On Profit Maximization in Mechanism Design

Matthew Cary

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

Doctor of Philosophy

University of Washington

2007

Program Authorized to Offer Degree: Computer Science & Engineering

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Matthew Cary

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of the Supervisory Committee:

Anna Karlin

Reading Committee:

Anna Karlin

Paul Beame

Oren Etzioni

Date: 6-4-2007

University of Washington

Abstract

On Profit Maximization in Mechanism Design

Matthew Cary

Chair of the Supervisory Committee:
Professor Anna Karlin
Computer Science & Engineering

Mechanism design is a subfield of game theory and microeconomics focused on *incentive engineering*. A mechanism is a protocol, typically taking the form of an auction, that is explicitly designed so that rational but non-cooperative agents, motivated solely by their self-interest, end up achieving the designer's goals. The challenge of mechanism design is to apply these methods to traditional computer science goals such as worst-case or competitive analysis. We consider two challenging problems related to mechanism design for profit maximization: the analysis of natural bidding strategies used by participants in sponsored search auctions, and the design of a mechanism for matroid procurement with provable performance guarantees.

The sponsored search auctions of web search engines such as Google or Yahoo! use the *generalized second-price* (GSP) mechanism, in which bidders do not have a dominant strategy. We develop a framework for studying a variety of greedy bidding strategies and analyze their revenue, convergence and robustness properties. We compare the performance of greedy bidding strategies to that of a Nash equilibrium, quantifying how close these bidding strategies are to one of the most natural and rational stable points of the system. ([9])

In the *procurement* problem a buyer is given a set of agents with values, along with a family of *feasible sets* over the agents. The goal is to procure a feasible set of maximum value, for minimum cost. Assuming that a buyer obtains a decreasing marginal benefit per

feasible set procured, the problem is to determine the optimal number of feasible sets to procure in order to maximize the buyer's profit. We develop a mechanism that approximates the optimal profit to within a constant factor, when the set system is a matroid. Matroids are important structures in combinatorial optimization: for example, minimum spanning trees and node-weighted maximum matchings are both matroid problems. We also show that the well-known cost sharing revenue extraction mechanism is only truthful for matroid set systems, so that procurement problems over non-matroid set systems are not likely to be solved with current techniques. ([10])

TABLE OF CONTENTS

LIST OF FIGURES

Figure Number Page

ACKNOWLEDGMENTS

The joy I have had in my graduate career is due to the guidance and inspiration from those who have advised me. I am privileged to have worked with and learned from them. I would like to particularly acknowledge Paul Beame, Ramarathnam Venkatesan, and especially Anna Karlin, who finally got me out the door.

DEDICATION

This dissertation is dedicated to my family, to whom I owe everything: my grandparents, James and Blanche Campbell, and Maynard and Marjorie Cary; my parents, John and Susan Cary, and sister Margaret; and most of all my wife, Sara Shaw.

Chapter 1
INTRODUCTION

The opening of the twenty-first century has witnessed a major change in the computer industry as the rise of mass-market auctions has changed economic fundamentals. Where once a self-respecting software company would earn its keep by selling word processor and spreadsheet applications, branching into the operating systems market, and then ruthlessly outmaneuvering and suppressing all competition until it had a stranglehold over the industry, now a software company can *give away* its products and instead earn billions of dollars selling advertisements, 98% of which is accounted for by advertising auctions.[1] Simultaneously, the rise of markets like eBay have transformed electronic commerce from niche business-to-business applications to a truly massive scale: according to Google finance and Reuters, eBay's total revenue for 2006 was $5.97 billion.

Alongside these changes, research has blossomed at the boundaries of game theory, economics, and the theory of computation. The scale of the new markets has meant that classical assumptions of cooperative users no longer approximate the situation, and instead the dynamics of interaction and competition must be explicitly studied. These areas of research include both the design of new protocols adapted to this new environment, as well as the study of existing systems using technology from game theory and economics in addition to the theory of computation.

Examples of the design of new protocols include problems of resource allocation. Early work in this area [37, 38] focused on producing efficient output at minimal cost to the central coordinator running the auction. This led to work analyzing the costs of achieving efficiency over sets of selfish agents in natural games and new concepts such as the prices of anarchy

[1] According to Edelman, Ostrovsky, and Schwarz [16], quoting the figure for 2005 when Google's total revenue was $6.14 billion.

and stability [42, 44, 2]. In tandem these tools were brought to bear on existing problems on the Internet like congestion control [30, 40].

We can illustrate the two approaches of analysis of existing systems and design of new protocols by considering the *sponsored search keyword auctions* that have had such an impact on the contemporary computer industry. The auction begins when a user submits a query to a search engine. On the received page, along with links that are relevant to the search, a small number of advertisements known as *sponsored links* are shown. The advertisements are chosen by auction, so that advertisers submit bids for particular keywords, and certain advertisers are chosen to be displayed with search results. Figure 1.1 shows this process in the context of a graduate student using Google search to look for a little bit of help on his dissertation. The query was "thesis research"; the search results occupy most of the page on the left, and the sponsored links appear on the right. Several different advertisers have bid on several different combinations of the search terms. If the graduate student were to have clicked on any of those links (of course, he did not), he would have been transferred to the advertiser's web site, and the advertiser would have been charged by Google for the click.

Given such an auction system, what is an advertiser to do? Bidding in auctions for single items, such as how art or livestock is sold, is a very well understood problem. The optimal bidding strategies for all common auction formats are known in detail (see Klemperer [33] or Milgrom [35]). In contrast, the sponsored search auctions involve bidding for the placement of an advertisement in one of several different locations on the web page, and users may react to each location in different ways. These auctions are run continuously, and advertisers can update their bids at any time. As a result, bidding strategies must dynamically respond to all other advertisers participating in the auction. In addition, advertisers are usually bidding for multiple search terms simultaneously and must balance their participation in each auction with budget constraints.

The complexity of this situation has created an active industry in *autobidders*, software robots that update an advertiser's auction portfolio in real-time. Autobidders from companies such as Atlas Search [6] and iProspect [27] can attack economic goals such as maximizing return on investment as well as behave strategically with anti-competitive and

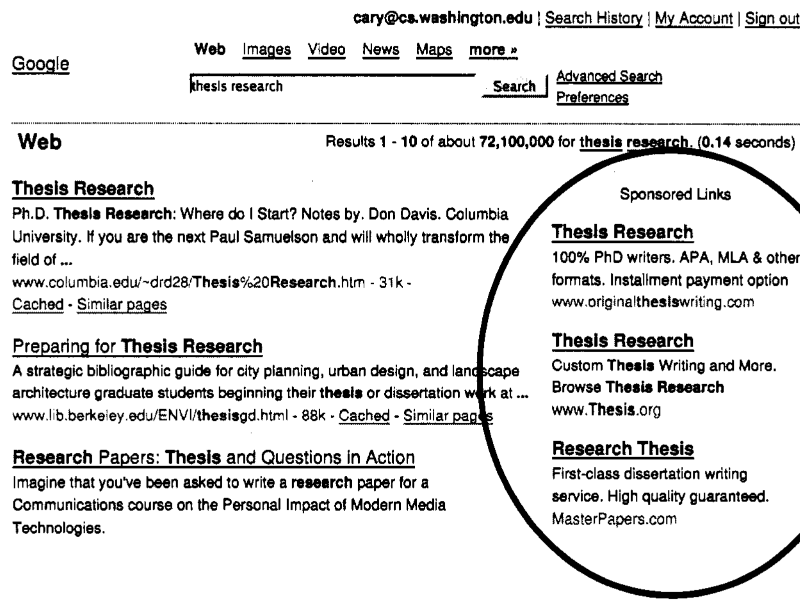

Figure 1.1: A sample search query result from Google with sponsored links circled

vindictive algorithms marketed under names like "max bid gap jammer" and "bidding war eliminator".

The current state of knowledge of bidding strategies for these dynamic auctions is bleak. Even when restricted to the problem of a single auction examined in isolation, very little is known about how to maximize an advertiser's profit when the auction is run repeatedly. Some strategies have been shown to work well or in practice [31], but only against a static backdrop of other advertisers who are not bidding reactively to the strategy in question. Other strategies have been shown to exhibit unstable behavior [50] but it is not clear what the effect of this instability is on either advertiser or search engine revenue.

An equally difficult question has been how the search engines should organize these auctions. This is the problem of *mechanism design.* Initial sponsored search auctions selected advertisers strictly in order of their bids. As sponsored search auctions developed, different mechanisms were created. Google, for example, selects advertisers based both on their bids and the rates at which their advertisements are clicked by users. This mechanism avoids the problem raised by some advertisers, such as eBay, who place advertisements simply for branding and not to elicit clicks. As these ads tend not to be clicked on very frequently, the advertisers can bid very high, knowing that they won't have to pay for too many clicks. As a result, advertisers whose ads may be quite relevant to the search query are forced out of the auction. The search engine loses revenue and users' search experiences are degraded: a lose-lose situation. Google's simple change of ranking by bid and advertisement quality reportedly increases revenue by half.[2]

Just as participants in an auction may balance several different goals when deciding on their bids, a mechanism designer has several different goals. Two of the most important are *efficiency* and *revenue*. The goal of revenue maximization is quite obvious in these sponsored search auctions: those search results don't appear by themselves, and Google engineers demand attractive paychecks. The goal of efficiency is most clear in resource allocation problems, such as the FCC spectrum auctions. These were run in order to distribute licenses in such a way to promote a healthy and active wireless market. Efficiency also appears as a secondary goal in many other contexts. The sponsored search auctions raise money for Google only to the extent that advertisers are attracted to them; advertisers are attracted only when there are lots of users; the users are attracted to Google by the quality of their search results. Google is thus very concerned about users' search experiences. Displaying advertisements with queries sharing little relevance degrades this experience, driving away users and discouraging advertisers. An efficient sponsored search auction will assign the advertisers who most value a search term to those queries. It is in Google's interest to find those advertisers even if it means sacrificing short-term revenue.

Mechanism design is also known as *inverse game theory* as it must consider all possible strategies that players could employ and select those that achieve the auctioneer's goal. This naturally places an emphasis on mechanisms where rational players will have a single strategy that is always best: a *dominant strategy*. The simplest strategy a player can use is to bid their actual valuation for the auction; mechanisms for which this is a dominant

---

[2] A New York Times article from January 24, 2007 stated that when Yahoo! switched ranking mechanisms for pure rank-by-bid to the method employed by Google, search revenue was projected to increase by 50%.

strategy are known as *truthful.* Research attention on such mechanisms over the last half a century has resulted in many beautiful theorems as well as Nobel prizes. The strongest line of work has characterized *truthful* and *efficient* auctions by the so-called Vickery-Clarke-Groves (VCG) mechanism [46, 11, 24]. While this has given mechanism designers a powerful tool, the trade-off between efficiency and revenue is not yet fully understood. Some examples may be found in *procurement* auctions, also known as reverse auctions, where instead of selling something, the auctioneer is trying to hire a team of agents to perform some task. If the auctioneer is trying to procure a path in a graph, it has been shown that no truthful mechanism will be able to avoid large overpayments compared to the true costs of the agents [4, 17]. Are such overpayments necessary for any mechanism for path procurement? To what extent can truthful mechanisms be used and avoid overpayments? These questions remain unanswered.

While much of the theory of truthful mechanisms was derived under the goal of efficiency, techniques have recently been developed extending it to a wider domain. An example of this, one that has appeared only with the advent of electronic commerce, is a *digital goods auction.* In this setting the auctioneer is selling off an item that can be reproduced at zero cost. Consider a cable company that wishes to find a set of viewers willing to pay for video-on-demand. In this case the efficient solution distributes the video across the entire network, because of the negligible cost of reproduction, and so the machinery of efficient auctions is not directly of much use to the cable company who wishes to identify the subset of viewers willing to show them the (most) money.

A series of works [22, 19] has developed a novel mechanism drawn from existing work on truthful and efficient auctions along with the economic study of *cost-sharing.* In true computer science tradition, they reduce the problem to a decision problem through a technique known as *profit extraction.* The result is a truthful mechanism that achieves in revenue a large fraction of the bidders' valuations.

In this dissertation we approach this area from both directions, studying bidder behavior and the design of optimal mechanisms. In the first half we study the sponsored search keyword auctions already mentioned, finding bidder strategies that are effective in practical auction settings. In the second half, we illustrate this by looking at procurement in an unlimited goods setting in which the auctioneer is able to resell the teams procured and hence may choose to buy any number of disjoint teams. In this setting, techniques from the study of digital goods auctions are applied to find near-optimal costs. Before describing our auction framework formally in the next chapter, we briefly outline the specific contributions of this dissertation.

1.1 Contributions

1.1.1 Keyword Auctions[3]

The preeminent mechanism used in practice for keyword auctions is generalized second-price (GSP), which has no dominant strategy. Hence this raises the question of what strategy should advertisers take, and what revenue can the search engine expect to get?

In the absence of budget constraints, there is one and only one *truthful* auction that can in principle be used: the Vickrey-Clarke-Groves or VCG mechanism [46, 11, 24]. This mechanism has the property that it is in the best interest of the participating advertisers to bid their true valuation of a click. Despite this appealing property of the VCG mechanism, for a number of reasons, no search engine uses the VCG mechanism. Rather, the most widely used auction mechanism is the non-truthful *Generalized Second Price* or GSP auction (described in Section 3.2).

