allocation constraint \hat{x} that any representative is served. Consider the distribution of the second order statistic of values and let $\nu_{(2)}(q)$ be the value that the q quantile of this random variable takes on. The optimal revenue obtainable via a second price auction with allocation constraint \hat{x} is $\mathbf{E}_q[\nu_{(2)}(q)\hat{x}(q)]$. To obtain this revenue, conditioning on the second highest value being v , with probability $\hat{x}(\nu_{(2)}^{-1}(v))$ we serve the highest valued representative and charge her v (only when we serve her). This auction is incentive compatible and revenue optimal (in expectation) among all second-price procedures that meet the allocation constraint. Therefore, it more than covers the expected deficit of the L mimicking auction.

We have given two incentive compatible auctions for the representative environment with combined expected revenue exceeding the revenue of the lottery pricing L . Therefore, twice the optimal representative revenue is at least the optimal unit-demand revenue. \square

Lemma 39. *The pseudo revenue curve $\tilde{R}(\cdot)$ from uniform virtual pricings for a unit-demand agent 2-approximates the optimal representative revenue curve (as a function of \hat{q} for any \hat{q} -step constraint).*

Proof. The proof closely follows the standard prophet inequality proofs (see for example Chawla et al. (2010a)). As in the proof to the previous lemma, we may view the representatives as one entity and consider its optimal revenue under ex ante constraint L on serving. Denote the optimal representative revenue for the \hat{q} -step constraint as a function of \hat{q} by the revenue curve $\text{ORR}(\hat{q})$. Consider the outcome of the optimal auction for the representative environment with ex ante service constraint \hat{q} . It sets a uniform virtual price (denoted $\psi(\hat{q})$) and serves the agent with the highest virtual value strictly bigger than $\psi(\hat{q})$ with probability one. If the probability that the largest virtual value is equal to $\psi(\hat{q})$ is strictly positive (which might happen if any

virtual value function is constant on an interval, e.g., from ironing), it probabilistically accepts or rejects the maximum virtual value when it is equal to $\psi(\hat{q})$ so as to serve with the desired ex ante probability \hat{q} . The optimal representative revenue can thus be calculated and bounded as follows. Let (ψ_1, \dots, ψ_m) denote the profile of virtual values of the representatives.

$$\begin{aligned} \text{ORR}(\hat{q}) &= \hat{q} \cdot \psi(\hat{q}) + \mathbf{E}[\max_i(\psi_i - \psi(\hat{q}))_+] \\ &\leq \hat{q} \cdot \psi(\hat{q}) + \sum_i \mathbf{E}[(\psi_i - \psi(\hat{q}))_+]. \end{aligned}$$

Above, the notation $(\psi_i - \psi(\hat{q}))_+$ is short-hand for $\max(0, \psi_i - \psi(\hat{q}))$.

Now we show a lower bound on $\tilde{R}(\hat{q})$ for \hat{q} that does not require probabilistic acceptance in the optimal representative auction described above; denote by $Q \subset [0, 1]$ the set of all such quantiles. Let \mathcal{E}_i denote the event that $\psi_j < \psi(\hat{q})$ for all $j \neq i$; our lower bound on the \hat{q} ex ante pseudo pricing revenue will ignore contributions to the virtual surplus from the case that more than one representative has virtual value at least $\psi(\hat{q})$.

$$\begin{aligned} \tilde{R}(\hat{q}) &\geq \hat{q} \cdot \psi(\hat{q}) + \sum_i \mathbf{E}[(\psi_i - \psi(\hat{q}))_+ \mid \mathcal{E}_i] \cdot \mathbf{Pr}[\mathcal{E}_i] \\ &\geq \hat{q} \cdot \psi(\hat{q}) + (1 - \hat{q}) \cdot \sum_i \mathbf{E}[(\psi_i - \psi(\hat{q}))_+ \mid \mathcal{E}_i] \\ &= \hat{q} \cdot \psi(\hat{q}) + (1 - \hat{q}) \cdot \sum_i \mathbf{E}[(\psi_i - \psi(\hat{q}))_+]. \end{aligned}$$

The second inequality followed because $\mathbf{Pr}[\mathcal{E}_i]$, the probability of the event that $\psi_j < \psi(\hat{q})$ for all $j \neq i$ is not less than the probability that $\psi_j < \psi(\hat{q})$ for all j , which is $(1 - \hat{q})$. To extend this lower bound on $\tilde{R}(\hat{q})$ from $\hat{q} \in Q$ to all $\hat{q} \in [0, 1]$, consider inserting a virtual value $\psi' = \psi(\hat{q}) + \epsilon$ with measure zero in the distribution. The \hat{q}' that corresponds to serving this virtual value or higher has revenue bounded by the formula above but $\psi' \approx \psi(\hat{q})$. Keeping the virtual value constant and varying \hat{q} in the formula interpolates a line between the two revenues. As the ex

ante pseudo pricings are closed under convex combination, this line gives a lower bound on the \hat{q} ex ante pseudo pricing. Therefore, the bound above on $\tilde{R}(\hat{q})$ holds for all \hat{q} .

