Computing with Strategic Agents

by

Nicole Immorlica

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2005

Author ....
Department of Electrical Engineering and Computer Science
April 29, 2005

Certified by ....
Erik D. Demaine
Associate Professor
Thesis Supervisor

Certified by ....
David Karger
Associate Professor
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Computing with Strategic Agents

by

Nicole Immorlica

Submitted to the Department of Electrical Engineering and Computer Science
on April 29, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This dissertation studies mechanism design for various combinatorial problems in the presence of strategic agents. A *mechanism* is an algorithm for allocating a resource among a group of participants, each of which has a privately-known value for any particular allocation. A mechanism is *truthful* if it is in each participant's best interest to reveal his private information truthfully regardless of the strategies of the other participants.

First, we explore a *competitive auction framework* for truthful mechanism design in the setting of *multi-unit auctions*, or auctions which sell multiple identical copies of a good. In this framework, the goal is to design a truthful auction whose revenue approximates that of an omniscient auction for any set of bids. We focus on two natural settings — the *limited demand* setting where bidders desire at most a fixed number of copies and the *limited budget* setting where bidders can spend at most a fixed amount of money. In the limit demand setting, all prior auctions employed the use of randomization in the computation of the allocation and prices. Randomization in truthful mechanism design is undesirable because, in arguing the truthfulness of the mechanism, we employ an underlying assumption that the bidders trust the random coin flips of the auctioneer. Despite conjectures to the contrary, we are able to design a technique to derandomize any multi-unit auction in the limited demand case without losing much of the revenue guarantees. We then consider the limited budget case and provide the first competitive auction for this setting, although our auction is randomized.

Next, we consider abandoning truthfulness in order to improve the revenue properties of *procurement auctions*, or auctions that are used to hire a team of agents to complete a task. We study first-price procurement auctions and their variants and argue that in certain settings the payment is never significantly more than, and sometimes much less than, truthful mechanisms.

Then we consider the setting of *cost-sharing auctions.* In a cost-sharing auction, agents bid to receive some service, such as connectivity to the internet. A subset of agents is then selected for service and charged prices to approximately recover the cost of servicing them. We ask what can be achieved by cost-sharing auctions satisfying a

strengthening of truthfulness called *group-strategyproofness*. Group-strategyproofness requires that even coalitions of agents do not have an incentive to report bids other than their true values in the absence of side-payments. For a particular class of such mechanisms, we develop a novel technique based on the probabilistic method for proving bounds on their revenue and use this technique to derive tight or nearly-tight bounds for several combinatorial optimization games. Our results are quite pessimistic, suggesting that for many problems group-strategyproofness is incompatible with revenue goals.

Finally, we study *centralized two-sided markets*, or markets that form a matching between participants based on preference lists. We consider mechanisms that output matching which are *stable* with respect to the submitted preferences. A matching is stable if no two participants can jointly benefit by breaking away from the assigned matching to form a pair. For such mechanisms, we are able to prove that in a certain probabilistic setting each participant's best strategy is truthfulness with high probability (assuming other participants are truthful as well) even though in such markets in general there are provably no truthful mechanisms.

Thesis Supervisor: Erik D. Demaine
Title: Associate Professor

Thesis Supervisor: David Karger
Title: Associate Professor

Acknowledgments

I would like to thank my advisors, Erik Demaine and David Karger, for their guidance in my research and academic career. To my committee, Erik Demaine, David Karger, David Parkes, and Madhu Sudan, thanks for your time and interest. I am especially grateful to David Karger, David Parkes, and Mohammad Mahdian for copious comments that significantly improved the presentation of my results.

I am forever indebted to my numerous colleagues, past, present, and future, for promoting a stimulating research environment. The results in this thesis were based on joint work with Gagan Aggarwal, Christian Borgs, Jennifer Chayes, Brian Dean, Amos Fiat, Michel Goemans, Andrew Goldberg, Jason Hartline, David Karger, Mohammad Mahdian, Vahab Mirrokni, Evdokia Nikolova, Amin Saberi, Rahul Sami, and Madhu Sudan and discussions with Joan Feigenbaum, Kamal Jain, Anna Karlin, Bobby Kleinberg, Hervé Moulin, Martin Pál, Al Roth, Kunal Talwar, Eva Tardos, and Rakesh Vohra. I have also benefited greatly from collaborations with Yigal Bejerano, Randeep Bhatia, Steve Chien, Mayur Datar, Erik Demaine, MohammadTaghi Hajiaghayi, Piotr Indyk, Tracy Kimbrel, Maria Minkoff, Seffi Naor, Baruch Schieber, and Mark Smith. Of the many others who have been instrumental in my career, I mention just two: Peter Fisher, my master's thesis advisor, who told me the secret to success, and Joel Spencer, who gave me the courage to pursue it.

Essential to my well-being throughout my four long years in grad school were my friends and family. To my friends, Mohammad Mahdian, Vahab Mirrokni, and all the rest, my best "results" are the memories we created. I will carry them with me longer than anything else I learned in grad school! To my family, especially my parents Annette and Tony Immorlica who have stuck by me for the past 26 years, thanks for a wonderful home, for Tiki, Justin, Trond, and Minx, and most of all, for you. I could not have completed this work without your unfaltering love and support.

Finally, to all the wonderful coffeeshops, especially 1369 and Toscaninni's and Fremont Coffee Company (where the majority of this thesis was written while sipping a steaming latte), thanks.

Contents

1 Introduction **11**

1.1 Mechanism Design 12

1.2 Sample Applications 14

1.2.1 Multi-Object Auctions 14

1.2.2 Procurement Auctions 15

1.2.3 Cost-Sharing Auctions 15

1.2.4 Two-Sided Markets 16

1.3 Our Contributions 16

2 Multi-Unit Auctions **23**

2.1 Truthful Auction Design 25

2.2 Limited Demand 27

2.2.1 Setting 28

2.2.2 A Hat Puzzle 29

2.2.3 Auction Derandomization 33

2.2.4 A Polynomial-time Deterministic Auction 36

2.3 Budget Constraints 42

2.3.1 Setting 42

2.3.2 An asymptotically optimal auction 45

2.3.3 Impossibility Result 51

3 Procurement Auctions **61**

3.1 Setting 63

3.2 Solution Concepts 65

3.2.1 Nash Equilibria 66

3.2.2 ϵ-Nash Equilibria 67

3.2.3 Strong ϵ-Nash Equilibria 68

3.3 Approximate Efficiency of First-Price Combinatorial Auctions 73

3.4 Payment Bounds for Flow Auctions 74

3.4.1 Known Demand Path Auction 75

3.4.2 Unknown Demand Path Auction 84

4 Cost-Sharing Auctions **99**

4.1 Setting 103

4.2 Upper bounds for cross-monotonic cost sharing schemes 105

4.2.1 A simple example: the edge cover game 105

4.2.2 The set cover game 108

4.2.3 The vertex cover game 110

4.2.4 The metric facility location game 113

4.2.5 Other combinatorial optimization games 116

4.3 Implications for Cost-Sharing Mechanisms 119

4.3.1 Preliminaries 119

4.3.2 A discussion about the definition 121

4.3.3 A partial characterization of group-strategyproof mechanisms 123

5 Two-Sided Markets **131**

5.1 Setting 135

5.2 Economic implications 139

5.3 Proof of Theorem 5.1.5 142

5.3.1 Counting the number of stable husbands 142

5.3.2 Analyzing the expectation 145

5.3.3 Number of singles 149

5.4 Tighter analysis for the uniform distribution 152

6 Conclusion **159**

A Group-Strategyproof Mechanisms and Cost-Sharing Schemes **167**

A.1 A group-strategyproof mechanism with no cross-monotonic cost-sharing scheme 167

A.2 A semi-cross-monotonic cost-sharing scheme with no group-strategyproof mechanism 170

B Expected Number of Singles in Two-Sided Markets **171**

C Unsplittable Stable Marriage Problems **175**

Chapter 1

Introduction

As computer scientists, we have traditionally assumed that our algorithms operate in isolation from their environment. In our models, algorithms are presented with environmental inputs which they use to design a solution. This solution is then implemented in the environment, and in our standard analytic methods such as adversarial analysis, average-case analysis, or simulations, we assume that the choice of the algorithm does not affect the environmental input. Although this is a reasonable model in many settings, the advent of new technologies such as the internet has resulted in the growth of new computational problem spaces where the design of the algorithm affects the inputs and the behavior of the users. Prominent examples include the allocation of radio spectrum, privatization of public services, or, closer to home, peer-to-peer networks like Gnutella and electronic commerce sites like eBay. In these systems, participants act in their own self-interest and therefore declare inputs which, based on their knowledge of the other participants and the mechanism itself, they believe will maximize their own gain.

The field of *mechanism design* or *implementation theory* attempts to build systems taking into consideration the strategic behavior of the participants (see [62, 63] for a survey). The basic paradigm postulates that each individual maintains some private information relevant to the problem at hand. A system solicits from participants their private information and then computes a global solution to the problem. The properties of this global solution are analyzed in a game-theoretic *equilibrium* or a steady-state induced by the behavior of rational participants. It is the goal of the designer to introduce incentives in the system such that these equilibria result in globally optimal solutions. *Algorithmic mechanism design* focuses on the application of mechanism design to computationally intensive settings (see [89] for an introduction).

1.1 Mechanism Design

We consider a setting in which there is a set N of n *agents* denoted $\{1, \ldots, n\}$ which must collectively reach a common decision from a set of feasible decisions D. These agents might be, for example, people at an art auction, the decision being the price and allocation of the work of art. Each agent i maintains some private information relevant to the decision problem at hand. This private information, denoted θ_i, is called his *type*, and is an element of his *type space* Θ_i. In the art auction, for example, this private information might be the agent's value for the work of art, the type space being the set of non-negative real numbers $\mathbb{R}^+$.

Agents have preferences over decisions as represented by a utility function $u_i : D \times \Theta_i \to \mathbb{R}$ for each agent i.[1] A decision $d \in D$ for an agent i with type $\theta_i \in \Theta_i$ is said to have utility $u_i(d, \theta_i)$ for agent i. In the art auction example outlined above, the utility function for a bidding agent, or *bidder*, i might be his value for the allocation minus his price. Utility functions of this form are called *quasi-linear*. For example, the value of agent i for an allocation x might be $v_i(\theta_i, x) = \theta_i$ if he is allocated the work of art, and $v_i(\theta_i, x) = 0$ otherwise. His utility function is then $u_i(d, \theta_i) = v_i(\theta_i, x) - p$ if the decision d is to choose allocation x and charge agent i at price p. We assume that the value of the auctioneer for the work of art is 0, and so his utility for a decision d is p if the decision d is to allocate the work of art to a bidder at price p and 0 otherwise.

A mechanism M is a pair $(\times_i A_i, g)$ defining for each agent i a set of actions A_i and a decision function $g : \times_i A_i \to D$ mapping the agents' actions to a decision. If the set of actions of an agent A_i equals the type space of that agent Θ_i, then the mechanism

[1] Note that this definition requires an agent's utility to be independent of the other agents' types. This is called the *private value* setting, and will be the focus of this dissertation.

is said to be a *direct revelation* mechanism. The action space in a direct revelation art auction might be the sealed bid announcement of a value (that is, $A_i = \mathbb{R}^+$ for all i); the decision function might allocate the work of art to an arbitrary agent with a maximum announced value, charging him a price equal to the maximum announced value among the remaining agents.

We study the properties of a mechanism in an *equilibrium* state. A vector of actions is an equilibrium if each agent's action maximizes his (expected) utility given perhaps some information regarding the types of the others (for example, knowing their types, or knowing a probability distribution over their types). In some cases, it is possible to define the mechanism M in such a way that each agent i has a *dominant strategy* action a_i^*, or one which maximizes his utility regardless of the actions of the other players. In other words, for each possible type $\theta_i \in \Theta_i$ there is an action $a_i^*(\theta_i)$ such that

$$u_i\left(g(a_i^*(\theta_i), a_{-i}), \theta_i\right) \geq u_i\left(g(a), \theta_i\right)$$

for all profiles a of actions of the agents (where, as is standard, we have used the notation a_{-i} to denote the $(n-1)$-dimensional vector $(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ and the notation (a_i, a_{-i}) to denote the n-dimensional vector a). A vector of dominant strategy actions is called a *dominant strategy equilibrium.* A direct revelation mechanism in which revealing θ_i is a dominant strategy for any agent i with type θ_i is called a *truthful* or *incentive compatible* mechanism (we use these terms interchangeably). So far, we have implicitly assumed that the only actions available to agents are those defined by the mechanism. In other words, an agent does not have a choice about whether to participate in the mechanism. A direct revelation mechanism is *individually rational* if, in an equilibrium, no agent is worse off by participating. Specifically, let d be an equilibrium outcome of a mechanism M. We say M is individually rational if for all agents i and types $\theta_i \in \Theta_i$, $u_i(\theta_i, d) \geq 0$ (here we assume an agent can guarantee himself a utility of zero by not participating). Henceforth, all auctions we consider will be individually rational unless otherwise stated. By this definition, a truthful auction is individually rational if no agent is charged more than his bid.

The goal of the mechanism designer is, given a *social choice function* $f : \times_i \Theta_i \rightarrow 2^D$, to define M in such a way that $g(a) \in f(\theta)$ for all type vectors $\theta \in \times_i \Theta_i$ and all equilibria $a \in \times_i A_i$ of agents with types θ. The social choice function might be, for example, the set of decisions which maximize the social welfare or the revenue.

A prominent example of a truthful mechanism is the Vickrey-Clark-Groves mechanism, or *VCG mechanism*, developed in a series of papers by Vickrey [115], Clark [19], and Groves [52] for a general setting where agents have quasi-linear utility functions. The VCG mechanism is a direct revelation auction that chooses a decision that maximizes the *social welfare*, or the sum of values of the bidders (equivalently, the sum of utilities of all the agents including the auctioneer). It then charges each bidder a price equal to his value minus a bonus. The bonus is defined as the amount by which his presence increases the social welfare. More formally, given reported valuations θ of the bidding agents, define $W(\theta) = \max_{x \in D} \sum_{i=1}^{n} v_i(\theta_i, x)$. The VCG mechanism chooses a decision x which maximizes $\sum_{i=1}^{n} v_i(\theta_i, x)$ and then charges each agent i a price equal to $p_i = v_i(\theta_i, x) - (W(\theta) - W(\theta_{-i}))$. In the case of a single item auction, the VCG mechanism is precisely the second price auction described earlier. A classic economic result states that this mechanism is truthful and *efficient* (that is, maximizes the social welfare) when utility functions are quasi-linear [19, 52, 115].

1.2 Sample Applications

The fundamental scenario in mechanism design is applicable in a wide variety of settings, of which we highlight a few.

1.2.1 Multi-Object Auctions

Multi-object auctions concern the sale of multiple related objects by an auctioneer (the seller) to interested bidders (the buyers). The notion of an object is intentionally generic and encompasses anything from works of art or cut flowers to, in recent times, radio spectrum or advertisement slots on web search pages. A special case discussed in this dissertation is the *multi-unit* setting in which the auctioneer wishes to sell multiple identical copies of an object like two copies of the same photograph.

In multi-object and multi-unit auction settings, the set of agents is the bidders.[2] The type of an agent is the utility function of the agent. The feasible decisions are all possible allocations of the objects along with a price vector. There are a variety of natural utility functions one might assume in this setting, the most prevalent being quasi-linear utilities. One plausible goal of the mechanism designer is to maximize revenue of the auctioneer. Another option is to maximize the social welfare of the decision.

1.2.2 Procurement Auctions

This setting is quite similar to the last except that the auctioneer is a *buyer* and the bidders are *sellers*. Typically, the auctioneer wants to hire a team of agents to complete some task. Each agent, if selected, performs some fixed service which facilitates completion of the task. For example, the auctioneer might be an Internet Service Provider (ISP) who needs to enlist the services of several routing domains, or autonomous systems (ASs), in order to route packets in the internet.

Again, the set of agents is the bidders. The type of a bidder is the cost to the bidder of performing his service. The feasible decisions are a set of bidders capable of completing the task along with a payment vector. The utility of a bidder is his payment minus his cost (if selected). A goal in this setting might be to minimize the payment of the auctioneer (that is, the sum of payments made to the bidders).

1.2.3 Cost-Sharing Auctions

In cost-sharing auctions, agents bid for a service. For each subset of bidders, there is a cost associated with providing the service for that subset, and the objective of the auction is to determine which bidders receive the service and how much each of them has to pay to recover the cost of the service. For example, the residents of a town

[2]Alternatively, we could define the set of agents to include the auctioneer, but it will prove more convenient in subsequent discussions to exclude him.

might wish to build a new power generator. To finance the cost of the generator, they can run a cost-sharing auction to decide which residents will be serviced by the power generater and how much each of them must pay.

The type of each agent in a cost-sharing auction is the value to him of receiving the service. A feasible decision is a set of agents to be serviced, and a cost to charge each agent. An agent's utility is his value for his allocation minus his price. One goal here is to charge prices which exactly recover the cost of the production (called *budget balance*), or, if that is not possible, to recover an α fraction of the cost (called *α-budget balance*).

1.2.4 Two-Sided Markets

Two-sided markets refer to scenarios in which there are two sets of participants, like workers and firms or men and women, which must be paired together. A motivating example in this setting is the National Residency Matching Program (NRMP) which matches medical school graduates to internship programs in hospitals following their schooling.

In this setting, the set of agents is the two sets of participants, and the type of a participant is an ordered list of preferences over members of the opposite set. A feasible decision is a matching between members of opposite sets. Such a matching is called *stable* if there are no two participants who prefer each other to their respective assignments in the matching. A goal in a two-sided market might be to output a stable matching.

1.3 Our Contributions

When they exist, truthful mechanisms are very attractive. Their equilibria are highly stable and predictable, and furthermore, it is fairly simple for each participant to compute his optimal strategy (he need only be concerned about his own situation and is not required to guess the strategies of his fellow participants). One important area of research is to find the best truthful mechanism for a given problem. However, for some problems, by considering stronger equilibrium concepts, it is possible to design a mechanism that, in equilibrium, significantly improves the objective. In other problems, it may be theoretically impossible for a truthful mechanism to optimize the global objective. This dissertation studies the benefits, drawbacks, and necessity of truthful mechanism design for a variety of market settings.

The field of truthful mechanism design has a long history in the economics literature. Of this large body of work, the result most relevant to this dissertation is the Vickrey-Clark-Groves (VCG) mechanism [19, 52, 115]. Applicable to a wide array of settings, the VCG mechanism defines a truthful mechanism that maximizes the social welfare. However, this mechanism does not necessarily maximize the revenue.

Chapter 2: Truthfulness in Multi-Unit Auctions

We first study revenue-maximizing truthful mechanism design in multi-unit auctions, that is an auction for multiple copies of a single good. This setting is an important special case of the multi-object (combinatorial) auction setting (for a discussion of this setting, see [22]). We consider the revenue maximization problem under various assumptions regarding the form of the bidders' utility functions. In particular, we assume bidders either have *limited demand* (have utility for at most a fixed and publicly known quantity of the good) or *limited budgets* (have a fixed privately known budget to spend). We further assume that their utility increases linearly with their allocation up to their budget or demand constraint.

The limited demand case is fairly standard and many of the known results from the economics literature apply. The VCG mechanism may be applied in this setting, but its revenue is quite low when the supply is large. Alternatively, if one assumes that the valuations of bidders are drawn independently from a known probability distribution, then a result of Myerson [88] yields a truthful mechanism with maximum expected revenue. However, assuming a probability distribution for the bidders' valuations is often unsatisfactory. In an effort to avoid any such assumption, Goldberg, Hartline, and Wright [50] proposed a framework called *competitive auctions* for designing high-revenue auctions. This framework seeks to maximize the ratio between the revenue of an auction to the revenue of an omniscient auctioneer (that is, an auctioneer who knows the valuations of the bidders). In a sequence of works, Goldberg et al. [49, 48, 50] and Hartline and McGrew [54] use this framework to design high-revenue truthful auctions without any assumption on the bidders' valuations.

All of these results employ randomization in the computation of the allocation. Use of randomization in truthful mechanism design is unsatisfactory from two viewpoints: first, the resulting revenue guarantees hold only with some probability, and second, the property of truthfulness holds only if the participants trust that the coin flips of the auctioneer are indeed fair.[3] In Chapter 2, we introduce a technique to "derandomize" truthful mechanisms without sacrificing much of the expected revenue guarantee, thus providing *deterministic* high-revenue auctions for the limited demand setting.

The limited budget setting has received much less attention in the literature. The utility functions of bidders in this setting are not quasi-linear (see Section 1.1). Thus, many classic results such as the VCG mechanism are not even well-defined in this setting. In fact, in Chapter 2, we show that, modulo some technical assumptions, there is *no* truthful mechanism for this setting. On the positive side, we drop some of the technical assumptions in order to provide the first (randomized) truthful mechanism for this setting with high revenue guarantees. Unfortunately, our derandomization techniques do not extend to this setting, and so the existence of a deterministic mechanism remains an open question.

The results of this chapter regarding limited demand are based on joint work with Aggarwal, Fiat, Goldberg, Hartline, and Sudan [2]. The results regarding budget-constrained bidders are based on joint work with Borgs, Chayes, Mahdian, and Saberi [14].

[3]It is standard to assume that bidders trust the auctioneer to implement the announced mechanism and so one might argue that trusting the auctioneer's random source is not unreasonable. However, an auctioneer can easily provide a certificate to prove the first fact (by, for example, announcing the inputs). It is much harder to prove that the coin flips are truly random.

Chapter 3: Weaker Notions than Truthfulness in Procurement Auctions

We then turn our attention to the setting of procurement auctions with a particular focus on *path* and *flow* auctions. In path auctions, the auctioneer, a buyer, seeks to buy a path of edges between a specified source and destination in a network while spending as little as possible. Each edge of the network is owned by an agent. Agents (network edges) have a privately known cost for transmitting traffic, and bid to attract traffic. A flow auction is a generalization of a path auction in which the auctioneer has a demand that might exceed the capacity of a single path and therefore needs to buy a set of edges, called a flow, capable of routing his demand between the source and the sink.

Nisan and Ronen [89] proposed applying the VCG mechanism to path auctions; Hershberger and Suri [55] and Feigenbaum et al. [35] study methods to make this mechanism efficiently computable and practical in this setting. However, as observed by Archer and Tardos [6], the VCG mechanism (and, in fact, all *min function* mechanisms) can force the auctioneer to pay far more than the true cost of the cheapest path. The tendency to overpay is exaggerated in path auctions because a bonus needs to be paid to every agent on the path. Thus, the payment to the lowest-cost path may even greatly exceed the cost of the second-cheapest path. Elkind, Sahai, and Steiglitz [29] generalized the result of Archer and Tardos [6] to prove that *all* truthful mechanisms have high overpayments.

As motivated by the work of Elkind, Sahai, and Steiglitz [29], to reduce the payment we must turn our attention to weakened equilibrium concepts. These equilibrium concepts require that the participants have more information about the setting (for example, that they know the others' types or a probability distribution over their types). Thus they are arguably less predictable and stable than dominant strategy equilibria.[4] Nonetheless, they promise substantial savings and so merit attention.

Elkind, Sahai, and Steiglitz [29] present and analyze an optimal *Bayesian-Nash* mechanism. Czumaj and Ronen [20] propose a mechanism which combines dominant

[4]Substantial effort has been made to understand when we might expect mechanisms to exhibit these weaker equilibria (see, for example, [38]).

and non-dominant strategy mechanisms; however they show that it has an arbitrary ratio between the payment of different equilibria and say that overall, "finding a natural and tractable measure of [non-dominant strategy] protocols seems challenging and important." In Chapter 3, we analyze the payment properties of a *first-price auction*, or one in which the cheapest feasible set is selected and each agent in this set is paid a price equal to his bid. We show that the payment of such an auction in a *strong ϵ-Nash equilibrium* is never more than, and often much less than, the payment in a truthful auction.

The results of this chapter are based on joint work with Karger, Nikolova, and Sami [59].

Chapter 4: Stronger Notions than Truthfulness in Cost-Sharing Auctions

In Chapter 4, we consider cost-sharing auctions based on combinatorial optimization games. In our setting, the "good" that is auctioned is a service, like connectivity to the internet. The cost of the service for a particular subset of agents can be computed by solving a combinatorial optimization problem, such as Steiner tree in the internet connectivity example. The allocation rule selects a set of agents to service. Thus, the allocation to an agent is a binary decision – he either receives service or does not.

In this setting, we explore a strengthened equilibrium concept called *group strategyproofness* intended to rule out collusion among participants. This equilibrium concept requires that no subset of participants has an incentive to collectively deviate from their truthful strategies. We assume there are no side payments in the system, and so a coalition has incentive to deviate only if no member of the coalition is worse off and some member is strictly better off after the deviation.

Moulin [85] presents a method for constructing group-strategyproof mechanisms based on combinatorial constructions called *cross-monotonic cost-sharing schemes*. This results together with cost-sharing schemes such as those of Shapley [107] or Dutta and Ray [26] imply budget-balanced group-strategyproof mechanisms for any submodular cost function. However, many interesting cost functions based on combinatorial optimization problems are not submodular. It is known that no cross-monotonic cost-sharing scheme for some problems can be perfectly budget-balanced. Accordingly, approximately budget-balanced schemes have been proposed for many combinatorial optimization problems including minimum spanning tree [64, 68], Steiner tree [64], Steiner forest [72], facility location [90], and connected facility location [78].

In Chapter 4, we provide a general methodology to prove upper bounds on the fraction of cost recoverable by cross-monotonic cost-sharing schemes. We apply our bounds to several cost functions to obtain significantly stronger bounds than previously known. In many cases, the bounds we get are tight (for example, for facility location), or nearly tight. We also provide a partial characterization of group strategyproof mechanisms in terms of cross-monotonic cost-sharing schemes and are thus able to argue that our bounds hold for all group strategyproof mechanisms that satisfy additional properties. Our results are quite pessimistic, suggesting that often group strategyproofness may be too strong a goal.

The results in this chapter are based on joint work with Mahdian and Mirrokni [61].

Chapter 5: Truth in Two-Sided Matching Markets

Finally, in Chapter 5, we study an existing marketplace in which truthful mechanisms are provably limited and introduce a new analysis technique to explain observed behavior in this marketplace. Specifically, we consider centralized two-sided marketplaces such as the National Residency Matching Program (NRMP). Ideally, a centralized mechanism should output a stable solution in order to prevent participants from forming matchings outside the market and decentralizing it. However, it is well-known that in a stable matching mechanism, it is not always in a participant's self-interest to announce their true preference list. Yet Roth and Peranson [99] observe that in practice in the NRMP, very few participants have incentives to lie. They suggest that this phenomenon is due to the fact that the length of the medical students' preference lists is necessarily quite short. We provide a theoretical justification of this observation by showing that in a reasonable probabilistic setting, the expected number of participants who can improve their match by submitting a false preference list is vanishingly small. This proves a conjecture of Roth and Peranson [99], and implies that, with high probability, a participant's best strategy is truthfulness when other participants are truthful. Thus, even a dishonest participant is incentivized to be truthful if he believes in the honesty of others. Furthermore, this result implies that the NRMP mechanism has an equilibrium in which most participants are truthful and an approximate equilibrium in which all participants are truthful. This proves a conjecture of Roth and Peranson [99] regarding the NRMP marketplace.

The results of this chapter are based on joint work with Mahdian [60].

Chapter 2

Multi-Unit Auctions

In this chapter, we focus on markets in which there are multiple *identical* units of a good for sale. In these markets, it is natural to consider various forms of the utility function for the buyers. As units of an identical good are perfect substitutes, like tickets to a concert or cut flowers, in general each additional unit has less value to a bidder than the last. Here, we concern ourselves only with the special case where this marginal return is initially constant and then zero so that every additional unit is equally valuable up to some limit. For instance, it is likely that buyers have limited demand; perhaps they desire at most one unit of the good as might be the case in the auction of a limited addition photograph. We refer to this general setting as the *limited demand case*, and the setting of unit demands as the *single-unit demand case.* On the other hand, as is the case for an advertisement slot on a web page, a buyer (advertiser) might have a desire to be shown during multiple viewings of the web page and so have a constant marginal return but be constrained by a limited advertising budget. These buyers are called *budget-constrained* buyers.

There are many plausible formats for the sale of identical goods. They may be offered at a fixed price, auctioned one after the other in multiple rounds, or sold in a single auction. This last format is known as a *multi-unit auction*, and is used extensively in settings such as the sale of government securities [113], the sale of advertising slots on web pages, and, even, the initial public offering of Google stock. Several auction formats have been proposed for these settings, the most abundant being the *discriminatory* auction and the *uniform-price* auction (see the book by Krishna [74] for a survey on multi-unit auctions).

A standard goal in the design of multi-unit auctions, and auction design in general, is to maximize revenue. One approach to maximizing revenue, traditionally applied in the limited supply and/or limited demand case, is to determine an optimal *reserve price* based on assumptions about the distribution of values of the bidders [16, 88]. The budget-constrained case has also been investigated using this approach in several recent papers [11, 17, 18, 30, 31, 76, 79, 120] often in the context of privatization of high-value public goods, such as FCC auctions of telecommunications bands. Given an accurate model of these distributions, this approach is often optimal in the sense that it maximizes expected revenue. However, a misconception regarding these distributions can significantly reduce the revenue of the auction.

Another approach, pioneered by Goldberg, Hartline, and Wright [50], makes no assumptions regarding the value distribution and still approximately maximizes revenue even in the worst-case scenario. A worst-case competitive analysis framework is employed to compare the revenue of the proposed auction to that of an optimal auction run by an omniscient auctioneer, or one who knows the private values of all the bidders. The minimum ratio, over all possible input values, of the mechanism's revenue to the optimal revenue is called the *competitive ratio* of the mechanism, and the goal is to maximize this ratio. This framework has been used to design truthful auctions for the single-unit demand case. [37, 46, 48, 50, 54].[1]

This chapter follows the competitive framework introduced by Goldberg, Hartline, and Wright [50]. We design truthful mechanisms for profit maximization in markets in both the limited demand case and the budget-constrained case. We begin with a folklore theorem which characterizes the set of truthful auctions (Section 2.1). In the limited demand case, all prior results in the competitive framework employed the use of randomization in the allocation and pricing scheme. We present the first derandomized mechanism for this setting, along with our derandomization technique

[1]All of this work is extendable to multi-unit demands when the demands are public knowledge and the marginal returns are constant.

(Section 2.2). We also give the first (randomized) mechanism in the competitive framework for the limited demand case (Section 2.3).

The results of Section 2.2 are from a joint work with Aggarwal, Fiat, Goldberg, Hartline, and Sudan [2]. The results of Section 2.3 are from a joint work with Borgs, Chayes, Mahdian, and Saberi [14].

2.1 Truthful Auction Design

The space of truthful auctions is characterized by the following theorem which states that the price charged to a bidder must be independent of his bid. Statements similar to this one have appeared in numerous places and date back to at least the 1970s [82]. For simplicity, we state the characterization for the special case of a single-item auction in which each bidder's utility is quasi-linear (see Proposition 2.3.1 for a generalization to a non-quasi-linear multi-unit setting).

Definition 1 (Bid-independent Auction) *Suppose there are n bidders. Let f_i be a function from bid vectors of dimension $n-1$ to $\mathbb{R}^+ \cup \{0\} \cup \{\infty\}$ (where ∞ is a number larger than any bid) and $\mathbf{f}$ be a set of n such functions. The* deterministic bid-independent auction *defined by $\mathbf{f}$ works as follows. For each bidder i:*

1. *Set $t_i = f_i(\mathbf{b}_{-i})$.*
2. *If $t_i < b_i$, bidder i wins at price t_i*
3. *If $t_i > b_i$, bidder i loses.*
4. *Otherwise, ($t_i = b_i$) the auction can either accept the bid at price t_i (in which case bidder i is a winner) or reject it.*

Theorem 2.1.1 *An (individually rational) deterministic auction is truthful if and only if it computes the same allocation and prices as a bid-independent auction.*

Proof. It is clear that a bid-independent auction is truthful. Consider a truthful auction A and define $f_i(\mathbf{b}_{-i})$ as the minimum bid b_i for which bidder i wins the good when the vector of bids is $\mathbf{b}$. Consider a bid vector $\mathbf{b}$ and let $t_i = f_i(\mathbf{b}_{-i})$. We consider the three cases $t_i < b_i$, $t_i = b_i$, and $t_i > b_i$.

First suppose $t_i > b_i$. Then, by individual rationality, A does not allocate the good to i.

Next, if $t_i = b_i$ and A does not allocate i the good then the payment is zero by individual rationality. If A allocates the good to i, then the price p is at most t_i by individual rationality. Suppose $p < t_i$. Then when the true value of i is in the interval (p, t_i), he has an incentive to report t_i, contradicting truthfulness. Therefore, $p = t_i$.

Finally suppose $t_i < b_i$. Then we claim A must allocate the good to i at price t_i. If not, suppose p is the price at which A sells the good to i ($p = \infty$ if i is not allocated the good). If $p > t_i$, then i has an incentive to report t_i when his true value is b_i in order to receive the good at a lower price p' where $p' \leq t_i$ by individual rationality. Similarly, if $p < t_i$, then i has an incentive to report b_i when his true value is t_i (the above argument proved that when i bids t_i he is either allocated the good at price t_i or not allocated the good).

Thus, A is equivalent to the bid-independent auction defined by $\mathbf{f}$ in which the allocation when $t_i = b_i$ is defined by calling A. □

As an example, consider the sealed-bid second-price auction for a single good. As described in Chapter 1, this auction allocates the good to the highest bidder and charges him a price equal to the second-highest bid. A simple thought experiment shows that this auction is truthful when the utility of an agent is his value for his allocation minus his price. Consider a bidder who submits his true value as his bid and suppose he is the highest bidder. Raising his bid has no effect on his price and allocation. On the other hand, by decreasing his bid, he runs the risk of undercutting the second highest bidder and losing the good even though the price was favorable to him. Similarly, if he is not the highest bidder, lowering his bid has no effect on his allocation. By increasing his bid, he runs the risk of overshooting the highest bidder at a price unfavorable to him. Indeed, the auction is truthful, as can be seen by defining the bid-independent functions $f_i(\mathbf{b}_{-i}) = \max_{j \neq i}\{b_j\}$ for all i.

2.2 Limited Demand

In many multi-unit markets, buyers have utility for at most a fixed number of units of a good. An important special case is when buyers have *unit* demand for the good, and this is the focus of the competitive auction literature.[2] In a sequence of papers, Goldberg et al. [37, 46, 48, 50] and Hartline and McGrew [54], gave truthful *randomized* auctions for the single-unit demand case that achieve a constant fraction of the optimal single price revenue with high probability. Our goal is to design a truthful auction that guarantees a constant fraction of the optimal revenue and is *deterministic*. Unfortunately, Goldberg, Hartline, and Wright [50] prove that randomization is necessary in order for truthful *symmetric* auctions, or ones whose outcome is not a function of the order of the input bids, to achieve a revenue comparable to the optimal revenue.

It was conjectured that this impossibility result holds for truthful *asymmetric* auctions as well. An asymmetric auction considers the order of bids in the allocation and pricing algorithms and can therefore produce outcomes which, for example, offer the i'th bidder a price equal to the average of the first $(i-1)$ bids. We assume bidders can not affect their position in the ordering. In this section, we show that truthful deterministic *asymmetric* auctions can generate revenue which is a constant fraction of optimal, thus disproving the conjectured impossibility. In fact, we show that *any* truthful randomized auction has a truthful deterministic counterpart with approximately the same revenue guarantees, and so asymmetry is, in some sense, as powerful as randomization. Specifically, for any truthful randomized auction with expected revenue R, we construct a truthful deterministic auction with revenue $R/4 - 2h$ where h is the highest bid in the instance. Combined with the best known truthful randomized auctions, this implies the existence of a truthful deterministic auction with revenue at least $\mathrm{OPT}/13 - 2h$ where OPT is the optimal single price revenue. The main contribution here is a derandomization technique that preserves truthfulness (standard algorithmic derandomization techniques do not have this property). Unfortunately,

[2]However, most results can be generalized to larger demands so long as they are publicly known and the marginal return is constant.

our reduction is exponential. Therefore, we also give a specific polynomial-time truthful deterministic auction whose revenue is always at least $\mathrm{OPT}/4 - h$.

2.2.1 Setting

We focus on the single-unit demand case where buyers have utility for at most *one* unit of the good. In this case, the utility function of a bidder can be specified by a single parameter u (his value for one unit): the utility for allocation x is then $ux - p$ where p is the price charged to the bidder and x is an indicator variable for the event that the bidder was allocated a unit.

Furthermore, we assume that the auctioneer has access to an unlimited supply of the good. This assumption is merely for ease of exposition. In fact, any algorithm for the unlimited supply case can be easily extended to a situation in which just k units are available by restricting allocations to the k highest bidders and increasing the minimum price to the $(k+1)$'st highest bid.

Clearly, our results depend on the definition of the optimal revenue. With no restrictions, an omniscient auctioneer can extract a revenue equal to the sum of the bidders' values $\sum_{i=1}^{n} u_i$. However, one can not hope to compete with this truthfully, and so the following definition, which essentially restricts the omniscient auctioneer by forcing him to offer a single price to every bidder, proves to be a more useful measure of comparison [50].[3]

Definition 2 *Given the values* $\mathbf{u} = (u_1, \ldots, u_n)$ *of the agents, assume* $u_1 \geq \ldots \geq u_n$ *and let* k *be the index at which* ku_k *is maximized. Then* u_k *is the* optimum price *and* ku_k *is the* optimum revenue*, denoted* $\mathrm{OPT}(\mathbf{u})$*, or simply* OPT *when* $\mathbf{u}$ *is clear from the context.*

Unfortunately, this goal is impossible to approach with a truthful auction. The problem is that one bidder may dominate the market with a very high value. In this case, the optimal revenue extracts all its profit from this single bidder while an

[3] In fact, the two measures are also within a factor of $\log h$ of one another where $h = \max_i u_i$ is the highest bid in the input.

auction can not extract even a fraction of this profit bid-independently. By Theorem 2.1.1, bid-independent computations are required for truthfulness, and so this line of reasoning implies that the revenue of any truthful auction can not compete with the optimal revenue. Instead, we allow ourselves to lose a constant fraction of the highest bid h in our revenue approximation: we look for auctions that obtain a profit of at least $\mathrm{OPT}/\beta - \gamma h$ for small constants β and γ. We refer to β as the *approximation ratio* and γh as the *additive loss.*

Definition 3 *We say an auction is* (β, γ)-approximately optimal *if its expected profit on any input,* $\mathbf{b} \in [1, h]^n$, *is at least* $\mathrm{OPT}(\mathbf{b})/\beta - \gamma h$ *for fixed constants* β *and* γ.

We can equivalently formulate any approximately optimal result as a multiplicative approximation in which the approximation factor is a function of the *market dominance parameter.* Intuitively, the market dominance parameter bounds the amount a single bidder contributes to the optimal revenue. In the single-unit demand case, we define the market dominance parameter as the ratio of the maximum bid h to the optimal single price revenue OPT. If we are promised that the market dominance parameter is at most ϵ, then a (β, γ)-approximately optimal auction has revenue at least $(\beta - \epsilon\gamma)\mathrm{OPT}$. While in this section we state our results according in the form of Definition 3, in Section 2.3, we will find it more convenient to prove bounds which are a function of the market dominance parameter.

2.2.2 A Hat Puzzle

As indicated by the impossibility result of Goldberg, Hartline, and Wright [50], our auctions will need to compute the allocation and prices *asymmetrically* and, in order to be truthful, *bid-independently.* Still, we want the sum of the sale prices to approximate a global optimum, namely the optimal revenue. This raises the question of how a group of people may use asymmetry to coordinate convergence on a global objective when they have access to only partial information.

We study the problem of asymmetric coordination through a *hat puzzle.* In a hat puzzle, n players enter a room wearing hats. Each hat has one of k colors. No one can see his own hat or communicate with the other players, but each player can observe the colors of other players' hats. The objective varies depending on the version of the puzzle, but typically requires some players to guess the colors of their hats correctly. In the version of the puzzle we study here, the *everywhere-balanced k-color hat puzzle* (or *balanced k-hats* for short), the objective is for at least a $1/k$ fraction (rounded down) of the players wearing each color to correctly guess their hat color. For example, if Adam and Eve are wearing red hats, and Cain, Abel, and Seth are wearing blue hats, they would win if Adam or Eve guesses red and one of the remaining three guesses blue.

Hat puzzles of various forms have been contemplated in the mathematics community, partially due to their connections to coding theory, discrepancy problems, and autoreducibility of random sequences, and often simply because they make interesting brain-teasers [27, 32, 93, 117]. In this chapter, we exhibit a connection between hat puzzles and truthful mechanism design. We draw an analogy between agents' bids and players' hat colors, and then use solutions to the hat puzzle to compute an offering price for each agent by setting his price equal to the corresponding player's guess. As a player must guess his hat color without observing his own hat, the mechanism we design will be bid-independent and hence truthful.

Balanced k-hats is easy to solve in expectation if the players have access to randomization. Each player simply guesses each of the colors with equal probability $1/k$. We call this set of strategies the *randomized hat guessing algorithm.* In order to use a hats solution in our auction derandomization technique, we need to find a *deterministic* hat guessing algorithm. Notice that standard algorithmic derandomization techniques such as trying all possible coin flips and selecting the best, can not be implemented within the rules of the puzzle. Instead, we give a flow-based technique that uses the ordering (or names) of the players in place of randomness to solve the puzzle deterministically. It is instructive to view this technique as a derandomization of the simple randomized hat guessing algorithm proposed above. Feige [32] independently derived a similar flow-based construction.

First we define a bipartite graph representing the game (see Figure 2.2.2). The

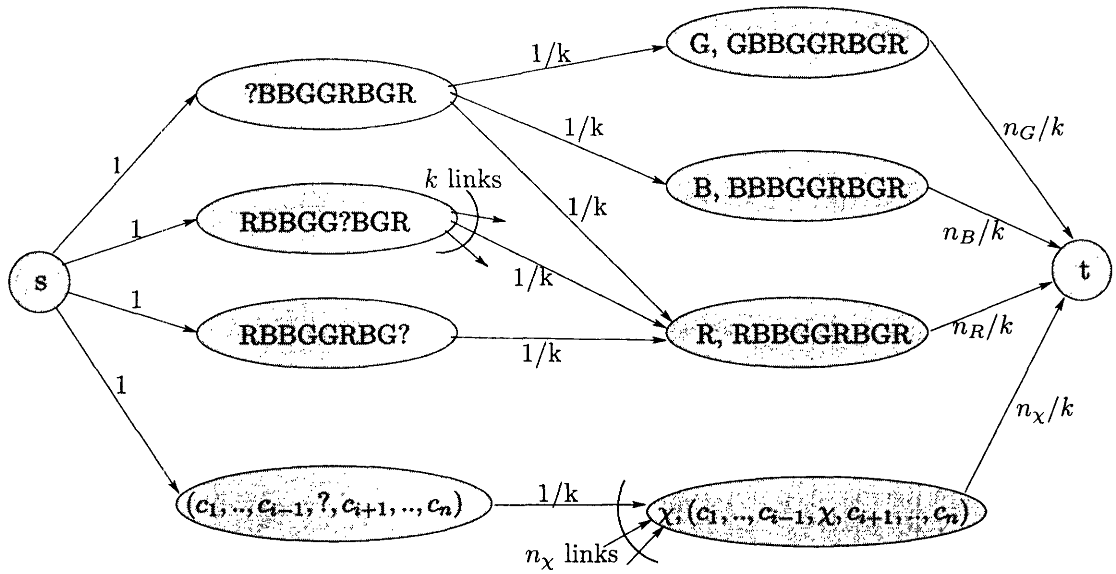

Figure 2-1: Hat puzzle graph

nodes are defined as follows. Each node on the left-hand side represents a possible viewpoint of player i. Let $\mathbf{c} = (c_1, \ldots, c_n)$ represent the array of colors and for any index i, $\mathbf{c}_{-i}$ represent the array of colors with the i'th color hidden. That is, $\mathbf{c}_{-i} = (c_1, \ldots, c_{i-1}, ?, c_{i+1}, \ldots, c_n)$ (by this definition $c_{-i} \neq c_{-j}$ for $i \neq j$). Note that $\mathbf{c}_{-i}$ is precisely the view of player i. For each of the nk^{n-1} possible values of $\mathbf{c}_{-i}$, we have a vertex $v_{\mathbf{c}_{-i}}$ on the left-hand side. Each node on the right-hand side represents a possible scenario (a hat color for each player) and a corresponding guess of some player. Let χ be one of the k colors. Then a scenario and corresponding guess is a pair $(\chi, \mathbf{c})$ (note this pair does not specify which player guesses χ). For each of the k^{n+1} possible values of the pair $(\chi, \mathbf{c})$, we have a vertex $v_{\chi,\mathbf{c}}$ on the right-hand side. We also include a source vertex s and sink vertex t. The arc set is defined as follows. We place an arc from the source s to each vertex on the left-hand side, and another arc from each vertex on the right-hand side to t. We also add an arc between $v_{\mathbf{c}_{-i}}$ and $v_{c_i,\mathbf{c}}$ signifying that we get $\mathbf{c}$ when we reveal that at position i in $\mathbf{c}_{-i}$ is a hat with color c_i. Notice that the in-degree due to such arcs of a vertex $v_{\chi,\mathbf{c}}$ is precisely the number $n_\chi(\mathbf{c})$ of hats of color χ in $\mathbf{c}$. The out-degree of a vertex $v_{\mathbf{c}_{-i}}$ is exactly k, one for each possible color of the hat at position i. The structure of the graph is sketched in Figure 2.2.2 (all arcs are directed from left to right in this figure).