The fact that the GSP mechanism is not truthful means that the participating advertisers are forced to undertake the complicated task of choosing a bidding strategy. Asdemir [5] and Edelman and Ostrovsky [15] observe that instability and bidding wars can result from the use of the GSP mechanism. To make matters worse, on a typical search page, there is room for multiple sponsored links. The positioning of these sponsored links affects the chances that a sponsored link will be clicked on and thus these advertising slots have varying desirability from the perspective of advertisers. This makes the advertisers' utilities a discontinuous function of their bids. Overall, the resulting bidding is sufficiently complex that many advertisers hire consultants or intermediaries to do their bidding for them.

---

[3] Joint work with Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna Karlin, Claire Mathieu and Micheal Schwarz. Portions of this work appear in Cary et al. [9].

We formalize a class that we call *greedy bidding strategies* that contains the first natural strategy known to converge to an equilibrium. Moreover, the equilibrium is revenue-minimizing over the class of symmetric Nash equilibria, where all advertisers are in the slot that would maximize their utility if all slot prices were fixed. This means that this equilibrium minimizes the overall payments that the advertisers make, making it a desirable strategy from the advertisers' point of view. In addition, the equilibrium is efficient, placing the highest-valued advertisers in the most desirable slots.

A greedy bidding strategy in each auction round identifies the optimal slot to bid for, assuming the bids of advertisers in previous rounds stay the same. This assumes a total information model, where each advertiser sees all the bids made in the previous round. The selection of the optimal slot determines a range of bids that are able to win that slot. This framework describes strategies that have been observed in practice. For example, the choice of the highest bid in the range implements *competitor busting* (also known as *gap jamming*), where an advertiser tries to maximize the costs to a competitor. Another greedy strategy, which we call *balanced bidding* and is based on strategies studied previously, chooses the highest bid that guarantees at least the same utility even if a competitor subsequently tries to underbid the advertiser.

The balanced bidding strategy converges to the equilibrium mentioned above, where the payments are the same as if the VCG mechanism were used. We prove this convergence in two settings, one where bidders change their bids *asynchronously* over a continuous series of auctions, and another when bidders bid *synchronously* at the beginning of each auction. We prove convergence in the asynchronous case, and show experimentally that the convergence is rapid. In the synchronous case, the unmodified balanced bidding strategy does not always converge, and if click-through rates on the slots are near each other, this non-convergence is quite common. However, a slightly modified strategy where an advertiser does not bid for high slots in an unrestricted fashion does provably converge.

Turning to the question of revenue to the search engine, we demonstrate general Nash equilibria for GSP with revenue both arbitrarily higher and lower than revenue at balanced bidding convergence. However, we also show that if advertisers are debt-averse in the sense of never bidding in a way that exposes them to payments greater than their value, then the maximum GSP revenue in equilibrium is at most a constant factor greater than the balanced bidding equilibrium, where the constant depends on the slot click-through rates and increases as the click-through rates become closer. This shows that under the reasonable assumption of debt-averse advertisers, balanced bidding does produce reasonable revenue. We confirm this analysis experimentally by calculating maximum equilibrium revenue from normally distributed bidders.

We also experimentally study the revenue produced by other greedy bidding strategies, including competitor busting. We observe that the competitor busting strategy rarely converges to an equilibrium, and that in experimental simulation it produces much more revenue than balanced bidding. In fact, if the click-through rate of all slots is significant, then competitor busting produces more revenue than the maximum equilibrium of debt-averse advertisers, even though competitor busting is itself a debt-averse strategy. Thus, the GSP mechanism can take advantage of a cut-throat advertiser market to raise more revenue for the search engine owner.

1.1.2 Set System Procurement[4]

In the set system procurement problem, an auctioneer is trying to hire teams of agents, such as a network provider who must buy spanning trees from a network of agents, each of whom owns a single link. In general, the possible teams are given as a system of feasible sets over the agent set; in this case the feasible sets are spanning trees in a graph. Another important procurement problem, also in the graph domain, is the problem of buying disjoint paths between two specified nodes when again each edge is owned by a single agent. Mechanism design for this problem tries to limit the cost of overpayment for the set procured.

In this dissertation we design a mechanism, in which the auctioneer chooses the number of sets to procure after viewing the bids, that is constant-competitive with the optimal auction for the spanning tree problem and other matroid set systems. In addition, we show that current techniques cannot be extended to procurement of disjoint paths and other

[4]Joint work with Abie Flaxman, Jason Hartline and Anna Karlin. Portions of this were submitted to FOCS2007.

non-matroid systems.

The consensus of current research on this problem is that when procuring a fixed number of sets the overpayment must be very large. The negative results known to current research on this problem are centered on path auctions. In particular there are graphs where any truthful mechanism must overpay by a large amount: $\Omega(n)$, where n is the number of edges. The situation was somewhat clarified for general set systems when it was shown that the VCG mechanism pays the same amount as a first-price auction for the cheapest feasible set if and only if the set system is a matroid, a class which includes the spanning tree setting mentioned above. This in turn implies that for any non-matroid set system, either the VCG mechanism pays the same as the first price auction, or the first price auction will not always select the cheapest feasible set. This is strong evidence that overpayment for a fixed number of sets from a non-matroid set systems cannot be avoided.

Recent work in the design of digital auctions has shown that when the auctioneer is free to choose the number of items sold in addition to the price, revenue can be constant-competitive with the optimal. We show that this intuition carries over to matroid procurement. Our near-optimal mechanism for the multi-unit procurement problem exploits this freedom, using techniques from the partition and cost-share mechanism for digital goods auctions. Furthermore, we show that this cost-sharing is not truthful on non-matroids set systems, continuing the trend of research from procurement of a single set that suggests that procurement is fundamentally different, and much more difficult, over non-matroids.

Our proof of the competitiveness of our mechanism for matroids uses a random-sampling result due to Karger along with novel structural results to show that a random partition of a matroid scales the costs for disjoint sets, that is, the cost for $k/2$ disjoint sets in one of the partitions is with high probability about half the cost for k disjoint sets in the original matroid. This enables us to estimate the value of the set system in such a way that the cost-sharing mechanism can be used to extract at least half the revenue of the optimal set.

Chapter 2
AUCTION THEORY

> ...as long as an extra penny can be gained, a profit-maximizing producer will leap to garner it. —Theodore Groves [24]

The keyword auction and procurement settings we study in this dissertation share a common framework, which we formalize in this chapter. After defining the framework, we apply it to well-known single-item auctions. These will serve as examples for the concepts of strategies, equilibria and truthful mechanisms whose definitions will finish the chapter.

Definition 2.0.1. An *auction* with players labeled $1, \ldots, n$ is described by

- A vector $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{R}^n$, corresponding to the player values,

and the following two functions, each of which map a vector of *bid values* to a vector of real numbers, one for each player:

- an *allocation* function $\mathbf{x} : \mathbb{R}^n \to \mathbb{R}^n$, and
- a *payment* function $\mathbf{p} : \mathbb{R}^n \to \mathbb{R}^n$. $\mathbf{p}$ is constrained to never force a player to pay more than her bid, that is, $p_i(\mathbf{b}) \leq b_i$ for all i.

A *round* of the auction consists of the players choosing *bids* $\mathbf{b} = (b_1, \ldots, b_n)$. Then

- The *utility* of a player is $u_i = x_i(\mathbf{b}) \cdot v_i - p_i(\mathbf{b})$, and
- the *revenue* to the auctioneer is equal to $\sum_i p_i(\mathbf{b})$.

In the case that the allocation or payment functions are random, the utility and revenue will be taken in expectation over the random choices made by the mechanism. When the bid vector $\mathbf{b}$ is understood we will write x_i for $\mathbf{x}(\mathbf{b})_i$ and p_i for $\mathbf{p}(\mathbf{b})_i$.

There are two important points to this formulation of utility. The first is that we are restricting our attention to a *private-value* model, that is, the utility of a player depends only on her value, the allocation, and her payment, and not on the values of any other players. This is contrasted with a *common-value* model, such as an oil-field auction where players bid for drilling rights for a tract holding an unknown, but estimable, quantity of oil. The second is that we are restricting ourselves to *quasi-linear* utilities, rather than more general functions of the allocation and the payment. In particular, this utility function depends only on the player's allocation, and not the allocation of any other player, so that it does not model situations where one player cares about what another player receives in the auction.

It will be useful to talk about how an auction changes if a single player changes her bid. Given a vector $\mathbf{a} = (a_1, \ldots, a_n)$, we will will use $\mathbf{a}_{-i}$ to denote the vector

$$\mathbf{a}_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$$

of all elements of $\mathbf{a}$ except the i-th. We let $(a', \mathbf{a}_{-i})$ denote the vector

$$(a', \mathbf{a}_{-i}) = (a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_n)$$

which is equal to $\mathbf{a}$ except for the i-th coordinate being replaced with a'.

2.1 Single-Item Auctions

Examples of the single-item auction include the selling of art by dedicated auction houses such as Sotheby's and the Department of Transportation putting a road-paving project out to bid. The first example shows something being sold, the second the procurement of a service. In the case of the painting, bidder values might be determined by combining resale value of the painting and prestige for owning it along with some intrinsic value related to the aesthetic effect the work produces on the potential bidder. In the procurement case each contractor knows or can estimate the cost for her to perform the paving job. In both cases when the auction is finished only one bidder is in possession of the object, so that the allocation function chooses the winning bidder i^*, and sets $x_{i^*} = 1$ and $x_j = 0$ for all $j \neq i^*$. The common mechanisms described below all choose the winner with the highest bid, so that $i^* = \operatorname{argmax} b_i$.

First-price sealed bid In this mechanism, all bidders submit a sealed bid. The winning bidder pays the price she bid. In the procurement auction, the lowest bidder receives the object and is paid her bid. Alternatively, for the procurement auction we may think of the bidders as submitting negative bids, so that the highest bidder wins and pays the (negative) amount of her bid. In this way the procurement auction is exactly parallel to the forward auction. In either case, we set the payment function as follows.

$$p_j = \begin{cases} b_{i^*} & \text{if } j = i^*, \\ 0 & \text{otherwise.} \end{cases}$$

A variation on this format is the *descending*, or *Dutch*, auction. No bids are submitted; instead, the auctioneer announces a very large bid, and successively lowers it until some player accepts it. This player receives the object, and pays the currently announced price.

Second-price sealed bid Like the first-price auction, all bidders submit a sealed bid to the auctioneer, and the highest bidder receives the object. In this case the winner pays the *second-highest* bid. If we let i_2 be this second-highest bidder, then the payments are as follows.

$$p_j = \begin{cases} b_{i_2} & \text{if } j = i^*, \\ 0 & \text{otherwise.} \end{cases}$$

While this might seem to be an artificial mechanism, it is equivalent to one of the most famous auction formats, the *ascending*, or *English*, auction. The auctioneer begins by announcing a low bid, successively raising it. All bidders begin the auction eligible to receive the good, and as the price rises, bidders decide to drop out of the bidding. The last player remaining receives the item at the currently announced price. This can be seen as the reverse of the Dutch auction. Several instances of ascending auctions, such as those found in online settings like eBay [14] or Google [23], are nominally ascending, but actually conducted through auctioneer-provided robots in such a way that the actual mechanism is second-price.

The iterative nature of the ascending and descending auctions mean that they do not immediately fit into our framework. However, as we are assuming a private-value model where bidder values are uncorrelated, hearing the bids of other players should not affect the strategy of a particular player. Hence in the ascending (English) auction, we may as well describe a player's strategy as fixing a maximum bid b_i before the auction commences, which is the price at which the player will drop out of the auction. The winner will be the player with the highest bid target, but that player's payment will be the second-highest target from the last player who drops out. Thus the payments are identical to those of a sealed-bid second-price auction. Similarly, in a descending (Dutch) auction, b_i will be the price at which the player is resolved to buy the item. The price will drop to the highest target, and that player will pay her target.

2.2 Strategy and Equilibrium

A *strategy* is how a player involved in an auction determines her bid. In many simple auctions, a player's bid may simply be a function of her value; in more complicated situations, like a repeated keyword auction, a strategy will take into account the bids of players from previous rounds as well as the click-through-rates of various slots, all compared with a player's value, in order to determine the bid. A strategy may be deterministic or random; like with a randomized mechanism the utility from a random strategy is usually taken in expectation over its random choices.

A central tenant of game theory is that rational players will choose strategies to maximize their utility as much as possible and, as suggested by the quote opening this chapter, are sensitive to even small changes in this utility. What exactly is possible depends on the situation. For example, one may assume that a player uses the *ex ante* best strategy, which maximizes the expected revenue over the choices from distribution of values for the other bidders. Alternatively, one may require an *ex post* optimal strategy, where even after seeing all bids, the player does not want to change her bid. This notion is most useful either when the auction is repeated for multiple rounds, or when the model implicitly assumes multiple rounds of the auction, so that *ex post* optimality approximates the situation by overstating the information bidders learn about each other. Following the modus operandi in algorithmic game theory we concentrate on *ex post* analysis for both player strategies and mechanisms. This is in contrast to much of the research from the economics community, which assumes a known prior distribution of bidder values and then searches for *ex ante* optimality.