To bound $\text{ORR}(\hat{q})$ in terms of $\tilde{R}(\hat{q})$ we consider two cases. When $\hat{q} \leq 1/2$ these terms can be directly bounded as the first terms in both bounds are the same and the second terms are within a factor of two of each other (by assumption $1 - \hat{q} \geq 1/2$). To show the claim for $\hat{q} > 1/2$ notice that

$$\begin{aligned} \text{ORR}(1) &= \mathbf{E}[\max_i(\psi_i)_+] \\ &= \hat{v} + \mathbf{E}[\max_i(\psi_i)_+ - \hat{v}] \\ &\leq \hat{v} + \mathbf{E}[(\max_i(\psi_i)_+ - \hat{v})_+] \\ &= \hat{v} + \mathbf{E}[(\max_i \psi_i - \hat{v})_+] \\ &\leq \hat{v} + \sum_i \mathbf{E}[(\psi_i - \hat{v})_+], \end{aligned}$$

for any \hat{v} . As a result, by setting $\hat{v} = \psi(1/2)$ we get

$$\begin{aligned} \text{ORR}(1) &\leq \psi(1/2) + \sum_i \mathbf{E}[(\psi_i - \psi(1/2))_+] \\ &\leq 2\tilde{R}(1/2). \end{aligned}$$

From monotonicity of ORR and \tilde{R} we then conclude that for any $\hat{q} > 1/2$, $\text{ORR}(\hat{q}) \leq \text{ORR}(1) \leq 2\tilde{R}(1/2) \leq 2\tilde{R}(\hat{q})$.

□

B.4. Revenue Linearity for Unit Demand Valuations Uniform on Hypercubes

In this section we show that unit-demand quasi-linear-utility agents whose values for m alternatives are i.i.d. drawn from $U[0, 1]$ are revenue linear. Recall from Section B.3 that an incentive compatible mechanism offers a menu of lotteries to the agent. Each lottery takes the form of $(p(t), \pi^1(t), \dots, \pi^m(t))$, where $\sum_j \pi^j(t) \leq 1$, with p denoting the price of the lottery and π_j the probability with which alternative j is allocated to the agent. We sometimes write π as the vector (π^1, \dots, π^m) . In this section we abuse the notation and use u to denote a mapping that maps a type $t \in T = [0, 1]^m$ to the expected utility of this type in an incentive compatible mechanism. We use the following lemma first noted by Rochet (1985).

Lemma 40. *For a quasi-linear-utility agent, a utility function u corresponds to an incentive compatible mechanism if and only if it is convex. In this case, $p(t) = \nabla u \cdot t - u(t)$, and $\pi(t) = \nabla u(t)$.*

In the above lemma $\nabla u(t)$ is the gradient of the function u . Since selling any alternative accounts as a service, by Lemma 40 the allocation of a type t is $\|\nabla u(t)\|_1$, the L_1 norm of the vector $\nabla u(t)$. Let W be the space of convex utility functions u , and c the cost of producing an alternative. Using Lemma 40, we can reformulate the problem of revenue maximization under allocation constraint \hat{x} as follows:

$$\begin{aligned} & \text{maximize } \int_T [\nabla u(t) \cdot t - u(t)] f(t) dt - c \int_T \vec{1} \cdot \nabla u(t) f(t) dt \\ & \text{s.t. } \quad u \in W \\ & \quad \forall S \subseteq T, \int_S \|\nabla u(t)\|_1 dt \leq \hat{X}(f(S)). \end{aligned}$$

Recall from Section 2.2 the definition of the cumulative allocation constraint \hat{X} . Note also that the second constraint automatically guarantees the feasibility constraint: for all but a measure zero set of types, $\|\nabla u(t)\|_1 \leq 1$. By our assumption, $f(t)$ is 1 everywhere on $[0, 1]^m$.

For any $t \in T$, define a scaling function $r_t : [0, 1] \rightarrow T$ as $r_t(\alpha) = \alpha t$. Then $r_t(0) = \vec{0}$, and $r_t(1) = t$, for any t . We now use the gradient theorem and write

$$\forall t, u(t) - u(0) = \int_0^1 \nabla u(r(\alpha)) \cdot r'_t(\alpha) \, d\alpha.$$

In a revenue optimal mechanism, $u(0) = 0$. Also, by definition of r , $r'(\alpha) = t$. Therefore,

$$u(t) = \int_0^1 \nabla u(\alpha t) \cdot t \, d\alpha, \quad \forall t \in T.$$

Using this, we can rewrite the objective function as

$$\begin{aligned} & \int_T \left[\nabla u(t) \cdot (t - c\vec{1}) - \int_0^1 \nabla u(\alpha t) \cdot t \, d\alpha \right] dt \\ &= \int_T \nabla u(t) \cdot (t - c\vec{1}) \, dt - \int_0^1 \int_T \nabla u(\alpha t) \cdot t \, dt \, d\alpha. \end{aligned}$$