Next, we set upper and lower capacities on the arcs as follows. For each χ and $\mathbf{c}$, we lower bound the flow of the arc $(v_{\chi,\mathbf{c}}, t)$ to $\lfloor n_\chi(\mathbf{c})/k \rfloor$ (recall $n_\chi(\mathbf{c})$ represents the number of hats in $\mathbf{c}$ that are colored χ). This represents the objective that at least $n_\chi(\mathbf{c})$ players should guess χ in scenario χ. For every other arc, we upper bound its flow by 1. This represents the requirement that each player can submit at most one guess.

We can interpret the randomized hat-guessing algorithm as a feasible flow on this graph. Between s and each $v_{\mathbf{c}_{-i}}$ place a flow of 1. This corresponds to the randomized algorithm, upon seeing $\mathbf{c}_{-i}$, having a total probability of 1 to spend on guessing a color for the i'th player's hat. On each of the outgoing arcs from $v_{\mathbf{c}_{-i}}$ we place a flow of $1/k$ corresponding to the probability with which the randomized algorithm picks each color (this is possible since each $v_{\mathbf{c}_{-i}}$ has k outgoing arcs). Now notice that the incoming flow to $v_{\chi,\mathbf{c}}$ is precisely $1/k$ times the number $n_\chi(\mathbf{c})$ of hats in $\mathbf{c}$ from color class χ. Thus by sending all of this flow on the outgoing arc to t, we satisfy all capacities. The flow is sketched in Figure 2.2.2 (the labels on the arcs represent the amount of flow on that arc).

Similarly, we can interpret an integral flow in the this graph as a deterministic hat-guessing algorithm: each player observes a vector $\mathbf{c}_{-i}$ of hat colors and identifies the corresponding node $v_{\mathbf{c}_{-i}}$ of the graph. If the integral flow sends flow from this node to $v_{\chi,\mathbf{c}}$, then the player guesses χ as his hat color (if the integral flow does not send flow through this node, then we let the player make a default guess). We now analyze the performance of such an algorithm on $\mathbf{c}$. Given $n_\chi(\mathbf{c})$ hats with color χ in $\mathbf{c}$, the lower bound on the capacity of the outgoing arc from $v_{\chi,\mathbf{c}}$ to t is $\lfloor n_\chi(\mathbf{c})/k \rfloor$. Therefore, it must be that $\lfloor n_\chi(\mathbf{c})/k \rfloor$ of the $n_\chi(\mathbf{c})$ incoming arcs have one unit of flow on them. For each such arc $(v_{\mathbf{c}_{-i}}, v_{\chi,\mathbf{c}})$, player i correctly guesses χ in the game setting $\mathbf{c}$. Thus, players guess the correct color χ for $\lfloor n_\chi(\mathbf{c})/k \rfloor$ positions out of a total of $n_\chi(\mathbf{c})$ such positions. This holds true for all colors χ, and so the players have solved the puzzle.

A classic result on integrality of network flows (see, for example, Hoffman's circulation theorem in the book by Schrijver [104, Theorem 11.2]) states that in a graph with integral capacities, if there is a feasible fractional flow, then there is a feasible integral flow. Therefore, the existence of the randomized hat-guessing algorithm implies the existence of a deterministic one. Unfortunately, it takes exponential time to construct an integral flow from a fractional one, and so our reduction is not polynomial-time constructive.

2.2.3 Auction Derandomization

By analogy to the hat-guessing technique of Section 2.2.2, we can show that *any* randomized auction has a deterministic counterpart that achieves approximately the same profit.

Theorem 2.2.1 *Corresponding to any single-round sealed-bid truthful auction $\mathcal{A}$ with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector $\mathbf{b}$, there is a deterministic truthful auction $\mathcal{A}'$ whose expected profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{A}(\mathbf{b})]/4 - 2h$, where $h = \max_i b_i$ is the highest bid.*

The proof of Theorem 2.2.1 reduces any (randomized) auction to a special type of (randomized) auction that we define, called a *guessing auction*, and then uses a flow-based construction similar to that in Section 2.2.2 to derandomize the guessing auction. The proof follows directly from Lemmas 2.2.1 and 2.2.2. Our proof

Guessing Auctions

The flow-based construction for balanced k-hats in Section 2.2.2 tries to reconstruct c_i from the vector $\mathbf{c}_{-i}$. In order to use this construction in the auction setting, we would like to draw an analogy between a player's hat color and an agent's bid, and between a player's guess for his hat color and an agent's price. However, an auction gets revenue not only when it charges a price equal to the bid value, but also when it charges a price below a bid value. In order to resolve this discrepancy, we define the notion of a *guessing auction* that uses only powers of two as prices and receives profit from a bidder only when it offers him a price equal to his bid rounded down to the nearest power of two (note that to preserve truthfulness a guessing auction may charge a bidder another price; however, the revenue generated from such bidders is not counted towards the profit of the guessing auction).

Guessing Auction $\mathcal{G}_{\mathcal{A}}$:

> Given an auction $\mathcal{A}$, simulate $\mathcal{A}$ on input bid vector $\mathbf{b}$. Suppose $\mathcal{A}$ offers bidder i price q_i and let 2^k be the largest power of two less than q_i. Then for integer $j \geq 0$, offer bidder i price p_i equal to 2^{k+j} with probability 2^{-j-1}.

Definition 4 *The* profit *of a guessing auction is the sum of those offering prices p_i for which* $\log p_i = \lfloor \log b_i \rfloor$.

It is possible to convert any truthful auction into a truthful guessing auction while only losing a factor of four from the profit. Denote the profit of a truthful auction $\mathcal{A}$ on input $\mathbf{b}$ as $\mathcal{A}(\mathbf{b})$. This profit is given by the sum of the prices charged to the winning bidders. For a randomized bid-independent auction $\mathcal{A}(\mathbf{b})$ is a random variable. The profit of a guessing auction is as defined above.

Lemma 2.2.1 *For any truthful auction $\mathcal{A}$ with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector $\mathbf{b}$, there is a corresponding truthful guessing auction $\mathcal{G}_{\mathcal{A}}$ whose expected profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{A}(\mathbf{b})]/4$.*

Proof. Given a bid-independent auction A, $\mathcal{G}_{\mathcal{A}}$ is bid-independent as well and therefore truthful by Theorem 2.1.1. We now bound the expected profit (in the sense of Definition 4) of $\mathcal{G}_{\mathcal{A}}$. Consider a bidder i who was offered price $q_i \leq b_i$ by $\mathcal{A}$. Thus the profit of $\mathcal{A}$ from i is q_i and the profit from all other bidders is zero. Thus, it is enough to show that the profit (in the sense of Definition 4) of $\mathcal{G}_{\mathcal{A}}$ from i is at least $q_i/4$. Let k be such that $2^k \leq q_i < 2^{k+1}$. Suppose bidder i's bid b_i is in the interval $[2^{k+j}, 2^{k+j+1})$. Then the probability that $\mathcal{G}_{\mathcal{A}}$ offers i price 2^{k+j} is 2^{-j-1}, and the profit from this offer is 2^{k+j}. Thus, the expected profit extracted from bidder i is $2^{j-1} \cdot 2^{k+j} = 2^{k-1}$. As $q_i < 2^{k+1}$ by assumption, the profit extracted from i is at least a quarter of his price q_i in $\mathcal{A}$. This completes the proof. $\square$

We note that if auction $\mathcal{A}$ only uses prices that are powers of two, then the profit of the corresponding guessing auction $\mathcal{G}_{\mathcal{A}}$ is actually within a factor of two of the profit of $\mathcal{A}$ instead of a factor of four.

The Flow Construction

We now show how to derandomize any guessing auction $\mathcal{G}_{\mathcal{A}}$. Our derandomization draws an analogy between bids and hat colors to deterministically compute prices for the guessing auction using the technique of Section 2.2.2.

Lemma 2.2.2 *Corresponding to any truthful guessing auction $\mathcal{G}_{\mathcal{A}}$ with expected profit $\mathbf{E}[\mathcal{G}_{\mathcal{A}}(\mathbf{b})]$ on bid vector $\mathbf{b}$, there is a truthful deterministic auction whose profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{G}_{\mathcal{A}}(\mathbf{b})] - 2h$, where h is the highest bid value in $\mathbf{b}$.*

Proof. First, round all bid values down to the nearest power of two. We draw an analogy between the k colors in the balanced k-hats puzzle and the $\log h$ powers of two that are the possible (rounded) values of bids. Set up a flow construction identical to that for the balanced k-hats puzzle, except that the fractional flow on an arc from $v_{\mathbf{b}_{-i}}$ to $v_{2^j,\mathbf{b}}$ is the probability that $\mathcal{G}_{\mathcal{A}}$ on seeing $\mathbf{b}_{-i}$ guesses 2^j. Furthermore, the flow from $v_{2^j,\mathbf{b}}$ to t is the expected number of times $\mathcal{G}_{\mathcal{A}}$ offers one of the bidders with (rounded) bid 2^j a price equal to 2^j. We represent this quantity by $E_j(\mathbf{b})$. We then set the capacities as before: we require flow on an arc between $v_{2^j,\mathbf{b}}$ and t to be at least $\lfloor E_j(\mathbf{b}) \rfloor$ and all other arc flows to be at most 1. Once again, the above fractional flow implies the existence of an integer-valued flow [104, Theorem 11.2], and this integer-valued flow corresponds to an auction making a deterministic bid-independent offer upon seeing $\mathbf{b}_{-i}$. The flow out of $v_{2^j,\mathbf{b}}$ is precisely the number of indices i such that the auction, upon seeing $\mathbf{b}_{-i}$, correctly guesses 2^j; since the flow is feasible, the flow out of $v_{2^j,\mathbf{b}}$ is at least $\lfloor E_j(\mathbf{b}) \rfloor$. Thus, considering a bid vector $\mathbf{b}$ where the expected profit of $\mathcal{G}_{\mathcal{A}}$ is $\mathbf{E}[\mathcal{G}_{\mathcal{A}}] = \sum_{j=1}^{\log h} 2^j E_j(\mathbf{b})$, the deterministic auction obtains $\sum_{j=1}^{\log h} 2^j \lfloor E_j(\mathbf{b}) \rfloor \geq \sum_{j=1}^{\log h} [2^j E_j(\mathbf{b}) - 2^j] \geq \mathbf{E}[\mathcal{G}_{\mathcal{A}}] - 2h$. □

As a corollary of Theorem 2.2.1, known approximately-optimal randomized auctions [37, 46, 50, 54] imply the existence of approximately-optimal deterministic auctions. Using a the best known randomized auction from Hartline and McGrew [54], we obtain the following result.

Theorem 2.2.2 *There is a deterministic auction whose revenue is at least* $\mathrm{OPT}/13 - 2h$ *in the worst case.*

Proof. The theorem follows from Theorem 2.2.1 and the Hartline-McGrew auction whose competitive ratio is 3.25. □

In this construction, we assumed that the range of bid values $[1, h]$ is known. This assumption is not necessary. When considering $\mathbf{b}_{-i}$, we can compute h, which is the maximum bid value scaled such that the minimum bid value is 1 on the new scale, correctly for all but the minimum and maximum bid value. Assuming the worst, that is the auction fails to get any profit from the highest and lowest bid, we only lose an additional additive $h + 1$.

2.2.4 A Polynomial-time Deterministic Auction

Unfortunately, the flow construction used to derandomize auctions in Section 2.2.3 had exponential size. Therefore the derandomized auctions of Section 2.2.3 are not efficiently computable. In this section, we describe a competitive deterministic asymmetric auction, the outcome of which can be computed in polynomial time. In particular, our deterministic auction guarantees a revenue which is at least $\mathrm{OPT}/4 - h$ where OPT the optimal single price revenue and h is the highest bid.

There are three key ingredients in this auction: (a) a method, called a *profit extractor*, for extracting a given feasible target revenue truthfully; (b) a method, called a *consensus estimator*, for each bidder to bid-independently compute the same feasible target revenue; and (c) a deterministic coin-flipping algorithm. To see how these pieces fit together, first suppose we knew the optimal revenue R. Could we then design an auction to truthfully recover revenue R? The goal of the profit extractor is to do just that: given bids $\mathbf{b}$, a profit extractor truthfully extracts a target revenue R from some subset of the bidders. Although this mechanism is truthful and extracts the optimal revenue deterministically, it requires the optimal revenue as an input. We can not hope to compute the optimal revenue bid-independently; rather we compute n bid-independent estimates of the optimal revenue, one for each bidder. If these estimates are coordinated appropriately (namely, if sufficiently many bidders compute the same estimate), then we can use these estimates as inputs to our profit extractor and generate sufficient revenue.

The profit extractor and consensus estimator were used previously by Goldberg and Hartline [46] along with a single random coin flip to get an approximately optimal randomized auction. Even though their consensus estimator uses just one bit of randomness, it is difficult to derandomize using standard techniques from randomized algorithms as they tend not to be bid-independent. For example, one might consider derandomizing the auction of Goldberg and Hartline [46] by running it twice – once with "heads" and once with "tails" – and outputting the higher-revenue solution. However, this clearly can not be represented bid-independently and therefore is not truthful. Instead, as the final ingredient of our construction, we design a deterministic coin flip which can be computed bid-independently and use this to derandomize the auction of Goldberg and Hartline [46].

Profit Extractor

The profit extractor we present here is a special case of a general cost-sharing scheme due to Moulin [85] (see Chapter 4).

Mechanism ProfitExtract_R:

> Given bids $\mathbf{b}$, find the largest k such that the highest k bidders can equally share the cost R (that is, each of their bids is at least R/k). These bidders are the *winners* and the rest are *losers*. Charge each of the winners R/k and the losers 0. If no such k exists, then all bidders are *losers* and are charged price 0.

As we will base our deterministic auction on this mechanism, it is important to note that it is truthful and actually extracts revenue R.

Lemma 2.2.3 ProfitExtract$_R$ *is truthful.*

Proof. We define a bid-independent implementation of ProfitExtract$_R$. Let the bid-independent function $\text{pe}_{i,R}(\mathbf{b}_{-i})$ equal $\frac{R}{l+1}$ where l is the largest number such that the highest l bidders in $\mathbf{b}_{-i}$ all have a bid at least $\frac{R}{l+1}$. If no such l exists, let $\text{pe}_{i,R}(\mathbf{b}_{-i})$ be ∞, a number larger than any bid. Let k be the number of winners in ProfitExtract$_R$. Then, by definition, each of these k winners has a bid at least $\frac{R}{k}$ while each of the losers in ProfitExtract$_R$ has bid strictly less than $\frac{R}{k+1}$. Thus, for a winner i, $\text{pe}_{i,R}(\mathbf{b}_{-i})$ equals $\frac{R}{k}$, i's price in ProfitExtract$_R$. For a loser j, $\text{pe}_{i,R}(\mathbf{b}_{-j})$ equals $\frac{R}{k+1}$, implying that j is a loser in the bid-independent implementation as well. □

Lemma 2.2.4 *If* $R \leq \text{OPT}(\mathbf{b})$, ProfitExtract$_R(\mathbf{b}) = R$*; otherwise it has no winners and no revenue.*

Proof. The proof is immediate from the definition of the mechanism and the optimal price OPT. □

Consensus Estimator

The goal of the consensus estimator is to compute an estimate of the optimal revenue for each bidder bid-independently. A *pair of consensus estimators* is a pair of functions, r_0 and r_1, having the following properties:

1. For any real number V, there exists an $r \in \{r_0, r_1\}$ such that for all $v \in [V/2, V]$, $r(v) = r(V)$. This r is called a *consensus* on V.

2. For any real number V and $r \in \{r_0, r_1\}$ that is a consensus on V, $r(V) \in [V/2, V]$. In this case, r is said to *estimate* V.

It is easy to see that the following functions form such a pair of consensus estimators [46].

$r_0(v) = 2v$ rounded down to the nearest *even* power of two.

$r_1(v) = 2v$ rounded down to the nearest *odd* power of two.

We will apply these consensus estimators to the value $\mathrm{OPT}(\mathbf{b}_{-i})$ in order to obtain a consensus on an approximate value for $\mathrm{OPT}(\mathbf{b})$.

Deterministic Coin-Flipping Algorithm

Our deterministic coin-flipping algorithm is best described via an analogy to another hat puzzle. This time, each of the hats is colored a distinct shade of red.[4] We would like each of the players to simulate a coin flip with the collective property that, for any particular shade of red, at least half the players with darker hats choose heads and at least half choose tails (rounding down). We call such an algorithm *balanced.*

Clearly, the randomized algorithm that instructs every player to simply flip a coin is balanced in expectation. The algorithm we are about to present achieves this property deterministically. In fact, our solution satisfies the following stronger property: lining the players up from darkest hat to lightest hat, the sequence of coin choices alternates. We call such a set of choices *perfectly alternating.* Our algorithm is based on the notion of the *sign of a permutation.*

Definition 5 *Given a vector of n hat shades, $\mathbf{c}$, the* sign of $\mathbf{c}$ *(shorthand for "the sign of the permutation of the ordering of hats") is the parity of the number of transpositions of adjacent hats it takes to sort $\mathbf{c}$, notated* $sign(\mathbf{c})$.[5]

The *deterministic coin flip algorithm,* ϕ works as follows. Each player has an identity i – this defines the fixed ordering used to compute the coin flips. Given $\mathbf{c}_{-i}$ as the shades of the hats that player i sees, player i computes his coin, $\phi(\mathbf{c}_{-i})$, by

[4] If hats are not distinct shades, then we can break ties according to the identities of the players.

[5] While the number of transpositions performed in sorting $\mathbf{c}$ is not unique, the parity of the number of transpositions is. See, for example, [58]

imagining that his own hat is the darkest shade, denoted ∞. As his coin flip, he chooses the sign of his imagined vector of hat colors, denoted $(\mathbf{c}_{-i}, \infty)$, whose i'th entry is ∞ and j'th entry for $j \neq i$ is c_j. To prove that this algorithm is balanced, we will use the notion of the *rank* of a player.

Definition 6 *Given a vector of n hat shades, $\mathbf{c}$, the* rank of i, *denoted $rank(\mathbf{c}, i)$, is the number hats in $\mathbf{c}$ that are darker than c_i.*

Lemma 2.2.5 *The deterministic coin flip algorithm, ϕ, is perfectly alternating with the shades of the hats' colors.*

Proof. This result is implied by the fact that

$$\phi(\mathbf{c}_{-i}) \equiv sign(\mathbf{c}) + rank(\mathbf{c}, i) \pmod 2,$$

which is evident because one way to sort $(\mathbf{c}_{-i}, \infty)$ would be to first sort $\mathbf{c}$ and then replace hat i with ∞ which would require $rank(\mathbf{c}, i)$ additional transpositions to move ∞ to the front of the vector. As the parity of the ranks alternate, this implies the lemma. □

In this solution to the deterministic coin-flipping problem, each player can compute his own coin by simply executing ϕ; however, no player can compute the coin of anyone else as he does not know his own hat shade. Clearly, each player can compute his coin in $O(n \log n)$ time; furthermore, as is evident from the above proof, the coins of all the players can be computed in $O(n \log n)$ time.

Remark 2.2.1 *Note that we can reduce the balanced k-hats puzzle introduced in Section 2.2.2 for $k = 2$ to the problem of deterministic coin flipping as follows. Run the algorithm with the two colors – light red and dark red. Interpret a* heads *coin as "light red" and a* tails *coin as "dark red". The resulting algorithm, modulo rounding, guesses half of the light red hats and half of the dark red hats correctly.*

Deterministic Auction

Finally, we have developed the tools necessary to describe our efficient auction, the deterministic consensus revenue estimate (DCORE) auction. This auction is built from the three components discussed above – the profit extractor, the pair of consensus estimators, and the deterministic coin flipping algorithm. The auction uses the deterministic coin flipping algorithm to pick a consensus estimator and corresponding estimate R_i for each bidder i. It then offers bidder i a price equal to the price computed for i by the profit extractor on input R_i.

Definition 7 *DCORE is the bid-independent auction implemented by the following functions, f_i:*

$$f(\mathbf{b}_{-i}) = \mathrm{pe}_{i,R_i}(\mathbf{b}_{-i}),$$

where $R_i = r_{\phi(\mathbf{b}_{-i})}(\mathrm{OPT}(\mathbf{b}_{-i}))$ and pe_{i,R_i} is the bid-independent function for the mechanism $\mathrm{ProfitExtract}_R$ defined in the proof of Lemma 2.2.3.

DCORE is bid-independent and therefore truthful. We now show that DCORE is approximately optimal. Our proof is similar to the proof of the revenue of the CORE auction in Goldberg and Hartline [46].

Theorem 2.2.3 *The profit of DCORE is at least $\mathrm{OPT}/4 - h$.*

Proof. If the optimal single price sale has exactly one winner, then the optimal revenue is h and approximating it within an additive h is trivial. Otherwise, let $\mathrm{OPT} = \mathrm{OPT}(\mathbf{b})$ be the revenue from the optimal single price sale. Then, for every i, we have $\mathrm{OPT}/2 \leq \mathrm{OPT}(\mathbf{b}_{-i}) \leq \mathrm{OPT}$. Since r_1 and r_0 are a pair of consensus estimates, one of them is a consensus on OPT. Suppose, without loss of generality, that it is r_0. Then $r_0(\mathrm{OPT}(\mathbf{b}_{-i})) = r_0(\mathrm{OPT})$ for all i. Now consider the following thought experiment. Suppose we had chosen consensus function r_0 for all i and so $R_i = r_0(\mathrm{OPT}) \in [\mathrm{OPT}/2, \mathrm{OPT}]$ for all i. In this case, our auction is equivalent to the profit extraction mechanism on input $r_0(\mathrm{OPT})$. Let p be the price charged to the k winning bidders in this thought experiment. Note, as $r_0(\mathrm{OPT}) \leq \mathrm{OPT}$, the total profit, pk, must be $r_0(\mathrm{OPT}) \in [\mathrm{OPT}/2, \mathrm{OPT}]$ by Lemma 2.2.4. In reality, in the deterministic coin-flipping procedure, at least $k/2 - 1$ of these k bidders computed $\phi(\mathbf{b}_{-i}) = 0$ and thus these bidders all pay p, exactly as they would have in the thought experiment. The total profit accounted for is $pk/2 - p \geq \mathrm{OPT}/4 - h$, which proves the theorem. □

2.3 Budget Constraints

In Section 2.2, we considered buyers who have limited demand. Here, we consider buyers with budget constraints. Budget constraints are a central feature of many real auctions. In the context of e-commerce, there is a great deal of interest in multi-unit auctions of relatively low-value goods, such as the auction of Internet ads for search terms and content pages on MSN, Google, Yahoo, etc., to bidders with budget constraints.

The theoretical framework of budget-constrained auctions is currently substantially less well-developed than that of unconstrained auctions – which is unsatisfactory both from a theoretical viewpoint, and from a practical viewpoint, where the absence of an appropriate framework leads to losses in revenue and efficiency. It is therefore of tremendous interest to design truthful mechanisms for budget-constrained auctions. Existing mechanisms for the budget-constrained case typically assume a distribution on the budgets and values of bidders and use these assumptions to compute high revenue auctions [11, 17, 18, 30, 31, 76, 79, 120]. In this section, we instead follow the framework of Goldberg, Hartline, and Wright [50] outlined in Section 2.2 to design a (randomized) auction with high revenue in the worst case. It is unknown whether the results of the last section can be modified to derandomize this auction. We also partially characterize the conditions under which deterministic truthful auctions exist.

2.3.1 Setting

Although the setting in this section is quite similar to that of Section 2.2, it differs in several significant respects. First, whereas in the last section solutions to the unlimited supply case implied solutions to the limited supply case, in this setting, maximizing revenue with unlimited supply is trivial: simply offer each unit for a very small fixed price and every buyer will exhaust his budget. Thus, we consider the *limited supply* case.

Second, whereas in the last section each agent submitted a single parameter to the mechanism, namely his value, now our mechanism must now solicit a *two*-parameter bid from each agent. The first parameter is interpreted as that agent's announced value per unit and the second parameter is that agent's announced budget. In order for our mechanism to be truthful, it must be the case that the agent's utility is maximized by reporting *both* parameters truthfully.

Finally, our utility function is no longer quasi-linear. Instead, each agent i has a private value $u_i \in \mathbb{R}^+ \cup \{0\}$ per unit of the good and a private budget constraint $b_i \in \mathbb{R}^+ \cup \{0\}$. The budget constraint is a *hard* constraint, that is the agent cannot spend more than his budget under any circumstances. In other words, the total utility $u_i(j, p)$ that agent i derives from an allocation of j units at a total price of p is:

$$u_i(j,p) = \begin{cases} ju_i - p & \text{if } p \leq b_i, \\ -\infty & \text{if } p > b_i. \end{cases}$$

The value $-\infty$ in the above definition means that this agent prefers receiving nothing and paying nothing than participating in any lottery with a non-zero risk of going over the budget.

This final distinction summarizes the most significant departure of this setting from the one in Section 2.2. Many of the results in the auction theory literature are not applicable when utility functions fail to be quasi-linear. In particular, the classical Vickrey-Clarke-Groves (VCG) mechanisms [115, 19, 52], which in the limited demand case yielded truthful mechanisms but with low revenue, are no longer even truthful. This fact is illustrated in the following example.

Example 2.3.1 One plausible mechanism for auctioning m units of a good to budget-constrained buyers is to apply the VCG mechanism assuming that the value of agent i for j units of the good is $\min(b_i, ju_i)$. A common mistake is to assume that since this mechanism is based on VCG, it is truthful. The following example shows that this is not the case: assume we have two units of the good to sell to two agents, and the truthful bids of these agents are given by $(u_1, b_1) = (10, 10)$ and $(u_2, b_2) = (1, 10)$. The above mechanism assumes that the value of the first agent for either one or two units of the good is 10, and therefore allocates one unit to each agent to maximize the total value (which is $10 + 1$). The payment charged to the agents by this mechanism is 1 and 0, respectively. Therefore, the utility of the first agent is 9. However, if the first agent announces the bid $(5, 10)$, then the mechanism will allocate both units to this agent at a total price of 2. Thus, the first agent would achieve a utility of 18 by bidding untruthfully. This example shows that the above VCG-based mechanism is not truthful even if the agents are not allowed to lie about their budget.

In fact, it is easy to observe that in this setting, no truthful mechanism can always produce an *efficient* allocation, that is an allocation that maximizes the social welfare, even when there is only one unit. The reason for this is that an efficient mechanism should always allocate the good to the bidder with the highest u_i, even if such a bidder has a zero budget and therefore cannot be charged any positive amount. Therefore, any agent can bid a high value and zero budget to get the good for free. This simple impossibility result shows that we cannot require efficiency from a truthful mechanism. In contrast, in the limited demand case, the classic VCG mechanism calculates an efficient allocation truthfully.

Despite the significant differences between this setting and the one in Section 2.2, we are still able to employ a framework similar to that of Goldberg, Hartline, and Wright [50] to design a revenue-maximizing auction. As in the last section, we compare our revenue to the revenue of the optimal single price sale. For any price p, we denote by $r(p, k)$ the revenue of allocating *at most* k units at price p.

Definition 8 *Suppose there are m units of the good for sale. Given the values and budgets of the agents, let p^* be the price at which $r(p, m)$ is maximized. Then p^* is* *the* optimum price *and* $r(p^*, m)$ *is the* optimum revenue, *denoted* OPT.

As before, a particular bidder may dominate the market, making it impossible to design a truthful auction which always receives a constant fraction of the optimum revenue. However, now a bidder dominates the market by having a high *budget* as opposed to a high value. Thus, the market dominance parameter is now defined as the ratio of the maximum budget of all agents, b_{max}, to the optimum revenue OPT.

The mechanism we design has the property that its revenue approaches that of the optimum single-price auction as the market dominance parameter tends to 0. In particular, we will prove that for all $0 < \delta < 1$, the revenue of our mechanism is at least a $(1-\delta)$ fraction of the optimum posted-price revenue with probability at least $1 - O(e^{-c\delta^2/\epsilon})$ where c is some constant and ϵ is the market dominance parameter.

2.3.2 An asymptotically optimal auction

Our mechanism is quite natural: similar to the random sampling optimal threshold auction of [49, 50], we divide bidders into two random subsets, compute the optimal price for each subset, and offer that price to the other subset. In order to guarantee that our auction doesn't oversell the good, we sell at most half the available units to each subset, greedily allocating units to interested agents arranged in an arbitrary order.

Note, although our units are indivisible, we can assume that fractional allocations are possible by using the proper randomization: whenever the algorithm asks us to allocate a fraction c of a unit to an agent, we instead charge the agent c times the offering price for participation in a lottery that offers him a full unit with probability c. Thus an agent's payment is deterministic and less than his budget, and his expected utility is constant. Only his allocation is randomized. For the remainder of this section, we assume without loss of generality that our units are divisible.

Let n be the number of agents and m be the number of available units of a good. Each agent i submits his value u_i per unit of the good and his maximum budget b_i.

Mechanism

- Partition the agents randomly into two sets A and B by independently putting each agent into either set uniformly at random with probability $\frac{1}{2}$.
- From the set of values u_i of agents $i \in A$, choose p_A to be the price which maximizes the revenue of selling at most $m/2$ units in A. In other words, if the u_i's are sorted in decreasing order so that $u_1 \geq u_2 \geq \cdots \geq u_n$, then

$$i_0 = \min \left\{ i : \sum_{j=1}^{i-1} b_j \geq \frac{u_i m}{2} \right\},$$

  define $p_A = u_{i_0-1}$. Compute p_B analogously.
- Consider the agents in A in an arbitrary order and allocate at most $\frac{m}{2}$ units to them as follows. In every step, if the per-unit value of agent i satisfies $u_i \geq p_B$, allocate $\frac{b_i}{p_B}$ units to i, or all remaining units if less than $\frac{b_i}{p_B}$ units remain. Charge i a price of p_B per unit. Apply the same procedure to the set B using the threshold value p_A.

First, we give a simple proof of the truthfulness of the mechanism.

Lemma 2.3.1 *The above mechanism is truthful, that is for every agent it is a dominant strategy to report the correct value and budget.*

Proof. Consider an agent i in A. First we argue that agent i does not have any incentive to misreport his value. We know that agent i receives a unit only if $u_i \geq p_B$, and that he pays p_B for a unit if he receives it. The two key observations are that (1) the threshold p_B is determined independently of all u_j and b_j, $j \in A$, including $j = i$; and (2) when the supply of units in A is inadequate to meet the demands of all agents in A whose values exceed p_B, then the allocation of units to those agents is done in an arbitrary order, again independently of all u_j and b_j.

Finally, by reporting a budget below b_i, agent i would potentially decrease his allocation and hence his total utility. By reporting a budget above b_i, if he was previously saturating his budget, then his allocation might increase causing him to be charged more than his budget and decreasing his total utility to negative infinity. Otherwise, if he was not previously saturating his budget, his allocation will not change. In either case, he has no incentive to misreport. □

The truthfulness of this mechanism is straightforward, but providing a revenue guarantee requires a more careful analysis. Our proof is fairly natural. We first show that, using a price that is at least the optimum posted price and disregarding supply limits, the revenue extracted from each random subset of bidders is approximately equal. We will use this to claim that the revenue extracted by our mechanism from each subset is almost half the optimum.

First, we introduce some notation. Recall that $r(p, k)$ is the revenue received by allocating at most k units at price p. We now extend this definition to subsets of bidders: for any price p, we denote by $r_S(p, k)$ the revenue of allocating at most k units to a set S of agents at price p:

$$r_S(p, k) = \min \left\{ kp, \sum_{j \in S | u_j \geq p} b_j \right\}.$$

Finally, we also define $r(p) = r(p, \infty) = \sum_{u_i \geq p} b_i$.

In our argument, we will use the following properties of an optimum posted-price auction for allocating at most k units, for any k.

1. There exists an agent i such that selling the units at price $p = u_i$ results in the optimum revenue.
2. For any k, if p is the optimum price for allocating at most k units, then $r(p, k) \leq r(p) \leq r(p, k) + b_{max}$ where $b_{max} = \max_i b_i$. In particular, OPT $\leq r(p^*) \leq$ OPT $+ b_{max}$.

Let ϵ denote the *market dominance parameter* (the ratio of the maximum budget of all agents, b_{max}, to the optimum revenue OPT). As we will show, the probability of success of our algorithm is asymptotically controlled by ϵ. The next lemma shows that the revenue extracted from each subset at or above the optimum price is approximately equal, disregarding supply limits, with probability approaching 1 as ϵ approaches zero.

Lemma 2.3.2 *Let* $\delta > 0$. *Then the probability that*

$$\Big|r_A(u_l) - r_B(u_l)\Big| < \delta\mathrm{OPT} \quad \text{for all} \quad l \quad \text{with} \quad u_l \geq p^*$$

is at least $1 - 2e^{-\delta^2/(4\epsilon)}$.

Proof. Define α_i to be a random variable indicating whether agent i is in A, with $\alpha_i = 1$ when $i \in A$ and $\alpha_i = -1$ when $i \in B$. Let $S_i = \sum_{j \leq i} \alpha_j b_j$. Then $|r_A(u_l) - r_B(u_l)| = |S_l|$. Thus we need to bound the probability that the random variable S_i deviates by more than $\delta\mathrm{OPT}$ from its expectation 0.

Let $\tau(\delta) = \min\{i : |S_i| \geq \delta\mathrm{OPT}\}$. We define the following martingale:

$$\tilde{S}_i = \begin{cases} S_i & \text{if } i \leq \tau(\delta) \\ S_{\tau(\delta)} & \text{otherwise} \end{cases}$$

Let k be such that $u_k = p^*$. Then we have

$$\begin{aligned} 1 - \mathbf{Pr}\Big(|r_A(u_i) - r_B(u_i)| < \delta\mathrm{OPT},\ \forall i \leq k\Big) \\ &= 1 - \mathbf{Pr}\Big(|S_i| < \delta\mathrm{OPT},\ \forall i \leq k\Big) \\ &= \mathbf{Pr}\Big(\exists\, i \leq k \ : \ |S_i| \geq \delta\mathrm{OPT}\Big) \\ &= \mathbf{Pr}\Big(\tau(\delta) \leq k\Big) \\ &= \mathbf{Pr}\Big(|\tilde{S}_k| \geq \delta\mathrm{OPT}\Big) \end{aligned}$$

Now since $\tilde{S}_i$ is a martingale, by the Azuma-Hoeffding inequality we have:

$$\mathbf{Pr}(|\tilde{S}_k| \geq \delta\mathrm{OPT}) \leq 2\exp(\frac{-\delta^2\mathrm{OPT}^2}{2\sum_{i \leq k} b_i^2}).$$

Bounding the sum $\sum_{i \leq k} b_i^2$ by $b_{max} r(p^*) \leq \epsilon\mathrm{OPT} r(p^*)$ and using the inequality $r(p^*) \leq \mathrm{OPT}(1 + \epsilon) \leq 2\mathrm{OPT}$, we obtain the lemma. □

From now on, we will say that an event happens with high probability if its probability is at least $1 - 2e^{-\delta^2/(4\epsilon)}$. From the previous lemma, it is clear that the revenue of each subset at the optimum price is almost half the optimum revenue with high probability, disregarding supply limits. In fact, it is not hard to see that this statement holds even *observing supply limits.*

Corollary 2.3.1 *With high probability, we have*

$$r_A(p^*, \frac{m}{2}) \geq \frac{1-\delta}{2}\text{OPT}.$$

Proof. First note $\text{OPT} = r(p^*, m) = \min\{p^* m, r(p^*)\}$, implying $mp^* \geq \text{OPT}$ and $r(p^*) \geq \text{OPT}$. From Lemma 2.3.2, we have with high probability, $r_A(p^*) \geq r(p^*) - \frac{\delta}{2}\text{OPT} \geq \frac{1-\delta}{2}\text{OPT}$. Furthermore, $\frac{m}{2}p^* \geq \frac{1-\delta}{2}\text{OPT}$ by definition of OPT, so $r_A(p^*, \frac{m}{2}) = \min(\frac{m}{2}p^*, r_A(p^*)) \geq \frac{1-\delta}{2}\text{OPT}$. □

At this point we would be done if our mechanism computed the offering price p^*. Unfortunately, we can not compute this price. Instead we compute prices p_A and p_B, the optimal prices for subsets A and B, and offer these prices to the opposing set. Thus we need to a prove statement similar to that of Lemma 2.3.2 for all offering prices. The following corollary states that for *any* offering price, the revenue of one subset is either close to the revenue of the other subset or close to half the optimum revenue, disregarding supply limits.

Corollary 2.3.2 *With high probability, we have that*

$$r_B(u_k) \geq \min\{r_A(u_k) - \delta\text{OPT}, \frac{1-\delta}{2}\text{OPT}\} \qquad \text{for all} \quad k.$$

Proof. By Lemma 2.3.2, we have that with high probability,

$$r_B(u_k) \geq r_A(u_k) - \delta\text{OPT} \quad \text{for all } k \text{ with } u_k \geq p^*. \tag{2.1}$$

Recalling that $r_A(p^*) + r_B(p^*) = r(p^*) \geq \text{OPT}$, we conclude that with high probability, both (2.1) and

$$r_B(p^*) \geq \frac{1-\delta}{2}\text{OPT}$$

hold simultaneously. By monotonicity, this implies the statement of the lemma. Indeed, either $u_k \geq p^*$ so that $r_B(u_k) \geq r_A(u_k) - \delta\text{OPT}$, or $u_k \leq p^*$ and $r_B(u_k) \geq r_B(p^*) \geq \frac{1-\delta}{2}\text{OPT}$, which gives the lemma. □

Finally, we have developed all the necessary machinery to prove the main theorem.

Theorem 2.3.1 *The mechanism described in the previous section is truthful. Furthermore, for all* $0 < \delta < 1$*, the algorithm has revenue at least* $(1-\delta)\text{OPT}$ *with probability* $1 - O(e^{-c\delta^2/\epsilon})$ *for some constant* c *and* $\epsilon = b_{max}/\text{OPT}$.

Proof.

Recall that p_A is the price which maximizes the revenue of selling at most $\frac{m}{2}$ units in A. Thus, for all p, we have $r_A(p, \frac{m}{2}) \leq r_A(p_A, \frac{m}{2})$, so in particular $r_A(p_A, \frac{m}{2}) \geq r_A(p^*, \frac{m}{2})$. Combined with Corollary 2.3.1, we conclude that with high probability, $r_A(p_A, \frac{m}{2}) \geq \frac{1-\delta}{2}\text{OPT}$, which in turn implies that

$$p_A \frac{m}{2} \geq \frac{1-\delta}{2}\text{OPT} \tag{2.2}$$

and

$$r_A(p_A) \geq \frac{1-\delta}{2}\text{OPT}. \tag{2.3}$$

Combined with Corollary 2.3.2, inequality 2.3 gives

$$r_B(p_A) \geq \frac{1-3\delta}{2}\text{OPT}, \tag{2.4}$$

again with high probability. Inequalities 2.2 and 2.4 together with the definition of $r_B(p_A, \frac{m}{2})$ imply that with high probability,

$$r_B(p_A, \frac{m}{2}) \geq \frac{1-3\delta}{2}\text{OPT}.$$

Exchanging the roles of A and B, we get the same result for $r_A(p_B, \frac{m}{2})$. Since

$r_B(p_A, \frac{m}{2}) + r_A(p_B, \frac{m}{2})$ is the revenue of the algorithm, this completes the proof of the theorem. $\square$

2.3.3 Impossibility Result

Although the auction given in Section 2.3.2 is truthful and asymptotically revenue-maximizing, it behaves strangely in some scenarios. First, it does not allow any bidder to receive more than $m/2$ units of the good. Ideally, if a bidder bids high enough he will be able to receive all the available units. Second, the allocation of an individual is affected by the bids of the losing agents. If losing agents ceased to bid, or submitted negligible bids, the allocation of other agents might change drastically. It is not surprising that the prices change with the bids of the losing agents, but ideally the allocation should be independent of their bids. In this section, we explore the space of truthful mechanisms proving that, in some sense, these violations were essential.

In particular, we prove that there is *no* truthful mechanism satisfying three properties defined below, even if there are only two buyers and two units of the good. This result automatically generalizes to auctions with more buyers, by considering the situation where all but two of the buyers bid zero. Our result also extends to randomized auctions that are strategyproof in the following stronger sense: no matter what the outcome of the coin flips are, it is a dominant strategy for the participants to reveal their true type. The randomized algorithm given in Section 2.3.2 is strategyproof in this sense only if the good is assumed to be divisible.

The first property is the following. This is similar to a property with the same name defined by Moulin [85] in the context of group-strategyproof mechanisms for cost sharing problems.

- *consumer sovereignty* – For any agent i and any vector of bids (u_{-i}, b_{-i}) for other agents, there is a bid (u_i, b_i) such that if agents bid according to (u, b), then agent i receives all units of the good.

Intuitively, consumer sovereignty requires that each agent must be able to win all units if he bids *high enough.* This precludes trivial mechanisms that for example sell at most one unit to each bidder.

The second property, which we call the *independence of irrelevant alternatives* (IIA), is a much weaker version of a property of the same name in Lavi et al. [77]. This property is defined as follows.

- *independence of irrelevant alternatives (IIA)* – For any agent i and a bid vector $(\mathbf{u}, \mathbf{b})$, if i receives no units at $(\mathbf{u}, \mathbf{b})$, then the allocation when every agent bids according to $(\mathbf{u}, \mathbf{b})$ is the same as the allocation when agent i bids $(0, 0)$ and others bid according to $(\mathbf{u}_{-i}, \mathbf{b}_{-i})$.

Intuitively, the above property states that if an agent who does not win the auction leaves, the allocation to other agents should not change (their payment, however, might change). As we will see in the proof of Theorem 2.3.2, in the case of two buyers and two units, IIA is equivalent to the property that if bids of both agents are large enough (both the value and the budget), then both units are allocated.