2.2.1 Nash Equilibria

An *equilibrium* is a stable set of bids, according to some restriction on player strategies. The most important is a *Nash equilibrium*, capturing the intuitive idea of rational players mentioned above. While it is a very general kind of equilibrium, it has the disadvantage that there frequently occur multiple Nash equilibria for a given set of players. This means that characterizing the Nash equilibria may not give very good bounds on revenue or player behavior.

Definition 2.2.1 (The Nash Equilbirium). Given a mechanism $(\mathbf{x}, \mathbf{p})$, a vector of bids $\mathbf{b}$ is said to be in *Nash equilibrium* if for each player i,

$$x_i(\mathbf{b})v_i - p_i(\mathbf{b}) = \max_{b'} x_i\big((b', \mathbf{b}_{-i})\big)v_i - p_i\big((b', \mathbf{b}_{-i})\big).$$

That is, there is no bid b' that player i can make that will increase her utility, assuming all other player's bids remain the same.

In many combinatorial situations, the auction outcome may depend critically on the tiebreaking rule. For example, Immorlica, Karger, Nikolova, and Sami [26] show examples where the tiebreaking rule can determine whether a Nash equilibrium exists at all. In most cases, these problems are technical in nature and distract from the intuition of the situation. Therefore we will frequently use an approximate version of the Nash equilibrium, which is more realistic practically as bids are usually drawn from a discrete space rather than the continuum.

Definition 2.2.2 (The ε-Nash Equilbirium). Given a mechanism $(\mathbf{x}, \mathbf{p})$, a vector of bids $\mathbf{b}$ is said to be in *ε-Nash equilibrium* if for each player i,

$$x_i(\mathbf{b})v_i - p_i(\mathbf{b}) \geq \max_{b'} x_i\big((b', \mathbf{b}_{-i})\big)v_i - p_i\big((b', \mathbf{b}_{-i})\big) - \varepsilon.$$

That is, there is no bid b' that player i can make that will increase her utility by more than ε, assuming all other player's bids remain the same.

Finally, note that the various Nash equilibria operate under the implicit assumption that no other players change their bids. While this is a useful myopic approximation, it does mean that in, for example, a second-price auction, where a player's payment is less than her bid, an equilibrium may entail a player bidding above her value. Consider two players, with values 10 and 5 in a second price auction. The first player bidding 4 and the second 12 forms a Nash equilibrium, as neither player will improve her utility by changing her bid—the second player is already winning, and the first player will pay more than her value if she bids to win. On the other hand, the first player can change her bid to be her value, without changing her utility, but now the second player will be forced to pay more than she thinks the item is worth. This motivates the following definition.

Definition 2.2.3. A player i is *debt-averse* if for any bid b_i, for all possible $\mathbf{b}_{-i}$, $p(b_i, \mathbf{b}_{-i}) \leq v_i$. That is, a debt-averse player never exposes herself to paying more than her value.

2.3 Dominant Strategies and Truthful Mechanisms

We now define our notion of a dominant strategy, and the link between so-called *truthful* mechanisms and dominance demonstrated by the revelation principle. In accordance with our *ex post* approach to analysis, we define dominance with respect to all possible bids of the other players.

Definition 2.3.1. Consider a deterministic auction with mechanism $\mathcal{M}$ and bidder i with value v_i. Let $\mathbf{b}_{-i}$ be bid vector from all other players, and let $u_i^{\mathcal{M}}(b, \mathbf{b}_{-i})$ be the utility of player i when bidding b under this mechanism. Then $\mathbf{b}$ is a *dominant strategy* for $\mathcal{M}$ if for all i, all choices of $\tilde{\mathbf{b}}_{-i}$, and all $\tilde{b}$, it holds that

$$u_i(\tilde{b}, \tilde{\mathbf{b}}_{-i}) \leq u_i(b_i, \tilde{\mathbf{b}}_{-i}).$$

We say that a randomized mechanism $\mathcal{M}$ has a dominant strategy if it is a distribution of deterministic auctions that have the same dominant strategy.

A mechanism is *truthful* if it has a dominant strategy $\mathbf{b}$ with $b_i = v_i$. In other words, the dominant strategy in a truthful mechanism is for each player to bid her value.

The *revelation principle* says that for any mechanism $\mathcal{M}$ with a dominant strategy, there is a truthful mechanism with the same allocation and payoffs. To see this, suppose that $\mathcal{M}$ has allocation and payment functions $\mathbf{x}(\cdot)$ and $\mathbf{p}(\cdot)$, respectively, and that for player i with value v_i, her bid for the dominant strategy is given by $b_i = f_i(v_i)$. Then let $\mathcal{M}'$ be the mechanism whose allocation and payment functions are

$$\mathbf{x}'(\mathbf{b}) = \mathbf{x}\big(f_1(b_1), \ldots, f_n(b_n)\big), \quad \text{and} \quad \mathbf{p}'(\mathbf{b}) = \mathbf{p}\big(f_1(b_1), \ldots, f_n(b_n)\big),$$

respectively. Then since $\big(f_1(v_1), \ldots, f_n(v_n)\big)$ is a dominant strategy for $\mathcal{M}$, $(v_1, \ldots, v_n)$ is a dominant strategy for $\mathcal{M}'$ and hence $\mathcal{M}'$ is truthful, and produces the same allocation and payments as $\mathcal{M}$ at the dominant strategy.

2.3.1 The VCG Mechanism

Here we will present the VCG mechanism as a natural development of an efficient truthful mechanism. Efficiency means that the outcome maximizes player value; that is, the allocation $\mathbf{x}$ maximizes $\sum_{i \in N} x_i \cdot v_i$. This is a natural requirement in many settings, in particular when the auction is being used as a protocol to solve an optimization problem over the players [37, 38]. The VCG mechanism has a much longer history in the economics literature in the context of maximizing social welfare. This arose while studying how the government can improve market function through "counter-speculation" (by Vickrey [46] and generalized by Clarke [11]) as well as how to induce truth-telling by subordinates reporting to a centralized bureau in either a large company or a government with a centralized economy, by Groves [24]. Classical results known as *revenue equivalence* imply that the VCG mechanism is the only efficient, truthful mechanism (up to a fixed participation price), and so it has naturally risen to preeminence in mechanism design.

Our requirement for the VCG mechanism is that it be truthful, which fixes the bids of the dominant strategy. The efficiency requirement fixes the allocation choice, as the allocation can simply be computed from the bids assuming that they are the actual player values. Hence the only thing left to describe is the payment structure, which must guarantee the assumed dominant strategy. The essential intuition behind the VCG mechanism is that each player pays the incremental cost to the remaining community of her presence, that is the difference between the value for all other players from the efficient allocation, and the total value of the efficient allocation were she not there.

For example, consider a single-item auction and suppose that we assume for simplicity that all bidder values are distinct, and that we have ordered them so that $v_1 > v_2 > \cdots > v_n$. The efficient allocation for the single-item has $x_1 = 1$ and $x_i = 0$ for all $i > 1$. The total value of the auction is just $V(\mathbf{x}, \mathbf{v}) = v_1$. Similarly, if player 1 does not participate the allocation $\mathbf{x}'$ has $x'_2 = 1$ so that $V(\mathbf{x}', \mathbf{v}_{-1}) = v_2$. As the value from the first allocation $\mathbf{x}$, not including the first player, is 0, we have that the difference for the community is v_2, which is the winner's price. Thus the VCG mechanism for a single-item auction is just the second-price auction. We formally define the VCG mechanism as follows.

Definition 2.3.2. Let $V(\mathbf{x}, \mathbf{v}) = \sum_i x_i \cdot v_i$, and let $\hat{\mathbf{x}}(\mathbf{v})$ be an efficient allocation over $\mathbf{v}$, that is it satisfies

$$V(\hat{\mathbf{x}}(\mathbf{v}), \mathbf{v}) = \max_{\mathbf{x}} V(\mathbf{x}, \mathbf{v}).$$

The VCG mechanism is that with allocation $\hat{\mathbf{x}}(\mathbf{b})$ and payment p where

$$p_i(\mathbf{b}) = V(\hat{\mathbf{x}}_{-i}(\mathbf{b}_{-i}), \mathbf{v}_{-i}) - \sum_{j \neq i} \hat{x}_i(\mathbf{b}) \cdot v_i.$$

Note that if $x_i = 0$, then the VCG mechanism specifies that i makes no payment.

The VCG mechanism is not universally truthful. For example, if there is one player who is always in the efficient allocation, that player will have no incentive to bid truthfully. This player can bid anything, and the mechanism, constrained as it is to choose an optimal set, must charge that player whatever she desires. Hence no mechanism will be able to guarantee truthful bidding by that player.

As it turns out, this example actually characterizes when the VCG mechanism is effective. Let us call a player i *pivotal* for $\mathbf{v}$ if $\mathbf{x}(\mathbf{v}) \neq \mathbf{x}(\mathbf{v}_{-i})$. A player with value v_i is *potentially pivotal* if there is some $\mathbf{v}_{-i}$ such that i is pivotal for $(v_i, \mathbf{v}_{-i})$. That is, a player is pivotal if her absence changes the auction allocation. An auction with all players potentially pivotal is called *monopoly-free.* The following theorem demonstrates this characterization.

Theorem 2.3.3 ([35]). *If all players are potentially pivotal, so that the auction is monopoly-free, then truthful reporting is a dominant strategy for the VCG mechanism.*

Chapter 3

SPONSORED SEARCH KEYWORD AUCTIONS

For our first application of auction theory, we turn to keyword auctions. In this chapter we give formal background into the variety of keyword auction formats explored both academically and in practice by search engine companies such as Yahoo! and Google. We will first describe concepts shared by all keyword auctions before comparing the various formats and specifying the one used in this dissertation. We close by describing related work on bidding strategies from both the theoretical and practical worlds.

3.1 Auction Format

Recall that the goal of a a keyword auction is to select a number of advertisements to display alongside search results for a keyword query. Each advertiser chooses a set of search terms and submits a bid for each one. When a user enters a search query, all advertisers with terms matching keywords in the query are entered into an auction using their submitted bids. Multiple advertisements are shown along with the search results, usually arranged in a column on the right of the page, as shown before in Figure 1.1. The displayed ads contain links to the advertiser's web site where presumably the user can purchase goods. There are three steps to this transaction: *impression*, when the advertisement is viewed, *clicking*, when the user selects a link and is transferred to the advertiser's web site, and *conversion*, when the user buys something from the advertiser.

Older forms of Internet advertising such as banner ads charged rates according to the number of impressions. In contrast, sponsored search auctions only charge advertisers at the second step, only when their ads are actually clicked. Thus, an important quantity in determining the effectiveness of an advertisement is its *click through rate* (CTR), the ratio of the number of clicks to the number of impressions. The CTR is a number between zero and one, as an ad can be clicked at most once per user impression.

The CTR is primarily determined by two factors: advertisement quality and position effects. Advertisement quality results from the design of the advertisement and how well it is related to the keyword. Position effects result from where the advertisement appears on the source page; each possible location for an advertisement is called a *slot*. For example, when the ads are arranged in a column on one side of the search result page, users will first view the ad in the top slot and, most importantly, are more likely to click that ad than any other [18, 34, 6]. For simplicity, these two factors are frequently assumed to be *separable*; that is, each advertiser i, has a *quality rating* α_i, and each possible slot j has a *click through rate* θ_j. The final CTR for advertiser i appearing in slot j is determined by the product $\beta_{ij} = \alpha_i \cdot \theta_j$.

An advertiser's benefit from this auction also depends on the conversion rate. This is much more difficult to quantify, as advertisers do not generally share information about conversion rates. What limited evidence there is suggests that the conversion rate does not depend on slot location [32], contradicting anecdotal claims that clicks on lower slots are more likely to lead to conversions because such users are more serious about buying than those clicking on higher slots. This allows us to assume that any conversion rate is accounted for in an advertiser's value per click, and so does not need to be modeled explicitly.

With the above observations in mind, we formally define our auctions as follows.

Definition 3.1.1. A *keyword auction* is defined by:

- A set of k slots with click-through rates (CTRs) $\theta_1 > \ldots > \theta_k$, where θ_i is the probability that the user will click on the advertisement in slot i.

- A set of n players (advertisers) participating in the auction, each one having a private valuation v_i for a click, and a quality rating α_i.

- Based on knowledge of the auction mechanism and their own private valuations, each player submits a bid to the auction. We denote by b_i the bid submitted by player i.

- The auction mechanism then

- computes a ranking σ. Players $\sigma(1), \ldots, \sigma(k)$ are assigned to slots $1, \ldots, k$, respectively. The allocation $\mathbf{x}$ computes the expected number of clicks received by each player using the CTR, so that for $s = 1, \ldots, k$, $x_{\sigma(s)} = \alpha_{\sigma(s)}\theta_s$.
- charges a price p_s to the player $\sigma(s)$ for each click on his advertisement, so that the expected total payment is $\alpha_{\sigma(s)}\theta_s p_s$.

Note that if player i is assigned slot s, the utility $u_i = x_i v_i - \alpha_i \theta_s p_s$ may be written as $\alpha_i \theta_s (v_i - p_s)$.

We also define a simplified setting that we will use in our analysis.

Definition 3.1.2. A *uniform* keyword auction is one where $\alpha_i = 1$ for all i, that is, all advertisers have identical quality ratings.