In the second term, change variables by defining $v = \alpha t \in [0, 1]^m$. Notice that $t = v/\alpha$, and $dv^j = \alpha \, dt^j$ for any $1 \leq j \leq m$. Therefore $dv = \alpha^m \, dt$. Define T_α to be the set of $t \in T$ such that $\max_j t^j \leq \alpha$. The objective is now rewritten as

$$\begin{aligned} & \int_T \nabla u(t) \cdot (t - c\vec{1}) \, dt - \frac{1}{\alpha^m} \int_0^1 \int_{v \in T_\alpha} \nabla u(v) \cdot (v/\alpha) \, dv \, d\alpha \\ &= \int_T \nabla u(t) \cdot (t - c\vec{1}) \, dt - \int_{v \in T} \nabla u(v) \cdot v \int_{\alpha = \max_j v^j}^1 \frac{1}{\alpha^{m+1}} \, d\alpha \, dv \\ &= \int_T \nabla u(t) \cdot (t - c\vec{1}) \, dt - \int_{v \in T} \nabla u(v) \cdot v \left[\frac{1}{m(\max_j v^j)^m} - \frac{1}{m} \right] dv \\ &= \int_T \nabla u(t) \cdot \left[t \left(\frac{m+1}{m} - \frac{1}{m(\max_j t^j)^m} \right) - c\vec{1} \right] dt. \end{aligned}$$

Now, if we relax the convexity constraint, the optimization problem is expressed solely in terms of the gradient of u . Next we argue that the optimal solution to this optimization problem takes a particularly simple form. First note that the function $t^j(\frac{m+1}{m} - \frac{1}{mt^j})$ is increasing in t^j . Consider any feasible solution ∇u to the program and its alteration $\nabla \tilde{u}$ in the following manner: at any type t where j^* is $\arg \max_j t^j$, let $\nabla^{j^*} \tilde{u}$ be $\sum_j \nabla^j u$, and $\nabla^j \tilde{u}$ be 0 for all $j \neq j^*$. Since this alteration keeps the L_1 -norm of ∇u , ∇u^* still satisfies all the constraints (except that we are relaxing the convexity constraint for now). But the objective function is pointwise better for $\nabla \tilde{u}$ than for ∇u . Therefore, it suffices to consider solution gradients whose coordinates at each type t are all zero except the one coordinate where the valuation is maximized (ties can be broken arbitrarily). But then the problem degenerates, and the optimal utility function is given by a simple greedy procedure, which grows, at each type, in the direction of the maximum valued alternative as much as allowed by the allocation constraint \hat{x} . Formally, the optimal utility function is given by

$$u^*(t) = \begin{cases} 0, & \max_j t^j \leq \hat{t}^c \\ \int_{\alpha=\hat{t}^c}^{\max_j t^j} \hat{x}(1 - \alpha^m) d\alpha, & \max_j t^j > \hat{t}^c, \end{cases}$$

where \hat{t}^c solves

$$\hat{t}^c \left(\frac{m+1}{m} - \frac{1}{m(\hat{t}^c)^m} \right) = c.$$

In particular, $\hat{t}^0 = \sqrt[m]{\frac{1}{m+1}}$. This utility function u specified above is convex and linear in \hat{x} . By Lemma 40, it is easy to see that u^* being linear implies that $\text{Rev}[\cdot]$ is also linear (noting that integral is a linear functional).

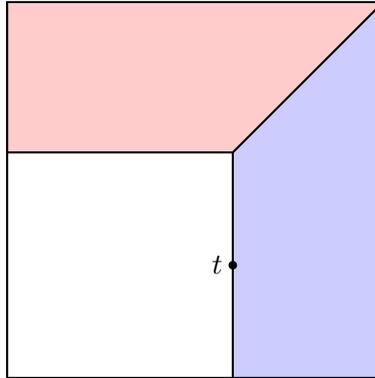


Figure B.1. The quantile mapping for the two dimensional uniform distribution. The quantile of type $t = (t^1, t^2)$ with $t^1 \geq t^2$ is $q = 1 - (t^1)^2$.

To summarize, we have shown that the ex ante optimal mechanism for constraint \hat{q} is to post a price of $\sqrt[m]{1 - \hat{q}}$ for any service. The quantile of each type $t = (t^1, \dots, t^m)$ is $q = 1 - (\max_i t^i)^m$ (see Figure B.1).

B.5. Reverse Auctions

Reverse auctions can naturally be modeled as service constrained environments. Different agents, here sellers, have different costs for providing different services, and the auctioneer has possibly different values for different services, and wishes to acquire at most one service, and to do so in order to maximize the value for the service acquired minus the payment for the service. Notice that the goal of maximizing value minus payment is equivalent to minimizing payment minus value. Such an objective can be modeled as a forward auction in which the seller has possibly different costs for selling items. In Section B.4 we solve the forward auction problem with uniform values and uniform costs, which implies the following results for the reverse auction problem.