As we will see at the end of this section, there are truthful mechanisms not satisfying the IIA. In fact, the following example shows that even with IIA, there are mechanisms that are truthful.

Example 2.3.2 *Bundling mechanism:* Consider the mechanism that always *bundles* the two units, that is it allocates both units to the agent i such that $\min(2u_i, b_i) > \min(2u_{3-i}, b_{3-i})$, and charges him $\min(2u_{3-i}, b_{3-i})$. It is easy to see that this mechanism is truthful and satisfies the IIA.

However, we conjecture that the bundling mechanism is essentially the only truthful mechanism satisfying the above properties. In other words, we would like to show that there is no truthful mechanism satisfying the above properties and the following.

- *non-bundling* – there is a bid vector (u, b) such that the mechanism allocates one unit of the good to each buyer.

Unfortunately, we do not know how to prove this conjecture. However, we can prove this statement under the following stronger condition.

- *strong non-bundling* – for any non-zero bid (u_1, b_1) of the first agent, there is a bid (u_2, b_2) for the second agent such that if both agents bid according to (u, b), the mechanism allocates one unit of the good to each buyer.

The following theorem is the main result of this section.

Theorem 2.3.2 *There is no deterministic truthful auction for two buyers and two units of a good that satisfies consumer sovereignty, IIA, and strong non-bundling.*

In the proof of our impossibility result, we will use the following simple characterization of truthful auctions. Similar to Theorem 2.1.1, this characterization essentially claims that any truthful auction determines the allocation and price for agent i by comparing his bid to thresholds computed from the other agents' bids.

Proposition 2.3.1 *For any deterministic truthful auction selling m units of a good to n agents in the budget-constrained case, there exist mn functions $p_i^1, \ldots, p_i^m : (\mathbb{R}^+ \cup \{0\})^{2(n-1)} \to \mathbb{R}^+ \cup \{0\} \cup \{\infty\}$ such that agent i receives j units at price $p_i^j(u_{-i}, b_{-i})$ where j maximizes $ju_i - p_i^j(u_{-i}, b_{-i})$ subject to $p_i^j(u_{-i}, b_{-i}) \leq b_i$.*

Proof. For any $(u_{-i}, b_{-i}) \in (\mathbb{R}^+ \cup \{0\})^{2(n-1)}$ and $j \in \{1, \ldots, m\}$, we define $p_i^j(u_{-i}, b_{-i})$ as the minimum, over the choice of (u_i, b_i) such that the auction allocates at least j units to i if agents bid (u, b), of the price that the auction charges to i at these bids. For any set of bids (u, b), let j^* be an index that maximizes $j^*u_i - p_i^{j^*}(u_{-i}, b_{-i})$ subject to $p_i^{j^*}(u_{-i}, b_{-i}) \leq b_i$. If when agents bid (u, b), the auction allocates j units to i at price p, then we must have $ju_i - p \geq j^*u_i - p_i^{j^*}(u_{-i}, b_{-i})$, since otherwise agent i would have an incentive to bid untruthfully to get j^* units at price $p_i^{j^*}(u_{-i}, b_{-i})$. The equality follows from the definition of j^*. □

By considering all cases for the relationship between the p_i^j's, the auction can be expressed as a concise set of inequalities. This is done for the case of two units of good and two buyers in the following corollary. We will use this corollary to prove that truthful mechanisms satisfying certain properties do not exist.

Corollary 2.3.3 *For any deterministic truthful auction selling 2 units of a good to 2 agents, there exist threshold functions* $p_i^j : (\mathbb{R}^+ \cup \{0\})^2 \to \mathbb{R}^+ \cup \{0\} \cup \{\infty\}$, $1 \leq i, j \leq 2$, *such that for* $i = 1, 2$, *the agent* i *receives*

- 2 *units at a total price of* $p_i^2(u_{3-i}, b_{3-i})$ *if*

$$b_i \geq p_i^2(u_{3-i}, b_{3-i})$$

  and

$$u_i > p_i^2(u_{3-i}, b_{3-i}) - \min(p_i^1(u_{3-i}, b_{3-i}), p_i^2(u_{3-i}, b_{3-i})/2)$$

  (or if the latter inequality holds with equality, the mechanism can choose to allocate 2 units to i);

- *else* 1 *unit at price* $p_i^1(u_{3-i}, b_{3-i})$ *if*

$$b_i \geq p_i^1(u_{3-i}, b_{3-i})$$

  and

$$u_i > p_i^1(u_{3-i}, b_{3-i})$$

  (or if the latter inequality holds with equality, the mechanism can choose to allocate 1 units to i);

- *else* 0 *units.*

Conversely, for any set of threshold function $p_i^j : (\mathbb{R}^+ \cup \{0\})^2 \to \mathbb{R}^+ \cup \{0\} \cup \{\infty\}$, $1 \leq i, j \leq 2$, *the mechanism defined above satisfies incentive compatibility and individual rationality.*

Proof. We prove the statement for $i = 1$ ($i = 2$ is analogous). Consider the threshold functions given by Proposition 2.3.1. Fix any bid (u_2, b_2) of the second agent. Suppose the true value and budget of the first agent is u_1 and b_1, respectively. For simplicity, we use the notation $p_1^1 := p_1^1(u_2, b_2)$ and $p_1^2 := p_1^2(u_2, b_2)$. Notice that by the definition of p_1^1 and p_1^2 in the proof of Proposition 2.3.1, $p_1^1 \leq p_1^2$. The first agent's utility for an allocation of 0 units is 0, 1 unit is $u_1 - p_1^1$ assuming $b_1 \geq p_1^1$, and 2 units is $2u_1 - p_1^2$ assuming $b_1 \geq p_1^2$. The first agent receives two units if and only if he has enough budget to pay for it (that is, $b_1 \geq p_1^2$), and his utility for receiving two units ($2u_1 - p_1^2$) is greater than or equal to his utility for receiving one unit ($u_1 - p_1^1$) and zero units (zero). This can be written as $u_1 \geq p_1^2 - p_1^1$ and $u_1 \geq p_1^2/2$, or equivalently, $u_1 \geq p_1^2 - \min(p_1^1, p_1^2/2)$. Otherwise, if the first agent does not receive two units, then he receives one unit if and only if he has the budget (that is, $b_1 \geq p_1^1$), and his utility for one unit ($u_1 - p_1^1$) is greater than or equal to his utility for zero units, or equivalently, $u_1 \geq p_1^1$. If these conditions do not hold, then the agent receives zero units. The converse follows easily from the definition of the mechanism. □

The proof of the impossibility result is based on examining functional relations imposed by our assumptions on the threshold functions of any truthful auction. We obtain the result by showing that this set of functional relations has no solution.

Proof of Theorem 2.3.2. The fact that our auction observes supply limits implies that whenever the threshold functions are such that the first (second) agent gets two units, then the second (first) agent must get zero units. The consumer sovereignty and IIA assumptions imply that these two situations are in fact equivalent in certain regions of the bid space, that is the mechanism always allocates all the units when the bids are large enough.

By consumer sovereignty, for each agent $i = 1, 2$, there is a bid (u_i^*, b_i^*) such that if i bids (u_i^*, b_i^*) and the other agent bids $(0, 0)$, then agent i wins both units. Furthermore, by Corollary 2.3.3, for every $u_i' \geq u_i^*$ and $b_i' \geq b_i^*$, if i bids (u_i', b_i') and the other agent bids $(0, 0)$, then i wins both units. Let $C = \max\{u_1^*, b_1^*, u_2^*, b_2^*\}$.

Claim 2.3.1 *For any set of bids (u_1, b_1) and (u_2, b_2) such that $u_1, b_1, u_2, b_2 \geq C$, the mechanism allocates both units when agents bid according to (u, b). Furthermore, the payment of any agent that receives at least one unit in this situation is non-zero.*

Proof. Assume, for contradiction, that for one such bid vector the mechanism allocates at most one unit of the good to the first agent and zero units to the second agent. Now, by IIA, if the second agent bids $(0, 0)$, the first agent must still receive at most one unit. This, contradicts the definition of C. Now, assume that an agent, say 1, receives at least one unit but has to pay 0. This means that if agent 1 bids $(0,0)$, he still wins at least one unit, and therefore agent 2 does not receive both units. This contradicts the definition of C. □

Immediate from Corollary 2.3.3 is the fact that the allocations and payments given bid (α_i, β_i) holding bid $(\alpha_{3-i}, \beta_{3-i})$ fixed is constant for all $\beta_i \geq 2\alpha_i$ and for all $\alpha_i \geq \beta_i$. We will use this observation to make statements about the properties of the threshold functions as one of the inputs becomes irrelevant (that is, sufficiently large). Let

$$\begin{aligned} r_i^j(x) &= p_i^j(x, 2x), \\ s_i^j(x) &= p_i^j(x, x) \end{aligned}$$

for $i, j = 1, 2$. By Corollary 2.3.3, all of the above functions are non-decreasing functions: holding the bid of bidder i fixed, as bidder $3-i$ increases his value and/or budget, his allocation can not decrease. Since there are a limited number of units available, this means the allocation of bidder i must not increase which means his thresholds p_i^j must be non-increasing.

As the r_i^j and s_i^j are non-decreasing, they can be discontinuous in at most a countable number of points [102]. Let T denote the set of points greater than C at which all of the above functions are continuous. Notice that since the number of discontinuity points of each of these functions is countable, the set T is dense in (C, ∞).

Claim 2.3.1 together with our characterization, Corollary 2.3.3, immediately imply the following functional relations:

Lemma 2.3.3 *For all $A, B \in T$,*

$$B < r_{3-i}^1(A) \Rightarrow A \geq (s_i^2 - \min(s_i^1, s_i^2/2))(B) \tag{2.5}$$

Proof. Suppose agent i bids $(A, 2A)$ and agent $(3-i)$ bids (B, B) and $B < r_{3-i}^1(A)$.

Then agent $(3-i)$ receives zero units, so agent i must receive two units. As agent i's budget is essentially unconstrained, this implies that his value is at least the threshold, or $A \geq (s_i^2 - \min(s_i^1, s_i^2/2))(B)$. □

Similarly, we can prove the following statements for every $A, B \in T$:

$$A > (s_i^2 - \min(s_i^1, s_i^2/2))(B) \Rightarrow B < r_{3-i}^1(A), \tag{2.6}$$

$$B \geq r_{3-i}^2(A) \Rightarrow A \leq \min(s_i^1, s_i^2/2)(B), \tag{2.7}$$

$$A < \min(s_i^1, s_i^2/2)(B) \Rightarrow B \geq r_{3-i}^2(A), \tag{2.8}$$

$$B > (r_{3-i}^2 - \min(r_{3-i}^1, r_{3-1}^2/2))(A) \Rightarrow A \leq \min(r_i^1, r_i^2/2)(B), \tag{2.9}$$

$$A < \min(r_i^1, r_i^2/2)(B) \Rightarrow B \geq (r_{3-i}^2 - \min(r_{3-i}^1, r_{3-1}^2/2))(A), \tag{2.10}$$

$$B \geq s_{3-i}^2(A) \iff A < s_i^1(B). \tag{2.11}$$

From these functional relations, we can derive the following inequalities.

Lemma 2.3.4 *For all $A \in T$,*

$$(r_2^2 - \min(r_2^1, r_2^2/2))(A) \geq (s_2^2 - \min(s_2^1, s_2^2/2))(A). \tag{2.12}$$

Proof. Choose $B \in T$, $B > (r_2^2 - \min(r_2^1, r_2^2/2))(A)$. Then relation 2.9 (with $i = 1$) implies $A \leq \min(r_1^1, r_1^2/2)(B) \leq r_1^1(B)$. Take $\epsilon > 0$ and note that relation 2.5 (with $i = 2$) implies $B > (s_2^2 - \min(s_2^1, s_2^2/2))(A - \epsilon)$. Taking the limit as ϵ goes to zero and using the continuity of s_2^1 and s_2^2 at A, we have that $B > (r_2^2 - \min(r_2^1, r_2^2/2))(A)$ implies $B \geq (s_2^2 - \min(s_2^1, s_2^2/2))(A)$. Since this statement holds for every $B \in T$ and T is dense, the lemma follows. □

Similarly, we prove the following lemma.

Lemma 2.3.5 *For all $A \in T$,*

$$\min(r_2^1, r_2^2/2)(A) \geq \min(s_2^1, s_2^2/2)(A). \tag{2.13}$$

Proof. Choose $B \in T$, $B > \min(r_2^1, r_2^2/2)(A)$. By the contrapositive of relation 2.9, $A \leq (r_1^2 - \min(r_1^1, r_1^2/2))(B)$. By Claim 2.3.1, $\min(r_1^1, r_1^2/2)(B) > 0$. Hence, $A < r_1^2(B)$. This, by the contrapositive of relation 2.8, implies that $B \geq \min(s_2^1, s_2^2/2)(A)$. Since this holds for every $B \in T$ and T is dense, the lemma follows. $\square$

Our non-bundling assumption implies that for all $Z \in T$ the interval $(r_2^1(Z), r_2^2(Z))$ is non-empty. Select a point t in this interval and observe that the contrapositive of relations 2.6 (with $i = 1$) implies

$$Z \leq (s_1^2 - \min(s_1^1, s_1^2/2))(t). \tag{2.14}$$

Let $\epsilon > 0$ and note that t is in the interval $(r_2^1(Z - \epsilon), r_2^2(Z - \epsilon))$ for small enough ϵ by continuity. Thus the contrapositive of relation 2.8 with $i = 1$ implies

$$Z > Z - \epsilon \geq \min(s_1^1, s_1^2/2)(t). \tag{2.15}$$

Combining inequalities 2.14 and 2.15, we get

$$(s_1^2 - \min(s_1^1, s_1^2/2))(t) > \min(s_1^1, s_1^2/2)(t)$$

and so

$$\min(s_1^1, s_1^2/2)(t) = s_1^1(t). \tag{2.16}$$

Equations 2.15 and 2.16 imply that for every $t \in (r_2^1(Z), r_2^2(Z))$, $Z > s_1^1(t)$. By Equation 2.11, this implies that $t < s_2^2(Z)$. Taking the limit of this equation as t tends to $r_2^2(Z)$, we obtain $r_2^2(Z) \leq s_2^2(Z)$. On the other hand, summing Equations 2.12 and 2.13 implies that $r_2^2(Z) \geq s_2^2(Z)$. Therefore, $r_2^2(Z) = s_2^2(Z)$. Thus, inequalities 2.12 and 2.13 must both attain equality at Z. Ranging over choice of $Z \in T$, we see that inequalities 2.12 and 2.13 must attain equality everywhere in T. Our contradiction arises from the observation that in fact for some $Z \in T$, inequality 2.13 is strict. By Claim 2.3.1, prices are always nonzero, and so $C < (r_1^2 - \min(r_1^1, r_1^2/2))(A) < r_1^2(A)$ for some $A \in T$. Select such an A and $Z \in ((r_1^2 - \min(r_1^1, r_1^2/2))(A), r_1^2(A)) \cap T$. Notice that since T is a dense set, this intersection is nonempty. Note that relation 2.8 with $i = 2$ implies that $A \geq \min(s_2^1, s_2^2/2)(Z)$. Therefore, for any small $\epsilon > 0$, $A + \epsilon > \min(s_2^1, s_2^2/2)(Z)$. Similarly, note that relation 2.9 with $i = 2$ implies $A + \epsilon \leq \min(r_2^1, r_2^2/2)(Z)$ for ϵ sufficiently small. But this means that, for this particular Z, $\min(r_2^1, r_2^2/2)(Z) > \min(s_2^1, s_2^2/2)(Z)$, yielding our contradiction. □

The following example shows that the IIA assumption in Theorem 2.3.2 is necessary, that is there are deterministic truthful mechanisms that satisfy consumer sovereignty and non-bundling, but not IIA.

Example 2.3.3 Consider an auction with 2 units and 2 bidders which uses the following rules for allocation to agent i ($i = 1, 2$):

- If $u_i > 2\min(u_{3-i}, b_{3-i})$ and $b_i > \frac{5}{2}\min(u_{3-i}, b_{3-i})$, then agent i gets 2 units and pays $\frac{5}{2}\min(u_{3-i}, b_{3-i})$;
- else if $u_i > \frac{1}{2}\min(u_{3-i}, b_{3-i})$ and $b_i > \frac{1}{2}\min(u_{3-i}, b_{3-i})$, then agent i gets 1 unit and pays $\frac{1}{2}\min(u_{3-i}, b_{3-i})$;
- else agent i receives nothing.

It is not hard to verify that this mechanism satisfies the characterization given in Corollary 2.3.3, and is therefore truthful. However, if, for example, agent 1 bids $(4a, 4a)$ and agent 2 bids $(9a, 9a)$ for any a, the mechanism allocates zero units to agent 1 and one unit to agent 2. Therefore, the mechanism does not allocate both units even if bids are sufficiently large, and hence it does not satisfy IIA.

Chapter 3

Procurement Auctions

In this chapter, we study markets in which an auctioneer wishes to assemble a team of agents to accomplish some task. These agents offer fixed services that incur some privately known cost. The auctioneer must select a team, or *feasible set* of agents, the combination of which is capable of performing the task. To this end, he designs a *procurement auction* in which he solicits bids from the agents and then selects some feasible set of agents (the *winners*) to perform the task at hand and pays them according to the rules of the auction.

Path and *flow* auctions are important special cases of procurement auctions. In path auctions, the auctioneer seeks to buy a path of edges of lowest price between a specified source and destination in a network. Sellers (network edges) have a privately known cost for transmitting traffic, and bid to attract traffic. Path auctions arise naturally in network routing – for example, an Internet Service Provider (ISP) might use a procurement auction to select autonomous systems (ASs) to route his demand. Flow auctions are a generalization of path auctions in which the demand of the auctioneer might exceed the capacity of any single source-destination path. In this case, the auctioneer must buy a set of edges capable of routing his demand.

One plausible mechanism for procurement auctions, proposed for use in path auctions by Nisan and Ronen [89], Hershberger and Suri [55], and Feigenbaum et al. [35], is the Vickrey-Clark-Groves (VCG) mechanism [19, 52, 115]. Roughly speaking, the VCG mechanism pays each winning agent the highest bid with which it could still have won, all other bids being unchanged. The utility of an agent is quasi-linear, and so the VCG mechanism is truthful (agents bid their true cost) and efficient (a feasible set of minimum total true cost is selected). However, as observed by Archer and Tardos [6], even in the special case of path auctions, the VCG mechanism (and, in fact, all *min function* mechanisms) can lead to the auctioneer paying far more than the true cost of completing the task at hand. In fact, the payment of the auctioneer may even greatly exceed the true cost of the second-cheapest feasible set. Elkind, Sahai, and Steiglitz [29] generalized the result of Archer and Tardos [6] to prove that *all* truthful mechanisms have high overpayments in general.

We are interested in reining in the cost to the auctioneer. There are two general approaches to this problem. One approach tries to characterize procurement settings in which the VCG mechanism has small overpayments. Tawlar [111] and Garg et al. [42] consider restricting the setting by imposing a structure on the collection of feasible sets of agents. Mihail, Papadimitriou, and Saberi [81] show that in a *random graph*, the expected payment of a VCG mechanism for a shortest path is small. Feigenbaum et al. [35] measure the average overpayment of the VCG mechanism for shortest path auctions in the Internet's autonomous systems (ASs) graph and conclude that it is relatively small. A second approach is to consider alternative solution concepts. Garg et al. [42] propose an *ascending price auction* format for procurement auctions that can perform well in settings of incomplete information. For the special case of path auctions, Elkind, Sahai, and Steiglitz [29] present and analyze an optimal *Bayesian-Nash* mechanism. Czumaj and Ronen [20] propose a mechanism that combines dominant and non-dominant strategy mechanisms and has small overpayments under certain assumptions. However they show that it has an arbitrary ratio between the payment of different equilibria and say that overall, "finding a natural and tractable measure of [non-dominant strategy] protocols seems challenging and important."

In this chapter, we follow the second of these approaches. We propose and analyze variants on *first-price* auctions, or auctions in which the team with the lowest bid is selected and paid their bid. First-price auctions are a natural class of auctions quite often implemented in practice. Therefore, it is interesting to ask if first-price auctions or their variants can reduce the payment of the auctioneer. These auctions are not truthful; instead, we motivate analyzing their properties in a *strong ϵ-Nash equilibrium* (see Definition 10). We show that in general procurement settings, strong ϵ-Nash exist, and the feasible set of agents selected in any strong ϵ-Nash equilibrium is approximately efficient. For path and flow auctions, we then bound the total payment to the winning agents by relating it to the true cost of routing one additional unit of demand (assuming all edges have unit capacity). Finally, we study the setting in which the demand of the auctioneer is not known, but rather the auctioneer and bidders share a *common prior belief* regarding the amount of demand. In other words, there is a publicly known *distribution* of possible demands. For this model, we design a first-price mechanism involving *two-parameter* bids and derive a bound on the payments of this mechanism similar to that of the known demand case.

In Section 3.1, we formalize the setting of procurement auctions and define the path and flow settings which we study later. In Section 3.2, we motivated the selection of strong ϵ-Nash equilibria as a solution concept for first-price auctions. In Section 3.3, we show that first-price auctions are approximately efficient in the general procurement setting. Finally, in Section 3.4, we show how to bound the payment of first-price auctions and their variants in the special case of path and flow settings.

The results of this chapter are based on joint work with Karger, Nikolova, and Sami [59].

3.1 Setting

Consider a *procurement setting* in which an auctioneer wishes to hire a team of agents to accomplish a particular task. There is a set U of n agents. Each agent is capable of performing a fixed service. In performing this service, an agent incurs a privately known cost $c_i \in \mathbb{R}^+ \cup \{0\}$. Some subsets of services can be combined to accomplish the auctioneer's task. We call a subset $S \subseteq U$ of agents a *feasible set* if their combined services can accomplish the task. The collection of feasible sets is denoted by $\mathcal{S}$.

The collection $\mathcal{S}$ could be publicly known to the auctioneer and all agents, or, more generally, they could share a *common prior* (a publicly known probability distribution over the collection of subsets of U) about $\mathcal{S}$.

A special case of the procurement setting is the *path* or *flow setting.* In this setting, there is a graph G. Each edge (u, v) is an agent capable of sending *one* unit of flow from u to v at a privately known cost $c_i \in \mathbb{R}^+ \cup \{0\}$. The auctioneer wants to route k units of demand from a known source node s and destination node t (in a *path* setting $k = 1$). Hence the collection $\mathcal{S}$ of feasible sets is the collection of all subgraphs in G that contain a k-flow from s to t. We assume that the structure of the graph G is public knowledge. The demand k could be publicly known to the auctioneer and all edges (the *known demand case*), or it might be drawn from a publicly known probability distribution (the *unknown demand case*).

The unknown demand case is modelled as follows: The demand can take any integral value in the range $[1, r]$, where r is a positive integer. Further, there is a known prior distribution on the demand values; say that the demand is k with probability p_k, for $k = 1, 2 \ldots, r$. We assume for simplicity that $p_k > 0$ for all k; our results easily extend to a situation in which $p_k = 0$ for some values of $k \in \{1, \ldots, r\}$.

In a *procurement auction* (similarly a *path auction* or *flow auction*), the auctioneer selects a feasible set by running an auction. He solicits from each agent a *bid* $b_i \in \mathbb{R}^+ \cup \{0\}$ which is supposed to represent the agent's true cost c_i. He then selects some feasible set S of agents and pays each agent $i \in S$ an amount $payment_i \in \mathbb{R}^+ \cup \{0\}$ and all other agents 0. The set S is called the *winning set.* Each agent $i \in S$ is a *winner*, and all other agents are *losers.* An agent's utility for the outcome is $payment_i - c_i x_i$, where $\mathbf{x}$ is the characteristic vector of S (that is, $x_i = 1$ if $i \in S$ and 0 otherwise).

We will focus on *first-price* auctions. In a first-price auction, the payment of every winner equals his bid. The auctioneer is restricted to select a *minimum price* feasible set S, or one which minimizes $\sum_{i \in S} b_i$. His only flexibility is in the definition of a *tie-breaking rule*, or method to select from among the collection of minimum price feasible sets. Thus, in specifying a first-price auction, we only need to specify a tie-breaking rule. We also consider variants of first-price auctions in which the minimum price feasible set is almost always selected and the winners are paid a quantity close to their bid.

To avoid confusion between the true costs and the prices of sets, we will adopt the following terminology: the *cost* of a set S is $\sum_{i \in S} c_i$, sometimes written $c(S)$. Similarly, the *price* of a set S is $\sum_{i \in S} b_i$, sometimes written $b(S)$. Additional notation will be introduced for the path and flow settings in Section 3.4.

3.2 Solution Concepts

First-price auctions are clearly not truthful. This raises the question of how we expect agents to bid. We want to retain the property that agents can see each others' bids, so that the bidding could be performed through posted prices. Thus, mixed-strategy equilibria are not very meaningful in our setting. Instead, we look for a pure strategy equilibrium solution concept which *always exists* and is *arguably reasonable* in that agents can be expected to reach that equilibrium. This section motivates the selection of *strong ϵ-Nash equilibria* (see Definition 10) as that solution concept through a series of examples. First we note that not every first-price procurement auction has a Nash equilibrium (Example 3.2.1), and those that do are impractical (Example 3.2.2). Both of these examples heavily rely on the continuity of the bid and payment space. In reality, bids and payments are restricted to a discrete space as they should be some multiple of a unit of money, like cents, for example. Thus it is simply not possible for agents to arbitrarily improve their payoffs, and so we suggest studying *ϵ-Nash equilibrium* (see Definition 9) where an agent deviates only if it improves his payoff by at least ϵ. Unfortunately, Example 3.2.3 shows that the overpayments in such an equilibrium can be quite high. In this example, however, if certain subsets of agents could arrange to jointly reduce their bids, all of them would benefit. This leads us to study *strong ϵ-Nash equilibria* (see Definition 10), or ϵ-Nash equilibria which are robust to such manipulations. As proven in Theorem 3.2.1, strong ϵ-Nash equilibria exist in all deterministic first-price procurement auctions (but may fail to exist in randomized ones as evidenced by Example 3.2.4). It remains to be seen if one can devise a bidding protocol that helps agents converge to a strong ϵ-Nash equilibrium.

3.2.1 Nash Equilibria

The most natural solution concept is that of a Nash equilibrium. Unfortunately, as the following example shows, not every first-price auction has a Nash equilibrium.

Example 3.2.1 *Suppose there are two agents A and B, either of whom forms a feasible set (that is, $S = \{\{A\}, \{B\}\}$). Consider any auction in which ties are broken by selecting agent B with probability p, $0 < p \leq 1$, independent of the bid values. Now suppose the costs of the agents are $c_A = 1$ and $c_B = 2$, and so in case of a tie the auction selects the higher-cost agent with positive probability.*

Suppose agent A bids x and B bids y. If $x \geq y$, then the expected payment of agent A is at most $(1-p)y$. As B has positive probability of winning, $y \geq c_B = 2$, and so the bid $y - \epsilon$ for $\epsilon < \min(yp, \frac{1}{2})$ is a better bid than x for agent A. If $x < y$, then the payment to A is x and so $x + (y-x)/2$ is a better strategy than x for A.

This example relies on the assumption that the tie-breaking rule is not a function of the bid values (otherwise we would have been unable to assume that the auction selects the higher-cost agent with positive probability). In fact, for a carefully chosen tie-breaking rule which is a function of the bid values, we can design first price auctions with pure strategy Nash equilibria, as the following example shows.

Example 3.2.2 *For ease of exposition, suppose all subsets 2^U of the set of agents U are feasible and index the subsets so $2^U = \{B_1, \ldots, B_{2^n}\}$. Partition the real numbers into 2^n subsets $S_1, \ldots, S_{2^n}$ such that each subset is dense in the reals. Let p be the price of the minimum price set and suppose $p \in S_k$. If B_k has price p, choose B_k. Otherwise, choose randomly among the collection of minimum price sets.*

We can construct a pure strategy Nash equilibrium for this tie-breaking rule as follows. If the minimum cost set is not unique, then it is a Nash equilibrium for all agents to bid $b_i = c_i$, their true cost. Otherwise, let B_1 be the minimum cost set and B_2 be the next cheapest set (in terms of true costs). Find a $p \in S_1 \cap [c(B_1), c(B_2))$

(thus B_1 wins in the case of a tie at price p). Consider a set of bids $\mathbf{b}$ *such that* $b(B_1) = b(B_2) = p$, $b_i \geq c_i$ *for* $i \in B_1$, $b_i \leq c_i$ *for* $i \in B_2$, *and* $b_i = c_i$ *for* $i \notin B_1 \cup B_2$. *Then* $\mathbf{b}$ *is a Nash equilibrium.*

However, this auction is arguably impractical as are the deviations discussed in the last example because they both assume that the bids and payments can be any real number. Yet, in many problems, payments are discrete, so it is simply not possible for agents to improve their utilities by arbitrarily small amounts. This motivates us to use the solution concept of ϵ-*Nash equilibrium.*

Remark 3.2.1 *The results in this chapter can be proved using tie-breaking rules such as that in Example 3.2.2 or using ϵ-equilibria concepts presented below. However, for clarity of presentation, we present our results in terms of ϵ-equilibria.*

3.2.2 ϵ-Nash Equilibria

In an ϵ-Nash equilibrium, we assume agents are indifferent to deviations that improve their payoff by a small amount.

Definition 9 *An* ϵ-Nash equilibrium *is a set of strategies, one for each agent, such that no agent can unilaterally deviate in a way that improves his payoff by at least ϵ.*

Unfortunately, there is a drawback to the ϵ-Nash solution concept as well. As the following example shows, when the winning set contains many agents, it may have a *price* higher than the *cost* of the best competing set.

Example 3.2.3 *Consider any first-price auction. Suppose there are four agents, A, B, C, and D with costs* 1, 2, 2, *and* 6 *respectively, and the collection $\mathcal{S}$ of feasible sets is* $\{\{A\}, \{B, C\}, \{D\}\}$. *Then it is an ϵ-Nash equilibrium for agent A to bid* $6 - \epsilon$, *and the rest to bid* 6. *In this case, the price to the auctioneer for the winning set* $\{A\}$ *is* $6 - \epsilon$ *which is higher than the cost,* 2, *of the best competing set* $\{B, C\}$.

This defeats our goal of reducing customer overpayment. We might argue that this solution would not be sustained in practice, since the agents in the second lowest-cost set are likely to each reduce their price. This leads us to explore the concept of *strong ϵ-Nash equilibria.*

3.2.3 Strong ϵ-Nash Equilibria

Strong ϵ-Nash equilibria, first introduced by Aumann [9] and used by Young [119], require that there is no *group* of agents who can deviate in a way that improves the payoff of each member by at least ϵ.

Definition 10 *A* strong ϵ-Nash equilibrium *is a set of strategies, one for each agent, such that no group of agents (called a* coalition*) can deviate in a way that improves the payoff of each member by at least ϵ.*

This definition captures the notion that agents might collude to win the auction if it is beneficial for each of them (for a discussion of stronger notions of collusion, see Chapter 4). For example, the bid vector in Example 3.2.3 is not a strong ϵ-Nash equilibrium as agents B and C could collude and bid $3 - \epsilon$, thus improving each of their payoffs by at least ϵ (assuming $\epsilon < \frac{1}{2}$).

Strong ϵ-Nash equilibria have several advantages over Nash and ϵ-Nash equilibria. First, although *randomized* first-price auctions may fail to have strong ϵ-Nash equilibria (see Example 3.2.4), Theorem 3.2.1 shows that every *deterministic* first-price auction has a strong ϵ-Nash equilibrium. Second, as demonstrated by Lemma 3.2.1, in a strong ϵ-Nash equilibrium of a determinist first-price auction, we can bound the bids of the winning agents by the true costs of the losing agents, furthering our goal of reducing payments and allowing us to prove that the winning set is approximately efficient (see Section 3.3). The rest of this section contains proofs of Theorem 3.2.1 and Lemma 3.2.1 and Example 3.2.4.

First, we show that any first-price auction with a *deterministic* tie-breaking rule has a strong ϵ-Nash equilibrium. Our proof is constructive. We consider the minimum cost feasible set and fix the bids of all items outside this set to be equal to their true cost. For the items in this set, we adjust their bids so that the price of the set is just less than the cost of the second-lowest cost feasible set.

Theorem 3.2.1 *Any first price auction with a deterministic tie-breaking rule has a strong ϵ-Nash equilibrium.*

Proof. Our proof is constructive. Let c_i be the cost of agent i, $\mathcal{S}$ be the collection of feasible sets, and S^* be the minimum cost feasible set selected by the auction under bid vector $\mathbf{c}$. Define a variable x_i for each $i \in S^*$ and consider the following linear program (LP for short):

$$\begin{array}{ll} \text{maximize} & \displaystyle\sum_{i \in S^*} x_i \\ \text{subject to} & \forall\, S \in \mathcal{S} : \displaystyle\sum_{i \in S^* - S} x_i \le \sum_{i \in S - S^*} c_i \\ & \forall\, i \in S^* : \; x_i \ge c_i \end{array}$$

The strong ϵ-Nash equilibrium that we construct will be a slightly modified optimal solution to this LP. The first constraint guarantees that S^* will be a minimum price set in this equilibrium, and the second that every agent has non-negative utility in this equilibrium. By setting $x_i = c_i$ for all i, we see that the LP is feasible.

Let x_i^* be an optimum solution of the LP, and define bid vector $\mathbf{b}$ where $b_i = \max\{c_i, x_i^* - \epsilon/(2n)\}$ for $i \in S^*$ and $b_i = c_i$ for all other i. Notice that our minimum cost set S^* is also a minimum price set with respect to bids $\mathbf{b}$.

We prove that $\mathbf{b}$ is a strong ϵ-Nash equilibrium. Note that only agents who are guaranteed winners (that is, agents in every minimum price set) are submitting a bid other than their true cost. For agents outside S^*, this is evident from the definition of $\mathbf{b}$. Consider an agent i in S^* that is not in every minimum price set, and let S be a minimum price set that does not contain i. Corresponding to this S is an inequality of type 1. This inequality together with those of type 2 for all $j \in S^* - S$ imply that $x_i = c_i$ and so $b_i = c_i$. Thus the bidders in a successful coalition can only increase their bids.

Let T be a successful coalition and $\mathbf{b}'$ be the bid vector when T deviates (so $b'_i = b_i$ for all $i \notin T$). Recall the notation $b(S) = \sum_{i \in S} b_i$. Then

$$b'(S^*) = b(S^*) + \sum_{i \in T \cap S^*} (b'_i - b_i). \tag{3.1}$$

In order for each member of the coalition to benefit by at least ϵ, he must increase his bid by at least ϵ, so

$$\forall\, i \in T,\ b'_i - b_i \geq \epsilon, \tag{3.2}$$

and T must be a subset of the selected minimum price set S'. Therefore,

$$b'(S') = b(S') + \sum_{i \in T} (b'_i - b_i). \tag{3.3}$$

Furthermore, as S^* is a minimum price set according to $\mathbf{b}$ and S' is a minimum price set according to $\mathbf{b}'$,

$$b(S^*) \leq b(S'), \tag{3.4}$$

and

$$b'(S') \leq b'(S^*). \tag{3.5}$$

Inequalities 3.1, 3.3, 3.4, and 3.5 imply

$$\sum_{i \in T \cap S^*} (b'_i - b_i) \geq \sum_{i \in T} (b'_i - b_i). \tag{3.6}$$

Together with inequality 3.2, inequality 3.6 implies $T \subseteq S^*$. As $T \subseteq S'$ as well, inequalities 3.4 and 3.5 imply

$$b'(S^*) = b'(S'). \tag{3.7}$$

Now consider the solution to the LP which sets each variable x_i to agent i's bid in bid vector $\mathbf{b}'$. By inequality 3.2,

$$
\begin{aligned}
\sum_{i\in S^*} x_i &= \sum_{i\in S^*} b'_i \\
&\geq \sum_{i\in S^*-T} b_i + \sum_{i\in T\cap S^*} (b_i+\epsilon) \\
&\geq \sum_{i\in S^*-T} (x_i^* - \frac{\epsilon}{2n}) + \sum_{i\in T\cap S^*} (x_i^* - \frac{\epsilon}{2n} + \epsilon) \\
&\geq \sum_{i\in S^*} x_i^* + \frac{\epsilon}{2}.
\end{aligned}
$$

By maximality of $\mathbf{x}^*$, this implies that $\mathbf{x}$ is not feasible. Since each $b'_i \geq c_i$ by construction, $\mathbf{x}$ must violate an inequality of type 1. Letting S be the set in the violating constraint, we see $b'(S) < b'(S^*)$ which, by equality 3.7, implies $b'(S) < b'(S')$, contradicting the optimality of S'. □

Next, we show that in a strong ϵ-Nash equilibrium, the price of the winning set can be bounded by the cost of losing feasible sets. The intuition for this proof is that if the winning agents are bidding significantly more than the losing agents, the losing agents can undercut the bidding agents and win at a profitable price. One powerful consequence of this definition is that, from the point of view of the total price, it lets us assume without loss of generality that items who are not winning in a strong ϵ-Nash equilibrium are bidding within ϵ of their cost. This notion is formalized in the following lemma.

Lemma 3.2.1 *Fix a strong ϵ-Nash equilibrium $\mathbf{b}$ and let S be the feasible set that wins with bids $\mathbf{b}$. Let T be any set (not necessarily feasible) such that $T \cap S = \emptyset$ and for all $i \in T$, $b_i > c_i + \epsilon$, where c_i is the true cost of item i. Consider the altered bid vector $\mathbf{b}'$ in which*

$$
b'_i = \begin{cases} c_i + \epsilon & \textit{for } i \in T, \\ b_i & \textit{otherwise.} \end{cases}
$$

Let S' be a minimum price feasible set with respect to bids $\mathbf{b}'$. Then $b'(S') = b(S)$.

Proof. Since $T \cap S = \emptyset$, $b(S) = b'(S)$, and so $b'(S') \leq b(S)$. Suppose $b'(S') < b(S)$. This means for all minimum price sets with respect to bids $\mathbf{b}'$, there are items in the set T. Let R' be a minimum price set with respect to bids $\mathbf{b}'$ which minimizes $|R' \cap T|$ (by the previous statement, this minimum is at least one). We will show that the agents in $R' \cap T$ form a coalition when the bids are $\mathbf{b}$, contradicting the assumption that $\mathbf{b}$ was a strong ϵ-Nash equilibrium.

Consider the bid vector $\mathbf{b}''$ constructed from $\mathbf{b}$ in which just the agents in $R' \cap T$ lower their bids to $c_i + \epsilon$:

$$b_i'' = \begin{cases} c_i + \epsilon & \text{for } i \in R' \cap T, \\ b_i & \text{otherwise.} \end{cases}$$

We will argue that the agents in $R' \cap T$ benefit by at least ϵ in this deviation. As $T \cap S = \emptyset$, all agents in $R' \cap T$ were losing agents with bid vector $\mathbf{b}$ and so their utility with bids $\mathbf{b}$ was zero. We argue that in bid vector $\mathbf{b}''$ the agents in $R' \cap T$ all win the auction and, therefore, as $b_i'' = c_i + \epsilon$ for $i \in R' \cap T$, increase their utility by ϵ. As first-price auctions choose a winning set from among the minimum price feasible sets, we must show that the agents in $R' \cap T$ are contained in *any* minimum price feasible set R'' with respect to bids $\mathbf{b}''$. As $b_i' \leq b_i''$ for all agents i,

$$b''(R'') \geq b'(R'') \geq b'(R') = b''(R') \geq b''(R''),$$

and so all statements hold with equality. Since items in $T - R'$ increased in price from $\mathbf{b}'$ to $\mathbf{b}''$, $b''(R'') = b'(R'')$ implies R'' does not contain any element of $T - R'$. Since $b'(R'') = b'(R')$, R'' is also a minimum price set with respect to bids $\mathbf{b}'$. As R' was chosen to minimize the intersection with T among all minimum price sets, this means R'' must contain $R' \cap T$. Therefore, the agents in $R' \cap T$ are winners in bid vector $\mathbf{b}''$ and so increase their utility by at least ϵ. Furthermore, as $b_i = b_i''$ for all agents outside of $R' \cap T$, the agents in $R' \cap T$ can form a successful coalition in $\mathbf{b}$, contradicting the assumption that $\mathbf{b}$ was a strong ϵ-Nash equilibrium. $\square$

Remark 3.2.2 *In the proof of Theorem 3.2.1, we used the determinism of the mecha-*

nism in assuming that there was a unique winner for every bid vector. As the following example shows, this assumption was necessary. Strong ϵ-Nash equilibria do not necessarily exist for randomized first-price auctions. Randomized tie-breaking rules pose a problem for the solution concept as a minor adjustment in bid value can drastically affected a bidder's expected utility.

Example 3.2.4 *Suppose there are two agents, A and B, either of whom forms a feasible set for the auctioneer (that is, $\mathcal{S} = \{\{A\}, \{B\}\}$). In the case of a tie, assume the auctioneer chooses uniformly at random between the two items. Suppose the cost of each agent is 0. Note for any set of bids $\{b_A, b_B\}$, the agents can form a coalition and each bid $2\max(b_A, b_B)+2\epsilon$. In this way they both profit by at least ϵ in expectation. Therefore no pure strategy bid vector forms a strong ϵ-Nash equilibrium.*

3.3 Approximate Efficiency of First-Price Combinatorial Auctions

It is often desirable to design auctions that choose *efficient* allocations. A procurement auction is *efficient* if it always select the minimum cost feasible set. The VCG mechanism guarantees that the set it selects is efficient. The strong ϵ-Nash equilibria of first-price procurement auctions are not necessarily efficient. For example, if the minimum cost and second-minimum cost feasible sets have costs within ϵ of one another, then it is a strong ϵ-Nash for the second-minimum cost set to bid truthfully and the minimum cost set to overbid by ϵ. In such a scenario, the first price procurement auction will select the second-minimum cost set. Still, the winning set is approximately efficient as its cost is within ϵ of the minimum cost set. In this section, we prove that this holds in general, that is the strong ϵ-Nash equilibria of first-price procurement auctions are approximately efficient.

Theorem 3.3.1 *Let* $\mathbf{b}$ *be a strong ϵ-Nash equilibrium of a deterministic first-price procurement auction. Then the cost $c(S)$ of the winning set S in a first-price procurement auction is at most the cost $c(S^*)$ of the minimum cost feasible set S^* plus* *an additive factor of ϵn:*

$$c(S) \leq c(S^*) + \epsilon n.$$

Proof. The proof is by contradiction. Assume the winning set S is not approximately efficient, that is, $c(S) > c(S^*) + \epsilon n$. Define a new bid vector $\mathbf{b}'$ in which the agents who are not winning but are in a minimum cost feasible set lower their bids to just above their cost:

$$b'_i = \begin{cases} \min\{b_i, c_i + \epsilon\} & \text{for } i \in S^* - S, \\ b_i & \text{otherwise.} \end{cases}$$

In this bid vector, S^* is cheaper than S:

$$\begin{aligned} b'(S) - b'(S^*) &= b(S - S^*) - b'(S^* - S) \\ &\geq c(S - S^*) - (c(S^* - S) + \epsilon n) \\ &= c(S) - c(S^*) - \epsilon n \\ &> 0. \end{aligned}$$

This contradicts Lemma 3.2.1 with $T = \{i \in S^* - S : b_i > c_i + \epsilon\}$. □

3.4 Payment Bounds for Flow Auctions

In this section, we bound the overpayments of first-price flow auctions. We assume that we have a deterministic tie-breaking rule so that if there is more than one cheapest feasible flow, we take the lexicographically first integral one. We consider two settings. In the *known demand path auction* studied in Section 3.4.1, the total demand of the auctioneer is known to the auctioneer and all the bidders at the time of the auction. It is easy to imagine that the assumptions of this model might be unrealistic in practice. Can the total demand really be known before it is realized? What if the auctioneer wishes to buy flow in advance? In our second model, the *unknown demand path auction* studied in Section 3.4.2, the auctioneer and bidders instead know a probability distribution over possible demand values.