The auction may be run in one of two formats. We refer to our standard format, where all players simultaneously update their bids on each round, as the *synchronous* model. In the *asynchronous* model, in each round, exactly one player updates her bid, while the other players merely repeat their previous bids. We consider both the case in which the player performing the update is arbitrary and the case where the player performing the update is chosen at random. This has been studied, for example in Zhang [49]. The asynchronous format is closer to realistic applications, although the synchronous model does apply in on-line settings where bids are updated in batches.

Equilibria are the same in the synchronous and asynchronous auction formats: if all players are happy with their bids, they will not change them whether or not they are bidding simultaneously with other players, or bidding by themselves. The different formats are significant when bidding strategies are considered; for example in Section 4.2 we show a bidding strategy that converges to its unique equilibrium in the asynchronous settings, but does not always converge in the synchronous format.

The auctioneer may also choose to work in an *open-cry* format, where all bids are public, or in a *sealed-bid* format, which keeps advertiser bids private. Historically, Yahoo! has used the former model whereas Google has used the latter. We assume an open-cry format in our analysis; in practice one can estimate the bids of other advertisers [31], so that our model can also apply to a setting where such estimation can be performed to an arbitrary degree of accuracy.

3.2 Ranking and Payment Schemes

Once the slots and the players are fixed, a keyword auction mechanism is defined by two things: a ranking scheme that assigns advertisers to slots based on their bids, and a payment scheme that determines the cost per bid. The two ranking schemes that have been used in practice are known as rank by bid (RBB) and rank by revenue (RBR). There are two payment schemes used in practice, first-price and generalized second-price (GSP). The VCG payment mechanism has not been used in practice but will be of theoretical interest. A company called Overture opened the field in 1997 with a first-price RBB scheme. The company Google followed with a generalized second-price RBR scheme in 2000, and in 2002 Yahoo!, which had acquired Overture, switched over to a generalized second-price auction while keeping the RBB scheme. In February 2007, Yahoo! joined Google in using RBR, and thus the generalized second-price RBR scheme has become the standard in practice.[1]

Recall that $\sigma(\cdot)$ is the ranking function, α_i is the quality rating of player i and b_i is that player's bid, and finally $\mathbf{p}$ is the vector of prices charged per click.

First-Price The first price payment scheme simply charges an advertiser her bid whenever her ad is clicked. Theoretical analysis by Lahaie [34] shows that a first-price scheme has no Nash equilibrium, and thus is more difficult to reason about. This, combined with the success of ascending auction formats used for example by eBay [14], may have contributed to the abandonment of the use of first-price scheme in practice.

Generalized Second-Price (GSP) The payment p_s is determined by the bid of the winner of slot $s+1$, by setting p_s equal to the minimum bid the advertiser $\sigma(s)$ winning slot s would have had to bid to out-rank the bid $b_{\sigma(s+1)}$ made by the advertiser winning

[1] This brief history was compiled from Lahaie [34] and Yahoo! [48]. As mentioned in the introduction, according to a New York Times article from January 24, 2007, the switch by Yahoo! from RBB to RBR was projected to increase revenue per search by 50%.

slot $s+1$. This will be made more clear as we examine the two ranking schemes, next.

Rank by Bid (RBB) In RBB the bids are ranked directly, so that $\sigma(\cdot)$ is chosen satisfying $b_{\sigma(1)} \geq b_{\sigma(2)} \geq \cdots \geq b_{\sigma(n)}$. Thus $p_s = b_{\sigma(s)}$ in a first-price auction, and $b_{\sigma(s+1)}$ under GSP.

Rank by Revenue (RBR) Bids are ranked according to the product of bid and advertisement quality, that is, $\alpha_{\sigma(1)}b_{\sigma(1)} \geq \alpha_{\sigma(2)}b_{\sigma(2)} \geq \cdots \geq \alpha_{\sigma(n)}b_{\sigma(n)}$. In a first-price auction the payments are the bids, as in RBB, which means that one advertiser may win a better slot then another advertiser, and yet have lower total payment, if her quality ranking is also lower. In contrast, the expected payment in a second-price auction satisfies $\alpha_{\sigma(j)}p_j = \alpha_{\sigma(j+1)}b_{\sigma(j+1)}$, or $p_j = \alpha_{\sigma(j+1)}b_{\sigma(j+1)}/\alpha_{\sigma(j)}$, so that expected revenue is decreasing with the ranking.

VCG Payments Finally, we mention the VCG mechanism, which is determined as in Section 2.3.1, each advertiser paying the difference in total value produced between her being in the auction or not. We will explain these in more detail in Section 3.5, below.

Lahaie [34] has examined the differences between allocating according to RBB versus RBR. He shows that while RBR and RBB behave similarly *ex post*, they have different behavior *ex ante*, when advertisers optimize against a known prior distribution of values and quality ratings. In particular, RBB is not *ex ante* efficient, while RBR is. If we turn to *ex post* analysis and examine what happens under specific advertiser values, neither ranking is always efficient, although the difference from efficiency can be bounded by a constant factor. This bound is the same in RBB and RBR, relative to the ranking, underlining the strategic similarity of the two schemes.

These formats have also been compared experimentally by Feng et al. [18]. They show that RBB and RBR behave very similarly if an advertiser's value is positively correlated with her quality rating, although in general RBR produces more revenue for the auctioneer than RBB. This confirms the universal adoption of RBR seen in practice that was mentioned above.

We next describe equilibria and VCG payments in the simplified uniform keyword auction, where all advertiser quality ratings are identically one. We will then justify this simplification through the use of *effective uniform auctions* in the following section.

3.3 Equilibrium and VCG Payments in a Uniform Keyword Auction

In a uniform keyword auction, the RBR and RBB schemes are identical. In this section, we show what the VCG payments are and characterize some natural equilibria.

Suppose that in a given allocation of advertisers to slots, we number the advertisers according to their assigned slot, so that b_1 is in slot 1, b_2 in slot 2, *etc.*. In this case the total expected value of the allocation from all advertisers except i is $\left(\sum_{1\leq j\leq k}\theta_j v_j\right) - \theta_i v_i$, and the expected VCG payment for advertiser i should be the difference between this value, and the total value without advertiser i. As the ordering between advertisers does not change when i is excluded, the effect of this is to bump all advertisers below i up one slot. Note that this is the payment *expected* from i, that is, after accounting for the click-through rates, and so gives us $\theta_i p_i^{\mathrm{VCG}}$ and not the payment itself. We summarize this as follows.

$$\begin{aligned}
\theta_i p_i^{\mathrm{VCG}} &= \sum_{j<i}\theta_j v_j + \sum_{j>i}\theta_{j-1}v_j - \sum_{j\neq i}\theta_j v_j \\
&= \sum_{j>i}(\theta_{j-1} - \theta_j)v_j.
\end{aligned}$$

A similar analysis occurs when we consider a Nash equilibrium for the GSP auction. A set of bids is in equilibrium if there is no advantage for any player to bid for a slot other than her current one. The prices of slots below a player would remain the same if she were to bid for them, but when bidding for a higher slot the player must account for the fact that the other players will move down into her vacant slot. Keeping the convention that advertisers are numbered according to their assigned slot, the result is that a Nash equilibrium satisfies the following constraints.

$$\theta_i(v_i - b_{i+1}) \geq \theta_j(v_i - b_{j+1}) \qquad \text{for } j > i, \text{ and} \tag{3.1}$$

$$\theta_i(v_i - b_{i+1}) \geq \theta_j(v_i - b_j) \qquad \text{for } j < i. \tag{3.2}$$

3.4 Symmetric Nash Equilibria

Varian [45] defines a natural subset of Nash equilibria he calls *symmetric Nash equilibria* that relates the VCG revenue to that from the GSP auction.

Definition 3.4.1. A set of bids in *symmetric Nash equilibrium* satisfy

$$\theta_i(v_i - b_{i+1}) \geq \theta_j(v_i - b_{j+1}) \qquad \text{for all } i \neq j.$$

Advertisers are in a symmetric Nash equilibrium when they prefer their slot to any other slot at the prices currently being paid. That is, when considering a higher slot ($j < i$), advertiser i does not account for the fact that slot prices are shifted down when she changes her bid, unlike the condition for a general Nash equilibrium. An advantage of a symmetric Nash equilibrium is that its bids are always in order, and it can be formulated by a local constraint.

Lemma 3.4.2 ([45]). *A set of bids is in a symmetric Nash equilibrium if and only if for each player i, $b_i \leq b_{i-1}$, and*

$$\theta_i(v_i - b_{i+1}) \geq \theta_{i-1}(v_i - b_i), \text{ and} \tag{3.3}$$

$$\theta_i(v_i - b_{i+1}) \geq \theta_{i+1}(v_i - b_{i+2}). \tag{3.4}$$

This local expression of the symmetric Nash equilibrium allows one to express the revenue explicitly, and show the following bounds.

Lemma 3.4.3 ([45]). *A bid vector* $\mathbf{b}$ *is in a symmetric Nash equilibrium if for all i,*

$$\frac{1}{\theta_{i-1}} \sum_{j>i} (\theta_{j-i} - \theta_j) v_j \leq b_i \leq \frac{1}{\theta_{i-1}} \sum_{j>i} (\theta_{j-i} - \theta_j) v_{j-1}.$$

Inspecting the lower bound of this recursion shows that it equals the VCG price for slot $i-1$. Hence we have the following corollary.

Corollary 3.4.4 ([45]). *For any set of player values $v_1, \ldots, v_n$, the minimum revenue $\min \sum_{k \leq ile1} \theta_i b_{i+1}$ over bids in a symmetric Nash equilibrium is equal to the VCG revenue obtained from $v_1, \ldots, v_n$.*

Thus, if advertisers are restricted to symmetric Nash equilibria, then they can do no better than paying the VCG revenue to the auctioneer. This helps justify using the VCG revenue as a benchmark when studying these auctions. We also have the following.

Lemma 3.4.5 ([45]). *Any symmetric Nash equilibrium is also a Nash equilibrium. The maximum revenue from a Nash equilibrium is equal to the maximum revenue from a symmetric Nash equilibrium.*

We complement this lemma in Section 4.3.3 by showing a Nash equilibrium with lower revenue than that from VCG. This equilibrium is not symmetric, showing that the symmetric Nash equilibria are a strict subest of Nash equilibria.

3.5 Effective Uniform Auctions

Here we justify the restriction to uniform auctions in the practically important RBR setting. This allows us to only consider click-through rates for slots by folding per-advertiser quality scores into their values. This model is also used by Edelman et al. [16] and Varian [45]. Let the *scaled* value of an advertiser i be $v_i' = \alpha_i v_i$ and the corresponding scaled bid be $b_i' = \alpha_i b_i$. Given a bid vector, renumber the players in order of their scaled bids, which is the RBR ranking. Then, as mentioned above, player i pays $\alpha_{i+1}b_{i+1}/\alpha_i = b_{i+1}'/\alpha_i$ per click. The expected utility of the advertiser is then $\alpha_i\theta_i(v_i - b_{i+1}'/\alpha_i) = \theta_i(v_i' - b_{i+1}')$. Similarly, the total expected revenue is $\sum \theta_i b_{i+1}'$. Compare this to the utilities of the RBB second price auction when the values are ordered by their bids: $\theta_i(v_i - b_{i+1})$ and $\sum \theta_i b_{i+1}$, respectively. Thus the RBR second-price auction where advertiser with value v_i and quality α_i makes a bid b_i is equivalent to a uniform auction with advertiser value v_i' and bid b_i'.

The analysis of VCG payments in the RBR auction is similar. For the same reasons as in the uniform auction, the total expected value for an allocation excepting that of advertiser i is

$$\sum_{j\neq i} \alpha_j\theta_j v_j,$$

and the total were i not included is

$$\sum_{j<i} \alpha_j\theta_j v_j + \sum_{j>i} \alpha_j\theta_{j-1} v_j.$$

Subtracting gives the expected payment from i is

$$\alpha_i \theta_i p_i = \sum_{j>i} (\theta_{j-1} - \theta_j) \alpha_j v_j.$$

Rewriting in terms of the scaled payments, we have $\theta_i p'_i = \sum_{j>i} (\theta_{j-1} - \theta_j) v'_j$, which is the VCG payment for a uniform auction with advertiser i's value v'_i.

3.6 Concrete Bidding Strategies

Observations of actual keyword auctions confirm that advertisers use complex and dynamic bidding strategies. Advertisers have been observed starting a bidding war to ramp up prices, followed by dropping their price to force competitors to overpay for a slot, strategic behavior which in addition can reduce search engine revenue [15, 5]. Even when not directly manipulating the outcome for other advertisers, advertisers tend to bid very actively, using different strategies depending on the type of goods they sell [1]. These studies also suggest widespread use of autobidding through online robots in addition to manual bid submission.

An industry selling such autobidding robots has been created in response to the complexity of bidding strategies. Companies such as Atlas Search [6] and iProspect [27] offer strategies with evocative names like "max bid gap jammer", "bidding war eliminator" and "timed caboose rule" that fulfill bidder priorities such as maximizing return on investment and consumption of competitors' budgets. There has been some theoretical analysis of anticompetitive behavior that confirms the lack of equilibria and suboptimal search engine revenue seen in practice [5, 50]. These analyses are for the most part limited to the case of two bidders.