More formally, we can transform a reverse auction problem into a forward auction as follows. Consider a single seller that can provide m services where each service i costs c_i . Assume that the cost of each service is drawn uniformly at random from the interval $[0, 1]$, and assume that

the value of each service for the auctioneer is 1 (the analysis generalizes to arbitrary distributions and valuations, but the general analysis is not required here). Let $\pi(c) = (\pi^1(c), \dots, \pi^m(c))$ be the vector of the probabilities of purchasing each service when the cost vector is c , and $p(c)$ the payment made to the seller by the auctioneer. Now the objective is to maximize

$$\int_{c \sim U[0,1]^m} \vec{1} \cdot \pi(c) - p(c) \, dc.$$

We can also write the incentive compatibility constraint as

$$p(c) - c \cdot \pi(c) \geq p(c') - c \cdot \pi(c')$$

for all cost vectors c and c' . Now define functions $\bar{\pi}$ and \bar{p} to be

$$\begin{aligned} \bar{\pi}(c) &= \pi(\vec{1} - c) \\ \bar{p}(c) &= \vec{1} \cdot \pi(\vec{1} - c) - p(\vec{1} - c). \end{aligned}$$

Using the above notation we can rewrite the objective to be

$$(B.10) \quad \int_{c \sim U[0,1]^m} \bar{p}(c) \, dc.$$

Also,

$$p(c) - c \cdot \pi(c) = (\vec{1} - c) \cdot \bar{\pi}(\vec{1} - c) - \bar{p}(\vec{1} - c).$$

Therefore, the incentive compatibility constraint is equal to

$$(B.11) \quad c \cdot \bar{\pi}(c) - \bar{p}(c) \geq c \cdot \bar{\pi}(c') - \bar{p}(c'), \quad \forall c, c'.$$

Now notice that the optimization problem given by (B.10) and (B.11) is equal to the standard formulation of a forward auction. We can therefore solve the reverse auction problems by transforming them into forward auction problems, solving the problem using our framework, and then transforming the solution back to the reverse auction setting.

In a reverse auction problem, classical auction theory says that (a) the optimal way to buy an object (henceforth: a bridge) with value 1 from a single agent with cost drawn from a uniform distribution on $[0, 1]$ is to offer a take-it-or-leave-it payment of $1/2$, (b) the optimal way to buy a bridge with value $1/2$ from one of multiple agents with uniformly distributed costs is to run a second-price reverse auction with reserve price $1/2$, in which the agent with the lowest cost (if it is less than $1/2$) constructs the bridge and is paid the minimum of the second lowest cost and $1/2$. The above interpretation of the marginal revenue mechanism in i.i.d. settings is one of the most important result in classical auction theory. Our theory generalizes this to multi-dimensional preferences as follows. Consider instead buying a bridge that can be built using technology 1 or technology 2. It says that (a) the optimal way to buy a bridge with value 1 from a single agent with costs for the different technologies each drawn independently and uniformly from $[0, 1]$ is to offer a take-it-or-leave-it payment of $1 - \sqrt{1/3}$ for either technology, (b) the optimal way to buy a bridge with value 1 from one of multiple agents each with i.i.d. uniform costs for each technology is to run the second-price reverse auction with reserve $1 - \sqrt{1/3}$ and allow the winning agent to choose her favorite technology to build the bridge.

APPENDIX C

Single-Agent to Multi-Agent Solutions: Computation

C.1. Proofs from Section 5.5.1

We first describe a network flow formulation of \mathbb{S} , which is used to prove Lemma 32 and Lemma 33.