Notation

For a graph G, let $\mathbf{c}$ be the vector of edge costs, $\mathbf{b}$ be the vector of edge bids, and $F_{\mathbf{w}}(k, G)$ be the set of edges in the winning k-flow[1] in G with respect to edge weights $\mathbf{w}$ (as we only consider deterministic first-price auctions, this is well-defined). We will refer to $F_{\mathbf{c}}(k, G)$ as the minimum cost k-flow and $F_{\mathbf{b}}(k, G)$ as the minimum price k-flow with respect to bid vector $\mathbf{b}$. When the bids, costs, or graph is clear from the context, we will sometimes drop them from the notation. As a shorthand, we sometimes write $c(k)$ for the (cost of the) lowest cost k-flow. Finally, to be consistent with the previous notation, we denote the number of agents, or edges in G, by n.

3.4.1 Known Demand Path Auction

In the known demand setting, we assume that the auctioneer has a publicly-known demand k. We will show that in such settings, the payments in a strong ϵ-Nash equilibrium of a deterministic first-price auctions is bounded. In particular, we show that the overpayment to each unit of flow is (approximately) at most the true marginal cost of sending an additional unit of flow (see Theorem 3.4.1). Together with the observation that the VCG mechanism pays each edge a bonus at least as large as this marginal (see Theorem 3.4.2), this shows that the payments in first-price auctions are (approximately) bounded by the payments in the VCG auction. We saw in Section 3.3 that the winning set in the first-price auction is also (approximately) efficient. These statements regarding the payments and efficiency of first-price auctions suggest that first-price auctions perform better than VCG auctions. However, first-price auctions have a significant drawback; it is not clear how agents might converge to a strong ϵ-Nash equilibrium. We partially address this concern by proposing another auction whose ϵ-Nash equilibria have the same properties as the strong ϵ-Nash equilibria of a first-price auction.

[1] The weight of this flow is equal to the weight of the minimum weight k-flow, that is requiring integrality doesn't change the value of the optimal solution.

Payment Bound

We first bound the payments in a strong ϵ-Nash equilibrium (see Definition 10) of a deterministic first-price auction. The edges announce bids and the auctioneer runs a first-price auction to select a cheapest k-flow according to the bid vector, paying each edge on the flow an amount equal to his bid. By Theorem 3.2.1, strong ϵ-Nash equilibria exist for such auctions. Given the existence of strong ϵ-Nash equilibria, we can bound the payments in any such equilibrium. In order to develop some intuition for the proof, it is useful to first consider sending 1 unit of flow in a graph consisting of just two parallel edges from the source s to the sink t of costs a and b, $a > b + \epsilon$. The lower-true-cost edge must be allocated the flow in equilibrium since he can bid just under the true cost of the higher cost edge and be guaranteed a profit of at least ϵ. Therefore, by the conditions of a strong ϵ-Nash equilibrium, we can assume that the bid of the higher cost edge is at most ϵ more than his true cost, and so the overpayment of any equilibrium will be at most $a + \epsilon - b$. The crux of this argument was to bound the bid of the winning path by the bid of an augmenting path. Since the augmenting path does not receive flow, Lemma 3.2.1 permitted us to assume, for the purposes of bounding the price, that the bid of this path was close to its true cost. This proof idea easily extends to auctions for k-flows in general graphs as can be seen below.

Theorem 3.4.1 *The total payment of the deterministic first price k-flow auction in a strong ϵ-Nash equilibrium is at most*

$$k\left[c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k))\right] + kn\epsilon,$$

where $\mathbf{c}$ *is the vector of true edge costs.*

Proof. Fix a strong ϵ-Nash equilibrium vector of bids $\mathbf{b}$ and define bid vector $\mathbf{b}'$ such that

$$b'_i = \begin{cases} b_i & \text{for } i \in F_{\mathbf{b}}(k), \\ \min\{b_i, c_i + \epsilon\} & \text{otherwise.} \end{cases}$$

By Lemma 3.2.1, $F_{\mathbf{b}}(k)$ is a minimum price k-flow with respect to $\mathbf{b}'$. Consider the (non-integral) flow $(k/(k+1))F_{\mathbf{c}}(k+1)$, that is the flow which sends $k/(k+1)$ units of flow along the flow paths determined by $F_{\mathbf{c}}(k+1)$. Since $F_{\mathbf{b}}(k)$ is a lowest-price k-flow with respect to $\mathbf{b}'$ and using the integrality of optimal network flows [104], we have

$$\left(\frac{k}{k+1}\right) b'(F_{\mathbf{c}}(k+1)) - b'(F_{\mathbf{b}}(k)) \geq 0. \tag{3.8}$$

Define edge sets

$$\begin{aligned}
E_+ &= \{e \in F_{\mathbf{c}}(k+1) - F_{\mathbf{b}}(k)\} \\
E_o &= \{e \in F_{\mathbf{c}}(k+1) \cap F_{\mathbf{b}}(k)\} \\
E_- &= \{e \in F_{\mathbf{b}}(k) - F_{\mathbf{c}}(k+1)\}
\end{aligned}$$

Then equation 3.8 reduces to

$$\left(\frac{k}{k+1}\right) b'(E_+) - \left(\frac{1}{k+1}\right) b'(E_o) - b'(E_-) \geq 0$$

which, solving for $b'(E_o) + b'(E_-)$, gives

$$\begin{aligned}
b(F_{\mathbf{b}}(k)) &= b'(E_o) + b'(E_-) \\
&\leq k(b'(E_+) - b'(E_-)) \\
&\leq k(c(E_+) + n\epsilon - c(E_-)) && (3.9) \\
&\leq k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{b}}(k)) + n\epsilon) \\
&\leq k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)) + n\epsilon) && (3.10)
\end{aligned}$$

where 3.9 follows from the fact that for any edge $b'_i \geq c_i$ and for all $i \in E_+$, $b'_i \leq c_i + \epsilon$; and 3.10 follows from the optimality of $F_{\mathbf{c}}(k)$ with respect to $\mathbf{c}$. □

In addition, it is easy to see that this bound is tight. Consider a graph with $(k+1)$ parallel edges where the cost of the bottom k edges is c and the cost of the remaining top edge is $c' > c$. Let all k lower cost edges bid $c' - \epsilon$ for a small $\epsilon > 0$, so their bid is less than the bid of the remaining higher cost edge (whose bid is at least c'). The minimum price k-flow with respect to this bid vector will use the bottom k edges for a total price of $k(c' - \epsilon)$ which approaches $k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)))$.

Finally, we emphasize that the total payment of our first price mechanism in a strong ϵ-Nash equilibrium is at most $kn\epsilon$ more than the VCG payment for the same graph in a Nash equilibrium.

Theorem 3.4.2 *Given a graph G with source s and sink t, the VCG payment for k units of flow from s to t is at least $k(c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k)))$.*

Proof. Let $P_1, \ldots, P_k$ be the k disjoint paths in the selected minimum cost k flow. Fix one path P_i with, say, l edges. We will prove that the sum of payments to edges on this path is at least $c(F_{\mathbf{c}}(k+1)) - c(F_{\mathbf{c}}(k))$. Recall that the VCG payment for an edge e on a minimum cost k-flow is

$$c_e + c(F_{\mathbf{c}}(k, G - \{e\})) - c(F_{\mathbf{c}}(k, G)). \tag{3.11}$$

We construct a new directed multi-graph on the same vertex set as G as follows. We use the term *forward* to mean an edge directed from s to t along the flow path and *backward* to mean an edge directed from t to s. For each edge e on path P_i, add a backward copy of each edge in $F_{\mathbf{c}}(k, G)$ and a forward copy of each edge in $F_{\mathbf{c}}(k, G - \{e\})$, retaining multiplicities. Now add a forward copy of the path P_i to the graph. Label each forward edge e with the cost c_e of the corresponding edge in G and each backward edge with the cost $-c_e$. Then, by equation 3.11, the sum of edge weights in this graph equals the VCG sum of payments to edges on path P_i. Note that this graph is a union of l $s-t$ flows, l $t-s$ flows, and one s-t path. Thus, the in-degree of every vertex except s and t is equal to its out-degree, and for s (t), the out-degree is one more (less) than its in-degree. For every pair of vertices, *cancel* the 2-edge cycles connecting them. That is, if the vertices are connected by k_1 forward edges and k_2 backward edges, replace the edges by $k_1 - k_2$ forward edges if $k_1 > k_2$, $k_2 - k_1$ backward edges if $k_1 < k_2$, or simply remove the edges if $k_1 = k_2$ (this does not change the degree or edge weight properties of the graph discussed above). Call the resulting graph G'. As the sum of edge weights in G' equals the sum of VCG payments to edges on path P_i, we can bound the sum of VCG payments to edges on path P_i by bounding the sum of edge weights in G'.

First note every edge of $F_{\mathbf{c}}(k,G)$ is either non-existent or directed backward in G': for edges $e \in F_{\mathbf{c}}(k,G) - P_i$, e is added exactly once in the backward direction and at most once in the forward direction by each of the l edges in P_i; for edges $e \in P_i$, e is added exactly once in the backward direction and at most once in the forward direction by each of the $l-1$ edges e' in P_i, $e' \neq e$. Furthermore, edge e is added once in the backward direction by itself and once in the forward direction in the last step of the construction of G'.

Select a path P from s to t in G' (such a path exists by the degree properties discussed above). As edges of $F_{\mathbf{c}}(k,G)$ exist only in the backward direction, our path P is a valid augmenting path in the original graph G, and so its weight is at least $c(F_{\mathbf{c}}(k+1,G) - c(F_{\mathbf{c}}(k,G))$ by minimality of $F_{\mathbf{c}}(k+1,G)$. We claim the weight of P is at most the sum of edge weights in G' (which equals the VCG payment), proving the result. This follows from the fact that, due to its degree properties, G' can be written as a union of P and a set of disjoint cycles, and, since $F_{\mathbf{c}}(k,G)$ is a minimum cost k-flow in G, the sum of edge weights on any cycle must be non-negative. Otherwise we could construct a cheaper k-flow in G by replacing the backward edges of a negative cycle with the forward edges in $F_{\mathbf{c}}(k,G)$: specifically, if C is a negative cycle in G' with backward edges A and forward edges B, then $F_{\mathbf{c}}(k,G) - A + B$ is a cheaper k-flow in G. □

Implementation in ϵ-Nash

The simple first-price auction may have costly ϵ-Nash equilibria, as shown in Example 3.2.3. In Section 3.4.1 we used the strong ϵ-Nash solution concept to get around this problem. However, assuming that the bidders will reach an strong ϵ-Nash equilibrium is perhaps too strong an assumption: it requires extensive coordination between agents. In this section, we present a variant of a first-price auction in which every ϵ-Nash equilibrium results in a low price.

One idea to achieve this is to pay edges a bonus that increases as their bid decreases. This encourages edges to submit low bids. However, this has the side-effect of giving edges incentives to bid even below their true cost, as long as they remain off the winning flow. This would make the bargaining problem that edges must solve much more complex, as it would include bargains between winning and losing edges. Alternatively, we could instead send flow on each edge with some probability that increases as the bid decreases. Thus an edge will not bid below its true cost, but it might have an incentive to bid very high. Using a combination of these two ideas, we can construct a payoff function such that an edge will bid close to its true cost if it is not on the lowest true cost flow. This is known as *virtual implementation* in the economics literature (see, for example, Jackson [62]). If the bonuses and probabilities are small enough, then the extra payment will not be very large in expectation, and we can prove a bound on the total payment of the mechanism similar to that in Theorem 3.4.1.

We describe the techniques in this section in the setting of path auctions, although they extend to more general settings as noted. Assume that there is a value B such that no edge bids more than B. (Alternatively, B can be the maximum amount that the buyer is willing to pay.) Further, we assume that the edges are risk-neutral. The mechanism starts by computing a collection of (not necessarily simple) paths $\{P_e\}$. The mechanism then solicits a bid b_e from each edge e. The lowest-price path is almost always picked; however, with a small probability, one of the paths from the collection is picked instead. In addition, each edge is paid a small bonus that depends on the bids. The selection probability and bonus are chosen to ensure that it is optimal for every edge that is *not* on the lowest-price path to bid its true cost. For simplicity, we present the mechanism and analysis for a single unit of flow; the results can easily be extended to any constant $k > 1$ units of flow.

Mechanism RandomPath: The parameters α and τ are selected to be small positive constants such that $\alpha < \min\{n^{-2}B^{-1}, \frac{2}{1+2n}\}$ and $\tau < \alpha n^{-1}B^{-1}$.

1. For each edge e, find P_e, a (not necessarily simple) path from s to t through e.

Let $\mathcal{P} = \{P_e\}_{e \in G}$. Note that an edge e may appear in multiple paths in $\mathcal{P}$.

2. Solicit bids $\mathbf{b} = (b_1, \ldots, b_e, \ldots, b_n)$ from the edges.

3. For each path $P \in \mathcal{P}$, compute

$$\sigma_P = \alpha - \tau \sum_{e \in P} b_e$$

4. Select each path $P \in \mathcal{P}$ with probability σ_P; with probability $(1 - \sum_{P \in \mathcal{P}} \sigma_P)$, select the lexicographically first lowest price path. Call the selected path P^*. Pay each edge $e \in P^*$ its bid b_e.

5. In addition to any payment edge e may get in step 4, pay each edge $e \in G$ the sum $f_e(\mathbf{b}) = \sum_{P \in \mathcal{P}, P \ni e} f_e^P(\mathbf{b})$, where

$$f_e^P(\mathbf{b}) = \alpha(B - b_e) + \tau b_e \sum_{j \in P} b_j - \tau \frac{b_e^2}{2}$$

Our payment rule is constructed in a way that encourages bidders not receiving flow to bid their true cost. Note that the bonus increases as the bid decreases, but the expected selection payment decreases as the bid decreases. Intuitively, we design the bonus and selection probabilities so that the total payoff function is maximized when $b_i = c_i$. Note that if an edge is selected, it incurs its true cost. In this way, the true cost automatically enters his expected payoff function—the mechanism does not need to know the cost in order to achieve the maximum at $b_i = c_i$.

Lemma 3.4.1 *For any edge e* not *on the lowest-price path with bids $\mathbf{b}$, if $b_e \notin [c_e - \sqrt{2\epsilon/\tau}, c_e + \sqrt{2\epsilon/\tau}]$, then $b_e = c_e$ will increase the expected payoff to e by at least ϵ.*

Proof. With the bid vector b, e's expected payoff is

$$\begin{aligned} f_e(\mathbf{b}) \;+\; \sum_{P \ni e} \sigma_P(b_e - c_e) \;&=\; \sum_{P \ni e} [f_e^P(\mathbf{b}) + \sigma_P(b_e - c_e)] \\ &= \sum_{P \ni e} [\alpha B - \tau \frac{b_e^2}{2} + \tau c_e \sum_{j \in P} b_j - \alpha c_e] \end{aligned}$$

Let $g(b_e) = [\alpha(B - c_e) - \tau\frac{b_e^2}{2} + \tau c_e \sum_{j \in P} b_j]$. Then, $g(b_e)$ is a quadratic function of b_e. Observe that $\frac{\partial g(b_e)}{\partial b_e} = -\tau b_e + \tau c_e = 0$ when $b_e = c_e$; at this point, $\frac{\partial^2 g(b_e)}{\partial^2 b_e} = -\tau < 0$. This is true for all paths P containing e. Further, for $\Delta > 0$,

$$g(c_e) - g(c_e + \Delta) = \tau c_e \Delta + \tau \Delta^2/2 - \tau c_e \Delta = \tau \Delta^2/2$$

Similarly, $g(c_e) - g(c_e - \Delta) = \tau\Delta^2/2$. Thus, if $b_e < c_e - \sqrt{2\epsilon/\tau}$, then edge e has incentive to raise his bid to $b_e = c_e$. Similarly, if $b_e > c_e + \sqrt{2\epsilon/\tau}$, then edge e has incentive to decrease his bid to $b_e = c_e$ (even if this puts him on the lowest-price path, then his payoff is still $g(c_e)$ per path so the above calculation still holds). □

Lemma 3.4.1 implies that if ϵ-Nash equilibria exist in mechanism RandomPath, then any edge not on the lowest-price flow must bid close to its true cost. This will help us bound the total expected payment in an ϵ-Nash, but first we must prove that ϵ-Nash equilibria exist in this mechanism. Indeed the same construction as in Theorem 3.2.1 yields an ϵ-Nash equilibrium.[2]

Lemma 3.4.2 *For any cost vector* $\mathbf{c}$ *and any* $\epsilon > 0$*, an* ϵ*-Nash equilibrium always exists in the mechanism RandomPath.*

Proof. Construct a bid vector $\mathbf{b}$ as in Theorem 3.2.1. By this construction, the lowest-cost path equals the lowest-price path. We have $b_e = c_e$ for any edge e that is *not* on the lowest-price path. Edges on the lowest-price path bid close to the maximum they can while still remaining on the lowest-price path (see the proof of Theorem 3.2.1 for the precise construction).

Following the analysis of $g(b_e)$, the expected payoff in Lemma 3.4.1, b_e maximizes e's payoff. (Note that e can only get onto the lowest-price path by bidding below its cost, which would result in a loss.) It remains to show that every edge i on the lowest-price path would not significantly benefit by changing it's bid. Note that, by construction of the bid vector, if i increased its bid by more than $\epsilon/2$, it would no longer be on the lowest-price path. Further, because of the shape of the bonus payoff

[2]Mechanism RandomPath can be extended to the general procurement setting. The proof of the following theorem can be generalized to prove existence of ϵ-Nash equilibria in this setting as well.

function, i's expected gain $g(b_e)$ from the bonus and probability of off-path selection would also drop. Thus, i cannot possibly gain more than ϵ by raising its bid. Consider the possibility that i lowers its bid by x. Then, i would still be on the lowest-price path. It would lose at least $(1 - n\alpha)x$ in profit from being on the lowest-price path, and gain at most $(g_e(b_e) - g_e(b_e - x)) = \tau(\frac{1}{2}x^2 + (c_e - b_e)x)$ in $g_e(b_e)$ per path. As $b_e \geq c_e$ in $\mathbf{b}$, its total gain is at most $\frac{n\tau}{2}x^2$. As $x \leq B$, the loss is more than the gain for any choice of τ less than $2(1 - n\alpha)/(nB)$ or, rewriting in terms of α, $\alpha < \frac{2}{1+2n}$. These conditions can be guaranteed by the choice of α and τ. □

Now, we observe that the values α and τ can be chosen small enough to make the probabilities $\{\sigma_P\}$ and bonuses $f_e^P(\mathbf{b})$ arbitrarily small. Thus, the total payment to edges not on the shortest path is very small. The bound on the payment of the mechanism RandomPath is more sensitive to the value of ϵ because edges not on the lowest-price path get very small payments in expectation. However, we can show that, in the limit as $\epsilon \to 0$, the maximum expected payment in any ϵ-Nash equilibrium is bounded. The following proof can be generalized to the flow setting to derive a bound similar to that in Theorem 3.4.1.

Theorem 3.4.3 *Choose any $\alpha < n^{-2}B^{-1}, \tau < \alpha n^{-1}B^{-1}$. For these values of α and τ,*

$$\lim_{\epsilon \to 0} \max_{\epsilon\text{-}NE\ \mathbf{b}} \{\textit{Total payments with bids } \mathbf{b}\} \to c(2) - c(1) + 3\alpha n^2 B.$$

Proof. Let $\mathbf{b}$ be an ϵ-Nash equilibrium bid vector, for sufficiently small ϵ. The total probability that the mechanism picks a path other than the lowest-price path is bounded by $n\alpha$. Any such path can have at most n edges on it, each with price at most B. Thus, the expected payment for using one of these paths is at most $\alpha n^2 B$. Similarly, we can bound the bonus $f_e(\mathbf{b})$ paid to any edge e: $f_e(\mathbf{b}) \leq n[\alpha B + \tau n B^2]$. This is always less than $2\alpha n B$.

Finally, using Lemma 3.4.1, we know that any edge not on the lowest-price path bids at most $c_e + \sqrt{2\epsilon/\tau}$. Combining this with a similar argument to Theorem 3.4.1, we can bound the total payment to edges on the lowest-price path by

$$b(F(1)) \leq c(2) - c(1) + n\sqrt{2\epsilon/\tau}$$

In the limit as $\epsilon \to 0$, the last term is negligible. Adding up all three sources of payment, we get the required result. □

Recall that mechanism RandomPath needs to compute a set of paths $\{P_e\}$, where P_e is a path from s to t that uses e. If e is to be relevant to the path auction, such a path must exist, however, it is not always straightforward to compute. In particular, if the network is a general *directed* graph, it is NP-hard to compute such a path, since it reduces to the two disjoint paths problem, which is NP-complete [41].

However, there are many interesting classes of graphs for which it is possible to compute such a path P_e in polynomial time, including undirected graphs and directed acyclic or planar graphs [41]. We can also modify the mechanism to ask each bidder to exhibit such a path, thus transferring the computational burden on to the bidders. Also, these paths may be precomputed and used in many executions of the mechanism—they do not depend on the costs or bids.

Another possibility is to use a set of covering paths that do not all terminate at t—this can be easily computed, even for general directed graphs. Then, if the path is picked, some arbitrary traffic is sent along this path. After this "audit" traffic has been delivered, the lowest-price path is used for the intended traffic from s to t. As long as the mechanism can verify that the traffic is correctly delivered, the edges would still have an incentive to bid as specified. Similarly, if we could verify the exact path that the traffic used, we could use non-simple paths to cover the edges; again, a set of non-simple covering paths can easily be found.

3.4.2 Unknown Demand Path Auction

In the previous sections, we studied first-price auctions to meet a known demand, argued that they had stable Nash equilibria, and showed how to adjust this auction so that the equilibria chosen by the auctioneer had relatively small overpayments.

In practice, however, it may not be possible to defer the setting of prices until the demand is known. In this section, we examine the problem of achieving stable prices without advance knowledge of the demand. Instead, the bidders and auctioneer share knowledge of a common prior or *probability distribution* over the possible demands.

Ideally, we would like our results for first-price auctions with known demand to carry over. For example, we proved in Section 3.4.1 that a first price auction for k units of demand led to a payment of $P_k = k[c(F_c(k+1)) - c(F_c(k))]$ in any strong ϵ-Nash equilibrium. It is thus natural to hope that the same auction operating over *random* k also has strong ϵ-Nash equilibria with expected payment $E_k[P_k]$. This turns out to be false—in fact, as we will show, a first-price auction might not even have ϵ-Nash equilibria (recall that strong ϵ-Nash equilibria are a subset of ϵ-Nash equilibria). As ϵ-Nash equilibria do not exist in first-price auctions, we turn to more complex auctions. We will exhibit an auction involving *two parameter bids* that, unlike the single-parameter first-price auction, *does* have ϵ-Nash equilibria. Furthermore, using an indifference-breaking technique similar to that of the mechanism RandomPath, we can restrict the set of equilibria in a variant of this auction to ones with bounded payments. The bound is not quite the $E_k[P_k]$ we hoped to achieve, but does bear a clear resemblance to it. Unfortunately, we are unable to prove that this auction is implementable in polynomial time as it involves solving an integer program. It remains to be seen if further modifications of this auction can result in a polynomial-time auction with bounded payments.

Definitions and Notation

The unknown demand case is modelled as follows: The demand can take any integral value in the range $[1, r]$, where r is a positive integer. Further, there is a known prior distribution on the demand values; say that the demand is k with probability p_k, for $k = 1, 2 \ldots, r$. We assume for simplicity that $p_k > 0$ for all k; our results easily extend to a situation in which $p_k = 0$ for some values of $k \in \{1, \ldots, r\}$.

An auction for the unknown demand case receives bids, and announces flows F_1, F_2, $\ldots$, F_r for each possible demand value. For a first-price auction in this setting, each $F_k \in \mathcal{F}$ must be a minimum price k-flow. We call the collection $\mathcal{F} = \{F_1, F_2 \ldots, F_r\}$ a candidate solution. We also identify a solution $\mathcal{F}$ with the set of edges in the union $F_1 \cup F_2 \cup \cdots \cup F_r$, and say that $i \in \mathcal{F}$ if $i \in F_k$ for some k.

As before, we use $c(\mathcal{F})$ to denote the total expected cost of a solution $\mathcal{F} = (F_1, \ldots, F_r)$ when the individual edge costs are c, and $\tilde{a}(\mathcal{F})$ to denote the price of the flow $\mathcal{F}$ when the bids are $\tilde{a}$. When the auction is clear from the context, we will denote the auction output by $\hat{\mathcal{F}}(\tilde{a})$.

Impossibility of ϵ-Nash Equilibria in First-Price Auctions

In this section, we show that a first-price auction may have *no* ϵ-Nash in the unknown demand case. Intuitively, this is because edges must tradeoff the *probability* of receiving flow with the *profit* of receiving flow. With a high bid, the profit is large, but the probability of winning the auction is low. If the other bids are also high, an edge will prefer to lower its bid to win with a higher probability. This will lead other edges to lower their bids so as to restore their high winning probability. Now, however, the first edge will increase its bid so as to increase its profit at the expense of its winning probability, and so a cycle emerges in the bidding strategies, as the following example shows.

Consider a graph with four parallel edges W, X, Y, and Z between the source and the sink, with true costs w, x, y, and z respectively. The demand is either 1, 2 or 3; for simplicity, let the probability of each demand value be $\frac{1}{3}$. Assign the costs such that $w + 50\epsilon < x + 42\epsilon = y + 12\epsilon = z$. Suppose there W, X, Y, Z bid a, b, c, d respectively. The proof repeatedly uses the ϵ-Nash conditions to show that one of the following must hold: (1) There is an agent who would gain by raising its bid, or, (2) There is an agent who would gain by undercutting another agent to win with a higher probability.

Theorem 3.4.4 *There is no pure-strategy ϵ-Nash equilibrium in the unknown demand first-price auction.*

Proof. First we prove a series of inequalities that the bids must satisfy in an ϵ-Nash equilibrium:

Claim 1: $a, b, c \leq d$.

Proof: First, suppose $d > y + 3\epsilon$, and $c > d$. Then, by changing its bid to $d - \delta$, for small enough δ, Y would be selected with probability 1/3 and so get utility greater than ϵ; thus, any solution in which Y had 0 expected payoff would not be an ϵ-Nash equilibrium. As the same is true for w and x, we must have $a, b, c \leq d$. Now, suppose $d \leq y + 3\epsilon$. Then, $d < z - 3\epsilon$, and as Z is selected with probability 1/3, its payoff is less than $-\epsilon$, which cannot be true in the equilibrium. Thus, in this case too, we have $d \geq a, b, c$.

Claim 2: $d > y - 3\epsilon$.

Proof: If $d \leq y - 3\epsilon$, then Y could not underbid Z without having expected utility less than $-\epsilon$. Hence, Z would be chosen with probability at least $\frac{1}{3}$ (if the demand was 3). But $d < z - 3\epsilon$, and hence Z's expected utility would be less than $-\epsilon$, and hence this cannot be true in an ϵ-Nash equilibrium.

Claim 3: $a, b, c > x + 21\epsilon$.

Proof: Suppose the order of the bids is $a < b < c$. Then, by Claim 1, W wins with probability 1, X with probability 2/3, and Y with probability 1/3. Thus, we must have $(b - a) \leq \epsilon$, $(c - b) \leq 2\epsilon$, and $(d - c) \leq 3\epsilon$, or else one of W, X, Y could increase her profit by ϵ. A similar argument holds if the bids are in a different order. Thus $a, b, c > d - 6\epsilon$. By Claim 2, this implies $a, b, c > y - 9\epsilon$ which equals $x + 21\epsilon$.

Claim 4: $b < c$.

Proof: By Claim 3, $c > x + 21\epsilon$. If we had $b \geq c$, then X could deviate by bidding $c - \delta$. This would involve a bid reduction of at most 6ϵ, but would enable X to win with a $\frac{1}{3}$ additional probability, leading to a net gain of at least ϵ.

Claim 5: $a < b$.

Proof: If $a \geq b$, W could deviate to $b - \epsilon$, resulting in a gain of at least ϵ, as above.

These claims imply that (a, b, c, d) is not an ϵ-Nash equilibrium: We have shown that $a < b < c \leq d$. It also must be true that $(c - b) < 2\epsilon$, and $c > y - 3\epsilon$. Thus, $b > x + 25\epsilon$. Further, $(b - a) < \epsilon$. Hence, X could deviate to $a - \delta$, resulting in a net gain of greater than ϵ. $\square$

Implementation in ϵ-Nash Using a 2-Parameter Bidding Scheme

In this section, we show that by allowing 2-parameter bids, we can define an auction with ϵ-Nash equilibria. Intuitively, a two-parameter auction gets around the problem of a single-parameter auction by letting the edges express their preferences over the entire price-probability space. It allows to an edge to bid a "price" such that the expected payment of any edge with a non-zero probability of winning is equal to its price. In particular, we will allow edges to report their cost along with a *demanded profit* and then guarantee that the *expected* payment of a winning edge is exactly its reported cost plus its demanded profit.

Auction 2-Parameter:

In the following auction, each edge i submits a pair $\tilde{a}_i = (\tilde{c}_i, \tilde{u}_i)$ as its bid, where $\tilde{c}_i$ is interpreted as the reported cost of edge i, and $\tilde{u}_i$ is interpreted as the profit that edge i demands.

1. Define an indicator variable x_{ik} for the event that edge i is on the selected flow F_k, and y_i for the event that edge i is selected to be on some flow. Also, for any node α in the network, let $\text{In}(\alpha)$ denote the set of incoming edges, and $\text{Out}(\alpha)$ denote the set of outgoing edges. Find an optimal solution to the following integer program (IP for short).

$$\begin{aligned}
\text{minimize} \quad & \sum_{k=1}^{r}\left[p_k \sum_{i\in E} \tilde{c}_i x_{ik}\right] + \sum_{i\in E} y_i \tilde{u}_i && (3.12)\\
\text{subject to} \quad & \forall\, \alpha \neq s,t,\ \forall\, 1 \le k \le r: \sum_{i\in \text{Out}(\alpha)} x_{ik} - \sum_{i \in \text{In}(\alpha)} x_{ik} = 0 && (3.13)\\
& \forall\, 1 \le k \le r: \sum_{i\in \text{Out}(s)} x_{ik} - \sum_{i\in \text{In}(s)} x_{ik} = k && (3.14)\\
& \forall\, 1 \le i \le n,\ \forall\, 1 \le k \le r:\ y_i - x_{ik} \ge 0 && (3.15)\\
& \forall\, 1 \le i \le n,\ \forall\, 1 \le k \le r:\ x_{ik} \in \{0,1\} \\
& \forall\, 1 \le i \le n,\ \forall\, 1 \le k \le r:\ y_i \in \{0,1\}
\end{aligned}$$

2. Set $F_k = \{i : x_{ik} = 1\}$ and $\mathcal{F} = \{F_1, \ldots, F_r\}$. For each $i \in \mathcal{F}$, calculate the probability $\rho_i = \sum_{\{k|i \in F_k\}} p_k$ that i wins. If the actual demand turns out to be r, use the edges in F_k to route the flow, and pay each edge $i \in F_k$ a sum $\tilde{c}_i + \frac{\tilde{u}_i}{\rho_i}$.

Remark 3.4.1 *Notice with these payments, IP 3.12 chooses a flow solution which minimizes the total expected payment for a fixed bid vector: constraints 3.13 and 3.14 guarantee that the set $F_k = \{i : x_{ik} = 1\}$ form a feasible k-flow and constraint 3.15 guarantees that edges selected to be on a flow are paid their reported cost.*

We now prove that this auction has ϵ-Nash equilibria. To develop some intuition for the proof, recall that in the known demand case, only bidders on the cheapest flow had the flexibility to submit a bid significantly more than their cost and still win the auction. A similar statement holds here when the first parameter of all bids are restricted to be equal to the cost. In particular, the following bid vector should intuitively be an ϵ-Nash equilibrium: for edges $i \notin \hat{\mathcal{F}}(\tilde{a})$, set $\tilde{a}_i = (c_i, 0)$; for edges $i \in \hat{\mathcal{F}}(\tilde{a})$, set $\tilde{a}_i = (c_i, \tilde{u}_i)$ where the $\tilde{u}_i$ divide up the available profit (the difference between the price of the cheapest and second cheapest flow). Edges $i \notin \hat{\mathcal{F}}(\tilde{a})$ can not afford to decrease their bids and have no chance of winning by increasing their bids, so they have no profitable deviation. As the expected payment of any edge $i \in \hat{\mathcal{F}}(\tilde{a})$ is the same regardless of their winning probability, these edges also have no incentive to decrease their bid. By an appropriate choice of $\{\tilde{u}_i\}$, we can arrange that if they increase their bid then they will drop out of the solution.

We formalize this argument by using a linear-programming technique similar to the proof of Theorem 3.2.1. The variables of the linear program (LP) are the profits demanded by the bidders (that is, the second parameter of the bid). The LP constrains the total profit demanded by a set of bidders to be at most the cost-savings induced by this set. Let $\mathcal{F}^*$ be the minimum cost solution and u_i be a variable corresponding to the profit demanded by bidder i. Consider the following linear program.

$$
\begin{aligned}
\text{maximize} \quad & \sum_{i=1}^{n} u_i && (3.16)\\
\text{subject to} \quad & \forall \text{ feasible solutions } \mathcal{F}: \ \sum_{i \notin \mathcal{F}} u_i \leq c(\mathcal{F}) - c(\mathcal{F}^*) && (3.17)\\
& \forall 1 \leq i \leq n: \ u_i \geq 0
\end{aligned}
$$

This LP is clearly feasible as $u_i = 0$ for all i satisfies all constraints. We will show that for an optimal solution $\{u_i^*\}$, the set of bids $\{(c_i, \max\{0, u_i^* - \epsilon/(2n)\})\}$ form an ϵ-Nash equilibrium.

Theorem 3.4.5 *For any $\epsilon > 0$, let $u_i = \max\{0, u_i^* - \frac{\epsilon}{2n}\}$ and consider the bid profile defined by $a_i^- = (c_i, u_i)$ for each edge i. Then a^- is an ϵ-Nash equilibrium.*

The proof uses three lemmas regarding the bids of the minimum price solution. The first lemma shows that edges $i \notin \mathcal{F}^*$ not in the minimum cost solution have zero demanded profit (that is, $\tilde{u}_i = 0$). This confirms the intuition that, as in the known-demand case, only edges in the minimum-cost solution can demand a payment significantly more than their cost.

Lemma 3.4.3 *The minimum cost solution includes all i with $u_i^* > 0$.*

Proof. Consider the inequality in LP 3.16 corresponding to solution $\mathcal{F}^*$. This inequality states that $\sum_{i \notin \mathcal{F}^*} u_i^* \leq 0$. Together with the non-negativity constraints, this implies that $u_i^* = 0$ for all edges i not in the minimum cost solution. Thus the minimum-cost solution includes all edges i with $u_i^* > 0$. □

The second lemma supports the intuition that the minimum cost solution $\mathcal{F}^*$ has minimum price.

Lemma 3.4.4 *The minimum cost solution is a minimum expected price solution with respect to bids $\tilde{a}_i = (c_i, \tilde{u}_i^*)$.*

Proof. As the first parameter of any bid $\tilde{a}_i$ is c_i, the expected price of any solution $\mathcal{F}$ is equal to its expected cost plus the sum of demanded profits of its edges. Since $u_i^* = 0$ for $i \notin \mathcal{F}^*$, we have

$$\tilde{a}(\mathcal{F}^*) = c(\mathcal{F}^*) + \sum_{i=1}^{n} u_i^*. \tag{3.18}$$

For any flow $\mathcal{F}$, the inequality 3.17 of LP 3.16 corresponding to $\mathcal{F}$ states that $c(\mathcal{F}^*) \leq c(\mathcal{F}) - \sum_{i \notin \mathcal{F}} u_i^*$. Adding $\sum_{i=1}^{n} u_i^*$ to both sides and using equation 3.18 gives $\tilde{a}(\mathcal{F}^*) \leq \tilde{a}(\mathcal{F})$. □

The third lemma argues that no single edge is essential to the minimum price solution. In other words, for each edge there is a minimum price solution that avoids that edge. Intuitively, if this were not the case, then the edge ought to be able to demand extra profit.

Lemma 3.4.5 *With bids $\tilde{a} = (c_i, \tilde{u}_i^*)$, for any edge i there is a minimum price solution $\mathcal{F}^{(i)}$ that does not contain i.*

Proof. Let $\mathcal{F}$ be a solution not containing i and suppose every minimum price solution contains i. Then, by Lemma 3.4.4, the inequality corresponding to $\mathcal{F}$ must be strict. As this holds for *any* solution $\mathcal{F}$ not containing i, *every* inequality containing u_i^* is strict. Therefore $u_i^* + \delta$ is a feasible solution for some $\delta > 0$, contradicting the optimality of solution u_i^*. □

Proof of Theorem 3.4.5. Suppose a^- is not an ϵ-Nash equilibrium. Then, there is some i which can change its bid to increase its payoff by ϵ. Let (c_i', u_i') be i's successful strategy, and let a' denote the bid profile given by $a_i' = (c_i', u_i')$ and $a_j' = a_j^-$ for all $j \neq i$. Let $\mathcal{F}$ be the solution output by the mechanism with bids a^- and $\mathcal{F}'$ be the solution output by the mechanism with bids a'. Note it must be the case that $i \in \mathcal{F}'$.

We observe that the change in expected price of $\mathcal{F}'$ from bid vector a' to a^- is at least ϵ. Let ρ_i be the probability (over the demand distribution) that i is in solution $\mathcal{F}'$. Then i's utility increases from u_i to $u_i' + (c_i' - c_i)\rho_i$, and so by assumption $u_i' + (c_i' - c_i)\rho_i - u_i \geq \epsilon$. Therefore, as only i's bid changes and $i \in \mathcal{F}'$,

$$a'(\mathcal{F}') - a^-(\mathcal{F}') = (u_i' + \rho_i c_i') - (u_i + \rho_i c_i) \geq \epsilon. \tag{3.19}$$

Now, by Lemma 3.4.5, there is a solution $\mathcal{F}^{(i)}$ not containing i which has minimum price with respect to bids $\tilde{a} = (c_i, \tilde{u}_i^*)$. Let $\mathcal{F}^{(i)}$ be that solution. Then $\tilde{a}(\mathcal{F}^{(i)}) \leq \tilde{a}(\mathcal{F}')$. Note that for any solution $\mathcal{F}$, the price with respect to bids a^- is within $\epsilon/2$ of the price with respect to bids $\tilde{a}$: $a^-(\mathcal{F}) \leq \tilde{a}(\mathcal{F}) \leq a^-(\mathcal{F}) + \epsilon/2$. Therefore

$$\begin{aligned} a'(\mathcal{F}^{(i)}) &= a^-(\mathcal{F}^{(i)}) \\ &\leq \tilde{a}(\mathcal{F}^{(i)}) \\ &\leq \tilde{a}(\mathcal{F}') \\ &\leq a^-(\mathcal{F}') + \epsilon/2 \\ &< a'(\mathcal{F}'), \end{aligned}$$

where the last inequality follows from inequality 3.19. This contradicts the optimality of $\mathcal{F}'$. □

Randomized 2-parameter Auction

The mechanism presented above has an ϵ-Nash equilibrium corresponding to every optimal solution to LP 3.16, but we cannot guarantee that there are no other ϵ-Nash equilibria. As a result, it was not possible to bound the total payoff to the edges. In this section, we consider a slightly modified mechanism in which we add a small random payment, as in the mechanism RandomPath. We prove that, with this modification, it is possible to bound the total payment. Our mechanism uses Auction 2-Parameter as a subroutine and therefore is not implementable in polynomial-time.

Randomized 2-parameter Auction: As before, each edge i bids a pair $\tilde{a}_i = (\tilde{c}_i, \tilde{u}_i)$ where $\tilde{c}_i$ is interpreted as i's reported cost, and $\tilde{u}_i$ is interpreted as i's demanded profit.

1. **The 2-parameter auction.** This step is conducted exactly as in Auction 2-Parameter by solving IP 3.12 to select the minimum price solution.

2. **Rejection.** If for any edge not in the selected solution $\tilde{u}_i \neq 0$, reject the bid profile. No edge is selected and no flow is sent.[3]

3. **The randomized audit.** For edges on a random source-destination path, the payoff is based entirely on the $\tilde{c}_i$ component of the bid, and is constructed as in the mechanism RandomPath. The parameters α, τ, and B are as defined in the mechanism RandomPath. If an edge has true cost c_i and bids $(\tilde{c}_i, \tilde{u}_i)$, its expected payoff from this component is $g(\tilde{c}_i) = \tau[c_i\tilde{c}_i - \frac{\tilde{c}_i^2}{2}]$. The exact form of the payoff was derived in the proof of Lemma 3.4.1.

The audit component of the auction encourages edges to submit bid vectors in which their costs are nearly truthful. The first two steps of the auction help guarantee that the demand profits form a nearly feasible solution to LP 3.16. These facts allow us to derive bounds on the expected payment as stated in the following theorem.

Theorem 3.4.6 *The total price paid by the auctioneer in the randomized 2-parameter auction is at most*

$$\left[\sum_{j=1}^{r} jp_j c(F_{r+1})\right] - rc(\mathcal{F}) + nr\sqrt{2\epsilon/\tau} + 3\alpha n^2 B.$$

The result of Theorem 3.4.6 stands in an interesting relation to that of Theorem 3.4.1. We do not achieve the intuitively appealing bound of the expectation of the bounds on the known demand auction in Section 3.4.1, *i.e.*, $E_j[P_j] = \sum_{j=1}^{r} jp_j(c(F_{j+1}) - c(F_j))$ but instead we achieve $\sum_{j=1}^{r} rp_j(c(F_{r+1})(j/r) - c(F_j))$. In other words, the external multiplier j is replaced by r (a larger quantity), while in the first term the quantity $c(F_{j+1})$ is replaced by $c(F_{r+1})(j/r)$, which can also be larger because the cost of j units of flow is a convex function of j. Our Theorem 3.4.1

[3]This step ensures that, for all edges i not in the winning solution, $\tilde{u}_i$ is 0. Alternatively, we could ensure that these $\tilde{u}_i$ are close to zero (which is enough for our purposes) by charging a small tax to all bidders who submit a positive $\tilde{u}_i$ component of the bid.

is therefore weaker in two important respects than Theorem 2, but it does have a similar overall structure.

To prove Theorem 3.4.6, we first show that *all* edges are nearly truthful about their costs in equilibrium:

Lemma 3.4.6 *Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be an ϵ-Nash equilibrium of the randomized 2-parameter auction. Then, for all i,*

$$c_i - \sqrt{2\epsilon/\tau} \leq \tilde{c}_i \leq c_i + \sqrt{2\epsilon/\tau}$$

Proof. We argue that player i can always do better by bidding his true cost; the bounds follow from the ϵ-Nash equilibrium condition and the expected-payoff graph of the randomized path audit. Let ρ_i be the probability of i being included in the lowest price solution in the ϵ-Nash equilibrium $\tilde{a}$. If $\rho_i = 0$, then i's entire expected payoff is due to her expectation of winning in the randomized path audit, and the bounds on $\tilde{c}_i$ follow directly. The same argument holds if $\rho_i > 0$ but i receives a negative expected payoff from the 2-parameter auction (because her bid $\tilde{c}_i$ was too low).