An important practical factor that we do not address in our keyword auction model is that advertisers are generally bidding simultaneously for several keywords under a budget constraint, so that instead of the luxury of considering each auction in isolation, advertisers must balance expenditures across all the auctions they are participating in.

Two approaches to bidding strategies for multiple auctions are explored by Kitts and Leblanc [31] and Borgs, Chayes, Etesami, Immorlica, Jain, and Mahdian [7]. In both settings, each advertiser participates in multiple auctions with a different valuation for receiving a click in each auction. The advertiser has a total budget which bounds the total payments made over some unit of time such as a day or a week. These budgets are usually soft, so that it is permissible to exceed it by a small amount in one time period, as long as spending is brought under control in the next.

Kitts and Leblanc [31] compute bids by solving an integer program that explicitly maximizes utility under budget constraints as a function of current prices and click-through rates. These variables are not directly known to the auctioneer, and the bulk of their analysis deals with finding good estimates of the prices and click-through rates. This is especially important in sealed-bid formats such as that used by Google [23], where competitors' bids are not directly revealed, forcing the bidding strategy to alternate between exploratory phases that bid in order to sample prices for different slots, and revenue-producing phases that use the estimated variables to solve the revenue-maximizing integer program. A notable feature of this work is that it does not explicitly take into account the dynamics of the auction, instead assuming that variable estimates are updated frequently enough to produce accurate bids. Thus it is unclear what would happen if, for example, several advertisers were all to use this strategy simultaneously. They do compare their strategies in practice against manual bidding and find that the automatic strategy both reduces expenditures as well as increasing click-through rates, making it an attractive tool for advertisers to use.

A more analytic approach to the problem of bidding for multiple auctions is taken by Borgs et al. [7]. Suppose that an advertiser with budget B per day is involved in m different auctions, and let her utility in auction j as function of her payment p_j be denoted $u_j(p_j)$; the payment, in turn, is a function of her bid. They analytically study the problem of maximizing $\sum u_j$ subject to the budget constraint $\sum p_j \leq B$, and note that Lagrangian relaxation implies that an optimal solution must satisfy $du_j/dp_j = \lambda$ for all j and some fixed constant λ. The quantity du_j/dp_j is known as *marginal return on investment* (marginal ROI), and hence this optimality condition intuitively chooses bids that balance this marginal ROI across all the auctions. Unfortunately, computing the marginal ROI may not be possible, as the fact that the utility and payment functions are in general not continuous means that the necessary derivative may not exist. They thus approximate the marginal return on investment by measuring value per unit payment, also known as *bang-per-buck*, which for quasi-linear utilities is simply $r_j = v_j/p_j$ where v_j is the advertiser's value for the j-th auction. Their bidding strategy begins with a target R and chooses bids so that $r_j = R$ for all auctions j in the first day. The advertiser then proceeds with a multiplicative weights-like algorithm, raising R if the advertiser has not used all her budget by the end of the day, and lowering R otherwise.

As this bidding strategy, known as the *ROI heuristic*, was developed by considering optimal bidding only with respect to a single advertiser, the effect of the interaction of multiple advertisers bidding according to it is not immediately clear. Borgs et al. [7] go on to prove that, under a first-price auction model where advertisers may randomly perturb their bids over the course of a day in order to share slots with other advertisers, a system with all players using the ROI heuristic will converge to the unique market equilibrium. They also present experimental evidence that the strategy converges in a second-price auction as well.

Chapter 4
GREEDY BIDDING STRATEGIES FOR KEYWORD AUCTIONS

4.1 Greedy Bidding Strategies

In the previous chapter we have described several natural equilibria for the repeated keyword auction. The existence of such equilibria, however comforting, does not necessarily translate into concrete bidding strategies for the participating advertisers. How should these players bid? Naturally, a player's main objective is to maximize his own expected utility over multiple rounds of the auction. However, without any real insight into the bidding strategies followed by the rest of the players, it is difficult for one player to make predictions about the future bids of other players and hence choose an optimal bidding strategy. Thus, a very natural approach is to assume that the immediate past is the best predictor of the future. This leads to a natural greedy-like bidding scheme where a player assumes that all the other bids will remain fixed in the next round. Under this assumption, the rational choice for a player j is to bid so as to win a slot s that maximizes his utility $u_j = \theta_s(v_j - p_s)$. This leads to the following definition.

Definition 4.1.1. A *greedy bidding strategy* for a player j is to choose a bid for the next round of a repeated keyword auction round so as to maximize her utility u_j, assuming the bids of all other players b_{-j} in the next round will remain fixed to their values in the previous round.

Given b_{-j}, denote by $p_{-j}(s)$ the payment player j would make if she bids so as to win slot s. Let s^* be the slot the greedy bidder j will target. Then if player j is greedy, she will bid $b' \in (p_{-j}(s^* - 1), p_{-j}(s^*))$. As b' is not fully specified, this defines a class of strategies that are distinguished by the choice of b' within the allowed range.

Since the advertisers participating in a keyword auction are often business competitors, one of the most common secondary objectives observed in practice, besides gaining the optimal slot, is the desire to "push" the prices paid by other advertisers higher. This is done by bidding in the high end of the range mentioned earlier. The inherent risk of this strategy is that a change in other players' bids could result in paying a higher price than expected. This naturally leads to the following bidding strategy.

Definition 4.1.2. Balanced Bidding

The *balanced* greedy strategy BB is the strategy for a player j that, given b_{-j}

- next targets the slot s_j^* which maximizes his utility (greedy bidding choice), that is,

$$s_j^* = \underset{s}{\operatorname{argmax}}\{\theta_s(v_j - p_{-j}(s))\}.$$

- chooses its bid b' by

$$b' = (1 - \gamma_{s_j^*})v_j + \gamma_{s_j^*} p_{-j}(s_j^*),$$

  where γ_i is defined to be θ_i/θ_{i-1}.

If s_j^* is the top slot, we choose $b' = (v_j + p_{-j}(1))/2$. We can thus deal with all slots uniformly by defining $\theta_0 = 2\theta_1$. Note that this choice of θ_0 is arbitrary; different values do not materially change our results.

Observe that in the definition b' is chosen so as to satisfy $\theta_{s_j^*}\big(v_j - p_{-j}(s_j^*)\big) = \theta_{s_j^*-1}\big(v_j - b'\big)$. That is, a player bids high enough to force the prices paid by her competitors to rise, but not so high that if one of her competitors were to just undercut her, she would mind getting a higher slot at a price just below her own bid of b'. Note that this rule is same as the local conditions which define a symmetric Nash equilibrium. It is also known in Edelman et al. [16] as a *locally envy-free* condition.

We also study a variant of the BB strategy, called RBB, according to which the players can only aim for their current slot or a slot of lower click-through rate than the one they currently have. The RBB strategy is designed so that it has the same unique fixed point as BB. However, by restricting the degree to which a player can be greedy, we will be able to show that even in the synchronous model, RBB always converges to the VCG equilibrium. This strategy also captures a model where the bids of all players are not publicly revealed.

Over the course of several rounds the payments made by a player reveal the bids of overbid competitors, but the player does not have any information about the bids of players who have outbid her. Hence a player is limited to targeting the slots at or below her current slot.

Definition 4.1.3. The *Restricted Balanced Bidding* (RBB) bidding strategy is the strategy where given b_{-j} from the previous round, player j

- targets the slot s_j^* which maximizes his utility among the slots with no higher click-through rate than his current slot s_j, that is,

$$s_j^* = \text{argmax}\{\theta_s(v_j - p_{-j}(s)) \; : \; s \geq s_j\}.$$

- chooses her bid b' for the next round as with the BB strategy: $b' = (1 - \gamma_{s_j^*})v_j + \gamma_{s_j^*} p_{-j}(s)$.

As before, this definition sets $\theta_0 = 2\theta_1$.

The two strategies are appealing for the following reason.

Theorem 4.1.4. *[16] If all players follow the* BB *strategy in an auction with all distinct θs, then the system has a unique fixed point. In this fixed point, the revenue of the auctioneer (and payments of each player) is equal to that of the VCG equilibrium. The equilibrium bids b_j^* of the players in the fixed point of* BB *satisfy the following equations:*

$$b_j^* = \begin{cases} v_j & \text{if } j \geq k+1 \text{ and} \\ \gamma_j b_{j+1}^* + (1-\gamma_j)v_j & \text{if } 2 \leq j \leq k. \end{cases} \tag{4.1}$$

The same set of bids is also a unique equilibrium when all players follow the RBB *strategy, under the same conditions.*

4.2 Convergence Properties of the BB Strategy

We study the convergence properties of the BB strategy in a repeated keyword auction under both the synchronous and asynchronous auction formats. We summarize the results of this section in the following theorem.

Theorem 4.2.1. *Consider a repeated keyword auction over k slots with click-through rates $\theta_1 > \cdots > \theta_k$ and n players with values $v_1, \ldots, v_n$ all using the* BB *strategy starting with arbitrary initial bids. Then the following hold.*

1. *A two-slot auction system always converges to its fixed point in both the synchronous and asynchronous models. The number of rounds until convergence in the synchronous model is $O(\log((v_2 - v_3)/v_3))$, where the constant depends on the click-through rates θ_1 and θ_2. The asynchronous model converges with the same bound on the expected number of rounds.*

2. *There exists a 3-slot auction system and set of initial bids that does not converge in the synchronous model.*

3. *There exists a 3-slot auction system, a set of initial bids, and an order in which all players update that does not converge in the asynchronous model.*

4. *In the synchronous model the* RBB *strategy always converges to its fixed point. The number of steps until convergence is 2^k times*
$$O\left(k + \frac{\log(1-\gamma^*)}{\log \gamma^*} + \log_{(1/\gamma^*)} \frac{v_1 - v_{k+1}}{\min_{1\le i\le k}(v_i - v_{i+1})}\right),$$
where $\gamma^* = \max_i \theta_i/\theta_{i-1}$.

5. *In the asynchronous model where players bid in random order, no matter how many slots there are, the system always converges to its fixed point.*

4.2.1 The Two-Slot Auction

Here we prove the first part of Theorem 4.2.1, the convergence of the two-slot auction in both the synchronous and asynchronous settings. We begin with a technical lemma.

Lemma 4.2.2. *At every round t such that $t > t_1 = 2 + \log_{\gamma^*}((1-\gamma^*)(v_2 - v_3)/v_3)$, where $\gamma^* = \max\{\theta_1/\theta_0, \theta_2/\theta_1\}$, we have:*

$$\begin{aligned} b_1, b_2 &> v_3, \\ b_3 &= v_3. \end{aligned}$$

Proof. Let b denote the third highest bid. By the second round, b can never be more than v_3. Suppose for some round that b is less than v_3. Consider a player i in $\{1,2,3\}$. In the next round, i will bid her value or target some slot $j \in \{1,2\}$ and bid $b' = (1-\gamma_j)v_i + \gamma_j p_j \geq (1-\gamma^*)v_3 + \gamma^* b = b + (1-\gamma^*)(v_3 - b)$. In either case,

$$(v_3 - b') \leq \gamma^*(v_3 - b).$$

Initially $v_3 - b \leq v_3$.

It takes at most $r = \log_{\gamma^*}((1-\gamma^*)(v_2 - v_3)/v_3)$ rounds before $v_3 - b < (1-\gamma^*)(v_2 - v_3)$. (Recall that $\gamma^* < 1$ as the θ_i's are decreasing.) In round $r+1$, bidders $i \in \{1,2\}$ will each bid either $v_i > v_3$ or target a slot $j \in \{1,2\}$ and bid

$$\begin{aligned} b' &= (1-\gamma_j)v_i + \gamma_j p_j \geq (1-\gamma^*)v_2 + \gamma^* b \\ &= b + (1-\gamma^*)(v_2 - b) > b + (1-\gamma^*)(v_2 - v_3) \\ &> v_3, \end{aligned}$$

hence in either case their bids are both above v_3. In round $r+2$, player 3 cannot profitably target either slot, and so will then bid v_3 while players 1 and 2 keep on bidding above v_3; this concludes the proof of the lemma. □

Once the conditions of Lemma 4.2.2 hold, the price of slot 2 is fixed at $p_2 = v_3$. Let $T_2 = b_2^* = (1 - \theta_2/\theta_1)v_2 + (\theta_2/\theta_1)p_2$ and $T_1 = (1 - \theta_2/\theta_1)v_1 + (\theta_2/\theta_1)p_2$ be the thresholds for each player to bid for the top slot. Then if the last bids of players 1 and 2 were b_1 and b_2, in the next round their bids will be:

$$b_1' = \begin{cases} T_1 & \text{if } b_2 > T_1, \\ (v_1 + b_2)/2 & \text{otherwise.} \end{cases} \tag{4.2}$$

$$b_2' = \begin{cases} T_2 & \text{if } b_1 > T_2, \\ (v_2 + b_1)/2 & \text{otherwise.} \end{cases}$$

Let $b_{\min} = \min(b_1, b_2)$ be the minimum of the two winning bids.