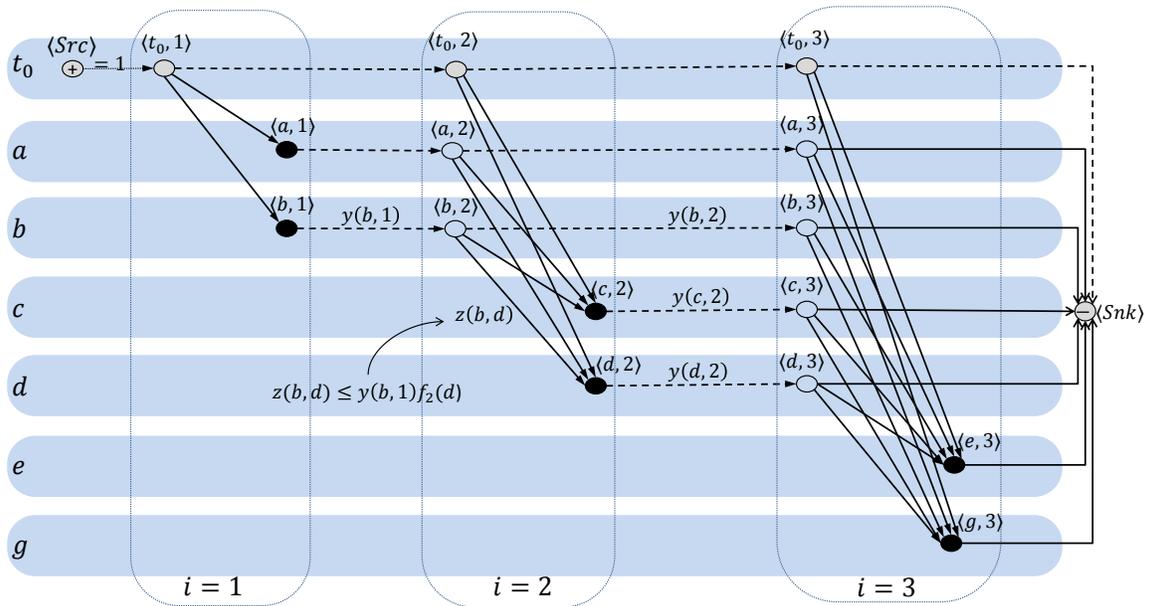


Figure C.1. The flow network corresponding to the SSA algorithm 31. In this instance, there are three agents with type spaces $T_1 = \{a, b\}$, $T_2 = \{c, d\}$, and $T_3 = \{e, g\}$. All nodes in the same row correspond to the same type. The diagonal edges have dynamic capacity constraints while all other edges have no capacity constraints. The flow going from $\langle t_{i'}, i \rangle$ to $\langle t_i, i \rangle$ corresponds to the ex-ante probability of t_i taking the token away from $t_{i'}$. The flow going from $\langle t_{i'}, i \rangle$ to $\langle t_{i'}, i + 1 \rangle$ corresponds to the ex-ante probability of $t_{i'}$ still holding the token after agent i is visited.

A network flow formulation of \mathbb{S} . We construct a network in which every feasible flow corresponds to some $(y, z) \in \mathbb{S}$. The network (see Figure C.1) has a source node $\langle \text{SRC} \rangle$, a sink node $\langle \text{SNK} \rangle$, and $n - i + 1$ nodes for every $t_i \in T_N$ labeled as $\langle t_i, i \rangle, \dots, \langle t_i, n \rangle$ where each node $\langle t_{i'}, i \rangle$ corresponds to the type $t_{i'}$ at the time SSA algorithm is visiting agent i . For each $t_{i'} \in T_N$ and for each $i \in \{i', \dots, n - 1\}$ there is an edge $(\langle t_{i'}, i \rangle, \langle t_{i'}, i + 1 \rangle)$ with infinite capacity whose flow is equal to $y(t_{i'}, i)$; we refer to these edges as “horizontal edges”. For every $t_{i'}$ and every t_i where $i' < i$ there is an edge $(\langle t_{i'}, i \rangle, \langle t_i, i \rangle)$ whose flow is equal to $z(t_{i'}, t_i)$ and whose capacity is equal to the total amount of flow that enters $\langle t_{i'}, i \rangle$ multiplied by $f_i(t_i)$, i.e., it has a dynamic capacity which is equal to $y(t_{i'}, i - 1)f_i(t_i)$; we refer to these edges as “diagonal edges”. There is an edge $(\langle \text{SRC} \rangle, t_0)$ through which the source node pushes exactly one unit of flow. Finally, for every $t_i \in T_N$, there is an edge $(\langle t_i, n \rangle, \langle \text{SNK} \rangle)$ with unlimited capacity whose flow is equal to $y(t_i, n)$. To simplify the proofs we sometimes use $\langle t_0, 0 \rangle$ as an alias for the source node $\langle \text{SRC} \rangle$ and $\langle t_i, n + 1 \rangle$ as aliases for the sink node $\langle \text{SNK} \rangle$. The network always has a feasible flow because all the flow can be routed along the path $\langle \text{SRC} \rangle, \langle t_0, 1 \rangle, \dots, \langle t_0, n \rangle, \langle \text{SNK} \rangle$.

We define the *residual capacity* between two types $t_{i'}, t_i \in T_N$ with respect to a given $(y, z) \in \mathbb{S}$ as follows.

$$(\text{RESCAP}) \quad \text{RESCAP}_{y,z}(t_{i'}, t_i) = \begin{cases} y(t_{i'}, i - 1)f_i(t_i) - z(t_{i'}, t_i) & i > i' \\ z(t_i, t_{i'}) & i < i' \\ 0 & \text{otherwise} \end{cases}$$

Due to dynamic capacity constraints, it is not possible to augment a flow along a path with positive residual capacity by simply changing the amount of the flow along the edges of the path, because reducing the total flow entering a node also decreases the capacity of the diagonal edges leaving that node, which could potentially violate their capacity constraints. Therefore,

we introduce an operator $\text{REROUTE}(t_{i'}, t_i, \rho)$ (algorithm 1 and Figure C.2) which modifies an existing $(y, z) \in \mathbb{S}$, while maintaining its feasibility, to transfer a ρ -fraction of $y(t_i, n)$ to $y(t_{i'}, n)$ by changing the flow along the cycle

$$\langle \text{SNK} \rangle, \langle t_{i'}, n \rangle, \langle t_{i'}, n-1 \rangle, \dots, \langle t_{i'}, \max(i', i) \rangle, \langle t_i, \max(i', i) \rangle, \dots, \langle t_i, n-1 \rangle, \langle t_i, n \rangle, \langle \text{SNK} \rangle$$

and adjusting the flow of the the diagonal edges which leave this cycle. More formally, the operator $\text{REROUTE}(t_{i'}, t_i, \rho)$ takes out a ρ -fraction of the flow going through the subtree rooted at $\langle t_{i'}, \max(i', i) \rangle$ ¹ and reassigns it to the subtree rooted at $\langle t_i, \max(i', i) \rangle$ (see Figure C.2).

Algorithm 1 $\text{REROUTE}(t_{i'}, t_i, \rho)$.