Now, suppose $\rho_i > 0$, and, further, i receives a positive payoff from the 2-parameter auction in the ϵ-Nash equilibrium. Consider the strategy $a'_i = (c_i, u'_i)$ with $u'_i = \tilde{u}_i + \rho_i[\tilde{c}_i - c_i]$. ($i$ received a non-negative profit under $\tilde{a}$, so it follows that u'_i is non-negative.) Let $\mathcal{F}$ be the solution chosen in the 2-parameter part of the mechanism when the bids are $\tilde{a}$. Note that if i were to deviate from $\tilde{a}_i$ to a'_i, the price of $\mathcal{F}$ would not change: the change in the utility component would exactly cancel the change in the cost component. Also, for any other flow $\mathcal{F}'$ that did not use i, the price of $\mathcal{F}'$ would not change with i's deviation; thus, using the consistency of the tie-breaking rule, $\mathcal{F}'$ would not be chosen above $\mathcal{F}$. Thus, we conclude that i remains in the winning solution (which need not be $\mathcal{F}$) under the bids a'_i.

Next, observe that i's expected payoff from the 2-parameter auction (with bid a'_i) is u'_i, because i bids her cost truthfully and is in the winning solution. This is exactly the same as i's payoff $\rho_i[\tilde{c}_i - c_i] + \tilde{u}_i$ from the 2-parameter auction in the ϵ-Nash equilibrium $\tilde{a}$.

To prove the bounds on $\tilde{c}_i$, we compare i's payoff from the randomized part of the mechanism with bids $\tilde{a}_i$ and a'_i. The bounds follow directly from the form of the randomized audit payoffs. □

Using the fact that the costs are nearly truthful, we can show that the utility values are an (almost) feasible solution to LP 3.16, and hence, derive the following bound on the total payment. We use the linear programming formulation given in LP 3.16, only this time we define the LP with respect to the *reported* costs rather than the true costs. Rewriting, we get

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{n} u_i \\ \text{subject to} \quad & \forall \text{ feasible solutions } \mathcal{F}: \ \sum_{i \notin \mathcal{F}} u_i \leq \tilde{c}(\mathcal{F}) - \tilde{c}(\mathcal{F}^*) \\ & \forall\, 1 \leq i \leq n: \ u_i \geq 0 \end{aligned} \tag{3.20}$$

where $\mathcal{F}^*$ is now the minimum cost solution with respect to costs $\tilde{c}$.

Let $\tilde{a} = (\tilde{c}, \tilde{u})$ be any ϵ-Nash equilibrium of the Randomized 2-Parameter Auction. Let $\mathcal{F}^*$ be a minimum cost solution with respect to costs $\tilde{c}$, and let F_{r+1} be a minimum cost $(r+1)$-flow with respect to costs $\tilde{c}$.

Lemma 3.4.7 *Let u be any feasible solution to LP 3.20. Then for bids $\tilde{a} = \{(\tilde{c}_i, u_i^*)\}$, the minimum price solution $\mathcal{F}$ satisfies*

$$\tilde{a}(\mathcal{F}) \leq \tilde{c}(F_{r+1}) \sum_{j=1}^{r} j p_j \quad - \quad r\tilde{c}(\mathcal{F}^*).$$

Proof. Throughout this proof, minimum cost refers to minimum cost with respect to cost vector $\tilde{c}$. Consider an integral $(r+1)$-flow F_{r+1} minimizing $\tilde{c}(F_{r+1})$. Then F_{r+1} is a minimum cost $(r+1)$-flow and consists of $(r+1)$ disjoint paths $\{P_1, \cdots, P_{r+1}\}$ from s to t. For each $k \in \{1, 2, \cdots, r, r+1\}$, define $F_r^{-k} = F_{r+1} \backslash P_k$, that is, the r-flow obtained by dropping the k'th path. Extend F_r^{-k} to a collection of flows $\mathcal{F}^{-k} = \{F_1^{-k}, F_2^{-k}, \cdots, F_r^{-k}\}$, where F_j^{-k} consists of the j lowest-priced paths in F_r^{-k}.

Noting that F_j^{-k} has cost at most $\frac{j}{r}$ that of F_r^{-k},

$$\tilde{c}(\mathcal{F}^{-k}) \leq \tilde{c}(F_r^{-k}) \sum_{j=1}^{r} p_j \frac{j}{r}.$$

Now, summing the inequality corresponding to $\mathcal{F}^{-k}$ over all k, we get:

$$\sum_{k=1}^{r+1} \sum_{i \notin F_r^{-k}} u_i \leq \sum_{k=1}^{r+1} \left(\tilde{c}(F_r^{-k}) \sum_{j=1}^{r} p_j \frac{j}{r} - \tilde{c}(\mathcal{F}^*) \right).$$

Note that the left hand side includes each element of F_{r+1} exactly r times. Similarly, the flows F_r^{-k} in the right hand side cover F_{r+1} exactly r times. Thus,

$$(r+1)\sum_{i=1}^{n} u_i - \sum_{i \notin F_r^{-k}} u_i \leq r\tilde{c}(F_{r+1}) \sum_{j=1}^{r} p_j \frac{j}{r} - (r+1)\tilde{c}(\mathcal{F}^*),$$

and so,

$$\begin{aligned}
\tilde{a}(\mathcal{F}) &\leq \tilde{a}(\mathcal{F}^*) \\
&\leq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i \\
&\leq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i + r\sum_{i=1}^{n} u_i - \sum_{i \notin F_r^{-k}} u_i \\
&\leq \tilde{c}(F_{r+1}) \sum_{j=1}^{r} j p_j - r\tilde{c}(\mathcal{F}^*).
\end{aligned}$$

□

Now, to prove our main theorem, we simply need to prove that the bid profile is a feasible solution of the linear program.

Proof of Theorem 3.4.6. Similar to Theorem 3.4.3, the total probability that the mechanism picks a path in the randomized audit is bounded by $n\alpha$. Any such path can have at most n edges on it, each with price at most B. Thus, the expected payment for using one of these paths is at most $\alpha n^2 B$. Similarly, we can bound the bonus $f_e(\mathbf{b})$ paid to any edge e: $f_e(\mathbf{b}) \leq n[\alpha B + \tau n B^2]$. This is always less than $2\alpha n B$.

Now we show that vector $\tilde{u}$ of demanded profits in bid profile $\tilde{a}$ is a feasible solution to LP 3.20. By assumption, for all losers, the demanded profit is zero. Therefore,

$$\tilde{a}(\mathcal{F}) = \tilde{c}(\mathcal{F}) + \sum_{i=1}^{n} u_i \geq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i.$$

Consider any solution $\mathcal{F}'$ and note that

$$\tilde{c}(\mathcal{F}') + \sum_{i \in \mathcal{F}'} u_i = \tilde{a}(\mathcal{F}') \geq \tilde{a}(\mathcal{F}) \geq \tilde{c}(\mathcal{F}^*) + \sum_{i=1}^{n} u_i,$$

and so the constraint corresponding to $\mathcal{F}'$ is satisfied. Therefore $\tilde{u}$ is a feasible solution. Since the $\tilde{u}$ satisfy the conditions of Lemma 3.4.7, noting that for any set of edges F, $c(F) - n\sqrt{2\epsilon/\tau} \leq \tilde{c}(F) \leq c(F) + n\sqrt{2\epsilon/\tau}$, we can apply Lemma 3.4.7 to get the result. $\square$

Chapter 4

Cost-Sharing Auctions

Consider a situation where a group of customers (the agents) wish to buy a service such as connectivity to a network. The total cost of this service is a function of the group of customers that is serviced: a group of customers in distant towns might incur a larger cost than a group of customers in the same town. The service provider wants to run an auction in order to determine which subset of agents to service and at what cost. To this end, he might implement the Vickrey-Clark-Groves (VCG) mechanism, and thus have a strategyproof mechanism which services the most efficient group (that is, a group for which the sum of valuations minus cost is maximized). There are, however, several drawbacks to this solution. First, it is impossible for the VCG mechanism, or for that matter *most* efficient mechanisms, to be *budget-balanced* (that is, charge prices that exactly recover the cost of the service).[1] A mechanism which under-charges or over-charges customers is not economically viable. Second, although the mechanism is strategyproof, agents can benefit by lying if they coordinate their bidding strategies and so the mechanism is not *group strategyproof.*

In a group strategyproof mechanism, not even a group of agents should be able to benefit by cooperatively lying. This discourages complicated bidding strategies, and reduces concerns that the equilibrium of a truthful mechanism might not be stable if bidders collude. In the Vickrey-Clark-Groves (VCG) auction for a single item, for

[1]This impossibility result can be avoided by imposing a compatibility condition on individual beliefs and using a Bayesian model [21]. See Moulin and Shenker [87] for a discussion of the tradeoff between budget-balance and efficiency in our setting.

example, the two highest bidders can collude to bid just above the value of the third-highest bidder. With such a strategy, the highest bidder improves his utility without harming the second-highest bidder.[2] Group strategyproofness attempts to guard against these coalitions. In other words, group strategyproof mechanisms require that, even for groups of agents, truthfulness remains a dominant strategy.

The notion of group strategyproofness bears an interesting relationship to that of strong Nash equilibria introduced in Chapter 3. As strong Nash equilibria strengthened the notion of Nash equilibria by permitting collusion in the strategies, group strategyproofness strengthens the notion of strategyproofness in the same way. The key difference in this analogy is that in a coalition in a strong Nash setting, every member was required to strictly benefit. In the group strategyproof setting, we only require that no member is sacrificed to benefit the rest. An even stronger notion can be defined if side-payments are permitted. In this case, the only requirement for a coalition to form is that the total utility of the colluding group must strictly increase. Side-payments require a transfer of money between bidders which might be restricted in some settings either due to legal concerns or issues of trust, and so we do not consider side-payments in this chapter. For a discussion of collusion with side-payments, see Goldberg and Hartline [47].

In this chapter, we study the budget-balance properties of group strategyproof mechanisms. At the base of any mechanism for these problems lies a *cost-sharing* scheme, or method for sharing the cost of the service among the serviced customers. Cost-sharing schemes are of independent interest and have been studied extensively especially in the context of the allocation of a public good (see, for example, [86] and [118]). The question of what constitutes an equitable cost-sharing is difficult to define and has been the subject of centuries of thought, dating from Aristotle's proclamation of "equal treatment of equals and unequal treatment of unequals in proportion to their inequality" in his book on Nicomachean Ethics [7] through modern times. One plausible notion of equity is that of *cross-monotonicity* or *population*

[2]Note that this coalition was forbidden by the strong Nash equilibrium notion defined in Chapter 3 as the second-highest bidder's utility is not strictly increased. In fact, any dominant strategy equilibrium is also a strong Nash equilibrium but not necessarily group strategyproof.

monotonicity (see [112] for a survey). Intuitively, cross-monotonicity requires that the price charged to any individual in a group does not increase as the group expands. There is a large body of literature [25, 26, 57, 66, 85, 94, 107, 110] on cross-monotonic cost-sharing schemes for *submodular* cost functions[3], a subclass of cost functions of particular interest. Many mechanisms exist, prominent among them the Shapley value [107], which minimizes the worst-case efficiency loss, and the Dutta-Ray solution [26]. Both of these are budget-balanced and cross-monotonic for any submodular cost function, but not efficient.

As observed by Moulin and Shenker [87], cross-monotonic cost-sharing schemes can be used to construct *group strategyproof mechanisms*, or mechanisms that resist collusion among the agents. As submodular cost functions have cross-monotonic cost-sharing schemes that are budget-balanced, the group-strategyproof mechanisms derived from these schemes are themselves budget-balanced. These schemes have been applied to derive group-strategyproof mechanisms for important submodular cost functions such as multicast on a tree [36, 34, 33]. Unfortunately, many classes of important cost functions arise from (often NP-hard) optimization problems and fail to be submodular. For example, the cost of providing the service for a set S of agents could be expressed as the cost of building the cheapest Steiner tree that covers the elements of S, or the minimum cost of opening facilities and connecting each member of S to an open facility. These two games, and many others of practical import, are instances of covering games. For such games, it is usually impossible for a cross-monotonic cost sharing scheme to be budget-balanced. Moreover, even if a budget-balanced cross-monotonic cost sharing scheme exists, it might be hard to compute. Therefore, it is natural to consider cost sharing schemes that are *approximately budget balanced*, that is, they recover only a fraction of the cost of the service.[4] Approximately budget-balanced schemes have been proposed for minimum spanning

[3]Sometimes called *concave* games in the cooperative game theory literature.

[4]Alternatively, we can relax the definition of budget balance by allowing the scheme to recover at least the cost of the service and at most a small multiple of the cost of the service. This definition seems more reasonable, since a business usually needs to at least recover its costs. However, the two definitions are equivalent up to a constant multiple. To be consistent with other papers on this topic, we use the first definition in this chapter.

tree [64, 68], Steiner tree [64], Steiner forest [72], facility location [90], and connected facility location [78].

We can derive simple bounds on the budget-balance factor of combinatorial optimization games using the integrality gaps of the "natural" LP-relaxations. The cross-monotonicity of a cost sharing scheme implies that for every set of agents the cost shares form an allocation in the core of the game (see Section 4.1 for definitions). Therefore, the best budget-balance factor achievable by a cross-monotonic cost sharing scheme cannot be better than that of a cost sharing in the core. A simple extension of the classic Bondareva-Shapley theorem [13, 108] implies that the best budget-balance factor for a cost sharing in the core of integer covering games is equal to the integrality gap of the "natural" LP-relaxation of the problem (this fact was observed by Jain and Vazirani [64]). This line of reasoning proves bounds on cross-monotonic cost sharing schemes for many combinatorial optimization games. In particular, metric facility location, vertex cover, and set cover games cannot recover more than a $\frac{1}{1.463}$, $\frac{1}{2}$, and $\frac{1}{\ln n}$ fraction of the total cost, respectively. Prior to this work, this was the only method known for upper bounding the cross-monotonic cost sharing schemes. In this chapter, we show stronger upper bounds for several combinatorial optimization games using a novel technique based on the probabilistic method that will be explained in Section 4.2. In particular, we prove that the best budget-balance factor achievable for the facility location game is $\frac{1}{3}$, proving optimality of the scheme given by Pál and Tardos [90]. Also, for the vertex cover and set cover games, we show that no cross-monotonic cost sharing scheme can recover more than an $O(n^{-1/3})$ and $O(\frac{1}{n})$ fraction of the total cost, respectively. We also apply this technique to several other games including the maximum flow and the maximum matching games. In subsequent work, Könemann et al. [73] used our techniques to prove a tight bound of $\frac{1}{2}$ on the budget-balance factor of the Steiner tree game.

One might wonder if our negative results on cross-monotonic cost sharing schemes imply similar negative results for group-strategyproof mechanisms. As we know that there are group-strategyproof mechanisms that do not correspond to any cross-monotonic cost-sharing scheme, our negative results for cross-monotonic schemes do not immediately imply negative results for group-strategyproof mechanisms. However, we give a partial characterization of group-strategyproof mechanisms in terms of cost-sharing schemes that satisfy a condition weaker than cross-monotonicity, and use this characterization to prove that group-strategyproof mechanisms that satisfy an additional condition called *upper continuity* give rise to cross-monotonic cost-sharing schemes, and therefore our negative results apply to such mechanisms.

The rest of this chapter is organized as follows. In Section 4.1, we present the definitions of cross-monotonic cost sharing schemes. Section 4.2 contains a description of our upper bound technique, highlighted by the example of the edge cover game (Section 4.2.1), and proof of bounds for the set cover game (Section 4.2.2), the vertex cover game (Section 4.2.3), the facility location game (Section 4.2.4) , and several combinatorial profit-sharing games (Section 4.2.5). In Section 4.3 we define group-strategyproof mechanisms and prove several results relating such mechanisms to cost-sharing schemes.

The results of this chapter are based on joint work with Mahdian and Mirrokni [61].

4.1 Setting

Let $\mathscr{A}$ denote a set of n agents who are interested in a service. A *cost-sharing game* is defined by a function $C : 2^{\mathscr{A}} \mapsto \mathbb{R}^+ \cup \{0\}$ which for every set $S \subseteq \mathscr{A}$, gives the cost $C(S)$ of providing service to S.[5] A *cost allocation* for a set $S \subseteq \mathscr{A}$ is a function $\psi : S \mapsto \mathbb{R}^+ \cup \{0\}$, that for each agent $i \in S$, specifies the share $\psi(i)$ of i in the total cost of servicing S. A *cost-sharing scheme* is a collection of cost allocations for every $S \subseteq \mathscr{A}$.

Definition 11 *A* cost sharing scheme *is a function $\xi : \mathscr{A} \times 2^{\mathscr{A}} \mapsto \mathbb{R}^+ \cup \{0\}$ such that, for every $S \subset \mathscr{A}$ and every $i \notin S$, $\xi(i, S) = 0$.*

[5]This is similar to the notion of a *coalitional game with transferable payoff*, where the cost function is replaced by a function that gives the value, or the worth of each set. This notion was first defined by von Neumann and Morgenstern [116].

Intuitively, we think of $\xi(i,S)$ as the share of i in the total cost if S is the set of agents receiving the service.

Ideally, we want cost sharing schemes (and cost allocations) to be *budget-balanced*, that is, for every $S \subseteq \mathscr{A}$, $\sum_{i \in S} \xi(i,S) = C(S)$. Budget-balance is desirable as it guarantees economic viability of the auction. However, it is not always possible to achieve budget balance in combination with other properties, or even if it is possible, it might be computationally hard to compute the cost shares. Therefore, we relax this notion to the notion of *α-budget balance* (for some $\alpha \leq 1$).

Definition 12 *A cost sharing scheme ξ is* α-budget-balanced *if, for every $S \subseteq \mathscr{A}$, $\alpha C(S) \leq \sum_{i \in S} \xi(i,S) \leq C(S)$.*

This definition guarantees that the mechanism does not over-charge agents, but it may under-charge them. Alternatively, one could define α-budget balance as $C(S) \leq \sum_{i \in S} \xi(i,S) \leq \frac{1}{\alpha} C(S)$ and equivalently relax the notion of α-core (see Definition 13). All negative results hold without modification in this alternative framework as well; the positive results extend by multiplying each $\xi(i,S)$ by $\frac{1}{\alpha}$. To be consistent with other papers, we use the first definition in this chapter.

In addition to budget balance, we usually require cost allocations and cost-sharing schemes to satisfy additional properties. One property that is extensively studied in the classic cooperative game theory literature [8, 13, 43, 103, 108, 109] is the property of being in the *core*, first suggested by Edgeworth [28] in 1881. This property intuitively says that no subset of agents should be overcharged for the service.

Definition 13 *A cost allocation ψ for a set $S \subseteq \mathscr{A}$ is in the* α-core *if and only if it is α-budget balanced and for every $T \subseteq S$, $\sum_{i \in T} \psi(i) \leq C(T)$. A cost-sharing scheme ξ is in the α-core if and only if for every S, $\xi(\cdot, S)$ is in the α-core.*

Another property, which was studied by Moulin [85] and Moulin and Shenker [87] in order to design *group-strategyproof mechanisms* (see Section 4.3), and has recently received considerable attention in the computer science literature (see, for example, [64, 66, 68, 90]), is *cross-monotonicity* (or *population monotonicity*). This property captures the notion that agents should not be penalized as the serviced set grows. Namely,

Definition 14 *A cost sharing scheme ξ is* cross-monotone *if for all $S, T \subseteq \mathscr{A}$ and $i \in S$, $\xi(i, S) \geq \xi(i, S \cup T)$.*

It is a simple exercise to show that every α-budget-balanced cross-monotonic cost sharing scheme is in the α-core, but the converse need not hold. Therefore, cross-monotonicity is strictly stronger than the core condition. Using this fact and a simple extension of the classic Bondareva-Shapley theorem [13, 108] (see Jain and Vazirani [64]), one can derive upper bounds on the budget-balance factor of cross-monotonic cost-sharing schemes for covering games in terms of the integrality gap of their LP formulation. In the next section, we derive a technique based on the probabilistic method which yields stronger bounds.

4.2 Upper bounds for cross-monotonic cost sharing schemes

In this section we present the main idea behind our upper bound technique and prove upper bounds for several games defined based on combinatorial optimization problems. We explain the technique in Section 4.2.1 with a simple example of the edge cover game and then extend it to the set cover game in Section 4.2.2. Sections 4.2.3, 4.2.4, and 4.2.5 contain the proofs of our bounds for the vertex cover, facility location, and several other games.

4.2.1 A simple example: the edge cover game

In this section, we explain our technique using the edge cover game as a guiding example. The edge cover game is defined as follows.

Definition 15 *Let $G = (V, E)$ be a graph with no isolated vertices. The set of agents in the edge cover game on G is the set of vertices of G. Given a subset S of vertices,* *the cost of S is the minimum size of a set $F \subseteq E$ of edges such that for every $v \in S$, at least one of the edges incident to v is in F. Such a set F is called an* edge cover *for S.*

It is easy to see that for every set S, one can obtain a minimum edge cover of S by taking a maximum matching on S and adding one edge for every vertex that is not covered by the maximum matching (see [23]). Using this fact, we can give a cost-sharing scheme that is in the $\frac{2}{3}$-core of the game: charge each vertex that is covered by the maximum matching $\frac{1}{3}$, and other vertices $\frac{2}{3}$. Since there is no edge between two vertices that are not covered by the maximum matching, this cost-sharing scheme satisfies the core property (but not cross-monotonicity). Furthermore, it is easy to see that the sum of the cost shares is always equal to $\frac{2}{3}$ times the edge cover for S. Therefore, there is a cost-sharing scheme satisfying the core property with a budget-balance factor of $\frac{2}{3}$. In fact, Goemans [44] showed that for every graph there is a cost sharing scheme in the $\frac{3}{4}$-core. However, in the following, we show that no cross-monotonic cost-sharing scheme can achieve a budget-balance factor better than $\frac{1}{2}$.

Theorem 4.2.1 *For every $\varepsilon > 0$, there is no $(\frac{1}{2}+\varepsilon)$-budget balanced cross-monotonic cost sharing scheme for the edge cover problem.*

Here is the high-level idea of the proof: We assume, for contradiction, that there is a cross-monotonic cost sharing scheme that always recovers at least a $(\frac{1}{2} + \varepsilon)$ fraction of the total cost. We explicitly construct a graph G (or in general the set of agents $\mathscr{A}$ and the structure based on which the cost function is defined), and look at the cost-sharing scheme on this graph. For edge cover, this graph is simply a complete bipartite graph $K_{n,n}$, with n large enough. Then, we need to argue that there is a set S of agents such that the total cost shares of the elements of S is less than $\frac{1}{2} + \varepsilon$ times the size of the minimum edge-cover for S. This is done using the probabilistic method: we pick a subset S at random from a certain distribution and show that in expectation, the ratio of the recovered cost to the cost of S is low. Therefore, there is a manifestation of S for which this ratio is low. In the edge-cover example, we pick one vertex v of G uniformly at random and let S be the union of v and the set of vertices adjacent to v. We now need to bound the expected value of the sum of cost shares of the elements of S. We do this by using cross-monotonicity and bounding the cost share of each vertex $u \in S$ by the cost share of u in a substructure T_u of S. Bounding the expected cost share of u in T_u is done by showing that for every substructure T, every $u \in T$ has the same probability of occurring in a structure S in which $T_u = T$. This implies that the expected cost share of u in T_u (where the expectation is over the choice of S) is at most the cost of T_u divided by the number of agents in T_u. Summing up these values for all u gives us the desired contradiction.

Proof of Theorem 4.2.1. Assume that there is a $(\frac{1}{2} + \varepsilon)$-budget-balanced cross-monotonic cost sharing scheme ξ. Let G be the complete bipartite graph $K_{n,n}$, where n will be fixed later, and consider ξ on G. For every $v \in V(G)$, we let S_v be the union of v and the set of vertices adjacent to v (that is, all vertices of the other part). We pick a set S of agents by picking v uniformly at random from $V(G)$ and letting $S = S_v$. By the definition of the edge cover game,

$$C(S_v) = n \qquad \text{for every } v. \tag{4.1}$$

On the other hand,

$$\begin{aligned} \mathrm{E}_S\Big[\sum_{i\in S}\xi(i,S)\Big] &= \mathrm{E}_v\left[\xi(v,S_v)\right] + \mathrm{E}_v\Big[\sum_{u\in S_v\setminus\{v\}}\xi(u,S_v)\Big] \\ &\leq 1 + \mathrm{E}_v\Big[\sum_{u\in S_v\setminus\{v\}}\xi(u,\{u,v\})\Big], \end{aligned} \tag{4.2}$$

where the last inequality follows from the facts that for every vertex u and every set S, $\xi(u,S) \leq 1$, and that for every $v \in V(G)$ and $u \in S_v \setminus \{v\}$, $\xi(u,S_v) \leq \xi(u,\{u,v\})$. Both of these facts are consequences of the cross-monotonicity of ξ. By the definition of expected values, we have

$$\mathrm{E}_v\Big[\sum_{u\in S_v\setminus\{v\}}\xi(u,\{u,v\})\Big] = n\,\mathrm{E}_{v,u}\left[\xi(u,\{u,v\})\right], \tag{4.3}$$

where the second expectation is over the choice of v from $V(G)$ and u in $S_v \setminus \{v\}$. However, choosing a vertex v and then a neighbor u of v at random is equivalent to choosing a random edge e in G at random, and letting u be a random endpoint of e and v be the other one. By the budget-balance condition, the sum of the cost shares of the endpoints of e is at most one. Therefore, for every e, if u is a random endpoint of e and v is the other endpoint, $\mathrm{E}[\xi(u, \{u, v\})] \leq \frac{1}{2}$. Thus, the right-hand side of Equation 4.3 is at most $\frac{n}{2}$. Therefore, by Equations 4.1 and 4.2, we have

$$\mathrm{E}_S\left[\frac{\sum_{i\in S}\xi(i,S)}{C(S)}\right] \leq \frac{1+\frac{n}{2}}{n} < \frac{1}{2} + \varepsilon$$

for $n > 1/\varepsilon$. Therefore, there is a set S satisfying $\frac{\sum_{i\in S}\xi(i,S)}{C(S)} < \frac{1}{2} + \varepsilon$, which is a contradiction with the assumption that ξ is $(\frac{1}{2} + \varepsilon)$-budget balanced. □

It is not difficult to see that the cost-sharing scheme ξ satisfying $\xi(i, S) = \frac{1}{2}$ for every $i \in S$ is cross-monotonic and $\frac{1}{2}$-budget balanced. Therefore, the bound given in the above theorem is tight.

4.2.2 The set cover game

The set cover game is defined as follows.

Definition 16 *Let $\mathscr{A}$ be a set of agents and $\mathscr{E}$ be a collection of subsets of $\mathscr{A}$ such that every element of $\mathscr{A}$ is contained in at least one set in $\mathscr{E}$. For every $S \subseteq \mathscr{A}$, the cost of S in the set cover game is the minimum size of a subcollection $\mathscr{F} \subseteq \mathscr{E}$ such that every $x \in S$ is contained in at least one set in $\mathscr{F}$. Such a collection $\mathscr{F}$ is called a* set cover *for S.*

One can think of the edge-cover problem as a special case of the set cover problem in which the size of each set is 2. It is not difficult to generalize Theorem 4.2.1 to the special case of set cover in which the size of each set is k, and prove that for k constant, no cross-monotonic cost-sharing scheme for this problem can recover more than a $\frac{1}{k}$ fraction of the cost. Using a similar argument, the next theorem shows that for the general case of the set cover game, no cross-monotonic cost-sharing scheme can recover more than a $O(\frac{1}{n})$ of the total cost.

Theorem 4.2.2 *There is no cross-monotonic cost-sharing scheme ξ for the set cover game such that for every set $S \subseteq \mathscr{A}$, ξ recovers more than a $O(\frac{1}{|S|})$ fraction of the cost of S.*

Proof. Assume that there is such a cross-monotonic cost sharing scheme ξ. Consider the following set cover game. Let $\mathscr{A}$ be a set of n^2 agents that can be partitioned as $\mathscr{A} = A_1 \cup A_2 \cup \cdots \cup A_n$, where A_i's are disjoint sets each of size n. Define $\mathscr{E}$ as the collection of all sets $S \subset \mathscr{A}$ such that $|S \cap A_i| = 1$ for every $i = 1, \ldots, n$. An alternative way to look at this is that $\mathscr{A}$ and $\mathscr{E}$ are sets of vertices and edges of an n-uniform n-partite complete hypergraph.

We pick a random set S of agents in the above game as follows: Pick a random i from $\{1, \ldots, n\}$, and for every $j \neq i$, pick an agent a_j uniformly at random from A_j. Let $T = \{a_j : j \neq i\}$ and $S = A_i \cup T$. The cost of the optimal set cover solution on S is always at least n, since no set in $\mathscr{E}$ contains two distinct elements of A_i, and therefore each element of A_i must be covered with a distinct set in $\mathscr{E}$.

We now bound the average recovered cost over the random choice of S.

$$\begin{aligned}
\mathrm{E}_S\Big[\sum_{x\in S}\xi(x,S)\Big] &= \mathrm{E}\Big[\sum_{x\in A_i}\xi(x,S)\Big] + \mathrm{E}\Big[\sum_{j\neq i}\xi(a_j,S)\Big] \\
&\leq \mathrm{E}\Big[\sum_{x\in A_i}\xi(x,\{x\}\cup T)\Big] + \mathrm{E}\Big[\sum_{j\neq i}\xi(a_j,T)\Big]
\end{aligned}$$

Since all elements of T can be covered by one set, the second term in the above expression is at most 1. We write the first term as $n\mathrm{E}_{S,x}[\xi(x, \{x\} \cup T)]$ where the expectation is over the random choice of S and the random choice of x from A_i. As in the proof of Theorem 4.2.5, the expected value of $\xi(x, \{x\} \cup T)$ in this experiment is equal to the expected value of $\frac{1}{n}\sum_{j=1}^{n}\xi(a_j, \{a_1, \ldots, a_n\})$ in an experiment that consists of choosing an agent a_j from each A_j uniformly at random. By the budget-balance property, we always have $\sum_{j=1}^{n}\xi(a_j, \{a_1, \ldots, a_n\}) \leq C(\{a_1, \ldots, a_n\}) = 1$. Therefore, the first term in the left-hand side of the inequality (4.4) is at most one.

This means that the expected total cost share recovered from the set S is at most two. Therefore, the ratio of recovered cost to total cost of S is at most $2/n < 4/|S|$. $\square$

It is worth noting that the above proof shows that even for the *fractional* set cover game, no cross-monotonic cost-sharing scheme can achieve a budget-balance factor better than $O(1/n)$.[6] This is particularly interesting for the following reason: It is easy to show that if there is an α-budget balanced cross-monotonic cost-sharing scheme for the fractional set cover, then for any special case of the set cover problem of integrality gap at most μ, there is an $\alpha\mu$-budget balanced cross-monotonic cost-sharing scheme. For example, if we could find a constant-factor for fractional set cover, we would automatically get a constant-factor for metric facility location, generalized Steiner tree, and many other network design games. Unfortunately, the above theorem shows this approach for designing cross-monotonic cost-sharing schemes fails to recover much of the cost.

4.2.3 The vertex cover game

The vertex cover game is defined on a graph $G = (V, E)$. The set of agents is the set of edges of G, and the cost of serving a set $S \subseteq E$ is equal to the minimum size of a set A of vertices such that for each $e \in S$, at least one of the endpoints of e is in A. Such a set is called a *vertex cover* for the set S. It is well-known that the integrality gap of the LP relaxation of vertex cover is 2, and therefore no allocation in core can recover more than half the cost of the solution in the worst case [13, 108]. We show in the following theorem that if we require the cost-sharing scheme to be cross-monotonic, then no constant-factor budget balanced scheme exists.

Theorem 4.2.3 *For every $\varepsilon > 0$, there is no cross-monotonic cost sharing scheme for vertex cover that on every set S of n agents, recovers at least a $(2+\varepsilon)n^{-1/3}$ fraction of the cost of S.*

[6] Other bounds in the section also apply to the fractional variants of the corresponding games.

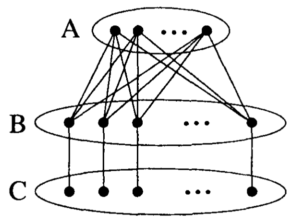

Figure 4-1: Vertex Cover Sample Distribution

Proof. Assume, for contradiction, that such a scheme ξ exists. We let G be a complete graph on $m + 2\ell$ vertices, where m and ℓ ($m < \ell$) are numbers that will be fixed later, and consider the cost-sharing scheme ξ on G. We show that there is some set S of edges of G for which ξ recovers at most a $|S|^{-1/3}$ fraction of the cost. We do this by picking S randomly from a distribution described below, and showing that the above statement holds in expectation, and therefore there should be a particular S satisfying the above statement.

Let π be a permutation of the $m + 2\ell$ vertices. Let A be the set of the first m vertices, B be the set of the next ℓ vertices, and C be the set of the remaining ℓ vertices. We denote the i'th vertices of B and C (based on the ordering given by π) by b_i and c_i. Let S_π denote the set of all $m\ell$ edges between A and B, together with the set of edges $b_i c_i$ for $i = 1, \ldots, \ell$. We pick S by picking the permutation π uniformly at random and letting $S = S_\pi$. See Figure 4-1 for an example.

If we denote the set of edges between A and B by T, we have

$$\mathrm{E}\left[\sum_{e \in T} \xi(e, S)\right] \leq \mathrm{E}\left[\sum_{e \in T} \xi(e, T)\right] \leq m, \tag{4.4}$$

where the first inequality follows from the cross-monotonicity of ξ and the second inequality is implied by the budget balance assumption and the fact that the cost of the minimum vertex cover in T is m. We also let T_i be the set of all $m + 1$ edges in S that have b_i as an endpoint (see Figure 4-1). Equation 4.4 and the cross-monotonicity of ξ imply the following.

$$\begin{aligned}
\mathrm{E}_S\left[\sum_{i\in S}\xi(i,S)\right] &= \mathrm{E}\left[\sum_{e\in T}\xi(e,S)\right]+\sum_{i=1}^{\ell}\mathrm{E}\left[\xi(b_ic_i,S)\right] \\
&\leq m+\sum_{i=1}^{\ell}\mathrm{E}\left[\xi(b_ic_i,T_i)\right], \qquad (4.5)
\end{aligned}$$

We now need to analyze the expectation of $\xi(b_ic_i,T_i)$ over the random choice of π. Notice that the only elements of π that are important in $\xi(b_ic_i,T_i)$ are the first m elements and the $m+i$'th and $m+\ell+i$'th elements (b_i and c_i). Therefore, the expectation of $\xi(b_ic_i,T_i)$ over the choice of π is equal to the expectation of $\xi(v_{m+2}v_{m+1},\{v_1v_{m+1},v_2v_{m+1},\ldots,v_mv_{m+1},v_{m+2}v_{m+1}\})$ over the random choice of an ordered list $v_1,v_2,\ldots,v_{m+2}$ of $m+2$ different vertices of G. However, in this experiment it is clear by symmetry that the expected cost share of v_iv_{m+1} is the same for $i=1,\ldots,m,m+2$, and therefore by the budget balance condition each of these expected cost shares is at most $\frac{1}{m+1}$. This, together with Equation 4.5 imply the following.

$$\mathrm{E}_S\left[\sum_{i\in S}\xi(i,S)\right]\leq m+\frac{\ell}{m+1}. \qquad (4.6)$$

On the other hand, the size of the minimum vertex cover in S is always ℓ. Therefore, the expected value of the ratio of $\sum_{i\in S}\xi(i,S)$ to $C(S)$ is at most $\frac{m}{\ell}+\frac{1}{m+1}$. Thus, there is a set S for which this ratio is at most $\frac{m}{\ell}+\frac{1}{m+1}$. Taking $m=\sqrt{\ell}$, we see that the allocation on S recovers at most a $\frac{2}{\sqrt{\ell}}<(2+\varepsilon)|S|^{-1/3}$ fraction of the cost. □

We can show the following positive result for cross-monotonic cost sharing schemes for the vertex cover which, together with a result of Moulin [85] (see Theorem 4.3.1), implies an approximately budget-balanced group-strategyproof mechanism for this problem (see Section 4.3). We do not know the right bound for the budget-balance factor of the vertex cover game.

Theorem 4.2.4 *For the vertex cover game, the cost sharing scheme that charges the edge uv in the set S an amount equal to $\min(1/\deg_S(u),1/\deg_S(v))$ is cross-monotonic and $\frac{1}{2\sqrt{n}}$-budget balanced.*

Proof. It is clear that this scheme is cross-monotone. We only need to verify the budget-balance factor. Consider a set S of n agents (edges), and the graph $G[S]$ induced on this set of edges. We prove that the total cost share of the agents in S is at least $\frac{1}{2\sqrt{n}}$ times the cost of a vertex cover for S.

Divide the set of vertices into two subsets L and H, where L is the set vertices of degree less than $\sqrt{n}$ in $G[S]$ and H is the rest of vertices ($H = V(G) - L$). As a vertex cover solution, select H and both endpoints of all edges (u, v) such that $u, v \in L$. We show that the cost shares of the edges in S sum to at least a $\frac{1}{2\sqrt{n}}$ fraction of the cost of this solution. First consider any edge e between vertices in L. The cost share of e is at least $\frac{1}{\sqrt{n}}$, thus its cost share covers $\frac{1}{\sqrt{n}}$ of the cost of picking both its endpoints. Now consider the vertices in H. Since the degree of each vertex $v \in H$ is greater than or equal to $\sqrt{n}$, the sum of the cost shares of the edges adjacent to v is at least $\frac{1}{n}\sqrt{n} = \frac{1}{\sqrt{n}}$. Each edge is included in at most two such summations (namely, when both its endpoints are in H), and thus the sum of the cost shares of edges adjacent to vertices in H is at least a $\frac{1}{2\sqrt{n}}$ fraction of the cost of H. Therefore, the sum of the cost shares of the agents in S is at least $\frac{1}{2\sqrt{n}}$ times the cost of the optimal vertex cover for S. □

4.2.4 The metric facility location game

Given a set of cities, facilities with opening costs, and metric connection costs between cities and facilities, the facility location problem seeks to open a subset of facilities and connect each city to a facility in a manner that minimizes the total cost. In the facility location game, each city is an agent. The cost of a subset of agents is the cost of the minimum facility location solution for that subset; a cross-monotonic cost-sharing scheme tries to share this cost among the agents. In this section, we prove that any cross-monotonic cost-sharing scheme for facility location is at best $\frac{1}{3}$-budget-balanced. This matches the budget-balance factor of the scheme given by Pál and Tardos [90].

We start by giving an example on which the scheme of Pál and Tardos [90] recovers

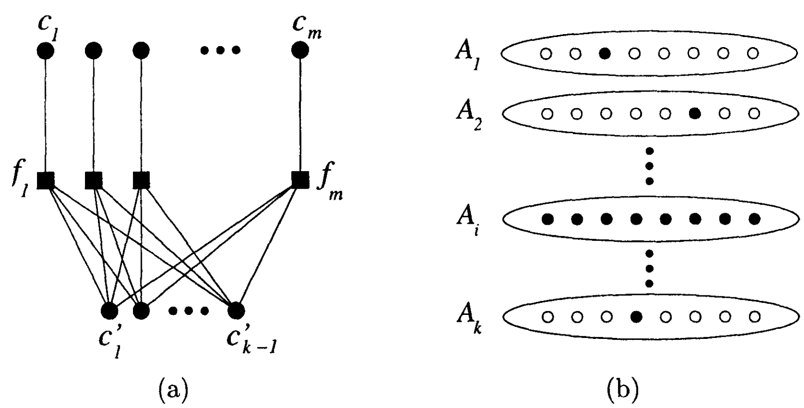

Figure 4-2: Facility Location Sample Distribution

only a third of the cost[7]. This example will be used as the randomly chosen structure in our proof.

Lemma 4.2.1 *Let $\mathcal{I}$ be an instance of the facility location problem consisting of $m + k - 1$ cities $c_1, \ldots, c_m$, $c'_1, \ldots, c'_{k-1}$ and m facilities $f_1, \ldots, f_m$ each of opening cost 3. For every i and j, the connection costs between f_i and c_i and between f_i and c'_j are all 1, and other connection costs are obtained by the triangle inequality. See Figure 4-2(a). Then if $m = \omega(k)$ and k tends to infinity, the optimal solution for $\mathcal{I}$ has cost $3m + o(m)$.*

Proof. The solution which opens just one facility, say f_1, has cost $3(m-1)+k+3 = 3m+o(m)$. We show that this solution is optimal. Consider any feasible solution which opens f facilities. The first opened facility can cover k clients with connection cost 1. Each additional facility can cover 1 additional client with connection cost 1. Thus, the number of clients with connection cost 1 is $k+f-1$. The remaining $m-f$ clients have connection cost 3. Therefore, the cost of the solution is $3f+k+f-1+3(m-f) = 3m+k+f-1$. As $f \geq 1$, this shows that any feasible solution costs at least as much as the solution we constructed. □

Theorem 4.2.5 *Any cross-monotonic cost-sharing scheme for the facility location game is at most 1/3-budget balanced.*

[7]This example also shows that the dual computed by the Jain-Vazirani facility location algorithm [65] can be a factor 3 away from the optimal dual.

Proof. Consider the following instance of the facility location problem. There are k sets $A_1, \ldots, A_k$ of m cities each, where $m = \omega(k)$ and $k = \omega(1)$. For every subset B of cities containing exactly one city from each A_i ($|B \cap A_i| = 1$ for all i), there is a facility f_B with connection cost 1 to each city in B. The remaining connection costs are defined by extending the metric, that is, the cost of connecting city i to facility f_B for $i \notin B$ is 3. The facility opening costs are all 3.

We pick a random set S of cities in the above instance as follows: Pick a random i from $\{1, \ldots, k\}$, and for every $j \neq i$, pick a city a_j uniformly at random from A_j. Let $T = \{a_j : j \neq i\}$ and $S = A_i \cup T$. See Figure 4-2(b) for an example. It is easy to see that the set S induces an instance of the facility location problem almost identical to the instance $\mathcal{I}$ in Lemma 4.2.1 (the only difference is that here we have more facilities, but it is easy to see that the only relevant facilities are the ones that are present in $\mathcal{I}$). Therefore, the cost of the optimal solution on S is $3m + o(m)$.

We show that for any cross-monotonic cost-sharing scheme ξ, the average recovered cost over the choice of S is at most $m + o(m)$ and thus conclude that there is some S whose recovered cost is at most $m + o(m)$. As in the previous proofs, we start bounding the expected total cost share by using the linearity of expectations and cross-monotonicity:

$$\begin{aligned} \mathrm{E}_S\Big[\sum_{c\in S}\xi(c,S)\Big] &= \mathrm{E}\Big[\sum_{c\in A_i}\xi(c,S)\Big] + \mathrm{E}\Big[\sum_{j\neq i}\xi(a_j,S)\Big] \\ &\leq \mathrm{E}\Big[\sum_{c\in A_i}\xi(c,\{c\}\cup T)\Big] + \mathrm{E}\Big[\sum_{j\neq i}\xi(a_j,T)\Big] \end{aligned}$$

Notice the set T has a facility location solution of cost $3 + k - 1$ and thus by the budget balance condition the second term in the above expression is at most $k + 2$. The first term in the above expression can be written as $m\mathrm{E}_{S,c}[\xi(c, \{c\} \cup T)]$ where the expectation is over the random choice of S and the random choice of c from A_i. However, it can be seen easily that this is equivalent to the following random experiment: From each A_j, pick a city a_j uniformly at random. Then pick i from $\{1, \ldots, k\}$ uniformly at random and let $c = a_i$ and $T = \{a_j : j \neq i\}$. From this description it is clear that the expected value of $\xi(c, \{c\} \cup T)$ is equal to $\frac{1}{k}\sum_{j=1}^{k}\xi(a_j, \{a_1, \ldots, a_k\})$. This, by the budget balance property and the fact that $\{a_1, \ldots, a_k\}$ has a solution of cost $k+3$, cannot be more than $\frac{k+3}{k}$. Therefore,

$$\mathrm{E}_S\left[\sum_{c \in S} \xi(c, S)\right] \leq m(\frac{k+3}{k}) + (k+2) = m + o(m), \tag{4.7}$$

when $m = \omega(k)$ and $k = \omega(1)$. Therefore, the expected value of the ratio of recovered cost to total cost tends to $1/3$. □

4.2.5 Other combinatorial optimization games

In this section we prove bounds for three other combinatorial optimization games (in particular, the ones considered by Deng, Ibaraki, and Nagamochi [23]). These problems are maximization problems; therefore instead of cost-sharing schemes, we consider *profit-sharing* schemes, as defined below.