Lemma 4.2.3. *After at most t_2 rounds, we have $b_{min} \geq T_2$, where*

$$t_2 \leq t_1 + 2\frac{\theta_1 - \theta_2}{\theta_2}.$$

Proof. Assume that at time $t > t_1$, we have $b_{\text{min}} < T_2$. Then both players will target the top slot, so in the next round we have $b'_{\text{min}} \geq (v_2 + b_{\text{min}})/2$. This implies

$$b'_{\text{min}} - b_{\text{min}} \geq \frac{v_2 - b_{\text{min}}}{2} \geq \frac{v_2 - T_2}{2} = \frac{\theta_2}{\theta_1}\left(\frac{v_2 - p_2}{2}\right) = \frac{\theta_2}{\theta_1}\left(\frac{v_2 - v_3}{2}\right).$$

Define δ to be $(\theta_2/\theta_1) \cdot \big((v_2 - p_2)/2\big)$. Since at time t_1 we have $b_{\text{min}} \geq v_3$, we can bound the number of rounds until $b'_{\text{min}} \geq T_2$ by

$$\begin{aligned}\frac{T_2 - v_3}{\delta} &= \frac{(1 - \theta_2/\theta_1)v_2 + (\theta_2/\theta_1)p_2 - v_3}{\delta} \\ &= \frac{(1 - \theta_2/\theta_1)(v_2 - v_3)}{\delta} \\ &= 2\frac{\theta_1(1 - \theta_2/\theta_1)}{\theta_2} = 2\frac{\theta_1 - \theta_2}{\theta_2}\end{aligned}$$

□

Finally, since at round t_2 we have $b_{\text{min}} \geq T_2$, at time $t_2 + 1$ we will have $b'_2 = T_2 < T_1$, and therefore at time $t_2 + 2$ we will have $b''_2 = T_2 = b^*_2$ and $b''_1 = (v_1 + T_2)/2 = b^*_1$; we have reached equilibrium. This proves Part 1 of the Theorem in the synchronous model.

The proof of convergence in the asynchronous model is similar, but complicated by the random order of bidding. We first mimic Lemma 4.2.2 to bound the time for the top two players to outbid the third. In any round before this holds, if the player with the minimum bid b is chosen to bid, then as before in the proof of Lemma 4.2.2, we have that her new bid b' satisfies $(v_3 - b') \leq \gamma^*(v_3 - b)$. Similar reasoning holds for the second part of the lemma, when player 3 is first outbid. As the lowest bidder will be chosen to bid in an expected constant number of rounds, we have that the conditions of Lemma 4.2.2 must hold in the asynchronous model after $O(t_1)$ rounds in expectation.

Once this occurs, whenever the player with the smaller winning bid is chosen to bid, the bid increases towards T_2 as in the proof of Lemma 4.2.3. Hence in $O(t_2)$ rounds in expectation, the condition of Lemma 4.2.3 holds, which implies convergence.

4.2.2 Nonconvergence

We now show parts 2 and 3 of Theorem 4.2.1, detailing when the balanced bidding strategy will not converge. First, in the synchronous setting a system of players all following the BB

	Round		
Bidder value	1	2	3
161	120.3	140.5	120.3
160	120	140.2	120
159	119.7	139.7	119.7
100	100	100	100

$\theta_1 = 1 \quad \theta_2 = 2/3 \quad \theta_3 = (2/3)^2$

Figure 4.1: Nonconvergent synchronous bidding

strategy is not guaranteed to converge when there are more than 2 slots. We prove this by counterexample. Let there be three slots with $\theta_1 = 1$, $\theta_2 = 2/3$, $\theta_3 = (2/3)^2$ and four players with values 161, 160, 159 and 100. Let the initial bids of 120.3, 120, 119.7 and 100, respectively. Then the bidding evolves as in Figure 4.1; in particular it is not convergent.

There is one inactive bidder, with the lowest value, who will never be able to bid for a slot as the prices are all above his utility. In the first round, the remaining bidders all target the lowest slot. The bid from the highest-valued player is low enough so that for the next round, all three top bidders target the first slot. With these high bids, the price for any player for the top slot is too high when compared with the price of the last slot. Thus for the third round all three top players target the last slot, and the cycle continues.

In this example the click-through rates follow a simple geometric sequence, similar to those observed in practice [18]. Note that in this example the bids are well-behaved, in the sense that the bids are in the same order as the players' values. Hence such regularity is not sufficient for convergence. Finally, even though we used our convention of $\theta_0 = 2$, a similar example can be constructed where the players are cycling while targeting intermediate slots.

Turning to the asynchronous setting, we show that a random bid ordering is essential by showing a three-slot example that does not converge even when all players appear in a fixed bidding sequence. Consider an auction where the slot click-through rates are given by $\theta_1 = 1$, $\theta_2 = 0.1$ and $\theta_3 = 0.09$ and the bidding is described in Figure 4.2.

	Round						
Bidder value	1	2	3	4	5	6	7
102	19.2	80.8	80.8	80.8	19.2	19.2	19.2
101	19.1	19.1	90.9	90.9	90.9	90.9	19.1
100	59.6	59.6	59.6	95.45	95.45	95.45	59.6
10	10	10	10	10	10	10	10

$\theta_1 = 1 \quad \theta_2 = 0.1 \quad \theta_3 = 0.09$

Figure 4.2: Nonconvergent asynchronous bidding under a fixed bidding sequence

In this example, the players bid cyclically in order from highest value to lowest. The pattern of bidding is the same as with the example for the synchronous case, where players all target the first slot, overpricing it and causing all players to next drop down to bid for lower slots.

4.2.3 Synchronous Convergence of RBB

To prove part 4 of Theorem 4.2.1, we first bound the number of steps until the price of slot k and the set of players who will be allocated slots have converged. As before, we define $\gamma_i = \theta_i/\theta_{i-1}$.

Lemma 4.2.4. *Player p prefers to target slot j rather than slot $j-1$ if and only if*

$$(1-\gamma_j)v_p + \gamma_j p_j < p_{j-1}.$$

Proof. By the balanced bidding rule, p prefers slot j rather than $j-1$ if and only if her utility in slot j is more than that in slot $j-1$, that is, $\theta_j(v_p - p_j) > \theta_{j-1}(v_p - p_{j-1})$. Rearranging gives the lemma. □

We next show a lemma analogous to the first part of the two slot proof, where bidding rises to a point where all but the k highest-valued players are priced out of a slot.

Lemma 4.2.5. *At every round t such that $t > t_1 = 2 + \log_{\gamma^*}((1-\gamma^*)(v_k - v_{k+1})/v_{k+1})$, where $\gamma^* = \max_{i>0} \theta_i/\theta_{i-1}$, we have:*

$$\begin{aligned} b_i &> v_{k+1} && \forall i \leq k, \\ b_i &= v_i && \forall i \geq k+1. \end{aligned}$$

Proof. Let b denote the $(k+1)$st highest bid. By definition, b can never be more than v_{k+1}. Suppose for some round that b is less than v_{k+1}. Take any player i in $\{1, 2, \ldots, k+1\}$. In the next round, i will either bid her value or target some slot $j \in \{1, \ldots k\}$ and bid $b_i' = (1-\gamma_j)v_i + \gamma_j p_j \geq (1-\gamma^*)v_{k+1} + \gamma^* b \geq b + (1-\gamma^*)(v_{k+1} - b)$. In either case,

$$(v_{k+1} - b_i') \leq \gamma^*(v_{k+1} - b).$$

As $b \geq 0$, we begin with $v_{k+1} - b \leq v_{k+1}$. The equation above says that each round this difference shrinks by a factor of at least γ^*, so that in $r \leq \log_{\gamma^*}((1-\gamma^*)(v_k - v_{k+1})/v_{k+1})$ rounds, we have $v_{k+1} - b < (1-\gamma^*)(v_k - v_{k+1})$. In round $r+1$, bidders $i \in \{1, \ldots k\}$ will each bid either $v_i > v_{k+1}$ or bid at least

$$\begin{aligned} b_i' &= (1-\gamma_j)v_i + \gamma_j p_j \geq (1-\gamma^*)v_k + \gamma^* b \\ &\geq b + (1-\gamma^*)(v_k - b) > b + (1-\gamma^*)(v_k - v_{k+1}) \\ &> v_{k+1}, \end{aligned}$$

hence in both cases their bids are above v_{k+1}. In round $r+2$, player $k+1$ will then bid v_{k+1} while players $1, 2, \ldots, k$ continue to bid above v_{k+1}; the other players don't get a slot, so they bid their value, which concludes the lemma. □

We now need to prove that the allocation of the k slots to these k players converges to a fixed point.

At any time, for any $i \in [1, k]$, consider the players allocated slots $[i+1, k]$. They are called *stable* if their bids and prices satisfy Equation (4.1), that is, the allocation is in order of decreasing values, and if $\pi(j)$ is the player currently allocated slot j, then the last bids of those players satisfied:

$$b_{\pi(j)} = \gamma_j b_{\pi(j+1)} + (1-\gamma_j)v_{\pi(j)},$$

for every $j \in [i+1, k]$. Note that a stable set may consist of a non-contiguous set of values. For example, we might have four players with values $v_1 > \cdots > v_4$ competing for three slots. If player 2 bids high, winning the top slot, players 1, 3 and 4 could form a stable set over the last two slots if they have bid appropriately.

If the players allocated all slots $[1, k]$ form a stable set, then we have reached the fixed point of the RBB strategy.

Assume that the current setting is not (yet) a fixed point of the RBB strategy. Let A be the maximum stable set, with associated $i \geq 2$, and let B be the set of players in slots $[1, i]$. Let $b_{\min}$ denote the minimum bid from players of B.

We define a partial order over sets of players which will be our measure of progress. We say that $A' \succ A$ if either $A \subset A'$ and $A \neq A'$, or if the player of minimum value in the symmetric difference of A and A' belongs to A'. This corresponds to a lexicographic ordering of the sets represented as bit-vectors ordered from minimum value. For example, if $v_6 < \cdots < v_1$, $\{6, 5, 3\} \succ \{6, 5, 2, 1\}$ as $110100 \succ 110011$ lexicographically.

In the next round, observe that, following the RBB strategy, players in A still bid in the same way as before. Let $b'_{\min}$ be the new minimum bid from players of B, and p be the player of B whose bid is $b'_{\min}$. There are three cases to consider.

1. p bids below p_i. Let $j \in [i+1, k]$ be the slot which was targeted by p. By definition, p prefers slot j to slot $j-1$, and so, by Lemma 4.2.4, the bid of p is less than p_{j-1}. By definition, the bid is $(1-\gamma_j)v_p + \gamma_j p_j > p_j$, thus it falls in the interval (p_{j-1}, p_j) and p will be allocated slot j. Recall that $\pi(j) \in A$ denotes the player who was in slot j in the previous round. Since set A is stable, by definition we have $p_{j-1} = (1-\gamma_j)v_{\pi(j)} + \gamma_j p_j$. Since this is greater than the bid of p, it follows that $v_{\pi(j)} > v_p$. Moreover, since p prefers slot j to slot $j+1$, by Lemma 4.2.4 again, we must have $(1-\gamma_{j+1})v_p + \gamma_{j+1}p_{j+1} > p_j = (1-\gamma_{j+1})v_{\pi(j+1)} + \gamma_{j+1}p_{j+1}$, and so $v_p > v_{\pi(j+1)}$. Thus $A' = \{p' \in A : v_{p'} < v_p\} \cup \{p\}$ is a stable set, and $A' \succ A$.

2. p targeted slot i. Then p is allocated slot i, $A' = A \cup \{p\}$ is a stable set, and $A' \succ A$.

3. p targeted some slot $j \leq i-1$. Then A is still a stable set, and $b_{\min} = p_{i-1}$ has increased: $b'_{\min} = (1-\gamma_j)v_p + \gamma_j p_j \geq b_{\min} + (1-\gamma^*)(v_p - b_{\min})$.

We will prove that Case 3 can only happen a bounded number x times before Case 1 or 2 must occur, where x depends on the θ_j's and the v_j's but not on the bids. Thus, the maximum stable set must change at least once every x rounds, and when that happens, it is replaced by a set which is larger in the $\succ$ ordering. This implies that the system converges to a fixed point and that the number of rounds until convergence is bounded by $2^k(x+1)$, hence the Theorem.

First, a useful technical lemma.