Input: An existing $(y, z) \in \mathbb{S}$ given implicitly, a source type $t_{i'} \in T_N$, a destination type $t_i \in T_N$ where $i' \neq i$, and a fraction $\rho \in [0, 1]$.

Output: Modify (y, z) to transfer a ρ -fraction of $y(t_{i'}, n)$ to $y(t_i, n)$ while ensuring that the modified assignment is still in \mathbb{S} .

- 1: **if** $i' < i$ **then**
 - 2: Increase $z(t_{i'}, t_i)$ by $\rho \cdot y(t_{i'}, i)$.
 - 3: **else**
 - 4: Decrease $z(t_i, t_{i'})$ by $\rho \cdot y(t_{i'}, i')$.
 - 5: **end if**
 - 6: **for** $i'' = \max(i', i)$ **to** n **do**
 - 7: Increase $y(t_i, i'')$ by $\rho \cdot y(t_{i'}, i'')$.
 - 8: Decrease $y(t_{i'}, i'')$ by $\rho \cdot y(t_{i'}, i'')$.
 - 9: **end for**
 - 10: **for** $t_{i''} \in T_{\{\max(i', i)+1, \dots, n\}}$ **do**
 - 11: Increase $z(t_i, t_{i''})$ by $\rho \cdot z(t_{i'}, t_{i''})$.
 - 12: Decrease $z(t_{i'}, t_{i''})$ by $\rho \cdot z(t_{i'}, t_{i''})$.
 - 13: **end for**
-

PROOF OF LEMMA 32. For any given $(y, z) \in \mathbb{S}$ we show that it is always possible to modify y and z to obtain a non-degenerate feasible assignment with the same induced interim allocation probabilities (i.e., the same $y(\cdot, n)$). Let d denote the number of degenerate types with respect

¹This subtree consists of the path $\langle t_{i'}, \max(i', i) \rangle, \dots, \langle t_{i'}, n \rangle, \langle \text{SNK} \rangle$ and all the diagonal edges leaving this path.

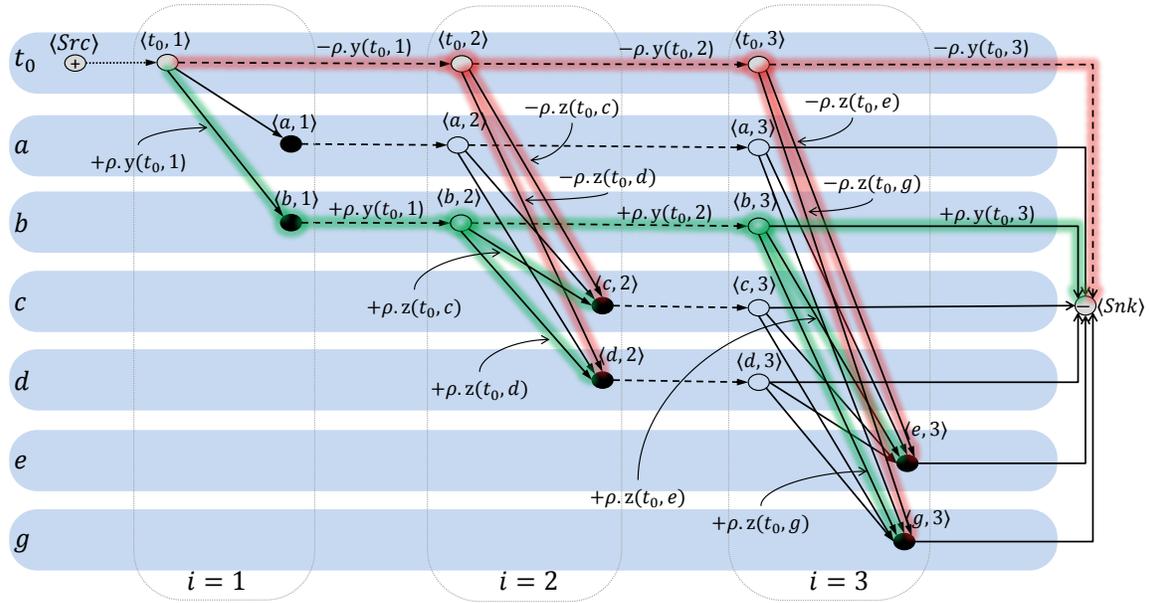


Figure C.2. Changes made by applying $\text{REROUTE}(t_0, b, \rho)$. A ρ -fraction of the red subtree rooted at t_0 is taken out and reassigned to the green subtree rooted at b . The exact amount of change is indicated for each green and each red edge. The flow along all other edges stays intact. The operator has the effect of reassigning ρ -fraction of ex-ante probability of allocation for type t_0 to type b .