Definition 17 *A* profit-sharing game *(or a* coalitional game with transferable utilities*) is defined by a set $\mathscr{A}$ of agents, and a function $v : 2^{\mathscr{A}} \mapsto \mathbb{R}^+ \cup \{0\}$ that for every set S, gives the value $v(S)$ of S (or the profit earned if agents in S collaborate). A profit-sharing scheme is a function $\xi : \mathscr{A} \times 2^{\mathscr{A}} \mapsto \mathbb{R}^+ \cup \{0\}$, such that for every $S \subseteq \mathscr{A}$ and every $i \notin S$, $\xi(i, S) = 0$. Such a scheme is called* α-budget-balanced *(for some $\alpha \geq 1$) if for every $S \subseteq \mathscr{A}$, $v(S) \leq \sum_{i \in S}\xi(i, S) \leq \alpha v(S)$. A profit-sharing scheme ξ is in the* α-core *if it is α-budget-balanced and for every S and $T \subseteq S$, $\sum_{i \in T}\xi(i, S) \geq v(T)$. A profit-sharing scheme ξ is* cross-monotone *if for all $S, T \subseteq \mathscr{A}$ and $i \in S$, $\xi(i, S) \leq \xi(i, S \cup T)$.*

In this section, we consider profit-sharing schemes for the games of maximum flow, maximum arborescence packing, and maximum matching, and derive lower bounds on the budget-balance factor of cross-monotonic profit-sharing schemes for these games.

The maximum flow game In the maximum flow game, we are given a directed graph $G = (V, E)$ with a source s and a sink t. Agents are directed edges of G. Given

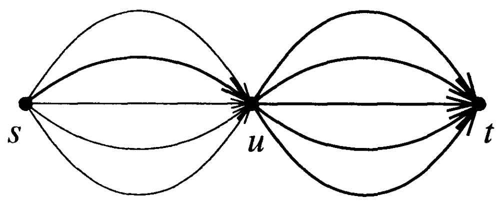

Figure 4-3: The graph G for the maximum flow game

a subset of edges, S, the value of S is the value of the maximum flow from s to t on the subgraph of G induced by the edges of S. It is known that the core of the maximum flow game is nonempty [23]. The situation is different for cross-monotonic profit-sharing schemes.

Theorem 4.2.6 *There is no $o(n)$-budget-balanced profit-sharing scheme for the maximum flow game where n is the number of agents in the set that receives the service.*

Proof. Let G be a graph consisting of three nodes named s, u, and t; $n-1$ edges from s to u; and $n-1$ edges from u to t. Let E_{su} and E_{ut} denote the set of edges from s to u and from u to t, respectively. See Figure 4-3. We pick a random set S of n agents as follows: With probability $1/2$, pick a random edge e from s to u, and let $S = \{e\} \cup E_{ut}$. With probability $1/2$, pick a random edge e from u to t, and let $S = \{e\} \cup E_{su}$. For example the set S could contain the thick edges in Figure 4-3.

Assume ξ is an $o(n)$-budget-balanced cross-monotonic profit-sharing scheme for G. We have

$$\begin{aligned}
\mathrm{E}_S\Big[\sum_{a\in S}\xi(a,S)\Big] &\geq \frac{1}{2}\mathrm{E}_{e\overset{R}{\leftarrow}E_{su}}\Big[\sum_{a\in E_{ut}}\xi(a,\{e\}\cup E_{ut})\Big] \\
&\quad + \frac{1}{2}\mathrm{E}_{e\overset{R}{\leftarrow}E_{ut}}\Big[\sum_{a\in E_{su}}\xi(a,\{e\}\cup E_{su})\Big] \\
&\geq \frac{1}{2}\mathrm{E}_{e\overset{R}{\leftarrow}E_{su}}\Big[\sum_{a\in E_{ut}}\xi(a,\{a,e\})\Big] + \frac{1}{2}\mathrm{E}_{e\overset{R}{\leftarrow}E_{ut}}\Big[\sum_{a\in E_{su}}\xi(a,\{a,e\})\Big] \\
&= (n-1)\mathrm{E}_{a\overset{R}{\leftarrow}E_{su},b\overset{R}{\leftarrow}E_{ut}}\Big[\frac{1}{2}\xi(a,\{a,b\}) + \frac{1}{2}\xi(b,\{a,b\})\Big] \\
&\geq \frac{n-1}{2}.
\end{aligned}$$

On the other hand, the value of every set S picked using the above procedure is one. Therefore, the expected ratio of the sum of profit shares to the value of S is at least $(n-1)/2$. $\square$

Remark 4.2.1 *It is easy to see that the above proof also works for the problems of packing the maximum number of arborescences in a digraph, and gives the same lower bound. An r-arborescence is a spanning tree rooted at r in which all edges are directed away from r. The maximum r-arborescence game is defined on a digraph $G = (V, E)$ with a root r where each edge is an agent. The value of a set S is the maximum number of edge-disjoint r-arborescences on the subgraph induced by S. One can think of the value of S as the maximum bandwidth for broadcasting messages from r to all vertices of the graph. It is known that the core of this game is nonempty [23].*

The maximum matching game As a last example, we consider the maximum matching game, in which the agents are vertices of a graph G, and the value of a subset of vertices S is the size of the maximum matching in the subgraph of G induced by S (denoted $G[S]$). One can show that there is a 2-budget-balanced profit-allocation in the core of this game.

Theorem 4.2.7 *There is no $o(n)$-budget-balanced profit-sharing scheme for the maximum matching game, where n is the set of agents that receive the service.*

Proof. We use the same construction that was used in the proof of Theorem 4.2.1. Let G be a complete bipartite graph with $n-1$ vertices in each part (here we use $n-1$ instead of n so that the size of S becomes n), and pick S by picking a random vertex in G and all vertices in the other part. Using an argument essentially the same as the one in the proof of Theorem 4.2.1, the expected sum of profit shares of the elements of S is at least $(n-1)/2$. On the other hand, the value of S is always one. Thus, there is an S on which the ratio between the total profit share and the value of S is at least $(n-1)/2$. $\square$

4.3 Implications for Cost-Sharing Mechanisms

One of the important applications of cross-monotonic cost-sharing schemes is in the construction of group-strategyproof cost-sharing mechanisms [85, 87]. In this section, we explore the connection between cross-monotonic cost-sharing schemes and group-strategyproof cost-sharing mechanisms, and implications of the upper bounds of the previous section on such mechanisms. In Section 4.3.1 we define the setting and present some preliminaries. In Section 4.3.2 we discuss an issue in the definition of group-strategyproof mechanisms, and note that in order to exclude a trivial mechanism, we need to use a stronger version of one of the axioms. In Section 4.3.3 we give a partial characterization of group-strategyproof mechanisms in terms of cost-sharing schemes satisfying a property weaker than cross-monotonicity. We then use this characterization to prove that group-strategyproof mechanisms that satisfy additional properties give rise to cross-monotonic cost-sharing schemes.

4.3.1 Preliminaries

Let $\mathscr{A}$ be a set of n agents interested in receiving a service. Each agent i has a value $u_i \in \mathbb{R}$ for receiving the service, that is, she is willing to pay at most u_i to get the service. We further assume that the utility of agent i is given by $u_i q_i - x_i$, where q_i is an indicator variable which indicates whether she has received the service or not, and x_i is the amount she has to pay. A *cost sharing mechanism* is an algorithm that elicits a bid $b_i \in \mathbb{R}$ from each agent, and based on these bids, decides which agents should receive the service and how much each of them has to pay. More formally, a cost sharing mechanism is a function that associates to each vector $\mathbf{b}$ of bids a set $Q(\mathbf{b}) \subseteq \mathscr{A}$ of agents to be serviced, and a vector $x(\mathbf{b}) \in \mathbb{R}^n$ of payments. When there is no ambiguity, we write Q and x instead of $Q(\mathbf{b})$ and $x(\mathbf{b})$, respectively. We assume that a mechanism satisfies the following conditions:[8]

- *No Positive Transfer (NPT)*: The payments are non-negative (that is, $x_i \geq 0$ for all i).

[8]For a discussion about these properties see Moulin [85] and Moulin and Shenker [87].

- *Voluntary Participation (VP)*: An agent who does not receive the service is not charged (that is, $x_i = 0$ for $i \notin Q$), and an agent who receives the service is not charged more than his bid (that is, $x_i \leq b_i$ for $i \in Q$).[9]

- *Consumer Sovereignty (CS)*: For each agent i, there is some bid b_i^* such that if i bids b_i^*, she will get the service, no matter what others bid.

Furthermore, we would like the mechanisms to be approximately budget balanced. Mimicking the definition for cost-sharing schemes, we call a mechanism α-budget balanced if the total amount the mechanism charges the agents is between $\alpha C(Q)$ and $C(Q)$. That is,

$$\alpha C(Q(\mathbf{b})) \leq \sum_{i \in Q(\mathbf{b})} x_i(\mathbf{b}) \leq C(Q(\mathbf{b})).$$

for every bid vector $\mathbf{b}$.

We look for mechanisms, called *group strategyproof mechanisms*, which satisfy the following property in addition to NPT, VP, and CS. Let $S \subseteq \mathscr{A}$ be a coalition of agents, and u, u' be two vectors of bids satisfying $u_i = u'_i$ for every $i \notin S$ (we think of u as the value of agents, and u' as a vector of strategically chosen bids). Let (Q, x) and (Q', x') denote the outputs of the mechanism when the bids are u and u', respectively. A mechanism is *group strategyproof* if for every coalition S of agents, if the inequality $u_i q'_i - x'_i \geq u_i q_i - x_i$ holds for every $i \in S$, then it holds with equality for every $i \in S$. In other words, there should not be any coalition S and vector u' of bids such that if members of S announce u' instead of u (their true value) as their bids, then every member of the coalition S is at least as happy as in the truthful scenario, and at least one person is strictly happier. We call such a coalition a *successful coalition.* [10]

Given a cross-monotonic cost-sharing scheme ξ, Moulin [85] defined a cost-sharing mechanism $\mathscr{M}_\xi$ as follows.

[9]This is equivalent to the condition of individual rationality as defined in Chapter 1.

[10]Notice that we do not allow members of the coalition to sacrifice their own utility to benefit the group's total utility, that is we disallow side-payments. Side-payments require a transfer of money between agents which might be restricted in some settings either due to legal concerns or issues of trust, and so we do not consider side-payments here. For a discussion of collusion with side-payments, see Goldberg and Hartline [47].

Mechanism $\mathscr{M}_\xi$:

> Initialize $S \leftarrow \mathscr{A}$.
>
> Repeat
>
> > Let $S \leftarrow \{i \in S : b_i \geq \xi(i,S)\}$.
>
> Until for all $i \in S$, $b_i \geq \xi(i,S)$.
>
> Return $Q = S$ and $x_i = \xi(i,S)$ for all i.

Notice that the mechanism $\mathscr{M}_\xi$ always services the maximal subset of agents whose bids are all at least as large as their cost shares in that set.[11] This mechanism is a generalization of the ProfitExtract$_R$ mechanism introduced in Chapter 2. There, the cost of every set was equal to a target revenue R, and the underlying cross-monotonic cost-sharing scheme set $\xi(i,S) = R/|S|$ for all $i \in S$. We saw in Chapter 2 that the ProfitExtract$_R$ mechanism was truthful. Moulin [85] proves a stronger result.

Theorem 4.3.1 (Moulin [85]) *If ξ is a cross-monotonic cost-sharing scheme, then $\mathscr{M}_\xi$ is group-strategyproof.*

4.3.2 A discussion about the definition

In the definition of group-strategyproof mechanisms in the paper by Moulin and Shenker [87] (which is the basis for the definition of this concept in most computer science papers), it is not required that an agent can bid in a way that guarantees her not to receive the service. In particular, it is assumed that the bids are non-negative, and an agent who bids zero can still be serviced, if her payment is also zero [87, page 517]. As we see in the following example, according to this definition, for every cost function there is a trivial budget-balanced group-strategyproof mechanism.

Example 4.3.1 *Arbitrarily order the agents from 1 to n. Then, find the first agent i in this order whose bid is at least $C(\{i,\ldots,n\})$. The set that will receive the service*

[11]Note that there is a unique maximal set as if two sets are feasible then, by cross-monotonicity, their union is as well.

is $Q = \{i, \ldots, n\}$, and the total cost of servicing this set is paid by the agent i. Other agents pay nothing.

Proposition 4.3.1 *Assuming non-negative bids, the mechanism in Example 4.3.1 is budget-balanced and group-strategyproof.*

Proof. It is not hard to see that this mechanism is budget-balanced and satisfies NPT, VP, and CS. To show that it is group-strategyproof, let i be the first agent to receive service when agents bid truthfully (or $n + 1$ if no agent receives service) and j be the first agent to receive service when a coalition deviates. If $j < i$, it must be that j is part of the coalition and raised his bid to a number greater than or equal to $C(\{j, \ldots, n\})$, but this decreases his utility. If $j = i$, then the outcome is identical to the truthful scenario and so no utility changes. If $j > i$, then the utility of any agent $k < j$ is now zero and so did not increase. The utility of any agent $k > j$ did not change as his allocation and payment remained the same. Finally, as the payment of j is at least his payment in the truthful scenario, the utility of agent j can not increase either. Thus the coalition can not be successful. □

Although it satisfies all of the axioms, this mechanism is unsatisfactory, since in practice a coalition can convince a member that has zero utility for receiving the service simply not to bid, thus reducing the cost to others. Furthermore, this mechanism fails to satisfy the axioms in the original paper of Moulin [85], where a stronger version of CS is assumed that guarantees that each agent can bid in a way that she does not receive the service, no matter how others bid.

In order to exclude mechanisms like the one in Example 4.3.1, we only consider mechanisms that satisfy the stronger definition of CS by Moulin [85]. To this end, we allow the utilities and bids to be *negative.* NPT and VP guarantee that any agent with negative bid will not receive the service. An alternative approach (adapted by Moulin [85]) is to assume that utilities, bids, and payments are all positive.[12] In many combinatorial games, the cost function is not strictly increasing and therefore it is reasonable to allow cost shares to be zero. Thus, we use negative bids to indicate

[12]This is equivalent to a property called *no free riders*, or *no free lunch.*

that an agent does not want to receive the service. However, it is easy to see that all our results hold in the setting considered by Moulin [85].

4.3.3 A partial characterization of group-strategyproof mechanisms

In Section 4.2, we proved that for certain games every cross-monotonic cost sharing scheme is poorly budget balanced. A natural question to ask is whether all group-strategyproof mechanisms for these games are so poorly budget balanced. Towards this aim, one might hope to show a converse to Theorem 4.3.1, namely that every group-strategyproof mechanism corresponds to a cross-monotonic cost sharing scheme. Unfortunately, this statement is not necessarily true (See, for example, Appendix A.1, or the incremental cost-sharing method for supermodular cost functions in the paper by Moulin [85]). In this section, we prove that for any group-strategyproof mechanism, we can construct a cost-sharing scheme that satisfies a weaker condition than cross-monotonicity. Then, we use this characterization to show that group-strategyproof mechanisms that satisfy certain additional properties correspond to cross-monotonic cost-sharing schemes.

We start by defining a property weaker than cross-monotonicity for cost-sharing schemes. Recall that a cost-sharing scheme is cross-monotonic, if the removal of each agent from the service set does not increase the cost to any other agent.

Definition 18 *Let $\xi : \mathscr{A} \times 2^{\mathscr{A}} \mapsto \mathbb{R}^+ \cup \{0\}$ be a cost-sharing scheme, $S \subseteq \mathscr{A}$, and $i \in S$. We say i is a* positive *element of S if for every $j \in S \setminus \{i\}$, $\xi(j, S \setminus \{i\}) \geq \xi(j, S)$ and for at least one such j a strict inequality holds; i is a* negative *element of S if for every $j \in S \setminus \{i\}$, $\xi(j, S \setminus \{i\}) \leq \xi(j, S)$ and for at least one such j a strict inequality holds. If for all $j \in S \setminus \{i\}$, $\xi(j, S \setminus \{i\}) = \xi(j, S)$, we say i is a* neutral *element of S. We say that ξ is* semi-cross-monotonic, *if every element of every set is either positive, negative, or neutral. In other words, ξ is semi-cross-monotonic if there is no set $S \subseteq \mathscr{A}$ and three distinct elements i, j_1, j_2 of S, such that $\xi(j_1, S \setminus \{i\}) < \xi(j_1, S)$*

and $\xi(j_2, S \setminus \{i\}) > \xi(j_2, S)$.[13]

Thus, cross-monotonicity is precisely a special case of semi-cross-monotonicity, when every element of every set is either positive or neutral. The results in this section are based on the following partial characterization of group-strategyproof mechanisms.

Theorem 4.3.2 *For every α-budget-balanced group-strategyproof cost-sharing mechanism $\mathscr{M}$ for a cost function C, there is a cost-sharing scheme $\xi_{\mathscr{M}}$ for C such that*

(a) $\xi_{\mathscr{M}}$ is α-budget-balanced and semi-cross-monotonic.

(b) for any set S and bid vector $\mathbf{b}$ such that $b_i = -1$ for $i \notin S$ and $b_i > \xi_{\mathscr{M}}(i, S)$ for $i \in S$, the mechanism $\mathscr{M}$ services the set S.

(c) for any bid vector $\mathbf{b}$, if the serviced set is S, then the payment of $i \in S$ is equal to $\xi_{\mathscr{M}}(i, S)$.

We note that this is not a complete characterization of group-strategyproof mechanisms, as there are semi-cross-monotonic cost-sharing schemes that do not correspond to any group-strategyproof mechanism (See Appendix A.2). Finding a complete characterization of cost-sharing schemes that give rise to group-strategyproof mechanisms is an interesting open direction.

Before proving the above theorem, we state two of the corollaries of this theorem. These results characterize group-strategyproof mechanisms that satisfy the following additional properties.

Definition 19 *A mechanism $\mathscr{M}$ is* upper continuous *if for every agent i, if i gets the service for every bid value greater than x holding other bids fixed, then i gets the service if he bids x.*

Definition 20 *A mechanism is* subsidy-free *if, for any bid vector, the total charge to any subset of agents is at most the cost of servicing that subset.*

[13]Notice that this definition allows sets that contain both negative and positive elements. Also, an element can be a positive element of one set and a negative element of another.

Although arguably not well-motivated, the condition of upper-continuity allows us to prove the following equivalence between cross-monotonic cost sharing schemes and group-strategyproof mechanisms satisfying this condition, hence implying that all the upper bounds on the budget-balance factor of cross-monotonic cost-sharing schemes proved in Section 4.2 apply to such mechanisms as well. This theorem can be viewed as guidance in the search for group-strategyproof mechanisms: in order to design a mechanism with better revenue properties than the best cross-monotonic cost-sharing schemes, one must build a mechanism which violates upper continuity.

Theorem 4.3.3 *The cost function C has an upper-continuous α-budget-balanced group-strategyproof cost-sharing mechanism if and only if it has an α-budget-balanced cross-monotonic cost-sharing scheme.*

The subsidy-freeness property was considered previously by Moulin [84]. This property parallels the core condition of cost-sharing games and is motivated by the argument that no subset of serviced agents should be over-charged to accommodate others. The following theorem shows the equivalence of group-strategyproof mechanisms satisfying this property and cross-monotonic cost-sharing schemes, in the case that the mechanism is perfectly budget balanced. We do not know if this theorem holds for budget-balance factors other than 1, and so the results of Section 4.2 only imply that the problems presented there do not have budget-balanced group-strategyproof mechanisms satisfying subsidy-freeness.

Theorem 4.3.4 *The cost function C has a budget-balanced group-strategyproof cost-sharing mechanism satisfying subsidy-freeness if and only if it has a budget-balanced cross-monotonic cost-sharing scheme.*

In the rest of this section, we present the proofs of Theorems 4.3.2, 4.3.3, and 4.3.4.

Proof of Theorem 4.3.2. **(a):** We start by defining the cost-sharing scheme $\xi_{\mathscr{M}}$. For an agent i, let b_i^* be a large enough value such that if agent i bids b_i^*, she will get the service, independent of other agents' bids (such a value exists by CS). For a set $S \subseteq \mathscr{A}$, consider the scenario where the agents in S bid their value in $\mathbf{b}^*$, and others bid -1. By CS and VP, the set of agents serviced by the mechanism in this scenario is precisely S. We define the cost share $\xi_{\mathscr{M}}(i, S)$ as the payment charged by the mechanism to the agent i in this scenario. By this definition and the fact that $\mathscr{M}$ is α-budget balanced, it is clear that $\xi_{\mathscr{M}}$ is also α-budget balanced.

Now, we prove that $\xi_{\mathscr{M}}$ is semi-cross-monotonic. Assume, for contradiction, that there is a set $S \subseteq \mathscr{A}$ and three distinct agents $i, j_1, j_2 \in S$ such that $\xi(j_1, S \setminus \{i\}) < \xi(j_1, S)$ and $\xi(j_2, S \setminus \{i\}) > \xi(j_2, S)$. Consider three bid vectors $\mathbf{b}^1, \mathbf{b}^2$, and $\mathbf{b}^3$ defined as follows: In all of these vectors, agents $j \in S \setminus \{i\}$ bid b_j^* and agents $j \in \mathscr{A} \setminus S$ bid -1. The bid of i in these vectors is $b_i^1 = b_i^*$, $b_i^2 = \xi_{\mathscr{M}}(i, S)$, and $b_i^3 = -1$. By VP and CS, the set of serviced agents at $\mathbf{b}^1$ is S, at $\mathbf{b}^3$ is $S \setminus \{i\}$, and at $\mathbf{b}^2$ is either S or $S \setminus \{i\}$. Furthermore, by the definition of $\xi_{\mathscr{M}}$, the payment of each agent j at the bid vectors $\mathbf{b}^1$ and $\mathbf{b}^3$ is $\xi(j, S)$ and $\xi(j, S \setminus \{i\})$, respectively. We consider two cases based on whether i is serviced at the bid vector $\mathbf{b}^2$:

Case 1: i is served at the bid vector $\mathbf{b}^2$. By VP, i's payment at $\mathbf{b}^2$ is at most $b_i^2 = \xi_{\mathscr{M}}(i, S)$. If i's payment is *strictly* less than $\xi_{\mathscr{M}}(i, S)$, then in a scenario where the utility of the agents is given by $\mathbf{b}^1$, i would have an incentive to announce a bid of b_i^2, contradicting the strategyproofness of the mechanism. Therefore, when all agents bid according to $\mathbf{b}^2$, the payment of i must be equal to $\xi_{\mathscr{M}}(i, S)$. Now consider the payment $x_{j_1}(\mathbf{b}^2)$ of j_1 when agents bid $\mathbf{b}^2$. If $x_{j_1}(\mathbf{b}^2) < \xi(j_1, S)$, then in the scenario where the utility of the agents is given by $\mathbf{b}^1$, $\{i, j_1\}$ can form a successful coalition: they can bid according to $\mathbf{b}^2$, thereby decreasing the payment of j_1, and not changing the payment of i. Also, if $x_{j_1}(\mathbf{b}^2) > \xi(j_1, S \setminus \{i\})$, then in the scenario where the utility of the agents is given by $\mathbf{b}^2$, $\{i, j_1\}$ can form a successful coalition: they can bid according to $\mathbf{b}^1$. This decreases the payment of j_1, and i is indifferent between the two situations, as her utility is zero in both. Thus, $\xi(j_1, S) \leq x_{j_1}(\mathbf{b}^2) \leq \xi(j_1, S \setminus \{i\})$, contradicting the definition of j_1.

Case 2: i is not served at the bid vector $\mathbf{b}^2$. Consider the payment $x_{j_2}(\mathbf{b}^2)$ of j_2 when agents bid $\mathbf{b}^2$. If $x_{j_2}(\mathbf{b}^2) < \xi(j_2, S \setminus \{i\})$, then if the true utility of the agents is given by $\mathbf{b}^3$, $\{i, j_2\}$ can form a coalition: they can bid according to b^2, thereby reducing j_2's payment while keeping the utility of i constant at zero. Also, if $x_{j_2}(\mathbf{b}^2) > \xi(j_2, S)$, then if the utility of the agents is given by b^2, $\{i, j_2\}$ can form a coalition and bid according to b^1, thereby reducing j_2's payment and keeping i's utility constant at zero. Therefore, $\xi(j_2, S\backslash\{i\}) \leq x_{j_2}(\mathbf{b}^2) \leq \xi(j_2, S)$, contradicting the definition of j_2.

The contradiction in both cases shows that $\xi_{\mathscr{M}}$ is semi-cross-monotonic.

(b): Index the agents such that $S = \{1, \ldots, k\}$. For $i = 0, \ldots, k$, define the bid vector $\mathbf{b}^{(i)}$ as follows: $b_j^{(i)} = b_j^*$ for $1 \leq j \leq k-i$, $b_j^{(i)} = b_j > \xi_{\mathscr{M}}(j, S)$ for $k - i < j \leq k$, and $b_j^{(i)} = -1$ for $j \in \mathscr{A} \setminus S$. We will prove by induction on i that if the agents bid $\mathbf{b}^{(i)}$, then the mechanism $\mathscr{M}$ will service the agents in S and charges $j \in S$ an amount equal to $\xi_{\mathscr{M}}(j, S)$. This statement for $i = k$ would imply (b). The induction basis ($i = 0$) is obvious from CS and the definition of $\xi_{\mathscr{M}}$. To show the induction step, we assume that the statement is true for i and prove it for $i + 1$. The only difference between the bid vectors $\mathbf{b}^{(i)}$ and $\mathbf{b}^{(i+1)}$ is the bid of the agent $k - i$. If at the bid vector $\mathbf{b}^{(i+1)}$ agent $k - i$ is either not serviced, or is charged an amount more than $\xi_{\mathscr{M}}(k-i, S)$, then this agent has an incentive to announce a bid of b_{k-i}^* when the true utilities of the agents is given by $\mathbf{b}^{(i+1)}$. Similarly, if $k - i$ is serviced and charged an amount less than $\xi_{\mathscr{M}}(k - i, S)$ when agents bid according to $\mathbf{b}^{(i+1)}$, then when the true utilities of the agents is given by $\mathbf{b}^{(i)}$, agent $k - i$ has an incentive to bid b_{k-i}. Therefore, at $\mathbf{b}^{(i+1)}$, $k - i$ gets serviced and pays $\xi_{\mathscr{M}}(k - i, S)$. This means that from the perspective of agent $k - i$, outcomes at $\mathbf{b}^{(i)}$ and $\mathbf{b}^{(i+1)}$ are the same. Therefore, for every other agent j, the agent j must be indifferent between these two outcomes as well, since otherwise $\{i, j\}$ can form a coalition at one of the two bid vectors $\mathbf{b}^{(i)}$ or $\mathbf{b}^{(i+1)}$. Therefore, by the induction hypothesis, at the bid vector $\mathbf{b}^{(i+1)}$, every agent $j \in S$ must receive the service and be charged $\xi(j, S)$.

(c): Let $S_1 = \{i \in S \mid b_i \leq \xi_{\mathscr{M}}(i, S)\}$, $S_2 = S \setminus S_1$, and $S_3 = \mathscr{A} \setminus S$. By VP, every $i \in S_1$ is not charged more than $\xi_{\mathscr{M}}(i, S)$ at $\mathbf{b}$. Suppose the price charged to some agent $i^* \in S_1$ is strictly less than $\xi(i^*, S)$. Consider a bid vector $\mathbf{b}'$ in which every agent $i \in S_1$ bids b_i^*, every $i \in S_2$ bids b_i (his bid in $\mathbf{b}$) and every $i \in S_3$ bids -1.

From part (b), at the bid vector $\mathbf{b}'$, set S will receive the service and $i \in S$ will pay $\xi_{\mathscr{M}}(i, S)$. Now, since the agent $i^* \in S_1$ is charged strictly less than $\xi_{\mathscr{M}}(i^*, S)$ at $\mathbf{b}$, then when the true utilities are given by $\mathbf{b}'$, i^* can form a coalition with the agents in $S_1 \cup S_3$ and submit the bid vector $\mathbf{b}$. As a result, i^* pays strictly less and no member of the coalition pays more, contradicting group-strategyproofness. Therefore the price of any agent $i \in S_1$ equals $\xi_{\mathscr{M}}(i, S)$ at the bid vector $\mathbf{b}$.

Now consider an agent $i \in S_2$. If his payment differs between $\mathbf{b}$ and $\mathbf{b}'$, then i can form a coalition with the agents in $S_1 \cup S_3$ and submit the bid vector in which he pays less. Agent i strictly benefits from this, while the situation of the agents in $S_1 \cup S_3$ does not change, again contradicting the group-strategyproofness of $\mathscr{M}$. Therefore the payment of every agent $i \in S_2$ also equals $\xi_{\mathscr{M}}(i, S)$. □

Proof of Theorem 4.3.3. The "if" part of this statement follows from Theorem 4.3.1 and the simple observation that the Moulin mechanism $\mathscr{M}_\xi$ is upper continuous.

Given an α-budget-balanced group-strategyproof mechanism $\mathscr{M}$, we show that the cost-sharing scheme $\xi_{\mathscr{M}}$ defined in the proof of Theorem 4.3.2 is cross-monotonic. In other words, we need to show that every element of every set is either positive or neutral. Define $\mathbf{b}^*$ as in the proof of Theorem 4.3.2. Consider a set $S \subseteq \mathscr{A}$ and an agent $i \in S$. Let $\mathbf{b}$ be a bid vector such that $b_j = b_j^*$ for every $j \in S \setminus \{i\}$, $b_j = -1$ for every $j \in \mathscr{A} \setminus S$, and b_i is any number greater than $\xi_{\mathscr{M}}(i, S)$. By part (b) of Theorem 4.3.2, at any such bid vector, the set S gets the service. Therefore, by the upper continuity of $\mathscr{M}$ and CS, the set S gets the service when i bids $\xi_{\mathscr{M}}(i, S)$ and every other agent bids according to $\mathbf{b}$. Call this bid vector $\mathbf{b}'$.

Now, assume, for contradiction, that $\xi_{\mathscr{M}}(j, S \setminus \{i\}) < \xi_{\mathscr{M}}(j, S)$ for some $j \in S \setminus \{i\}$. We argue that $\{i, j\}$ can form a successful coalition when the utilities of the agents is given by $\mathbf{b}'$. In this situation, if i bids -1 and j does not change her bid, then by Theorem 4.3.2 the set $S \setminus \{i\}$ receives the service and agent j pays $\xi_{\mathscr{M}}(j, S \setminus \{i\})$. This outcome makes the agent j strictly happier, and agent i is indifferent between the two outcomes. This contradicts the group-strategyproofness of $\mathscr{M}$. This contradiction shows that every element i of every set S is either positive or neutral, and hence $\xi_{\mathscr{M}}$ is cross-monotonic. $\square$

Proof of Theorem 4.3.4. As in the previous proof, the "if" direction is a direct corollary of Theorem 4.3.1 and the simple observation that $\mathscr{M}_\xi$ satisfies subsidy-freeness.

Given a subsidy-free 1-budget-balanced mechanism $\mathscr{M}$, we show that the cost-sharing scheme $\xi_{\mathscr{M}}$ defined in Theorem 4.3.2 is cross-monotonic. First, notice that by part (c) of Theorem 4.3.2, subsidy-freeness of $\mathscr{M}$ implies that $\xi_{\mathscr{M}}$ is in the 1-core of C, that is, for every $T \subseteq S \subseteq \mathscr{A}$, we have

$$\sum_{j \in T} \xi_{\mathscr{M}}(j, S) \leq C(T). \tag{4.8}$$

Now, consider a set $S \subseteq \mathscr{A}$ and an agent $i \in S$. If i is a negative element of S, then for every $j \in S \setminus \{i\}$, we have $\xi_{\mathscr{M}}(j, S) \geq \xi_{\mathscr{M}}(j, S \setminus \{i\})$, and at least for one j this inequality is strict. Therefore,

$$\sum_{j \in S \setminus \{i\}} \xi_{\mathscr{M}}(j, S) > \sum_{j \in S \setminus \{i\}} \xi_{\mathscr{M}}(j, S \setminus \{i\}) = C(S - \{i\}), \tag{4.9}$$

where the last equality follows from the fact that $\mathscr{M}$ is 1-budget-balanced. Equation 4.9 contradicts Equation 4.8 for $T = S \setminus \{i\}$. $\square$

Chapter 5

Two-Sided Markets

Suppose all the eligible bachelors and bachelorettes in a town confide in the town's matchmaker their ideal spouses. Each man submits an ordered preference list of the women he would like to marry. Similarly, each woman submits an ordered preference list of the men she would like to marry. The matchmaker must arrange marriages such that no one is tempted to ask for a divorce. In particular, the matchmaker must be sure that there is no pair of young lovers who prefer each other to their assigned spouses. Such a set of marriages is called *stable*, and finding a set of stable marriages is known as the *stable marriage problem*. Gale and Shapley [39] showed that the stable marriage problem always has a solution and developed an algorithm, called the *deferred acceptance algorithm*, to find it. Since the seminal work of Gale and Shapley, there has been a significant amount of work on the mathematical structure of stable marriages and related algorithmic questions. See, for example, the books by Knuth [69], Gusfield and Irving [53], or Roth and Sotomayoror [101].

The stable marriage problem has many promising applications in two-sided markets such as job markets [98], college admissions [98], sorority/fraternity rush [83], and assignment of graduating rabbis to their first congregation [12]. Since most applications of the stable marriage algorithm involve the participation of independent agents, it is natural to investigate how we should expect these agents to behave. In particular, we would like to know whether agents can benefit by being dishonest about their preference lists. Ideally, in economic settings such as job markets, we would like to design truthful mechanisms which always output a stable matching. Unfortunately, as shown by Roth [95], there is no mechanism for the stable marriage problem in which truth-telling is a dominant strategy for both men and women [101].

Nonetheless, stable matching algorithms have had spectacular success in practical applications. One particular job market — the medical residency market — has been using a centralized stable marriage market system called the National Residency Matching Program (NRMP) since the 1950s [96]. To this day, most medical residences are formed through an updated version of this centralized market system redesigned in 1998 by Roth [97].[1] It seems surprising that an algorithm like the one used by the NRMP which provably admits strategic behavior can be so successful. Roth and Peranson [99] noted that, in practice, very few students and hospitals could have benefited by submitting false preferences. They analyzed several years of data from the NRMP and calculated whether any applicant could improve his or her match (according to his submitted preference list) by altering his preference list. For example, in 1996, they calculated that out of 24,749 applicants, just 21 could have affected their match by changing their submitted preferences. One explanation for this observation is that the data did not in fact reflect the true preferences of the applicants but rather an equilibrium of the mechanism. Another clear factor that influences the medical market is the correlation between preference lists. Applicants share a general opinion of "desirable" and "undesirable" hospitals. Similarly, hospitals tend to agree on the "desirable" and "undesirable" applicants. Taken to the extreme where all preference lists are identical, this correlation induces a unique stable matching where no participant can benefit by altering their preference list. Conversely, Knuth, Motwani, and Pittel [70, 71] showed that in the general stable marriage setting, if preference lists are independent random permutations of all members of the opposite sex, then almost every person has more than one stable partner. However, Roth and Peranson [99] conjectured that the main factor influencing the medical market is its sheer size. In a small town, every man knows every woman, but in the medical market, a

[1]Although not addressed in this chapter, the algorithm currently used by the NRMP has the feature that it can accommodate married couples among students that submit joint preference lists. For a discussion about stable marriage with couples, see Appendix C.

student can not possibly interview at every hospital. In practice, the length of applicant preference lists is quite small, about 15, while the number of positions is large, about 20,000. Experimentally, Roth and Peranson [99] showed that size matters. They generated random preference lists of limited length and computed the resulting number of uniquely matched participants. Even though these randomly generated lists are, in a sense, the worst case (that is, there is no correlation between the lists), their experiments show that the number of participants with more than one stable partner (and therefore the number of those that can benefit by lying) is quite small when the length of the lists is sufficiently limited. This led them to conjecture that in this probabilistic setting, the fraction of such people tends to zero as the size of the market tends to infinity.[2]

In this chapter, we prove and generalize this conjecture. More precisely, we prove the following: Assume there are n men and n women in the town, and each woman has an *arbitrary* ordering of all men as her preference list. Each man independently picks a random preference list of a constant (that is, independent of n) number of women by choosing each woman independently according to an *arbitrary* distribution $\mathscr{D}$. These are the true preference lists. We show that in this setting the expected number of people with more than one stable spouse is vanishingly small. We use the following technique for our proof: First, we design an algorithm, based on an algorithm of Knuth, Motwani, and Pittel [70, 71], that for a given woman checks whether she has more than one stable husband in *one* run of proposals. Using this algorithm, we prove a relationship between the probability that a given woman has more than one stable husband and the number of single (that is, unmatched) women who are more popular than she. This relationship, essential to our main result, seems difficult to derive directly, without going through the algorithm. Given this relationship, we are able to derive our result by computing bounds on the expectation and variance of the

[2]Existing results in the literature study the core of markets and conclude that, under certain conditions, the size of the core shrinks as the size of the market grows (see, for example, the seminal paper of Aumann [10] or the book by Hildenbrand [56]). The set of stable matchings is the core of the stable marriage game, but our market setting is quite different from that in the literature. In fact, in our setting, the core is often large even though the fraction of people with more than one stable partner is small.

number of single popular women.

This result has a number of interesting economic implications. We can interpret the preference lists together with a stable marriage algorithm as a game G, in which everybody submits a preference list (not necessarily their true preference list) to the algorithm and receives a spouse. The goal for each player is to receive the best spouse possible according to their true preference list. First, we show that, with probability $1 - o(1)$ (as n approaches infinity), in any stable marriage mechanism, the truthful strategy is the best response for a given player when the other players are truthful. We also show that when a deferred acceptance mechanism is used, there is a Nash equilibrium of this game in which a majority of the players are truthful. Finally, we prove that in the more realistic setting of a game of incomplete information (where each player only knows the distribution of the preference lists), the set of truthful strategies in the game induced by the *women-proposing mechanism* form a $(1+o(1))$-approximate Bayesian-Nash equilibrium. In this ordinal setting, a $(1+\epsilon)$-approximate equilibrium is one in which no player can improve the *rank* of his allocation by more than a factor of $(1+\epsilon)$ in expectation. If the ratio of the largest cardinal preference to the smallest cardinal preference is bounded by a constant, our results carry over to the cardinal setting as well. It is important to note that our results hold for *any* distribution $\mathscr{D}$ of women. For the special case of uniform distributions (which includes the conjecture of Roth and Peranson), the $o(1)$ in the above bounds is roughly e^k/n, and thus the bounds converge quite quickly.

Mechanisms that are truthful in a randomized sense (that is, in expectation, or with high probability) have been a subject of research in theoretical computer science [4, 5]. These mechanisms seek to encourage truthfulness by introducing randomization into the mechanism. Our results are of a different flavor. We show that one can conclude statements regarding truthfulness by introducing randomization into the players' utility functions. To the best of our knowledge, our result is the first result of this type.

One can also view our results as an analysis of stable matching with random preferences. There has been a considerable amount of work in this area [70, 71, 91, 92], mostly assuming *complete* preference lists for participants, and none motivated by the economic aspects of the problem. We will use some of the techniques developed in these papers in our analysis. Sethuraman, Teo, and Tan [105, 106] have studied the stable matching game when participants are *required* to announce complete preference lists, and have given an optimal cheating algorithm and several experimental results regarding the chances that an agent can benefit by lying in this game.

The results of this chapter are based on joint work with Mahdian [60].

5.1 Setting

Consider a community consisting of a set $\mathscr{W}$ of n women and a set $\mathscr{M}$ of n men. Each person in this community has a *preference list*, which is a strictly ordered list of a subset of the members of the opposite sex. We assume that if a occurs before b on c's preference list, then c prefers a to b. A *matching* is a mapping μ from $\mathscr{M} \cup \mathscr{W}$ to $\mathscr{M} \cup \mathscr{W}$ in such a way that for every $x \in \mathscr{M}$, $\mu(x) \in \mathscr{W} \cup \{x\}$ and for every $x \in \mathscr{W}$, $\mu(x) \in \mathscr{M} \cup \{x\}$, and also for every $x, y \in \mathscr{M} \cup \mathscr{W}$, $x = \mu(y)$ if and only if $y = \mu(x)$. If for some $m \in \mathscr{M}$ and $w \in \mathscr{W}$, $\mu(m) = w$, we say that w is the wife of m and m is the husband of w in μ; or, if for some $x \in \mathscr{M} \cup \mathscr{W}$, $\mu(x) = x$, we say that x remains single in μ. A pair $m \in \mathscr{M}$, $w \in \mathscr{W}$ is called a *blocking pair* for a matching μ, if m prefers w to $\mu(m)$, and w prefers m to $\mu(w)$. A matching with no blocking pair is called a *stable matching.* If a man m and a woman w are a couple in *some* stable matching μ, we say that m is a *stable husband* of w, and w is a *stable wife* of m. Naturally, each person might have more than one stable partner. In the stable marriage problem, the objective is to find a stable matching given the preference lists of all men and women.

The stable marriage problem was first introduced and studied by Gale and Shapley [39] in 1962. They proved that a stable matching always exists, and a simple algorithm called the *deferred acceptance procedure* can find such a matching. This procedure iteratively selects an unmarried man m and creates a proposal from him to the next woman on his list. If this woman prefers m to her current assignment, then she tentatively accepts m's proposal, and rejects the man she was previously matched to (if any); otherwise, she rejects the proposal of m. The algorithm ends when every man either finds a wife that accepts him, or gets rejected by all the women on his list, in which case he remains single. This algorithm is sometimes called the *men-proposing algorithm*. Similarly, one can define the *women-proposing algorithm*. Gale and Shapley [39] proved the following.

Theorem 5.1.1 *The men-proposing algorithm always finds a stable matching μ. Furthermore, this stable matching is* men-optimal, *that is, for every man m and every stable wife w of m other than $\mu(m)$, m prefers $\mu(m)$ to w. At the same time, μ is the worst possible stable matching for women, that is, for any woman w and any stable husband m of w other than $\mu(w)$, w prefers m to $\mu(w)$.*

Notice that in the description of the men-proposing algorithm we did not specify the order in which single men propose. One might naturally think that choosing a different order for proposals might lead to a different stable matching. However, the above theorem together with the fact that the men-optimal stable matching is unique imply the following.

Theorem 5.1.2 *The men-proposing algorithm always finds the same stable matching, independent of the order in which the proposals are made.*

We will also need the following theorem of Roth [96] and McVitie and Wilson [80], which says that the choice of the stable matching algorithm does not affect the number of people who remain unmarried at the end of the algorithm.