Lemma 4.2.6. *Let $\epsilon = (1/2)\theta_k(1-\gamma^*)\min_{q\neq q'}|v_q - v_{q'}|/\theta_1$. If $p_{i-1} > v_p - \epsilon$ and $v_p > p_i$, then player p prefers slot i to any slot $j < i$.*

Proof. From player p's viewpoint, the utility of slot i is $\theta_i(v_p - p_i)$, the utility of slot $j < i$ is $\theta_j(v_p - p_j) < \theta_j(v_p - p_{i-1})$, and the ratio is

$$\frac{\theta_j(v_p - p_{i-1})}{\theta_i(v_p - p_i)} \leq \epsilon\frac{\theta_j}{\theta_i(v_p - p_i)} \leq \epsilon\frac{\theta_1}{\theta_k(v_p - p_i)}.$$

Now,

$$\begin{aligned} v_p - p_i &= v_p - ((1-\gamma_{i+1})v_{\pi(i+1)} + \gamma_{i+1}p_{i+1}) \\ &= (1-\gamma_{i+1})(v_p - v_{\pi(i+1)}) + \gamma_{i+1}(v_p - p_{i+1}), \end{aligned}$$

which is at least $(1-\gamma^*)\min_{q\neq q'}|v_q - v_{q'}|$. Plugging this into the previous expression proves the Lemma. $\square$

Now, let $x = \log_{1/\gamma^*}((v_1 - v_{k+1})/\epsilon)$. Assume that Case 3 happens for x consecutive rounds. Let $p_{\min}$ be the player in B whose value is minimum and $v_{\min}$ be its value. Let $b^{(t)}_{\min}$ be the minimum bid of players in B after t rounds, $0 \leq t \leq x$. If $p \in B$ is the player defining the minimum bid in round $t+1$, we have:

$$b^{(t+1)}_{\min} \geq (1-\gamma^*)v_p + \gamma^* b^{(t)}_{\min} \geq (1-\gamma^*)v_{\min} + \gamma^* b^{(t)}_{\min}.$$

After x rounds, we get $b^{(x)}_{\min} \geq v_{\min} - (\gamma^*)^x(v_{\min} - b^{(0)}_{\min})$, hence $b^{(x)}_{\min} \geq v_{\min} - (\gamma^*)^x(v_i - v_{k+1})$. Plugging in the value of x yields $b^{(x)}_{\min} \geq v_{\min} - \epsilon$. From Lemma 4.2.6, we know that $p_{\min}$

prefers slot i to any slot $j < i$. In the next round $p_{\min}$ targets slot i and has to be the minimum bidder from B, therefore we are now in Case 2. Thus there are at most x occurrences of Case 3 between any two occurrences of Case 1 or Case 2, and convergence is proved.

4.2.4 Asynchronous Convergence

Finally, we argue part 5 of Theorem 4.2.1, that there exists an asynchronous bidding sequence for the players that converges to the fixed point. At time t, we say that a player p is *activated* if p is the player who updates his bid while the other bidders repeat their previous bids.

Lemma 4.2.7. *There is an integer T depending on n, k, $\theta_1, \ldots, \theta_k$, and $v_1, \ldots, v_n$ such that for every starting configuration, there exists a sequence of at most T player activations, such that the resulting configuration is a fixed point.*

Proof. Consider an arbitrary starting configuration. To construct the sequence, the idea is to emulate the proof of the synchronous convergence of RBB shown in Section 4.2.3.

First, we emulate the proof of Lemma 4.2.5 as follows. Repeatedly activate all the players $1, \ldots, k+1$ until each of them bids at least v_{k+1}. In other words, if one of the players in $[1, k]$ has a current bid $< v_{k+1}$, we activate all the players $1, \ldots, k+1$ one at a time. By Lemma 4.2.5 after a total of at most $(k+1)t_1$ activations all players $[1, k+1]$ bid at least v_{k+1}, where t_1 is defined as in Lemma 4.2.5. Now if there are some players in $[k+2, n]$ who are not bidding their value, activate one of these players so that she will now bid her value. The conclusion of Lemma 4.2.5, that

$$\begin{aligned} b_i &> v_{k+1}, && \text{for all } i \leq k, \text{ and} \\ b_i &= v_i, && \text{for all } i \geq k+1 \end{aligned}$$

will hold after at most $(k+1)t_1 + n - k + 1$ activations. From that point on until the end of the sequence, players $[k+1, n]$ will not be activated again. The rest of the sequence is partitioned into phases, corresponding to stable sets, defined in Section 4.2.3. A stable set stays the same throughout a phase. To define the sequence during a phase, let A be the current stable set, $[i+1, k]$ be the slots occupied by the players of A, and B be the set of players occupying slots $[1, i]$. Let $p_{\min}$ be the player in B whose value is minimum, $v_{\min}$ be its value, and p_{i-1} be the price of slot $i-1$.

Consider the three cases enumerated in the Section 4.2.3. We repeatedly activate the player currently in slot i until either Case 1 or Case 2 occurs, or $p_{i-1} > v_{\min} - \epsilon$, where ϵ is defined as in Lemma 4.2.6. We then activate player $p_{\min}$. At this point, Case 1 or Case 2 must occur, a new stable set is defined, and the phase ends. This completes the definition of the sequence.

We now need to bound the length of the constructed sequence of player activations, independently of the starting configuration and bids.

In the initial part players $[1, k+1]$ are activated at most t_1 times and players $[k+2, n]$ are activated at most once. Thus this part has length at most $(k+1)t_1 + n - k + 1$.

It is also easy to see, as done in Section 4.2.3, that the stable set in the next phase will be larger in the $\succ$ order, hence there will be at most 2^k phases, ending with a fixed point.

To bound the length of a phase, it is again easy to see that during a phase, $b_{\min}$ can only increase; moreover, when a player p is activated for several times, the new value of $b'_{\min}$ must be larger than the value of $b_{\min}$ during the time of his previous activation by at least $b'_{\min} \geq b_{\min} + (1-\gamma^*)(v_p - b_{\min})$, precisely the same inequality as in the analysis of Case 3 in Section 4.2.3. Since we must have repetitions of choices of player at least once every k activations, if we define x as in Section 4.2.3, it follows that after at most xk activations we are ready to activate $p_{\min}$ and end the phase.

Hence the length T of the sequence is at most $n + 3k$ plus k times the bound x from the end of Section 4.2.3, which is independent of the starting configuration.

□

With this Lemma, it is easy to complete the proof of Theorem 4.2.1: In a random sequence, at every step we have probability at least $1/n$ to choose the next activation as in Lemma 4.2.7, hence the sequence will occur after about n^T steps on average, and in any case, it will occur with probability 1 after a finite time. This proves convergence.

4.2.5 Notes

The example above, showing that for 3 slots BB doesn't converge in the synchronous model, is not anomalous. In Section 4.3, we show experimentally that in a significant fraction of instances, BB does not converge in the synchronous model.

The bound we give on the time to convergence of BB in the random asynchronous model is extremely loose. In Section 4.3, we also show results of simulations demonstrating that convergence is actually quite fast, certainly no more than a polynomial in n, the number of bidders.

4.3 Empirical Evaluation

In Section 4.2, we showed how simple greedy bidding strategies can lead to the VCG equilibrium. Is this the "right" equilibrium for bidders to be shooting for? How desirable is this outcome for the bidders and for the search engine? In this section, we study these questions empirically.

We first compare the search engine revenue in equilibrium to the VCG revenue. We then compare these benchmarks with two alternative greedy bidding strategies, both from the perspective of search engine revenue and from the perspective of bidder utility.

4.3.1 Experimental Setup

Except as otherwise noted, in our experiments we used three slots and four players. Following the study of Feng et al. [18], we choose the click-through rates as a geometrically decreasing sequence by $\theta_i = \delta^{i-1}$ for some value of δ between 0 and 1, and plot our results as a function of δ. We take the average of 150 instances, where for each instance, the values of the four players are each independently chosen from a normal distribution with mean 500 and deviation 200.

To analyze a bidding strategy, for each of the 150 instances, the simulated strategy is run for 75 rounds from starting bids of 1 for all players, except for AB (defined later), which is run for 150 rounds from starting bids equal to the minimum player value. Some experiments, such as those comparing equilibrium revenue with VCG revenue, are not affected by the

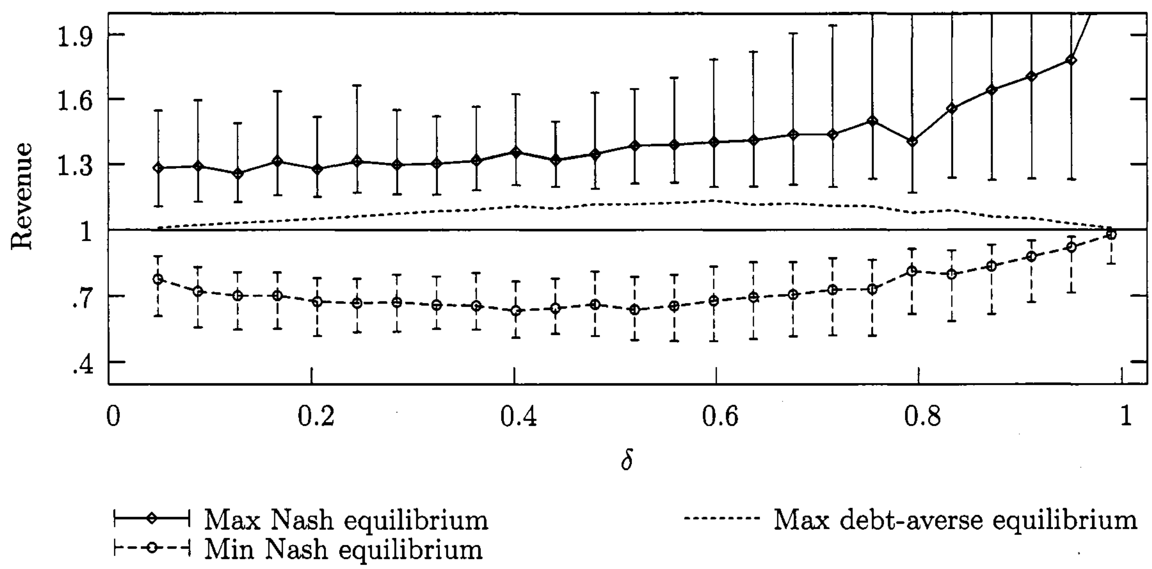

Figure 4.3: Comparing revenue from GSP equilibria with the VCG revenue (VCG=1)

choice of asynchronous or synchronous settings. Others, such as BB convergence, do depend on the setting, and in such cases we will discuss the differences.

4.3.2 Revenue in GSP Equilibria

How much revenue does GSP bring to search engines, compared to the revenue generated by the VCG mechanism?

Figure 4.3 plots the average, over instances, of the ratio between the revenue in some Nash equilibrium of GSP and the VCG revenue. Since GSP has many different Nash equilibria, we consider three extreme points: the maximum revenue Nash equilibrium, the minimum revenue Nash equilibrium, and the maximum revenue Nash equilibrium when the bidders are *debt-averse*, i.e. never bid above their value (see Definition 2.2.3). Kitts et al. [32] show evidence that bidders in keyword auctions usually follow a debt-averse bidding strategy.

The various Nash equilibria were computed by solving linear programs using the constraints (3.1) and (3.2) from Section 3.3, for each possible player ranking. While this method

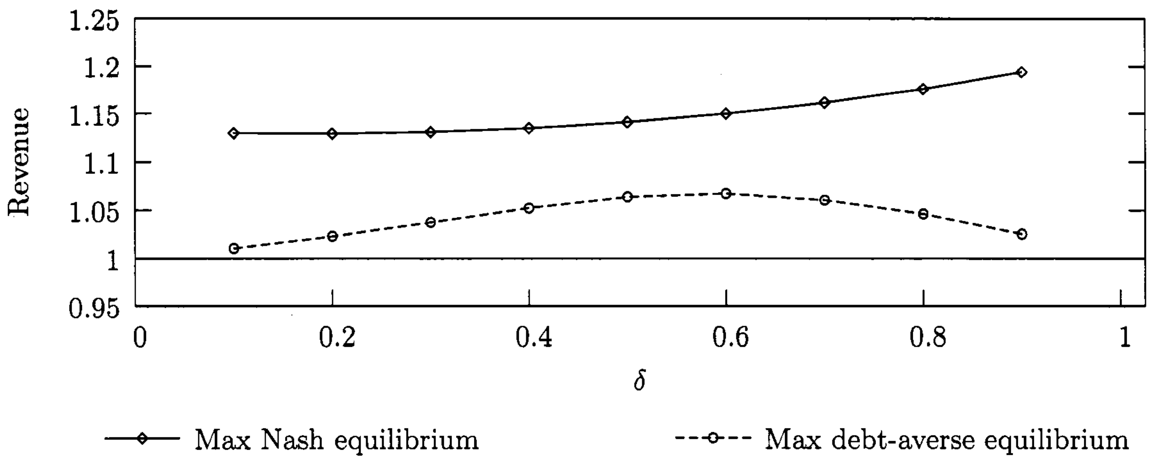

Figure 4.4: Revenue from six players competing for 5 slots (VCG=1)

takes exponential time in the number of players, the computations with only four players are quite manageable.

From these experiments, we conclude the following: *the equilibrium revenue of GSP is within* $\pm 40\%$ *of the VCG revenue, unless* δ *is close to 1 and bidders are willing to bid above their value.* This should be contrasted with Part 1 of Theorem 4.3.1, where we exhibit an example of a GSP equilibrium with revenue much smaller than the VCG revenue: our simulations indicate that such an example is an anomaly.

We also see the following interesting behavior: *If bidders are debt-averse, then in equilibrium GSP generates at most* 15% *more revenue than VCG.* Since being debt-averse seems a very likely behavior on the part of rational bidders, this indicates that choosing GSP over VCG may not have a huge impact on revenue. The influence of debt-averse behavior on the maximum revenue is rigorously justified in Part 3 of Theorem 4.3.1.