to (y, z) , i.e., define

$$d = \# \{t_i \in T_{\{1, \dots, n\}} \mid y(t_i, n) = 0, y(t_i, i) > 0\}$$

The proof is by induction on d . The base case is $d = 0$ which is trivial. We prove the claim for $d > 0$ by modifying y and z , reducing the number of degenerate types to $d - 1$, and then applying the induction hypothesis. Let t_i be a degenerate type. For each $t_{i'} \in T_{\{0, \dots, i-1\}}$, we apply the operator $\text{REROUTE}(t_i, t_{i'}, \frac{z(t_{i'}, t_i)}{y(t_i, i)})$ unless $y(t_i, i)$ has already reached 0. Applying this operator to each type $t_{i'}$ eliminates the flow from $\langle t_{i'}, i \rangle$ to $\langle t_i, i \rangle$, so eventually $y(t_i, i)$ reaches 0 and t_i is no longer degenerate and also no new degenerate type is introduced, so the number of degenerate types is reduced to $d - 1$. It is also easy to see that $y(t_{i'}, n)$ is not modified because $y(t_i, n) = 0$. That completes the proof. \square

PROOF OF LEMMA 33. To prove the lemma it is enough to show that for any augmentable type $t_{i'}$ and any non-augmentable type t_i , $\text{RESCAP}_{y,z}(t_{i'}, t_i) = 0$ which is equivalent to the statement of the lemma (the equivalence follows from the definition of RESCAP and equation (π)). The proof is by contradiction. Suppose $t_{i'}$ is augmentable and $\text{RESCAP}_{y,z}(t_{i'}, t_i) = \delta$ for some positive δ ; we show that t_i is also augmentable. Since $t_{i'}$ is augmentable, there exists a $(y', z') \in \mathbb{S}$ such that $y'(\tau, n) = y(\tau, n)$ for all $\tau \in T_N \setminus \{t_0, t_{i'}\}$ and $y'(t_{i'}, n) - y(t_{i'}, n) = \epsilon > 0$. Define

$$(y'', z'') = (1 - \alpha) \cdot (y, z) + \alpha \cdot (y', z')$$

where $\alpha \in [0, 1]$ is a parameter that we specify later. Note that in (y'', z'') , $t_{i'}$ is augmented by $\alpha\epsilon$, and $\text{RESCAP}_{y'',z''}(t_{i'}, t_i) \geq (1 - \alpha)\delta$, and $(y'', z'') \in \mathbb{S}$ because it is a convex combination of (y, z) and (y', z') . Consider applying $\text{REROUTE}(t_{i'}, t_i, \rho)$ to (y'', z'') for some parameter $\rho \in [0, 1]$. The idea is to choose α and ρ such that the exact amount, by which $t_{i'}$ was augmented, gets reassigned to t_i , by applying $\text{REROUTE}(t_{i'}, t_i, \rho)$; so that eventually t_i is augmented while every other type (except t_0) has the same allocation probabilities as they originally had in (y, z) . It is easy to verify that by setting

$$\alpha = \frac{y(t_{i'}, n)\delta}{2} \qquad \rho = \frac{\epsilon\delta}{2 + \epsilon\delta}$$

we get a feasible assignment in which the allocation probability of t_i is augmented by $\alpha\epsilon$ while every other type (except t_0) has the same allocation probabilities as in (y, z) . We still need to show that $\alpha > 0$. The proof is again by contradiction. Suppose $\alpha = 0$, so it must be $y(t_{i'}, n) = 0$, which would imply that $t_{i'}$ is a degenerate type because $y(t_{i'}, i') > 0$ (because $\text{RESCAP}_{y,z}(t_{i'}, t_i) > 0$), however (y, z) is a non-degenerate assignment by the hypothesis of the lemma, which is a contradiction. That completes the proof. \square

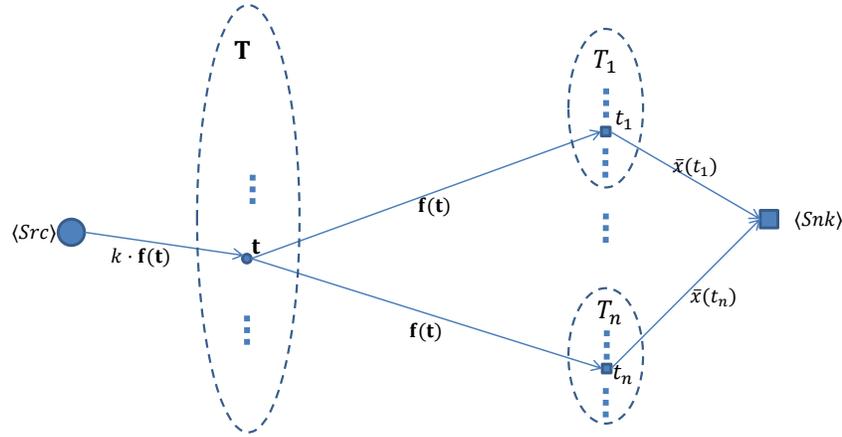


Figure C.3. The bipartite graph used in the max-flow/min-cut argument of the proof of Theorem 21.