Theorem 5.1.3 *In all stable matchings, the set of people who remain single is the same.*

A *stable matching mechanism* is an algorithm that elicits a preference list from each participant, and outputs a matching that is stable with respect to the announced preferences. Ideally, we would like to design mechanisms in which truthfulness (that is, announcing the true preference list to the mechanism) is a dominant strategy for all participants. However, Roth [95] proved that there is no such mechanism for the stable marriage problem. On the positive side, Gale and Sotomayor [40] show that in any stable marriage mechanism, each player has an optimal strategy which is simply a *truncation* (a prefix) of his true preference list. The following theorem (due to Roth [95] and Dubins and Freedman [24]) shows that in deferred acceptance mechanisms, truthfulness is a dominant strategy for half the population.

Theorem 5.1.4 *In the men-optimal stable marriage mechanism, truth-telling is a dominant strategy for men.*

Consider a situation where there are n men and n women. Assume the preference list of each man is chosen independently and uniformly at random from the set of all ordered lists of k women, and the preference list of each woman is picked independently and uniformly at random from the set of all orderings of all men. We want to bound the expected number of people who might be tempted to lie to the mechanism about their preferences when the other players are truthful. As we will show, only people who have more than one stable partner might be able to influence their final match by altering their preference lists. Therefore, we focus on bounding the expected number of women with more than one stable husband in this model. Notice that this number is equal to the expected number of men with more than one stable wife, since, in a market where the two sides are of equal size, the number of single and uniquely matched men must equal the number of single and uniquely matched women. Roth and Peranson [99] conjectured the following.

Conjecture 5.1.1 *Let $c_k(n)$ denote the expected number of women who have more than one stable husband in the above model. Then for all fixed k,*

$$\lim_{n\to\infty} \frac{c_k(n)}{n} = 0.$$

We prove this conjecture. In fact, we will prove the following stronger result. Let $\mathscr{D}$ be an *arbitrary* fixed distribution over the set of women such that the probability of each woman in $\mathscr{D}$ is nonzero.[3] Intuitively, having a high probability in $\mathscr{D}$ indicates

[3]This assumption is needed to make sure that the problem is well-defined.

that a woman is popular. The preference lists are constructed by picking each entry of the list according to $\mathscr{D}$, and removing the repetitions. More precisely, we construct a random list $(l_1, \ldots, l_k)$ of k women as follows. At step i, repeatedly select a women w independently according to $\mathscr{D}$ until $w \notin \{l_1, \ldots, l_{i-1}\}$ and then set $l_i = w$. Let $\mathscr{D}^k$ be the distribution over lists of size k produced by this process. Notice that if $\mathscr{D}$ is the uniform distribution, $\mathscr{D}^k$ is nothing but the uniform distribution over the set of all lists of size k of women. Therefore, the model of Roth and Peranson [99] is a special case of our model. We also generalize their model in another respect: we assume that women have *arbitrary* complete preference lists, as opposed to the assumption in [99] that they have random complete preference lists. Our main result is the following theorem.

Theorem 5.1.5 *Consider a situation where each woman has an* arbitrary *complete preference list, and each man has a preference list chosen independently at random according to $\mathscr{D}^k$. Let $c_k(n)$ denote the expected number of women who have more than one stable husband in this model. Then, for all fixed k,*

$$\lim_{n\to\infty} \frac{c_k(n)}{n} = 0.$$

Remark 5.1.1 *One might hope to further generalize this model to one where each man picks a random list from an arbitrary distribution over lists of size k. However, the following example shows that Theorem 5.1.5 is not true in this model: Assume women $1, \ldots, n/2$ rank men in the order $1, 2, \ldots, n$, and women $n/2+1, \ldots, n$ rank them in the reverse order. Each man picks a random $i \in \{1, \ldots, n/2\}$, and with probability $1/2$ picks preference list $(i, i+n/2)$ and otherwise picks preference list $(i+n/2, i)$. It is not difficult to see that with these preferences, at least a $1/(8e)$ fraction of women will have more than one stable husband.*

Even though we state and prove our results assuming that all preference lists are of size exactly k, it is straightforward to see that our proof carries over to the case where preference lists are of size at most k. For uniform distributions, we can prove a strong result on the rate of convergence of this limit.

Theorem 5.1.6 *Consider a situation where each woman has an* arbitrary *complete preference list, and each man has a preference list of k women chosen uniformly and independently. Then, the expected number of women who have more than one stable husband is bounded by $e^{k+1} + k^2$, a constant that only depends on k (and not on n).*

5.2 Economic implications

There are a number of interesting economic implications of this theorem. Our first result states that, with high probability, a given player's best strategy is truth-telling when the other players are truthful. Thus, a dishonest player who believes in the honesty of the other players has an economic incentive to be honest.

Corollary 5.2.1 *Fix any stable matching mechanism, and consider an instance with n women with arbitrary complete preference lists and n men with preference lists drawn from $\mathscr{D}^k$ (as in Theorem 5.1.5). Then, for any given person x, the probability (over the men's preference lists) that for x the truthful strategy is not the best response in a situation where the other players are truthful is $o(1)$ (at most $O(e^k/n)$ for uniform distributions).*

Proof. Fix a person, say a man named Adam, and suppose all other players are truthful. Theorem 5.1.5 implies that with probability at least $1 - o(1)$, Adam has at most one stable wife, Eve, with respect to the true preference lists of the players. Suppose all other players are truthful. We claim Adam's best response is truth-telling. Suppose not. Allow Adam to play his best response p and let μ be the matching that the stable matching mechanism outputs. Now run the men-optimal algorithm with the same preference lists (that is, p for Adam, and true preference lists for others) and let μ_M be the resulting matching. By Theorem 5.1.1, Adam must prefer his match in μ_M to his match in μ. However, by Theorem 5.1.4, in the men-proposing algorithm, Adam's dominant strategy is truth-telling and, by assumption, matches him to Eve. Therefore, Adam must prefer Eve to his match in μ_M and thus to his match in μ. But Eve is the woman that Adam would have been matched to in the original mechanism if he had been truthful (since it was his unique stable match), and so his altered strategy p was not his best response. □

The previous corollary states that a player can benefit by lying only with a vanishingly small probability when the other players are truthful. Now we turn to the situation in which the other players are not necessarily truthful, but are playing an equilibrium strategy of the game induced by the stable matching mechanism. There are two ways to interpret our stable marriage setting as a game. One way is to consider it as a game of complete information: Let P_m and P_w denote the preference lists of men and women. Knowing these preferences, each player chooses a strategy from the strategy space of all possible preference lists. The corresponding preference lists are submitted to a fixed stable marriage algorithm and a matching is returned. A player's goal is to choose the strategy that gets him/her a spouse as high on his/her preference list as possible. Let G_{P_m,P_w} denote this game.

Corollary 5.2.2 *Assume the preference lists P_w of women are arbitrary, and the preference lists P_m of men are drawn from $\mathscr{D}^k$ (as in Theorem 5.1.5). The game G_{P_m,P_w} induced by these preferences and the men-proposing (or women-proposing) mechanism has a Nash equilibrium in which, in expectation, a $(1-o(1))$ fraction of strategies are truthful.*

Proof. Suppose we are using the men-proposing mechanism (the women-proposing situation is analogous). We prove that the following set of strategies forms an equilibrium in the game G_{P_m,P_w}: all men announce their true preferences; all women who have at most one stable husband (with respect to P_m, P_w) announce their true preferences; and all women who have more than one stable husband truncate their preference lists just after their optimal stable husband. We denote the altered preference lists of women by P'_w. By Theorem 5.1.4, men cannot improve their situation by altering their strategy. Consider a woman, say Eve, and assume Eve will be assigned to Adam if the players use the strategies in (P_m, P'_w). It is easy to see that there is a unique stable matching with respect to (P_m, P'_w). Therefore, if we run the women-optimal mechanism on (P_m, P'_w), we get the same outcome as in the men-optimal mechanism. However, by Theorem 5.1.4 we know that no woman can benefit from altering her preferences in a women-optimal mechanism. Thus, if Eve changes her strategy from the one dictated by P'_w, then she gets a match, say Tom, that according to P'_w is not better than Adam. However, by the definition of P'_w, this implies that Tom is not better than Adam according to the true preferences of Eve. This shows that the set of strategies (P_m, P'_w) is an equilibrium. By Theorem 5.1.5, we know that all men and all but at most a $o(1)$ fraction of women are truthful in this equilibrium. $\square$

In the above setting, we assume that each player knows the preference lists of the other players when he/she is selecting a strategy, that is, we have a game of complete information. A more realistic assumption is that each player only knows the distribution of preference lists of the other players. Each player's goal is to alter his/her preference list and announce it to the mechanism in a way that the *expected* rank of his/her assigned spouse is as high as possible. A strategy for a player is a function that outputs an announced preference list for any input preference list. Hence the truthful strategy is the identity function. We wish to analyze the *Bayesian-Nash equilibria* in this incomplete information game. A $(1+\varepsilon)$-*approximate Bayesian-Nash equilibrium* for this game is a collection of strategies, one for each player, such that no single player can improve the rank (computed according to his/her true preference list) of his/her spouse by more than a multiplicative factor of $1+\varepsilon$ by deviating from his/her equilibrium strategy.

Corollary 5.2.3 *Consider the game described above with the women-optimal mechanism. Then for every $\varepsilon > 0$, if n is large enough, the above game has a $(1+\varepsilon)$-approximate Nash equilibrium in which everybody is truthful.*

Proof. Since the women-optimal mechanism is used, we know by Theorem 5.1.4 that truthfulness is a dominant strategy for women. It is enough to show that if all men and women are truthful, then no man can improve his match by more than a $(1+\varepsilon)$ factor if he uses a dishonest strategy. Fix a man, Charlie. With probability $1-o(1)$, preferences are such that Charlie does not have more than one stable wife.

In this case, the argument used in the proof of the previous two corollaries shows that Charlie cannot gain by being dishonest about his preferences. With probability $o(1)$, Charlie has more than one stable wife, and in that case, he might be able to improve his match from someone ranked at most k in his list to someone ranked first. However, k is a constant. Using this, it is easy to verify that on average, he can improve his match by at most a factor of $1 + k^2 \times o(1) = 1 + o(1)$. Thus, everyone being truthful is an approximate equilibrium in this game. □

Although we defined approximate equilibrium in Corollary 5.2.3 with respect to *ordinal* preferences, the result also holds in the following *cardinal* setting: Each player i has a *distinct* utility $u_{ij} \in \mathbb{R}$ for being matched to player j (hence the true preference list of i is $(j_1, \ldots, j_l)$ where $u_{ij_1} > u_{ij_2} > \ldots > u_{ij_l} \geq 0$), and the ratio $\frac{\max_j(u_{ij})}{\min_j(u_{ij})}$ of the maximum utility to the minimum utility is bounded by a constant for all i.

5.3 Proof of Theorem 5.1.5

In this section, we will prove our main technical result, Theorem 5.1.5. The proof consists of three main components. First, we present an algorithm that, given the preference lists, counts the number of stable husbands of a given woman (Section 5.3.1). We would like to analyze the probability that the output of this algorithm is more than one, over a distribution of inputs. In Section 5.3.2, we bound this probability assuming a lemma concerning the number of singles in a stable marriage. This lemma is proved in Section 5.3.3 by bounding the expectation of the number of singles and proving that it is concentrated around its expected value using the Chebyschev inequality.

5.3.1 Counting the number of stable husbands

The simplest way to check whether a woman g has more than one stable husband or not is to compute the men-optimal and the women-optimal stable matchings using the algorithm of Gale and Shapley (See Theorem 5.1.1) and then check if g has the same husband in both these matchings. However, analyzing the probability that g has more than one stable husband using this algorithm is not easy, since we will not be able to use the principle of deferred decisions (as described later in Section 5.3.2). In this section we present a different algorithm that outputs all stable husbands of a given woman in an arbitrary stable marriage problem in one run of a men-propose algorithm. This algorithm is a generalization of the algorithm of Knuth, Motwani, and Pittel [70, 71] to the case of incomplete preference lists.

Suppose we want the stable husbands of woman g. Initially all the people are unmarried (the matching is empty). The algorithm closely follows the man-proposing algorithm for finding a stable matching. However, g's objective is to explore all her options. Therefore, every time the men-proposing algorithm finds a stable marriage, g divorces her husband and lets the algorithm continue.

Algorithm A

1. Initialization: Run the man-proposing algorithm to find the men-optimal stable matching. If g is unmarried, output $\emptyset$.

2. Selection of the suitor: Output the husband m of g as one of her stable husbands. Remove the pair (m, g) from the matching (woman g and man m are now unmarried) and set $b = m$. (The variable b is the current proposing man.)

3. Selection of the courted: If b has already proposed to all the women on his preference list, terminate. Otherwise, let w be his favorite woman among those he hasn't proposed to yet.

4. The courtship:

   (a) If w has received a proposal from a man she likes better than b, she rejects b and the algorithm continues at the third step.

   (b) If not, w accepts b. If $w = g$, the algorithm continues at the second step. Otherwise, if w was previously married, her previous husband becomes the suitor b and the algorithm continues at the third step. If w was previously unmarried, terminate the algorithm.

Notice that in step 4(a) of the algorithm, w compares b to the best man who has proposed to her so far, and not to the man she is currently matched to. Therefore, after g divorces one of her stable husbands, she has a higher standard, and will not accept any man worse than the man she has divorced. For $w \neq g$, step 4(a) is equivalent to comparing b to the man w is matched to at the moment.

We must prove that this algorithm outputs all stable husbands of g. In fact, we will prove something slightly stronger.

Theorem 5.3.1 *Algorithm Aoutputs all stable husbands of g in order of her preference from her worst stable husband to her best stable husband.*

Proof. We prove the theorem by induction. As the man-proposing algorithm returns the worst possible matching for the women (by Theorem 5.1.1), the first output is g's worst stable husband. Now suppose the i'th output is g's i'th worst stable husband m_i. Consider running the man-proposing algorithm with g's preference list truncated just before man m_i (so that it includes all men she prefers to m_i but not m_i himself). As the order of proposals in the men-proposing algorithm do not affect the outcome (Theorem 5.1.2), let the order of proposals be the same as Algorithm A.Then, up until Algorithm Aoutputs the $i+1$'st output m_{i+1}, its tentative matching during the j'th proposal is the same as the tentative matching of the man-proposing algorithm during the j'th proposal (except, possibly, woman g is matched in Algorithm Aand unmatched in the man-proposing algorithm). Now since m_{i+1} was accepted, the fourth step guarantees that g preferred m_{i+1} to m_i. Thus m_{i+1} is on g's truncated preference list, and so the tentative matchings of the two algorithms are the same. Furthermore, m_{i+1} is the first proposal g has accepted in the man-proposing algorithm. All other women who get married in the set of stable matchings already have husbands since they have husbands in Algorithm A,and so the man-proposing algorithm terminates with the current matching. Thus, m_{i+1} is the worst possible stable husband for g that is better than m_i. □

5.3.2 Analyzing the expectation

We are interested in the expected number of women with more than one stable husband, or, equivalently, the probability that a fixed woman g has more than one stable husband. We can compute this probability by analyzing the output of Algorithm Afrom Section 5.3.1 on male preference lists drawn from the distribution $\mathscr{D}^k$. We simulate this experiment using the *principle of deferred decisions*: a man only needs to determine his i'th favorite woman when he makes his i'th proposal. If we make these deferred decisions independently according to $\mathscr{D}$, then the distribution of the output of this new algorithm over its coin flips will be exactly the same as the distribution of the output of the old algorithm over its input. This motivates the definition of the following algorithm which counts the number x_g of stable husbands of a girl g. At any point in this algorithm, the variable A_i denotes the set of women that man i has proposed to so far. Men and women are indexed by numbers between 1 and n.

Algorithm B

1. Initialization: Let $l = 1$, $\forall\ 1 \leq i \leq n,\ A_i = \emptyset$, $x_g = 0$. (The matching is empty and no men have made any proposals).

2. Selection of the suitor:

   (a) If $l \leq n$, let b be the l'th man and increase l by one.

   (b) Otherwise, we have found a stable matching. If g is single in this stable matching, then terminate. Otherwise, increment x_g, remove the pair (m, g) from the matching (man m and woman g who were previously married to each other are now unmarried) and set $b = m$.

3. Selection of the courted:

(a) If $|A_b| \geq k$, then do the following: If $x_g \geq 1$ (we have found a stable matching and a previously married man is now single), then terminate. Otherwise, return to step two.

(b) Repeatedly select w randomly according to distribution $\mathscr{D}$ from the set of all women until $w \notin A_b$. Add w to A_b.

4. The courtship:

(a) If w has received a proposal from a man she likes better than b, she rejects b and the algorithm continues at step 3.

(b) If not, w accepts b. If w was previously married, her previous husband becomes the suitor b and the algorithm continues at the third step. If w was previously single and $x_g = 0$ (we have not found a stable matching), the algorithm continues at the second step. If w was previously single and $x_g \geq 1$, the algorithm continues at the second step if $w = g$ and terminates if $w \neq g$.

Before giving a proof of Theorem 5.1.5, we introduce some notation. For every woman i, let p_i denote the probability of i in the distribution $\mathscr{D}$. We say that a woman i is *more popular* than another woman j, if $p_i \geq p_j$. Assume, without loss of generality, that women are ordered in the decreasing order of popularity, that is, $p_1 \geq p_2 \geq \cdots \geq p_n$.

Proof of Theorem 5.1.5. Recall that $c_k(n)$ is the expected number of women with more than one stable husband. We show that for every $\epsilon > 0$, if n is large enough, then $c_k(n)/n \leq \epsilon$. By linearity of expectation, $c_k(n) = \sum_{g \in \mathscr{W}} \mathbf{Pr}[g$ has more than one stable husband$]$. Fix a woman $g \in \mathscr{W}$. We want to bound the probability that g has more than one stable husband. By Theorem 5.3.1 and the principle of deferred decisions, this is the same as bounding the probability that the random variable x_g in Algorithm Bis more than one.

We divide the execution of Algorithm Binto two phases: the first phase is from the beginning of the algorithm until it finds the first stable matching, and the second phase is from that point until the algorithm terminates. Assume at the end of the first phase, Algorithm Bhas found the first stable matching μ. We bound the probability that $x_g > 1$ conditioned on the event that μ is the matching found at the end of the first phase (we denote this by $\mathbf{Pr}[x_g > 1 \mid \mu]$), and then take the expectation of this bound over μ.

Let the set $S_\mu(g)$ denote the set of women more popular than g that remain single in the stable matching μ and $X_\mu(g) = |S_\mu(g)|$. If g is single in μ, then $x_g = 0$ and therefore $\mathbf{Pr}[x_g > 1 \mid \mu] = 0$. Otherwise, $x_g > 1$ if only if woman g accepts another proposal before the algorithm terminates. We bound this by the probability that g receives another proposal before the end of the algorithm. The algorithm may terminate in several ways, but we will concentrate on the termination condition in step 4(b), that is, that some man proposes to a previously single woman. Thus, we are interested in the probability that in the second phase of Algorithm Bsome man proposes to a previously single woman before any man proposes to g.

Note that at the end of the first phase of the algorithm, all A_i's are disjoint from $S_\mu(g)$, since women have complete preference lists. Thus whenever the random oracle in step 3(b) outputs a woman from set $S_\mu(g)$, the algorithm will advance to step 4(b) and terminate. Thus, the probability $\mathbf{Pr}[x_g > 1 \mid \mu]$ is less than or equal to the probability that in a sequence whose elements are independently picked from the distribution $\mathscr{D}$, g appears before any woman in $S_\mu(g)$. By the definition of $S_\mu(g)$, for every $w \in S_\mu(g)$, every time we pick a woman randomly according to $\mathscr{D}$, the probability that w is picked is at least as large as the probability that g is picked. Therefore, the probability that g appears before all elements of $S_\mu(g)$ in a sequence whose elements are picked according to $\mathscr{D}$ is at most the probability the g appears first in a random permutation on the elements of $\{g\} \cup S_\mu(g)$, which is $1/(X_\mu(g)+1)$. Thus, for every μ,

$$\mathbf{Pr}[x_g > 1 \mid \mu] \leq \frac{1}{X_\mu(g) + 1} \tag{5.1}$$

Thus,

$$\begin{aligned}
\mathbf{Pr}[x_g > 1] &= \mathrm{E}_\mu \left[\mathbf{Pr}[x_g > 1 \mid \mu]\right] \\
&\le \mathrm{E}_\mu \left[\frac{1}{X_\mu(g)+1}\right].
\end{aligned} \tag{5.2}$$

We complete the proof assuming the following lemma, whose proof is given in Section 5.3.3.

Lemma 5.3.1 *For every $g > 4k$,*

$$\mathrm{E}\left[\frac{1}{X_\mu(g)+1}\right] \le \frac{12e^{8nk/g}}{g}.$$

Thus, using equation (5.2) and Lemma 5.3.1 for $g \ge \frac{16nk}{\ln(n)}$, and $\mathbf{Pr}[x_g > 1] \le 1$ for smaller g's, we obtain

$$\begin{aligned}
c_k(n) &\le \frac{16nk}{\ln(n)} + \sum_{g=\frac{16nk}{\ln(n)}}^{n} \frac{12e^{8nk/g}}{g} \\
&\le \frac{16nk}{\ln(n)} + \sum_{g=\frac{16nk}{\ln(n)}}^{n} \frac{3\ln(n)e^{\ln(n)/2}}{4nk} \\
&\le \frac{16nk}{\ln(n)} + 3\sqrt{n}\ln(n)/(4k) = o(n),
\end{aligned}$$

and so for every constant k, the fraction of women with more than one stable husband, $c_k(n)/n$, goes to zero as n tends to infinity. □

For the case of uniform distributions, since every woman is equally popular, $S_\mu(g) = S_\mu(g')$ is the set of all women, and so $\mathrm{E}[\frac{1}{X_\mu(g)+1}] = \mathrm{E}[\frac{1}{X_\mu(g')+1}] \le \frac{12e^{8nk/g}}{g}$ for $g > 4k$. Thus, $c_k(n) \le 4k + \sum_{g=4k}^{n} \frac{12e^{8k}}{n} \le 4k + 12e^{8k}$. We derive an even tighter bound in this case, as stated in Theorem 5.1.6, using a slightly different technique. This bound is proved in Section 5.4.

5.3.3 Number of singles

In this section we prove Lemma 5.3.1. This completes the proof of Theorem 5.1.5. We start with the following simple fact: the probability that a woman w remains single is greater than or equal to the probability that w does not appear on the preference list of any man. More precisely, let E_w denote the event that the woman w does not appear on the preference list of any man when these preferences are drawn from $\mathscr{D}^k$. Let Y_g denote the number of women $w \leq g$ for which the event E_w happens. Then we have the following lemma.

Lemma 5.3.2 *For every g, we always have $X_\mu(g) \geq Y_g$.*[4]

Proof. Every woman $w < g$ for which E_w happens is a woman who is at least as popular as g and will remain unmarried in any stable matching. □

We now bound the expectation of $1/(Y_g + 1)$ in a sequence of two lemmas. In Lemma 5.3.3 we bound the expectation of Y_g. Then, in Lemma 5.3.4 we show the variance of Y_g is small and therefore it does not deviate far from its mean.

Lemma 5.3.3 *For $g > 4k$, the expected number $\mathrm{E}[X_\mu(g)]$ of single women more popular than woman g is at least $\frac{g}{2}e^{-8nk/g}$.*

Proof. Let $Q = \sum_{j=1}^{k} p_j$ denote the total probability of the first k women under $\mathscr{D}$. The probability that a man m does not list a woman w as his i'th preference given that he picks $w_1, \ldots, w_{i-1}$ as his first $i-1$ women, is equal to

$$1 - \frac{p_w}{1 - \sum_{j=1}^{i-1} p_{w_j}} \geq 1 - \frac{p_w}{1 - Q}.$$

Thus the probability that m does not list w at all is at least $(1 - \frac{p_w}{1-Q})^k$, and so the probability that woman w is not listed by any man is at least $(1 - \frac{p_w}{1-Q})^{nk}$. If $w > k$, there are at least $w - k$ women who are at least as popular as w, but not among the k most popular women. Therefore, $p_w \leq \frac{1-Q}{w-k}$. By these two inequalities, for every $w > 2k$ we have

[4]In more mathematical terms, this means that $X_\mu(g)$ stochastically dominates Y_g.

$$\mathbf{Pr}[E_w] \geq (1 - \frac{1}{w-k})^{nk} \geq e^{-2nk/(w-k)} \geq e^{-4nk/w}.$$

Therefore, for every $g > 4k$, the expectation of Y_g is at least

$$\mathrm{E}[Y_g] = \sum_{w=1}^{g} \mathbf{Pr}[E_w] \geq \sum_{j=2k}^{g} e^{-4nk/j} \geq \sum_{j=g/2}^{g} e^{-8nk/g} = \frac{g}{2} e^{-8nk/g}, \tag{5.3}$$

yielding the result. □

Lemma 5.3.4 *The variance $\sigma^2(Y_g)$ of the random variable Y_g is at most its expectation $\mathrm{E}[Y_g]$.*

Proof. We show the events E_i are negatively correlated, that is, for every i and j, $\mathbf{Pr}[E_i \wedge E_j] \leq \mathbf{Pr}[E_i].\mathbf{Pr}[E_j]$. Let F_i be the event that a given man does not include woman i on his preference list. By the independence and symmetry of the preference lists of men, we have $\mathbf{Pr}[E_i] = (\mathbf{Pr}[F_i])^n$, and $\mathbf{Pr}[E_i \wedge E_j] = (\mathbf{Pr}[F_i \wedge F_j])^n$. Therefore, it is enough to show that for every i and j, $\mathbf{Pr}[F_i | F_j] \leq \mathbf{Pr}[F_i]$.

Let M be an arbitrarily large constant. The following process is one way to simulate the selection of one preference list $L = (l_1, \ldots, l_k)$: Consider the multiset Σ consisting of $\lfloor p_i M \rfloor$ copies of the name of woman i for each i. Pick a random permutation π of Σ. Let l_i be the i'th distinct name in π. It is not hard to see that as $M \to \infty$, the probability of a given list L in this process converges to its probability under distribution $\mathscr{D}^k$. Therefore, $\mathbf{Pr}[F_i]$ is equal to the limit as $M \to \infty$ of the probability that k distinct names occur before i in π. Similarly, if Σ' denotes the multiset obtained by removing all copies of woman j from Σ, then $\mathbf{Pr}[F_i | F_j]$ is equal to the limit as $M \to \infty$ of the probability that k distinct names occur before i in a random permutation of Σ'. However, this is precisely equal to the probability that k distinct names other than j occur before i in a random permutation π of Σ. But that certainly implies that k distinct names (including j) occur before i in π, and so for every π where $F_i | F_j$ happens, F_i also happens. Therefore, $\mathbf{Pr}[F_i | F_j] \leq \mathbf{Pr}[F_i]$. As argued above, this implies that $\mathbf{Pr}[E_i \wedge E_j] \leq \mathbf{Pr}[E_i] \cdot \mathbf{Pr}[E_j]$, and so the variance $\sigma^2(Y_g)$ is

$$\begin{aligned}
\sigma^2(Y_g) &= \mathrm{E}[Y_g^2] - \mathrm{E}[Y_g]^2 \\
&= \sum_{i=1}^{g} \mathbf{Pr}[E_i] + 2\sum_{1\leq i<j\leq g} \mathbf{Pr}[E_i \wedge E_j] - \sum_{i=1}^{g} \mathbf{Pr}[E_i]^2 - 2\sum_{1\leq i<j\leq g} \mathbf{Pr}[E_i]\cdot\mathbf{Pr}[E_j] \\
&\leq \sum_{i=1}^{g} \mathbf{Pr}[E_i] \\
&= \mathrm{E}[Y_g]
\end{aligned}$$

as required. □

Using the above three lemmas and the Chebyshev inequality (see the book by Alon and Spencer [3] for a discussion of this and related inequalities), we can easily conclude the statement of Lemma 5.3.1.

Proof of Lemma 5.3.1. Let q be the probability that $Y_g < \mathrm{E}[Y_g]/2$. By the Chebyshev inequality and Lemma 5.3.4,

$$\begin{aligned}
q &\leq \mathbf{Pr}\left[\left|Y_g - \mathrm{E}[Y_g]\right| > \mathrm{E}[Y_g]/2\right] \\
&\leq \frac{\sigma^2(Y_g)}{(\mathrm{E}[Y_g]/2)^2} \\
&\leq \frac{4}{\mathrm{E}[Y_g]}.
\end{aligned}$$

Thus, by Lemma 5.3.2 and the fact that $1/(Y_g+1)$ is always at most one, we have

$$\begin{aligned}
\mathrm{E}\left[\frac{1}{X_\mu(g)+1}\right] &\leq \mathrm{E}\left[\frac{1}{Y_g+1}\right] \\
&\leq (1-q)\frac{1}{\mathrm{E}[Y_g]/2+1} + q \\
&\leq \frac{6}{\mathrm{E}[Y_g]},
\end{aligned}$$

which together with Lemma 5.3.3 completes the proof. □

In this section, we analyzed the expected number of agents that remain single in a stable marriage mechanism, and used this lemma to prove our main result. Analyzing the expected number of singles in a probabilistic setting is of independent interest, and in Appendix B, we present a tighter analysis of the expected number of singles when men have random preference lists of size k and women have *random* complete preference lists. If, in addition to the results of this appendix, one could prove that the number of singles is concentrated around its expectation, then the bound for the setting in Conjecture 5.1.1 (proven to be $(e^k + k^2)/n$ in this chapter) would be improved.

5.4 Tighter analysis for the uniform distribution

For the case of uniform distributions (the setting in Theorem 5.1.6), it is possible to derive a much tighter bound on the expected number of women with more than one stable husband.

Recall that in the proof of Theorem 5.1.5, we bounded the probability that a fixed woman g is single by $\mathrm{E}_\mu[1/(X_\mu(g)+1)]$, where $X_\mu(g)$ is the number of women at least as popular as g that are single in matching μ. In the case of the uniform distribution, for every woman g, $X_\mu(g)$ is equal to the number of singles in μ. Therefore, if we define the random variable X as the number of women who remain unmarried in the men-optimal stable matching (recall that by Theorem 5.1.3, the set of unmarried women is independent of the choice of the stable marriage algorithm), then we have

$$c_k(n) \leq n\mathrm{E}\left[\frac{1}{X+1}\right].$$

Thus, the following lemma shows that if men have random preference lists of size k, then the expected number of women who have more than one stable partner is at most $e^{k+1} + k^2$. This completes the proof of Theorem 5.1.6.

Lemma 5.4.1 *Let X denote the random variable defined above. Then,*

$$\mathrm{E}\left[\frac{1}{X+1}\right] \leq \frac{e^{k+1}+k^2}{n}.$$

The proof of the above lemma is based on a connection between the stable marriage problem and the classical *occupancy problem.* In the occupancy problem, m balls are distributed amongst n bins. The distribution of the number of balls that end up in each bin has been studied extensively from the perspective of probability theory [67]. We denote the occupancy problem with m balls and n bins by the (m, n)-occupancy problem. The following lemma establishes the connection between the number of singles in the stable marriage game and the number of empty bins in the occupancy problem.

We use the techniques of *amnesia*, the *principle of deferred decisions*, and the *principle of negligible perturbations* used by Knuth [69] and Knuth, Motwani, and Pittel [70, 71]. These techniques allow us to show that our algorithm is almost equivalent to the following random experiment: every man names exactly $k+1$ (not necessarily different) women. Thus, there are $(k+1)n$ proposals which we will think of as balls. There are n women which we will think of as bins. The number of women who are not named in this experiment, denoted by X', is closely related to the number of singles, X, in the algorithm.

Lemma 5.4.2 *Let $Y_{m,n}$ denote the number of empty bins in the (m, n)-occupancy problem and X denote the random variable in Lemma 5.4.1. Then,*

$$\mathrm{E}\left[\frac{1}{X+1}\right] \leq \mathrm{E}\left[\frac{1}{Y_{(k+1)n,n}+1}\right] + \frac{k^2}{n}.$$

Proof. Assume every woman has an arbitrary ordering of all men. We define the following five random experiments:

- Experiment 1 is the experiment defined before Lemma 5.4.1: every man chooses a random list of k different women as his preference list. Then, we run the men-proposing stable marriage algorithm, and let the random variable $X_1 = X$ indicates the number of single women at the end of this experiment. Notice that in this experiment, as in Section 5.3.2, men do not have to select their entire preference list before running the algorithm. Instead, every time a man has to propose to the next woman on his list, he chooses a random woman among the women he has not proposed to so far, and proposes to that woman. It is clear that this does not change the experiment.

- In Experiment 2, each man names k different women at random. We let X_2 be the number of women that no man names in this game.

- Experiment 3 is the same as experiment 2, except here the men are *amnesiacs*. That is, every time a man wants to name a woman, he picks a woman at random from the set of *all* women. Therefore, there is a chance that he names a woman that he has already named. However, each man stops as soon as he names k different women. Let X_3 be the number of women who are not named in this process.

- In Experiment 4, we restrict every man to name at most $k+1$ women. Therefore, each man stops as soon as either he names k different women, or when he names $k+1$ women in total (counting repetitions). Let X_4 denote the number of women who are not named in this experiment.

- In Experiment 5 every man names exactly $k+1$ (not necessarily different) women. The number of women who are not named in this experiment is denoted by X_5. Clearly, $X_5 = Y_{(k+1)n,n}$.

Now, we show how the random variables X_1 through X_5 are related. It is easy to see that for any set of men's preference lists, the number of unmarried women in Experiment 1 is at least the number of women who are not named in Experiment 2. Therefore, $X_1 \geq X_2$. Also, it is clear from the description of Experiments 2 and 3 that $X_2 = X_3$.

In order to relate X_3 and X_4, we use the *principle of negligible perturbations*. Experiments 4 is essentially the same as Experiment 3, except in X_4 we only count women who are not named by any man as one of his first $k+1$ choices. Let E denote the event that no man names more than $k+1$ women in Experiment 3. We first show that $\mathbf{Pr}[\bar{E}] < k^2/n$. Fix a man, say Homer. We want to bound the probability that Homer names at least $k+2$ women before the number of different women he has named reaches k. By the union bound, this probability is at most the sum, over all pairs $\{i, j\} \subset \{1, \ldots, k+2\}$ that the i'th and j'th proposal of Homer are redundant.

It is easy to see that for any such pair, this probability is at most $1/n^2$. Therefore, the probability that Homer makes more than $k+1$ proposals is at most $\binom{k+2}{2}/n^2 < k^2/n^2$. Thus, by the union bound, the probability of this happens for at least one man is less than k^2/n. That is, $\mathbf{Pr}[\bar{E}] < k^2/n$. Now, notice that by the definition of X_3 and X_4, the random variables X_3 and X_4 are equal when conditioned on the occurrence of E. Therefore, $\mathrm{E}\left[\frac{1}{X_3+1}|E\right] = \mathrm{E}\left[\frac{1}{X_4+1}|E\right]$. Let $C = \left|\mathrm{E}\left[\frac{1}{X_3+1}\right] - \mathrm{E}\left[\frac{1}{X_4+1}\right]\right|$ be the unconditioned difference of these expectations. Then, letting $q = \mathbf{Pr}[E]$ and $\bar{q} = \mathbf{Pr}[\bar{E}]$,

$$
\begin{aligned}
C &= \left|q\mathrm{E}\left[\frac{1}{X_3+1}|E\right] + \bar{q}\mathrm{E}\left[\frac{1}{X_3+1}|\bar{E}\right] - q\mathrm{E}\left[\frac{1}{X_4+1}|E\right] - \bar{q}\mathrm{E}\left[\frac{1}{X_4+1}|\bar{E}\right]\right| \\
&= \bar{q}\left|\mathrm{E}\left[\frac{1}{X_3+1}|\bar{E}\right] - \mathrm{E}\left[\frac{1}{X_4+1}|\bar{E}\right]\right| \\
&\leq \bar{q} \\
&< \frac{k^2}{n}.
\end{aligned}
$$

Finally, we observe that by the definition of Experiments 4 and 5, we have $X_4 \geq X_5$. The above observations imply

$$
\begin{aligned}
\mathrm{E}\left[\frac{1}{X+1}\right] &\leq \mathrm{E}\left[\frac{1}{X_2+1}\right] \\
&= \mathrm{E}\left[\frac{1}{X_3+1}\right] \\
&\leq \mathrm{E}\left[\frac{1}{X_4+1}\right] + \frac{k^2}{n} \\
&\leq \mathrm{E}\left[\frac{1}{Y_{(k+1)n,n}+1}\right] + \frac{k^2}{n}.
\end{aligned}
$$

This completes the proof of the lemma. □

By the above lemma, the only thing we need to do is to analyze the expected value of $1/(Y_{m,n}+1)$ in the occupancy problem. We do this by writing the expected value of $1/(Y_{m,n}+1)$ as a summation and bounding this summation by comparing it term-by-term to another summation whose value is known.

Lemma 5.4.3 *Let $Y_{m,n}$ denote the number of empty bins in the (m,n)-occupancy* *problem. Then,*

$$\mathrm{E}\left[\frac{1}{Y_{m,n}+1}\right] \le \frac{e^{m/n}}{n}.$$

Proof. Let $P_r(m,n)$ be the probability that exactly r bins are empty in the (m,n)-occupancy problem. Then $P_0(m,n)$, the probability of no empty bin, can be written as the following summation by the principle of inclusion-exclusion.[5]

$$P_0(m,n) = \sum_{i=0}^{n}(-1)^i \binom{n}{i}\left(1-\frac{i}{n}\right)^m \tag{5.4}$$

The probability $P_r(m,n)$ of exactly r empty bins can be written in terms of the probability of no empty bin in the $(m,n-r)$-occupancy problem:

$$P_r(m,n) = \binom{n}{r}\left(1-\frac{r}{n}\right)^m P_0(m,n-r). \tag{5.5}$$

By equations 5.4 and 5.5,

$$P_r(m,n) = \sum_{i=0}^{n-r}(-1)^i \binom{n}{r,i}\left(1-\frac{r+i}{n}\right)^m, \tag{5.6}$$

where $\binom{n}{a,b}$ denotes the multinomial coefficient $\frac{n!}{a!\,b!\,(n-a-b)!}$. Using equation 5.6 and the definition of expected value we have,

$$\begin{aligned}
\mathrm{E}[\frac{1}{Y_{m,n}+1}] &= \sum_{r=0}^{n}\frac{1}{r+1}P_r(m,n) \qquad (5.7)\\
&= \sum_{r=0}^{n}\sum_{i=0}^{n-r}\frac{(-1)^i}{r+1}\binom{n}{r,i}\left(1-\frac{r+i}{n}\right)^m\\
&= \sum_{r=0}^{n}\sum_{i=0}^{n-r}\frac{(-1)^i}{n+1}\binom{n+1}{r+1,i}\left(1-\frac{r+i}{n}\right)^m\\
&= \sum_{r=1}^{n+1}\sum_{i=0}^{n+1-r}\frac{(-1)^i}{n+1}\binom{n+1}{r,i}\left(1-\frac{r+i-1}{n}\right)^m.
\end{aligned}$$

It is probably impossible to simplify the above summation as a closed-form formula.

[5] This can also be derived by dividing a well-known formula for Stirling numbers of the second kind (see, for example, [51, 114]) by n^m.

Therefore, we use the following trick: we consider another summation S with the same number of terms, and bound the ratio between the corresponding terms in these two summations. This gives us a bound on the ratio of the summation in equation 5.7 to the summation S. The value of S can be bounded easily using a combinatorial argument.

Consider the $(m, n+1)$-occupancy problem. The probability that at least one bin is empty is the sum, over $r = 1, \ldots, n+1$, of $P_r(m, n+1)$. We denote this probability by S. By equation 5.6 we have

$$S = \sum_{r=1}^{n+1} \sum_{i=0}^{n+1-r} (-1)^i \binom{n+1}{r, i} \left(1 - \frac{r+i}{n+1}\right)^m \leq 1,$$

where the inequality follows from the fact that S is the probability of an event. The summation in equation 5.7 and S have the same number of terms, and the ratio of each term in the summation in equation 5.7 to the corresponding term in S is equal to

$$\frac{(1 - \frac{r+i-1}{n})^m}{(n+1)(1 - \frac{r+i}{n+1})^m} = \frac{(\frac{n-r-i+1}{n})^m}{(n+1)(\frac{n+1-r-i}{n+1})^m} = \frac{(1 + \frac{1}{n})^m}{n+1}.$$

Therefore,

$$\mathrm{E}\left[\frac{1}{Y_{m,n}+1}\right] = \frac{1}{n+1}\left(1 + \frac{1}{n}\right)^m S < \frac{e^{m/n}}{n},$$

as desired. □

Lemma 5.4.1 immediately follows from Lemmas 5.4.2 and 5.4.3.

Chapter 6

Conclusion

In this dissertation, we studied mechanism design for various combinatorial optimization problems in the presence of strategic agents. We considered four important settings — limited demand and limited budget multi-unit auctions, procurement auctions, cost-sharing auctions, and two-sided markets. In each setting, we proposed and/or studied mechanisms to solve the allocation problem with respect to a particular equilibrium concept.

In the first three settings, our goal was to derive revenue-maximizing (or payment-minimizing) auctions. For multi-unit auctions (Chapter 2), we studied implementation in dominant strategies, following a competitive auction framework. In the limited demand setting, our auctions are the first provably competitive *deterministic* auctions. In the limited budget setting, we give the first competitive auction, although it is randomized. This chapter also contains a powerful auction derandomization technique that, given any truthful multi-unit auction for limited demands, computes a deterministic one with approximately the same revenue guarantee by using asymmetry.

For procurement auctions (Chapter 3), we studied implementation in variants of Nash equilibria. We proved that the strong ϵ-Nash equilibria of first-price flow auctions for publicly known demands are approximately efficient and produce payments which, in the limit as $\epsilon \rightarrow 0$, are at most (and sometimes much less than) those of the VCG mechanism. We then provided an implementation in ϵ-Nash equilibria with approximately the same payment properties. In the unknown-demand model, we designed an auction whose ϵ-Nash equilibria had expected payments similar to that of the known-demand model. However, we were unable to prove that this auction can be implemented in polynomial time.

For cost-sharing auctions (Chapter 4), we derived a general technique to bound the revenue properties of cross-monotonic cost-sharing schemes and used this technique to derive tight or nearly tight bounds for many combinatorial optimization games. We then explored the implications of our results on group-strategyproof mechanisms. We derived a partial characterization of group strategyproof mechanisms in terms of semi-cross-monotonic cost-sharing schemes. By imposing certain additional assumptions, we were able to prove a complete characterization in terms of cross-monotonic cost-sharing schemes. For mechanisms satisfying these additional assumptions, our bounds on the budget-balance factors of cross-monotonic cost-sharing schemes indicate that group-strategyproofness is incompatible with revenue goals for many combinatorial optimization games.

In the last setting — that of two-sided markets (Chapter 5) — we studied the stable marriage game in a probabilistic setting and showed that the expected fraction of singles tends to zero as the size of the market grows. In doing so, we answered a question asked by Roth and Peranson [99], and generalized their model to one where women have arbitrary preferences and each man independently picks each women on his preference list from an arbitrary fixed distribution. As discussed, this result has a number of economic implications which indicate that dishonesty almost surely does not benefit a player.

We conclude with some open questions related to each setting. Regarding limited-demand multi-unit auctions, there are many interesting open questions surrounding the issue of randomness (besides the obvious question of finding a polynomial-time derandomization technique).

1. **Revenue Guarantees in Mass Markets.** One issue with the deterministic auctions that we present is that their revenue is a constant factor away from the optimal revenue even for markets in which the number of winners is large. In such situations, the revenue of known randomized auctions asymptotically

approaches the optimal revenue [50]. Is it possible to design deterministic auctions that match this guarantee? More generally, in what sense are randomized auctions provably more powerful than deterministic ones?