C.2. Proofs from Section 5.5.2

REST OF THE PROOF OF THEOREM 21. We give a proof of $P_{g_k} \subseteq \tilde{\mathbf{X}}$ based on the min-cut/max-flow theorem. We start by constructing a directed bipartite graph as illustrated in Figure C.3. On one side we put a node $\langle \mathbf{t} \rangle$, for each type profile $\mathbf{t} \in \mathbf{T}$. On the other side we put a node $\langle t_i \rangle$, for each type $t_i \in T_{\{1, \dots, n\}}$. We also add a source node $\langle \text{SRC} \rangle$ and a sink node $\langle \text{SNK} \rangle$. We add a directed edge from $\langle \text{SRC} \rangle$ to the node $\langle \mathbf{t} \rangle$, for each $\mathbf{t} \in \mathbf{T}$ and set the capacity of this edge to $k \cdot \mathbf{f}(\mathbf{t})$. We also add n outgoing edges for every node $\langle \mathbf{t} \rangle$, each one going to one of the nodes $\langle t_1 \rangle, \dots, \langle t_n \rangle$ and with a capacity of $\mathbf{f}(\mathbf{t})$. Finally we add a directed edge from the node $\langle t_i \rangle$, for each $t_i \in T_{\{1, \dots, n\}}$, to $\langle \text{SNK} \rangle$ with capacity of $\tilde{x}(t_i)$. Consider a maximum flow from $\langle \text{SRC} \rangle$ to $\langle \text{SNK} \rangle$. It is easy to see that there exists a feasible ex post implementation for \tilde{x} if and only if all the edges to the sink node $\langle \text{SNK} \rangle$ are saturated. In particular, if $\rho(\mathbf{t}, t_i)$ denotes the amount of flow from $\langle \mathbf{t} \rangle$ to $\langle t_i \rangle$, a feasible ex post implementation can be obtained by allocating to each type t_i with probability $\rho(\mathbf{t}, t_i)/\mathbf{f}(\mathbf{t})$ when the type profile \mathbf{t} is reported by the agents.

We show that if a feasible ex post implementation does not exist, then $\tilde{x} \notin P_{g_k}$. Observe that if a feasible ex post implementation does not exist, then some of the incoming edges of

$\langle \text{SNK} \rangle$ are not saturated by the max-flow. Let (A, B) be a minimum cut such that $\langle \text{SRC} \rangle \in A$ and $\langle \text{SNK} \rangle \in B$. Let $B' = B \cap T_N$. We show that the polymatroid inequality

$$(C.1) \quad \tilde{x}(B') \leq g_k(B')$$

must have been violated. It is easy to see that the size of the cut is given by the following equation.

$$\text{CUT}(A, B) = \sum_{\mathbf{t} \in \mathbf{T} \cap A} \#\{i | t_i \in B\} \mathbf{f}(\mathbf{t}) + \sum_{\mathbf{t} \in \mathbf{T} \cap B} k \cdot \mathbf{f}(\mathbf{t}) + \sum_{\tau \in T_N \cap A} \tilde{x}(\tau)$$

Observe that for each $\mathbf{t} \in \mathbf{T} \cap A$, it must be that $\#\{i | t_i \in B\} \leq k$, otherwise moving $\langle \mathbf{t} \rangle$ to B would decrease the size of the cut. So the size of the minimum cut can be in simply written as:

$$\text{CUT}(A, B) = \sum_{\mathbf{t} \in \mathbf{T}} \min(\#\{i | t_i \in B\}, k) \mathbf{f}(\mathbf{t}) + \sum_{\tau \in T_N \cap A} \tilde{x}(\tau)$$

On the other hand, since some of the incoming edges of $\langle \text{SNK} \rangle$ are not saturated by the max-flow, it must be that

$$\sum_{\tau \in T_N} \tilde{x}(\tau) = \text{CUT}(A \cup B - \langle \text{SNK} \rangle, \langle \text{SNK} \rangle) > \text{CUT}(A, B),$$

so

$$\sum_{\tau \in T_N \cap B} \tilde{x}(\tau) > \sum_{\mathbf{t} \in \mathbf{T}} \min(\#\{i | t_i \in B\}, k) \mathbf{f}(\mathbf{t}).$$

The right hand side of the above inequality is the same as $\mathbf{E}_{\mathbf{t} \sim \mathbf{f}}[\min(\#\{i | t_i \in B\}, k)]$ which shows that polymatroid inequality (C.1) of P_{g_k} is violated so $\tilde{x} \notin P_{g_k}$. That completes the proof. \square

C.3. Proofs from Section 5.5.3

PROOF OF LEMMA 35. Assuming that agents are independent (i.e., assuming $\mathbf{f}(\cdot)$ is a product distribution), $g_k(S)$ can be computed in time $O((n + |S|) \cdot k)$ using the following dynamic program in which G_j^i denotes the probability of the event that $\min(|\mathbf{t} \cap S \cap T_{\{1, \dots, i\}}|, k) = j$.

$$g_k(S) = \sum_{j=1}^k j \cdot G_j^n$$

$$G_j^i = \begin{cases} G_k^{i-1} + (\sum_{t_i \in S \cap T_i} f_i(t_i)) \cdot G_{k-1}^{i-1} & 1 \leq i \leq n, j = k \\ G_j^{i-1} + (\sum_{t_i \in S \cap T_i} f_i(t_i)) \cdot (G_{j-1}^{i-1} - G_j^{i-1}) & 1 \leq i \leq n, 0 \leq j < k \\ 1 & i = 0, j = 0 \\ 0 & \text{otherwise} \end{cases}$$

□