2. **Derandomization in Multi-Parameter Settings.** Another promising direction for future research is an extension of the derandomization techniques presented in this chapter. Can these techniques be extended to settings with multi-parameter bids like the limited budgets setting of Section 2.3? What if the demand limit is a private value?

3. **Derandomization with Feasibility Constraints.** Alternatively, one could try to derandomize more general single-parameter auction settings with feasibility constraints. For example, consider a related machine scheduling auction, such as the setting studied by Archer and Tardos [5], where machines bid a processing speed, and further suppose that the set of feasible allocations is restricted (for example, perhaps jobs have precedence constraints). In settings like these, even if the randomized auction observes the feasibility constraints, the derandomized auction derived from the flow graph may violate these constraints as the set of allocations output by the derandomized auction is not necessarily contained in the set of allocations output by the randomized one. Perhaps the flow graph or the selection of the integral flow can be modified to accommodate these constraints.

4. **Hat Puzzles.** Finally, we have left unsolved many intriguing hat puzzles. Can the everywhere balanced k-coloring hat puzzle be solved in polynomial time for $k > 2$? A positive answer to this question might lead to a polynomial-time derandomization technique for auctions as well. Given the solution for $k = 2$ (that is, the deterministic coin flipping algorithm), one natural approach to this question is to try and find a "perfectly alternating" deterministic assignment of hat colors such that, when players are sorted according to their hat colors, their guesses cycle through the list of possible colors (for example, Red, Green, Blue, Red, Green, Blue, Red, Green, Blue).

For the limited budget setting, our results and those in the literature are still quite preliminary. In particular, the assumptions we make in deriving our impossibility result are somewhat questionable. Both the strong non-bundling and independence of irrelevant alternatives (IIA) assumptions are difficult to motivate in practice. Does a similar impossibility result hold under weaker assumptions?

5. **Ad Auction Design.** Internet ad auction design was a main motivation behind our study of limited-budget multi-unit auctions, but budget constraints are just one of a myriad of issues related to ad auction design. Other important considerations in this market include the matching algorithm (which advertisers are interested in a particular web surfer?), the ranking algorithm (who should be displayed and where?), click fraud (how can we discourage advertisers from clicking on their competitors' advertisements?), and reserve prices (how should reserve prices be set for different keywords and how do they impact the revenue?). Our results do not even fully address the issue of budget constraints in ad auctions, for ad auctions are combinatorial. How should an advertiser's budget be allocated across multiple keywords?

In the chapter on procurement auctions, we showed that first-price auctions entail potentially lower payments than VCG mechanisms. However, they suffer from one major drawback, in that the solution concept (strong ϵ-Nash equilibrium) requires agents to know all costs, and coordinate on the choice of an equilibrium. This is much more demanding than the dominant-strategy solution concept, and could lead to inefficiency and high payments in practice. Thus, the auction models analyzed here are not completely satisfying, as there is no mechanism prescribed for the agents bids to reach equilibrium. This is true even for the weaker concept of ϵ-Nash equilibrium.

6. **Convergence to Equilibria.** A promising direction for future research is to find bargaining mechanisms that enable the bidders to converge to an equilibrium. When the edges all know each others' costs, an n-party bargaining protocol, such as the one in the Krishna and Serrano [75], could be used. When there is uncertainty, the situation is more complex. Such a mechanism may

be subsidized; for example, the links may be given an additional payment that decays with time, to incentivize them to quickly reach an agreement. As long as the subsidy is smaller than the VCG premium, it may be a better alternative. See the book by Fudenberg and Levine [38] for a discussion of convergence problems in general settings.

In the setting of cost-sharing auctions, there are two main directions for future research — the study of cost-sharing schemes and the study of group-strategyproof auctions.

7. **Budget-Balance Factor of Other Combinatorial Optimization Games.** In this chapter, we presented a technique for proving bounds on the budget-balance factor of cross-monotonic cost-sharing schemes for a variety of combinatorial optimization games. Our technique was quite general and may prove applicable to a variety of other combinatorial games. For example, the facility location game restricted to a tree always has a budget-balanced cost allocation in the core [45], but we do not have a tight lower and upper bound on the budget-balance factor of the best cross-monotonic cost sharing schemes for this game. For the facility location game on the line, we have an upper bound of $\frac{6}{7}$.

8. **Characterization of Group-Strategyproof Mechanisms.** Another significant open question is to fully characterize cost-sharing schemes that can arise as $\xi_{\mathscr{M}}$ for some group-strategyproof mechanism $\mathscr{M}$. Our characterization in terms of cross-monotonic cost-sharing schemes imposed a technical condition on group-strategyproof mechanisms which we called upper-continuity. Is this essential to this equivalence result? Are there other well-motivated axioms that can be imposed on group-strategyproof mechanisms which would then imply the equivalence result?

9. **Budget-Balance versus Efficiency.** Moulin and Shenker [87] investigate the tradeoff between exact budget-balance and efficiency in group-strategyproof mechanisms. It would be interesting to extend their results by exploring the

possible budget-balance factor of group-strategyproof mechanisms that are in some sense close to efficient.

In the chapter on centralized two-sided markets, our main motivation was the National Residency Matching Program (NRMP). We studied the incentive issues facing participants in stable matching mechanisms. However, there are many other economic considerations surrounding the NRMP.

10. **When Doctors Get Married.** Since its conception, the NRMP market was redesigned to accommodate *couples* among students who want to live in the same city. We studied this problem in Appendix C when couples want to work at the same hospital. However, in general, such couples can submit a joint preference list of *pairs* of hospitals, and the algorithm has to match them to one of the pairs in their list. With this extra twist, there are instances for which no stable matching exists. However, so far every year the NRMP algorithm has been able to find a stable matching. A theoretical justification for this (in a reasonable probabilistic model), and a study of incentive properties in mechanisms with couples are interesting open directions for future research.

11. **When Hospitals Lie.** In this chapter, we proved that individual participants probably can not benefit by altering their preference lists. However, hospitals have another possible strategic manipulation; namely, they can withhold demand. Is there some way to quantify how much a hospitable can benefit from this manipulation? An answer to this question involves first proposing a model for comparing allocations of varying size.

12. **Salary Considerations in the NRMP.** The stable marriage mechanism, as we have presented it, does not have any provision for salary negotiations. In fact, salaries are announced by the hospital together with the position before the preferences are formed. In a recent lawsuit, medical students argue that this has the effect of decreasing their wages below the levels they could achieve in the absence of the match. It would be very interesting to provide a theoretical model

and analysis of salary considerations in stable matching markets and perhaps present mechanisms which admit salary negotiations. For more information, see the discussion paper by Bulow and Levin [15].

The techniques we developed to address incentive issues in the NRMP might prove useful in the study of matching markets in the presence other rules. Some matching markets restrict participants to submit complete preference lists and/or allow indifferences in the preference lists, as is the case in the New York City public school matching system [1]. One could also consider incentives in one-sided centralized markets like college dormitory assignments or the kidney exchange market [100].

Appendix A

Group-Strategyproof Mechanisms and Cost-Sharing Schemes

In this appendix, we show that not all group-strategyproof mechanisms are cross-monotonic and not all semi-cross-monotonic cost-sharing schemes give rise to group strategyproof mechanisms.

A.1 A group-strategyproof mechanism with no cross-monotonic cost-sharing scheme

As the following example shows, for some cost functions, group-strategyproof mechanisms do not correspond to cross-monotonic cost-sharing schemes.

Example A.1.1 *Suppose there are three agents,* 1, 2, *and* 3, *with a cost function given by*

$$C(S) = \begin{cases} 2 & \text{if } |S| = 3, \\ 1 & \text{otherwise.} \end{cases}$$

We consider the following mechanism for this cost function:

Mechanism $\mathscr{M}$:

If $b_1 \geq 1$ *then*

If $\min(b_2, b_3) > \frac{1}{2}$ *then* $Q = \{1, 2, 3\}$ *and* $x = (1, \frac{1}{2}, \frac{1}{2})$,

else if $\max(b_2, b_3) < \frac{1}{2}$ *then* $Q = \{1\}$ *and* $x = (1, 0, 0)$,

else if $b_2 \geq b_3$ *then* $Q = \{1, 2\}$ *and* $x = (\frac{1}{2}, \frac{1}{2}, 0)$,

else if $b_3 > b_2$ *then* $Q = \{1, 3\}$ *and* $x = (\frac{1}{2}, 0, \frac{1}{2})$,

else if $\frac{1}{2} \leq b_1 < 1$ *then*

If $\min(b_2, b_3) > \frac{1}{2}$ *then* $Q = \{2, 3\}$ *and* $x = (0, \frac{1}{2}, \frac{1}{2})$,

else if $\max(b_2, b_3) < \frac{1}{2}$ *then* $Q = \emptyset$ *and* $x = (0, 0, 0)$,

else if $b_2 \geq b_3$ *then* $Q = \{1, 2\}$ *and* $x = (\frac{1}{2}, \frac{1}{2}, 0)$,

else if $b_3 > b_2$ *then* $Q = \{1, 3\}$ *and* $x = (\frac{1}{2}, 0, \frac{1}{2})$,

else if $b_1 < \frac{1}{2}$ *then*

If $\min(b_2, b_3) \geq \frac{1}{2}$ *then* $Q = \{2, 3\}$ *and* $x = (0, \frac{1}{2}, \frac{1}{2})$,

else if $b_2 \geq 1$ *then* $Q = \{2\}$ *and* $x = (0, 1, 0)$,

else if $b_3 \geq 1$ *then* $Q = \{3\}$ *and* $x = (0, 0, 1)$,

else $Q = \emptyset$ *and* $x = (0, 0, 0)$.

The cost-sharing scheme $\xi_{\mathscr{M}}$ is not cross-monotonic since, for example, the cost-share $\xi_{\mathscr{M}}(1, \{1, 2, 3\})$ of agent 1 in the set $\{1, 2, 3\}$ is strictly greater than his cost-share $\xi_{\mathscr{M}}(1, \{1, 2\})$ in the subset $\{1, 2\}$. In fact, it is not hard to see that no cross-monotonic cost-sharing scheme for C exists. Still, as the following lemma shows, the mechanism $\mathscr{M}$ is group-strategyproof.

Proposition A.1.1 *The mechanism $\mathscr{M}$ in Example A.1.1 is group-strategyproof.*

Proof. Let u_i denote the true utility of i for receiving the service, b_i denote his bid, and $x_i(\mathbf{b})$ denote his payment when the bids are $\mathbf{b}$. Note $x_i(\mathbf{b}) = 0$ if and only if i does not receive the service.

We first prove by contradiction that any successful coalition must include 1. Suppose not (that is, $b_1 = u_1$). First consider the case $u_1 \geq \frac{1}{2}$. Note that for $i \in \{2, 3\}$, whenever i receives the service, he pays $\frac{1}{2}$. Therefore, i can benefit only if $u_i > \frac{1}{2}$ and he is not receiving service. However, in any input bid vector with $b_1 \geq \frac{1}{2}$, $b_i > \frac{1}{2}$ implies that i receives the service, so i can not benefit in any coalition. Next suppose $u_1 < \frac{1}{2}$. Consider the cross-monotonic cost-sharing scheme $\xi : \{2,3\} \to \mathbb{R}_+$ where for $i \in \{2,3\}$, $\xi(i, \{2,3\}) = \frac{1}{2}$ and $\xi(i, \{i\}) = 1$. The Moulin mechanism $\mathscr{M}_\xi$ is equivalent to $\mathscr{M}$ when $u_1 < \frac{1}{2}$ and so Theorem 4.3.1 implies that there is no subset of $\{2,3\}$ can form a successful coalition in this case.

Now consider any coalition including 1. Suppose $u_1 < \frac{1}{2}$. If $b_1 < \frac{1}{2}$, then the outcome does not change if we set $b_1 = u_1$. Thus, we only need to consider coalitions in which $b_1 \geq \frac{1}{2}$. As $u_1 < \frac{1}{2}$ and the minimum non-zero price of 1 is $\frac{1}{2}$, it must be that $1 \notin Q(\mathbf{b})$ even though $b_1 \geq \frac{1}{2}$. This happens only when $\frac{1}{2} \leq b_1 < 1$ and $\max(b_2, b_3) < \frac{1}{2}$ or $\min(b_2, b_3) > \frac{1}{2}$. In the first case, as no agent receives service, all utilities are zero and so no one can benefit. In the second case, for $i \in \{2,3\}$, the payment of i is $\frac{1}{2}$. Therefore, if i is in the coalition, it must be that $u_i \geq \frac{1}{2}$. If i is not in the coalition, then $u_i = b_i > \frac{1}{2}$ by assumption. Thus $\min(u_2, u_3) \geq \frac{1}{2}$. But then $x(\mathbf{b}) = x(\mathbf{u})$ and so no agent's utility for the outcome changes.

Next, suppose $u_1 \geq \frac{1}{2}$. For $i \in \{2,3\}$, in the truthful scenario i pays at most $\frac{1}{2}$. As i's payment is always at least $\frac{1}{2}$, i can not benefit from a decrease in price. Therefore i can benefit only if $u_i > \frac{1}{2}$ and $i \notin Q(\mathbf{u})$. But this is impossible for any vector with $u_1 \geq \frac{1}{2}$, so i can not benefit in any coalition. Therefore, 1 must be the agent that benefits from the coalition. As the minimum price for 1 is $\frac{1}{2}$, in order for 1 to benefit, it must be that $u_1 > \frac{1}{2}$ but either $1 \notin Q(\mathbf{u})$ or $x_1(\mathbf{u}) = 1$. This means that either $\min(u_2, u_3) > \frac{1}{2}$ (case one) or $\max(u_2, u_3) < \frac{1}{2}$ (case two). Furthermore, 1 can only benefit if $x_1(\mathbf{b}) = \frac{1}{2}$ since, when $u_1 \geq 1$, 1 is receiving the service at price 1 and so the price must decrease, and when $\frac{1}{2} \leq u_1 < 1$, 1 is not receiving the service but can not afford to pay 1 and so must receive the service at price $\frac{1}{2}$. Now, in case one ($\min(u_2, u_3) > \frac{1}{2}$), in the truthful scenario 2 and 3 have positive utility. In order for $x_1(\mathbf{b}) = \frac{1}{2}$, i for $i = 2$ or $i = 3$ must lower his bid to $b_i \leq \frac{1}{2}$. But then if the coalition consists of just i and 1, $i \notin Q(\mathbf{b})$ and so i's utility decreases. Similarly, if the coalition is $\{1,2,3\}$, then 1 only benefits if $\{2,3\} \not\subset Q(\mathbf{b})$ and so the utility of 2 or 3 decreases. In case two ($\max(u_2, u_3) < \frac{1}{2}$), 1 can only benefit if i for $i = 2$ or $i = 3$ raises his bid to $b_i \geq \frac{1}{2}$. But then if the coalition consists of just i and 1, $x_i(\mathbf{b}) = \frac{1}{2}$ and so i's utility becomes negative. Similarly, if the coalition is $\{1, 2, 3\}$, then at least one of 2 or 3 must pay $\frac{1}{2}$, and so his utility becomes negative. $\square$

A.2 A semi-cross-monotonic cost-sharing scheme with no group-strategyproof mechanism

Suppose there are just two agents, 1 and 2. The cost of servicing both agents is 6 while the cost of servicing either agent individually is 1. The following is a budget-balanced semi-cross-monotonic cost-sharing scheme:

$$\xi(1, \{1, 2\}) = \xi(2, \{1, 2\}) = 3, \quad \xi(1, \{1\}) = \xi(2, \{2\}) = 1$$

However, this scheme can not correspond to the payments in any group-strategyproof mechanism. First consider the bid vector $\mathbf{b}^1 = (3, 3)$. By group-strategyproofness, the mechanism must service exactly one of the agents; otherwise they could collude and bid either $(-1, 2)$ or $(2, -1)$. Without loss of generality, suppose it services agent 2. Now consider the bid vector $\mathbf{b}^2 = (3, 2)$. Again, the mechanism must service agent 2 since otherwise he could bid 3 and get the service at price 1. Finally, consider the bid vector $\mathbf{b}^3 = (b_1^*, 2)$, where b_1^* is as in the proof of Theorem 4.3.2. Now the mechanism must service just agent 1 at price 1. But this implies that in bid vector $\mathbf{b}^2$, agent 1 could have profitably deviated by bidding b_1^*.

Remark A.2.1 *Notice that in this cost-sharing scheme, removing either agent from the set $\{1, 2\}$ decreased the cost share of the other agent. This property allowed us to draw conclusions about the serviced set in bid vector $\mathbf{b}^1$ which led us to our contradiction. This highlights the following general fact: if two agents i and j are both negative in a set S, then either $\xi(i, S \setminus \{j\}) = \xi(i, S)$ or $\xi(j, S \setminus \{i\}) = \xi(j, S)$ (or both).*

Appendix B

Expected Number of Singles in Two-Sided Markets

In this appendix, we analyze the expected number of singles in a stable matching when men have random preference lists of size k and women have random complete preference lists.

Lemma B.0.1 *Consider a collection of n men and n women, each man having a random ordering of k random women, and each woman having a random ordering of men. Let $p_k(n)$ denote the probability that in a stable matching with respect to these preference lists a fixed man remains single. Then for $k \geq 2$, $p_k(n) \geq \frac{1}{k2^{k+2}}(1 - o(1))$.*

In order to prove the above lemma, we first generalize the scenario to a case where there are m men and n women ($m \leq n$). Let $p_k(m, n)$ denote the probability that a fixed man remains unmarried in this scenario. Therefore, $p_k(n) = p_k(n, n)$. We start by proving that if the population of women remains constant, an increase in the number of men can only make it harder for a man to find a stable wife.

Lemma B.0.2 *For every k, n, m_1, and m_2, if $m_1 \leq m_2$ then $p_k(m_1, n) \leq p_k(m_2, n)$.*

Proof. It is sufficient to prove that for every k, n, and m, $p_k(m, n) \leq p_k(m + 1, n)$. Consider a fixed man, Cain, in the scenario where there are $m + 1$ men. We want to compute the probability that after running the men-proposing algorithm, Cain remains single. By Theorem 5.1.2 we know that the order of proposals does not affect the outcome of the algorithm. Therefore, we can assume that one of the $m+1$ men, say Abel, starts proposing to women only after everyone else is done with his or her proposals. By definition, before Abel starts proposing, the probability that Cain is single is precisely $p_k(m,n)$. If Cain is married at this point, then there is a chance he becomes single after Abel starts proposing, since his wife might leave him. However, if he is single before Abel starts proposing, there is no chance that he gets married. Therefore, the probability that Cain remains single is at least $p_k(m,n)$. $\square$

Proof of Lemma B.0.1. Let $c < 1$ be a constant that will be fixed later. By Lemma B.0.2, we have $p_k(n) = p_k(n,n) \geq p_k(\lceil cn \rceil + 1, n)$, so it is enough to prove that $p_k(\lceil cn \rceil + 1, n) \geq \frac{1}{ek2^k}$. The proof of this is based on the following inequalities.

$$p_k(\lceil cn \rceil + 1, n) \geq \left(\frac{c}{2}\left(1 - p_k(\lceil cn \rceil, n)\right)\right)^k \tag{B.1}$$

$$p_k(\lceil cn \rceil, n) \leq c^k \tag{B.2}$$

We start by proving inequality B.2. Consider the situation where there are $\lceil cn \rceil$ men and n women. Fix a man, say Abel. The probability that Abel remains single is $p_k(\lceil cn \rceil, n)$. Now, consider the men-proposing algorithm. Since the order of proposals does not change the outcome, we can assume that Abel will wait until everyone else stops proposing, and then he will make his first proposal. Suppose there are s single women at this point and let S denote the set of single women. At this moment, there are at most $\lceil cn \rceil - 1 < cn$ women who are married, so $s \geq (1-c)n$. Since Abel's list consists of k randomly chosen women, the probability that his i'th choice is not in S given that his first $(i-1)$ choices are not in S is $\frac{n-s}{n-i+1}$. Therefore, the probability that none of his choices are in S is at most $\prod_{i=1}^{k}\left(\frac{n-s}{n-i+1}\right) \leq \left(\frac{n-s}{n}\right)^k \leq c^k$. We claim that if at least one of the women in Abel's list is in S then Abel will find a wife. The reason is that every time Abel makes a proposal, if he proposes to a single woman, the proposal will be accepted and the algorithm ends. But if he proposes to a married woman, he might start a chain of proposals that will either end at a single woman, in which case Abel ends up married, or gets back to Abel, in which case the set of single women does not change and we can repeat the same argument for the next proposal of Abel until he reaches a woman in his list that is in S. By this claim, the probability that Abel remains single is upper bounded by the probability that none of the women in his list are in S, which is at most c^k.

Now, we prove inequality B.1. Consider a situation where there are $\lceil cn \rceil + 1$ men and n women, and fix a man, say Cain. We bound the probability Cain remains single. Consider the men-proposing algorithm, and let everyone other than Cain make proposals. Let M denote the set of married women at this point and s denote its size. Then, let Cain enter and start proposing. The probability that Cain's i'th proposal is to a married woman given that his first $(i-1)$ choices were married is $\frac{s-i+1}{n-i+1}$. A married woman rejects a new proposal with probability at least $1/2$. Therefore, conditioning on the random choices of the other men, the probability Cain faces rejection immediately after each of his proposals and therefore ends up single is at least $\prod_{i=1}^{k} \frac{s-i+1}{2(n-i+1)} \geq \left(\frac{s-k+1}{2(n-k+1)}\right)^k \geq \left(\frac{s-k}{2n}\right)^k$. Removing the conditioning, this probability becomes the expectation $\mathrm{E}[(\frac{s-k}{2n})^k] \geq (\frac{\mathrm{E}[s]}{2n} - \frac{k}{2n})^k$ over the random choices of the other men. The expected size $\mathrm{E}[s]$ of M is the same as the expected number of married men, which, by the definition of $p_k(m,n)$, is $(1-p_k(\lceil cn \rceil, n))\lceil cn \rceil \geq c(1-p_k(\lceil cn \rceil, n))$. Thus the probability that Cain ends up single is at least $\left(\frac{c}{2}(1-p_k(\lceil cn \rceil, n)) - \frac{k}{2n}\right)^k$.

Inequalities B.1 and B.2 imply that $p_k(\lceil cn \rceil + 1, n) \geq \frac{1}{2^k}\left(c(1-c^k) - \frac{k}{n}\right)^k$. Choosing $c = k^{-1/k}$, we see that for $k \geq 2$

$$\begin{aligned}
p_k(\lceil cn \rceil + 1, n) &\geq \frac{1}{2^k}\left(k^{-1/k}(1-\frac{1}{k}) - \frac{k}{n}\right)^k \\
&\geq \frac{1}{4k2^k}\left(1 - \frac{k}{nk^{-1/k}(1-\frac{1}{k})}\right)^k \\
&\geq \frac{1}{k2^{k+2}}\left(1 - \frac{k^2}{nk^{-1/k}(1-\frac{1}{k})}\right) \\
&\geq \frac{1}{k2^{k+2}}\left(1 - o(1)\right),
\end{aligned}$$

as desired. □

Appendix C

Unsplittable Stable Marriage Problems

In this appendix, we study unsplittable stable marriage problems.[1] This problems address situations in which each side of the market has a "size" or "capacity". Instances of unsplittable stable marriage problems occur in situations like the assignment of couples in the National Resident Matching Program (NRMP), where each student submits a ranking over hospitals and each hospital submits a ranking over students as well as a capacity governing the number of students it can accept. Unsplittability constraints arise when pairs of students are married and wish to be similarly assigned (as opposed to the more general setting in which they wish to be assigned to certain pairs of hospitals).

We define the unsplittable stable marriage problem in terms of machine scheduling on unrelated parallel machines where (job, machine) preferences are specified in an ordinal setting, as in a stable marriage problem: rather than specifying costs for every possible (job, machine) assignment, jobs submit ranked preference lists over machines, and machines over jobs. In addition, each job i has a processing time p_{ij} on machine j. For a generalized notion of stability (defined below) similar to that from the classical stable marriage problem, we would like to compute a stable integral assignment in which machines are "congested" beyond the minimum stable fractional

[1]The results of this appendix are based on ongoing joint work with Dean and Goemans.

makespan M by at most the processing time of a single job. That is, we would like an integral assignment in which, for every machine j, the sum of the processing times of jobs i assigned to j is at most $M + \max p_{ij}$.[2]

Generalizing the stability property of one-to-one stable matchings in the natural fashion, we declare a fractional assignment to be *stable* if it admits no *blocking pair* (i, j), where either (i) both job i and machine j prefer each other more than some of their current partners, or (ii) job i prefers machine j to some of its current partners and j is not utilized fully up to the makespan of the schedule. The type (ii) constraints assume that machines have complete preference lists, and this is without loss of generality since we can transform an instance with incomplete preference lists into an equivalent instance with complete preference lists by appropriately introducing "dummy" jobs whose processing times fill the under-capacitated machines with incomplete preference lists.

Given a candidate makespan M, in finite (but not necessarily polynomial) time we can compute a fractional stable assignment (if it exists) using a natural generalization of the classical Gale-Shapley (GS) propose/reject algorithm [39]: jobs propose to machines in order of their preferences, and in each step a non-fully-assigned job i proposes all of its unassigned load to the next machine j on its list, which accepts only as much load as allowed by the makespan constraint and its preference list, possibly rejecting (fractionally) some of the jobs already assigned to it if they are less preferred than the proposing job. Whenever a job is "split" due to a fractional acceptance or rejection, it remains split into two "virtual jobs" for the remainder of the algorithm, each of which carries out independent sequences of proposals. Just as with the classical one-to-one stable marriage problem, one can show that order of proposals and rejections does not matter — we always obtain a "job-optimal" stable assignment in which for every job i simultaneously, the allocation of its load among the machines in i's preference list is lexicographically maximal among all fractional stable allocations.

[2]The nonuniform capacities of the couples problem in the NRMP can be accommodated in our setting by adjusting the p_{ij}'s appropriately. The target makespan M then becomes the normalized capacities as opposed to the minimum stable fractional makespan.

Our goal is to compute an assignment in which each job is integrally assigned. Assuming that a fractional stable solution of makespan M exists, we wish to round it to an unsplit solution satisfying condition (i) for stability, where each machine is "congested" by at most $\max p_{ij}$, and where each uncongested machine satisfies condition (ii). We treat congested machines as stable with respect to condition (ii), even though they may appear to have extra capacity from the perspective of an inflated makespan. This makes sense particularly if we think in terms of machines having specified capacities (as in the NRMP example mentioned above), since a machine congested beyond its capacity will not want any new jobs assigned to it. Henceforth, we refer to an unsplit rounding of a fractional stable solution satisfying stability conditions (i) universally and (ii) for all uncongested machines as a *proper* unsplit rounding.

Our algorithm, which we call the integral GS algorithm as it is a variant of the fractional GS algorithm, computes a proper unsplit stable assignment of makespan $M + \max p_{ij}$ in which every job is assigned to a machine that is at least as good as its best allocation in the fractional assignment. Assume the target makespan M is known.[3] Jobs issue proposals in sequence according to their preference lists. In each step of the algorithm, an arbitrary unassigned job issues a proposal (all proposals and rejections are integral this time) to the next machine on its preference list. The machine accepts, but then proceeds to reject in sequence the least favored jobs assigned to it (possibly including the proposing job) until the machine is overcongested by at most the processing time of a single job — that is, until rejecting the next job would leave the machine being utilized at or below M units of load. If each machine stores its accepted jobs in a heap based on preference list ranking, this variant of the GS algorithm runs in $O(mn \log n)$ time.

Theorem C.0.1 *Suppose that a fractional stable assignment of makespan M exists. Then the integral GS algorithm computes a proper unsplit assignment of makespan*

[3]Our integer GS algorithm does not know the minimum value of M for which a fractional stable assignment of makespan M exists; however, this is not a problem, as we can apply the algorithm inside a binary search on M (we know that a "guess" for M is too low if the integral GS algorithm terminates with unassigned jobs).

$M + \max p_{ij}$ in which every job is assigned to a machine that is at least as good as its best allocation in the fractional assignment.

Proof. Jobs are never fractionally assigned, and the rejection procedure ensures that the resulting makespan is at most $M + \max p_{ij}$. After the integral GS algorithm terminates, let each machine fractionally reject its least-preferred load until its utilization drops to M. This current state of assignment is one that would be reachable by the fractional GS algorithm via some sequence of proposals and rejections starting from an empty assignment. Utilizing the crucial property that the proposal/rejection sequence for the fractional algorithm does not matter, the fractional algorithm must be able to continue from this point and terminate with a job-optimal assignment (which we know exists by assumption), in which every job i's best fractional assignment is no better than its original integer assignment. Hence, every job i's integer assignment must be at least as good as its best fractional assignment in any fractional stable assignment of makespan M. $\square$

Remark C.0.2 *Given an instance of the couples problem in the NRMP as described above, the integral GS algorithm computes in $O(mn \log n)$ time a stable assignment such that every couple is assigned to the same hospital and hospitals are overcapacitated by at most one position.*

Bibliography

[1] A. Abdulkadiroglu, P.A. Pathak, and A.E. Roth. The New York City high school match. *American Economic Review.* forthcoming.

[2] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of auctions. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, 2005. forthcoming.

[3] N. Alon and J. Spencer. *The Probabilistic Method.* Wiley-Interscience, 2000.

[4] A. Archer, C. Papadimitriou, K. Talwar, and E. Tárdos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2003.

[5] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001.

[6] A. Archer and E. Tardos. Frugal path mechanisms. In *Proceedings of 13th ACM Symposium on Discrete Algorithms*, pages 991–999, 2002.

[7] Aristotle. Book V. In R. Crisp, K. Ameriks, and D.M. Clarke, editors, *Nicomachean Ethics*, pages 81–102. Cambridge University Press, 2000.

[8] R. J. Aumann. *Lectures on Game Theory.* Westview Press, 1989.

[9] R.J. Aumann. Acceptable points in general cooperative n-person games. In A.W. Tucker and R.D. Luce, editors, *Contribution to the Theory of Games: Annals of Mathematical Studies*, volume 40. Princeton University Press, 1959.

[10] R.J. Aumann. Markets with a continuum of traders. *Econometrica*, 32:39–50, 1964.

[11] J.P. Benoit and V. Krishna. Multiple-object auctions with budget constrained bidders. *Review of Economic Studies*, 68:155–179, 2001.

[12] L. Bodin and (Rabbi) A. Panken. High tech for a Higher authority: the placement of graduating rabbis from Hebrew Union College–Jewish Institute of Religion. *INTERFACES*, 33(3):1–11, May-June 2003.

[13] O.N. Bondareva. Some applications of linear programming to cooperative games. *Problemy Kibernetiki*, 1963.

[14] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, 2005. forthcoming.

[15] J. Bulow and J. Levin. Matching and price competition, December 2003. SIEPR Policy paper No. 03-006.

[16] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–1090, 1989.

[17] Y.K. Che and J. Gale. Expected revenue of the all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters*, 50:373–380, 1996.

[18] Y.K. Che and J. Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92:198–233, 2000.

[19] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

[20] A. Czumaj and A. Ronen. On the expected payment of mechanisms for task allocation. In *Proceedings of 23rd ACM Symposium on Principles of Distributed Computing*, pages 98–106, 2004.

[21] C. d'Aspremont and L.A. Gerard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[22] S. de Vries and R.V. Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, 15:284–309, 2003.

[23] X. Deng, T. Ibaraki, and H. Nagamochi. Algorithms and complexity in combinatorial optimization games. In *Proceedings of the 8th ACM Symposium on Discrete Algorithms*, 1997.

[24] L.E. Dubins and D.A. Freedman. Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, 88(7):485–494, 1981.

[25] B. Dutta. The egalitarian solution and reduced game properties in convex games. *International Journal of Game Theory*, 19:153–169, 1990.

[26] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. *Econometrica*, 57:615–635, 1989.

[27] T. Ebert, W. Merkle, and H. Vollmer. On the autoreducibility of random sequences. *SIAM Journal on Computing*, 32:1542–1569, 2003.

[28] F.Y. Edgeworth. *Mathematical Psychics.* Kegan Paul Publishers, 1881.

[29] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proceedings of 15th ACM Symposium on Discrete Algorithms*, pages 701–709, 2004.

[30] H. Fang and S.O. Parreiras. Equilibrium of affiliated value second price auctions with financially constrained bidders: The two-bidder case. *Games and Economic Behavior*, 39:215–236, 2002.

[31] H. Fang and S.O. Parreiras. On the failure of the linkage principle with financially constrained bidders. *Journal of Economic Theory*, 110:374–392, 2003.

[32] U. Feige. You Can Leave Your Hat On (if you guess its color). Technical Report MCS04-03, Computer Science and Applied Mathematics, The Weizmann Institute of Science, 2004.

[33] J. Feigenbaum, A. Archer, A. Krishnamurthy, R. Sami, and S. Shenker. Approximation and collusion in multicast cost sharing. *Games and Economic Behavior*, 2003.

[34] J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Hardness results for multicast cost sharing. *Theoretical Computer Science*, pages 215–236, 2003.

[35] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A bgp-based mechanism for lowest-cost routing. In *Proceedings of 21st Symposium on Principles of Distributed Computing*, pages 173–182, 2002.

[36] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmission. *Journal of Computer and System Sciences*, 63:21–41, 2001.

[37] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proceedings of 34th ACM Symposium on the Theory of Computing*, pages 72–81, 2002.

[38] D. Fudenberg and D.K. Levine. *The Theory of Learning in Games.* MIT Press, 1998.

[39] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

[40] D. Gale and M. Sotomayor. Ms. Machiavelli and the stable matching problem. *American Mathematical Monthly*, 92:261–268, 1985.

[41] M. Garey and D. Johnson. *Computers and Intractability.* Freeman Press, 1979.

[42] R. Garg, V. Kumar, A. Rudra, and A. Verma. Coalition games on graphs: core structure, substitutes and frugality. In *Proceedings of 4th ACM Conference on Electronic Commerce*, pages 248–249, 2003.

[43] D.B. Gillies. Solutions to general non-zero-sum games. In A.W. Tucker and R.D. Luce, editors, *Contributions to the Theory of Games, Volume IV (Annals of Mathematics Studies, 40)*, pages 47–85. Princeton University Press, 1959.

[44] M. Goemans. Personal Communication, 2004.

[45] M.X. Goemans and M. Skutella. Cooperative facility location games. *Proceedings of the 11th ACM Symposium on Discrete Algorithms*, 2000.

[46] A. Goldberg and J. Hartline. Competitiveness via Concensus. In *Proceedings of 14th ACM Symposium on Discrete Algorithms*, pages 215–222, 2003.

[47] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of 16th ACM Symposium on Discrete Algorithms*, pages 620–629, 2005.

[48] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior.* forthcoming.

[49] A.V. Goldberg, J.D. Hartline, A.R. Karlin, A. Wright, and M. Saks. Competitive auctions. *preprint*, 2003.

[50] A.V. Goldberg, J.D. Hartline, and A. Wright. Competitive auctions and digital goods. *Proceedings of 12th ACM Symposium on Discrete Algorithms*, pages 735–744, 2001.

[51] R.L. Graham, D.E. Knuth, and O. Patashnik. *Concrete Mathematics, 2nd ed.* Addison-Wesley, 1994.

[52] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[53] D. Gusfield and R.W. Irving. *The Stable Marriage Problem: Structure and Algorithms.* Foundations of Computing. MIT Press, 1989.

[54] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, 2005. forthcoming.

[55] J. Hershberger and S. Suri. Vickrey prices and shortest paths: what is an edge worth? In *Proceedings of 42nd IEEE Symposium on the Foundations of Computer Science*, pages 129–140, 2001.

[56] W. Hildenbrand. *Core and Equilibria of Large Economies.* Princeton University Press, 1974.

[57] T. Hokari. Population monotonic solutions on convex games. *International Journal of Game Theory*, 29:327–338, 2000.

[58] T.W. Hungerford. *Algebra.* Springer, 1997.

[59] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, 2005. forthcoming.

[60] N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proceedings of 16th ACM Symposium on Discrete Algorithms*, pages 53–62, 2005.

[61] N. Immorlica, M. Mahdian, and V.S. Mirrokni. Limitations on cross-monotonic cost-sharing schemes. In *Proceedings of 16th ACM Symposium on Discrete Algorithms*, pages 602–611, 2005.

[62] M. Jackson. A Crash Course in Implementation Theory. *Social Choice and Welfare*, 18:655–708, 2001.

[63] M. Jackson. Mechanism Theory. In U. Derigs, editor, *Optimization and Operations Research.* Encyclopedia of Life Support Systems, 2003. `http://www.eolss.net`.

[64] K. Jain and V.V. Vazirani. Applications of approximation algorithms to cooperative games. In *Proceedings of the 33rd ACM Symposium on Theory of Computing*, pages 364–372, 2001.

[65] K. Jain and V.V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *Journal of the ACM*, 48:274–296, 2001.

[66] K. Jain and V.V. Vazirani. Equitable cost allocations via primal-dual-type algorithms. In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pages 313–321, 2002.

[67] N. Johnson and S. Kotz. *Urn models and their application: an approach to modern discrete probability theory.* John Wiley & Sons, 1977.

[68] K. Kent and D. Skorin-Kapov. Population monotonic cost allocation on MST's. *In Operational Reearch Proceedings KOI*, pages 43–48, 1996.

[69] D.E. Knuth. *Marriages Stables et leurs relations avec d'autres problémes combinatoires.* Les Presses de l'Université de Montréal, 1976.

[70] D.E. Knuth, R. Motwani, and B. Pittel. Stable husbands. In *Proceedings of the 1st ACM Symposium on Discrete Algorithms*, pages 397–404, 1990.

[71] D.E. Knuth, R. Motwani, and B. Pittel. Stable husbands. *Random Structures and Algorithms*, 1:1–14, 1990.

[72] J. Könemann, S. Leonardi, and G. Schäfer. A group-strategyproof mechanism for steiner forests. In *Proceedings of 16th ACM Symposium on Discrete Algorithms*, pages 612–619, 2005.

[73] J. Könemann, S. Leonardi, G. Schäfer, and S. van Zwam. From primal-dual to cost shares and back: A stronger lp relaxation for the steiner forest problem. In *Proceedings of 32nd International Colloquium on Automata, Languages and Programming*, 2005. forthcoming.

[74] V. Krishna. *Auction Theory.* Academic Press, 2002.

[75] V. Krishna and R. Serrano. Multilateral bargaining. *Review of Economic Studies*, 63:61–80, 1998.

[76] J-J. Laffont and J. Robert. Optimal auctions with financially constrained buyers. *Economics Letters*, 52:181–186, 1996.

[77] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proceedings of the 44rd IEEE Symposium on Foundations of Computer Science*, pages 574–583, 2003.

[78] S. Leonardi and G. Schäfer. Cross-monotonic cost-sharing methods for connected facility location games. In *Proceedings of 5th ACM Conference on Electronic Commerce*, pages 242–243, 2004.

[79] E.S. Maskin. Auctions, development and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44:667–681, 2000.

[80] D. G. McVitie and L. B. Wilson. Stable marriage assignments for unequal sets. *BIT*, 10:295–309, 1970.

[81] M. Mihail, C. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. In *Proceedings of the 44rd IEEE Symposium on Foundations of Computer Science*, 2003.

[82] J. Mirrlees. An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies*, 38:175–208, April 1971.

[83] S.J. Mongell and A.E. Roth. Sorority rush as a two-sided matching mechanism. *American Economic Review*, 81:441–464, 1991.

[84] H. Moulin. *Axioms of cooperative decision making*, volume 1, chapter Cost sharing games and the core. Cambridge University Press, 1988.

[85] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16:279–320, 1999.

[86] H. Moulin. Axiomatic Cost and Surplus Sharing. In K.J. Arrow, A.K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, pages 289–357. Elseveir Science Publishers B.V., 2002.

[87] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance vs. Efficiency. *Economic Theory*, 18:511–533, 2001.

[88] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.

[89] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[90] M. Pál and E. Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proceedings of 44th IEEE Symposium on Foundations of Computer Science*, pages 584–593, 2003.

[91] B. Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, pages 530–549, 1989.

[92] B. Pittel. On likely solutions of a stable matching problem. In *Proceedings of the 3rd ACM Symposium on Discrete Algorithms*, pages 10–15, 1992.

[93] S. Robinson. Why mathematicians now care about their hat color. *New York Times*, page F5, April 10, 2001.

[94] E.C. Rosenthal. Monotonicity of the core and value in dynamic cooperative games. *International Journal of Game Theory*, 19:45–57, 1990.

[95] A.E. Roth. The economics of matching: stability and incentives. *Mathematics of Operations Research*, 7:617–628, 1982.

[96] A.E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016, 1984.

[97] A.E. Roth. The national residency matching program as a labor market. *Journal of the American Medical Association*, 275(13):1054–1056, 1996.

[98] A.E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. Fisher-Schultz Lecture. *Econometrica*, 70(4):1341–1378, 2002.

[99] A.E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89:748–780, 1999.

[100] A.E. Roth, T. Sonmez, and M.U. Unver. Kidney exchange. working paper.

[101] A.E. Roth and M.A.O. Sotomayor. *Two-sided Matching; A Study in Game-theoretic Modeling and Analysis.* Cambridge University Press, 1990.

[102] W. Rudin. *Principles of Mathematical Analysis.* McGraw-Hill, 1976.

[103] H.E. Scarf. The core of an n person game. *Econometrica*, 35(1):50–69, 1967.

[104] A. Schrijver. *Combinatorial Optimization.* Springer, 2003.

[105] J. Sethuraman, C-P. Teo, and W-P. Tan. The Gale-Shapley stable marriage problem revisited: strategic issues and applications. In *IPCO'99*, pages 429–438, 1999.

[106] J. Sethuraman, C-P. Teo, and W-P. Tan. The Gale-Shapley stable marriage problem revisited: strategic issues and applications. *Management Science*, 47(9):1252–1267, September 2001.

[107] L.S. Shapley. A value for n-person games. In H. Kuhn and A.W. Tucker, editors, *Contributions to the Theory of Games*, volume 2, pages 307–317. Princeton University Press, 1953.

[108] L.S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.

[109] M. Shubik. Edgeworth market games. In A.W. Tucker and R.D. Luce, editors, *Contributions to the Theory of Games, Volume IV (Annals of Mathematics Studies, 40)*, pages 267–278. Princeton University Press, 1959.

[110] Y. Sprumont. Population monotonic allocation schemes for cooperative games with transferable utility. *Games and Economic Behavior*, 2:378–394, 1990.

[111] Kunal Talwar. The price of truth: Frugality in truthful mechanisms. In *20th Annual Symposium on Theoretical Aspects of Computer Science*, pages 608–619, 2003.

[112] W. Thomson. Population-monotonic allocation rules. In W.A. Barnett, H. Moulin, M. Salles, and N.J. Schofield, editors, *Social Choice, Welfare and Ethics*, pages 79–124. Cambridge University Press, 1985.

[113] U.S. Securities & Exchange Commission and Board of Governors of the Federal Reserve System. *Joint Report on the Government Securities Market.* Government Printing Office, 1992.

[114] J. H. van Lint and R. M. Wilson. *A Course in Combinatorics.* Cambridge Univ. Press, 1992.

[115] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[116] J. von Newmann and O. Morgenstern. *Theory of Games and Economic Behavior.* John Wiley and Sons, 1944.

[117] P. Winkler. Games People don't Play. In D. Wolfe and T. Rogers, editors, *Puzzler's tribute. A feast for the mind.* A K Peters, 2001.

[118] H.P. Young. Cost Allocation. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 2, pages 1193–1235. Elseveir Science Publishers B.V., 1994.

[119] H.P. Young. Cost allocation, demand revelation, and core implementation. *Mathematical Social Sciences*, 36:213–228, 1998.

[120] C.Z. Zheng. High bids and broke winners. *Journal of Economic Theory*, 100:129–171, 2001.