Figure 4.4 shows maximum Nash and debt-averse equilibrium for six players and five slots. Revenues are generally closer to the VCG revenue than that from four players. Qualitative trends are the same between the six and four player instances, suggesting that our observations are independent of the number of slots used.

player	value	bid	slot won	GSP pmt.	VCG pmt.
1	$x+1$	$x+1$	1	2	$x+\frac{1}{2}$
2	x	2	2	$\frac{1}{2}$	$\frac{1}{2}$
3	1	1	–	0	0

$\theta_1 = 1 \qquad \theta_2 = 1/2$

Figure 4.5: A GSP equilibrium much smaller than VCG

4.3.3 Theoretical Revenue Bounds

In order to evaluate the experimental results presented above, we compare them with a theoretical analysis of GSP revenue in keyword auctions. Unlike single-item auction revenue equivalence and matroid procurement auctions, there is a multitude of Nash equilibria and they may differ considerably from the VCG revenue.

Theorem 4.3.1. *1. For every $K > 0$, there exists a keyword auction and a Nash equilibrium of the GSP mechanism whose revenue is at most $1/K$ times the revenue of the VCG mechanism. Moreover, the equilibrium is debt-averse, that is, every bidder i bids $b_i \leq v_i$.*

2. For every $K > 0$, there exists a keyword auction and a Nash equilibrium of the GSP mechanism whose revenue is at least K times the revenue of the VCG mechanism.

3. If bidders are debt-averse, then for every keyword auction and Nash equilibrium of the GSP mechanism, the revenue is at most α^* times the VCG revenue, where $\alpha^* = \max_i \theta_i/(\theta_i - \theta_{i+1})$.

Proof. Part 1: Figure 4.5 contains an example showing that the revenue to the auctioneer from a GSP equilibrium may be arbitrarily smaller than the revenue from the VCG equilibrium. The example is for a two-slot auction, where $\theta_1 = 1$ and $\theta_2 = 1/2$. The GSP revenue is $2\theta_1 + (1/2)\theta_2 = 2.25$ whereas the VCG revenue is $(x + 1/2)\theta_1 + (1/2)\theta_2 = x + (3/4)$.

player	value	bid	slot won	GSP pmt.	VCG pmt.
1	x	x	1	$(x+2)/2$	$\frac{3}{2}$
2	2	$\frac{(x+2)}{2}$	2	$\frac{1}{2}$	$\frac{1}{2}$
3	1	1	–	0	0

$\theta_1 = 1 \qquad \theta_2 = 1/2$

Figure 4.6: A GSP equilibrium much larger than VCG

Part 2: Figure 4.6 shows an example showing that the revenue to the auctioneer from a GSP equilibrium may be arbitrarily larger than the revenue from the VCG equilibrium. The example is again for a two-slot auction with $\theta_1 = 1$ and $\theta_2 = 1/2$.

These examples rely on player values that are arbitrarily separated. The second example was constructed using the following expression due to Varian [45], which gives the payments that achieve the maximum Nash revenue.

$$\theta_i p_i = \sum_{j>i} v_{j-1}(\theta_j - \theta_{j+1}) \tag{4.3}$$

Compare this with the VCG payments:

$$\theta_i p_i = \sum_{j>i} v_j(\theta_j - \theta_{j+1}) \tag{4.4}$$

Observe that if the values are close to each other, the maximum revenue from a Nash equilibrium is close to the VCG revenue. For example, if $v_{i+1} \geq \alpha v_i$, then the payments from (4.3) are at most a factor of $1/\alpha$ from the payments of (4.4).

Note that the second example showing a Nash revenue larger than the VCG revenue has the second player bidding much more than his value and so is not debt-averse. Under the more realistic debt-averse assumption we have the third part of this theorem.

Part 3: Let R^M be the maximum debt-averse Nash revenue. Note that the payment of i, p_i^M, is at most v_{i+1}. If $p_i^M > v_{i+1}$, as all players are bidding at most their values, there would not be enough winners to fill the top i slots. Thus the maximum risk-free Nash equilibrium revenue is $R^M \geq \sum_{1\leq i\leq k} \theta_i v_{i+1}$. On the other hand, the VCG payment

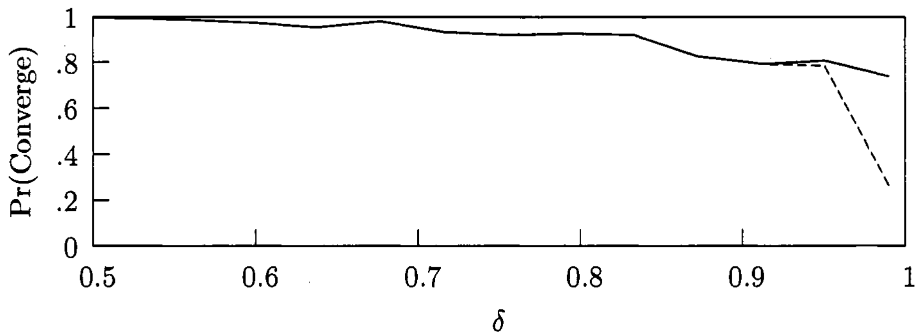

Figure 4.7: Synchronous convergence fraction of BB

of i is $\theta_i p_i^{\mathrm{VCG}} \geq (\theta_i - \theta_{i+1})v_{i+1}$. Thus the total VCG revenue is at most $\sum \theta_i p_i^{\mathrm{VCG}} \geq \sum(\theta_i - \theta_{i+1})v_{i+1} = \sum(\theta_i - \theta_{i+1})\theta_i v_{i+1}/\theta_i \geq R^M/\alpha^*$ as required. □

4.3.4 Balanced Bidding Strategies: BB and RBB

We next demonstrate the frequency of non-convergence of BB in the synchronous setting shown in Theorem 4.2.1 part 2. Figure 4.7 plots the probability that BB converges. Note that when $\delta > .95$, it takes much longer for the system to reach a fixed point or a cycle, and our simulations did not always enable us to reach that point, so that we plot two curves representing lower and upper bounds to the convergence probability. This figure indicates that the non-convergent example in Theorem 4.2.1 is not pathological and is a practically observed phenomenon which occurs for a significant fraction of instances.

On the other hand, while we have proved that the BB strategy always converges in an asynchronous setting, our theoretical bounds on the rate of convergence are quite weak. Figure 4.8 plots the average number of bids per player until convergence in this setting, as a function of the number of players. All runs used a conservative $\delta = 0.9$. As can be seen in the figure, the average number of bids per player increases linearly in the number of players. This suggests a $O(n^2)$ overall rate of convergence.

One contribution of this dissertation is to suggest that in a synchronous setting, bidders should use RBB rather than BB, since it gets to the same fixed point and has the advantage

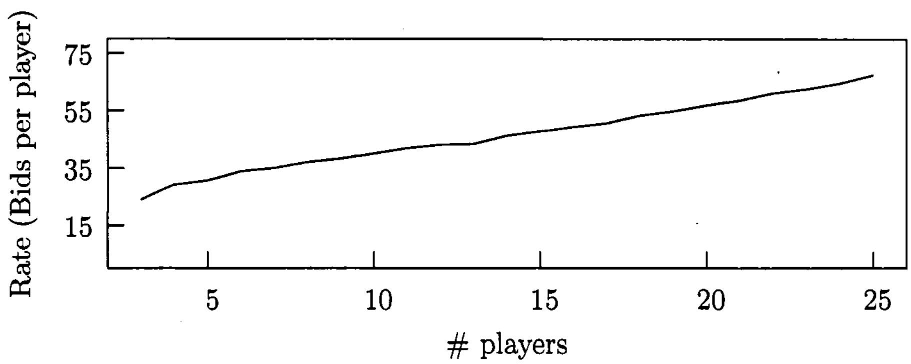

Figure 4.8: Asynchronous convergence rate of BB

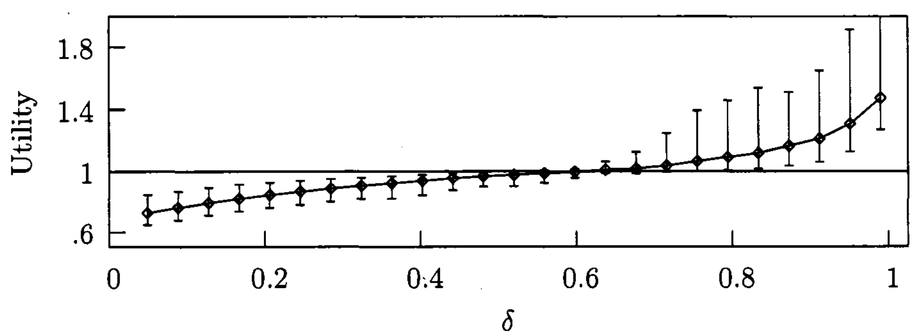

Figure 4.9: RBB vs. BB player utility (BB=1)

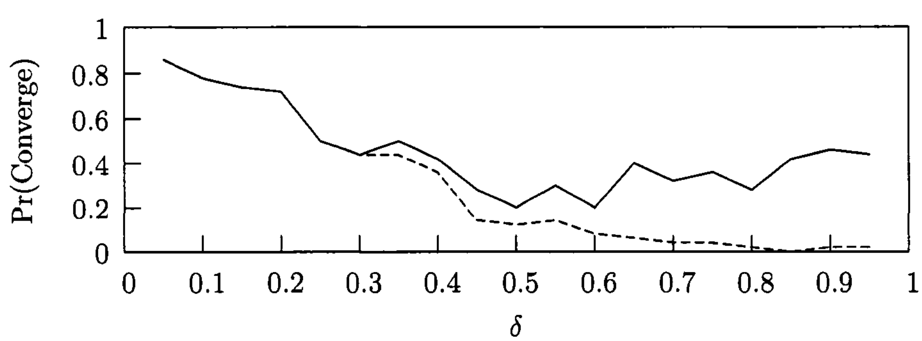

Figure 4.10: Convergence of CB

of always converging. However, players might wonder about the effect on their utility. To study this, we take one special player p and fix his value at 500, while the values of the other players are each independently chosen from a normal distribution with mean 500 and deviation 200. We then plot, as a function of δ, the ratio between player p's utility under RBB and his utility under BB.

We see in Figure 4.9, that when $\delta < .6$, player p's utility in RBB is only about 90% of the utility he would get in BB. Interestingly, as δ approaches 1, this is reversed: when δ is close to 1, RBB yields better utility than BB. We expect this to be due to the frequent non-convergence of BB in this regime. It is claimed [18] that in practice the click-through rates are fit well by a geometric sequence with δ about .7, and in that range RBB and BB yield about the same revenue.

Note that our result is robust: we have observed that changing the value of the player under focus from 500 to 400 or to 600 (so as to make him an either relatively lower valued player or relatively higher valued player) does not change the graph significantly.

4.3.5 Other Greedy Bidding Strategies: CB and AB

So far we have focused on BB and RBB, but it is conceivable that a bidder could obtain a higher revenue by following some other strategy. In this section we examine two other greedy bidding strategies, one that is often considered and another that is its natural complement. In the following we let $\varepsilon_{\text{price}}$ be a suitably small bid increment, for example 1¢.

Competitor Busting

A popular bidding strategy used in practice is known as *competitor busting* [47]. This strategy is also referred to as anti-social or vindictive bidding [8, 50], and may be used by as many as 40% of the bidders on Yahoo! [50]. The idea is that a player bids as high as possible while retaining her desired slot, in order to make competitors pay as much as possible and exhaust their advertising resources.

Definition 4.3.2. The *competitor-busting* greedy bidding strategy (CB) is the strategy for a player j that, given b_{-j}

- next targets the slot s_j^* which maximizes her utility (greedy bidding choice), that is,

$$s_j^* = \text{argmax}\{\theta_s(v_j - p_{-j}(s)) : s \geq s_j\}.$$

- chooses her next bid as $b' = \min\{v_j, p_{-j}(s_{j-1}^*) - \varepsilon_{\text{price}}\}$.

Convergence issues for CB are much more serious than for BB: in general, the CB strategy does not have a fixed point. Indeed, a player p in slot i will only have non-negligible utility if the player in slot $i+1$ cannot raise her bid any further, because it is equal to her value. In this case, unless the situation is degenerate, p will prefer to move to a different slot, which will in turn lead other players to want different slots.

Thus, the only fixed point for CB is when all players are bidding their values and those bids happen to be a Nash equilibrium for the GSP strategy. Bids equal to the values form a Nash equilibrium if and only if $\theta_i(v_i - v_{i+1}) \geq \theta_j(v_i - v_{j+1})$ for each i and $j > i$; for a random instance, these constraints are satisfied with significant probability when δ is close to 0.

Indeed, Figure 4.10 confirms this finding. As in Figure 4.7, we plot two curves, which are lower and upper bounds on the probability of convergence. The auction was run in the asynchronous setting. The space between the two curves corresponds to instances that have neither converged nor begun to cycle; because cycles are difficult to detect in the asynchronous setting the true convergence rate is likely much closer to the lower curve. We see that CB frequently converges when δ is small, and rarely converges when δ is large. In the synchronous setting the results are essentially identical, which shows that there is no cycling dependent on the format the way there is with BB.

Given that CB is a well-known strategy, the auctioneer might wonder how much revenue he obtains if every player followed the CB strategy. Figure 4.11 compares the revenue obtained from running GSP with the players following CB, with the revenue which would have been obtained from running VCG (with VCG revenue normalized to 1). This is good news for the auctioneer: if the players follow the CB strategy, then the auctioneer's revenue is higher than the VCG revenue. In the $\delta \approx 0.7$ setting seen in practice [18], auctioneer revenue is 10% higher than VCG revenue; as the CTR ratio for the slots becomes more

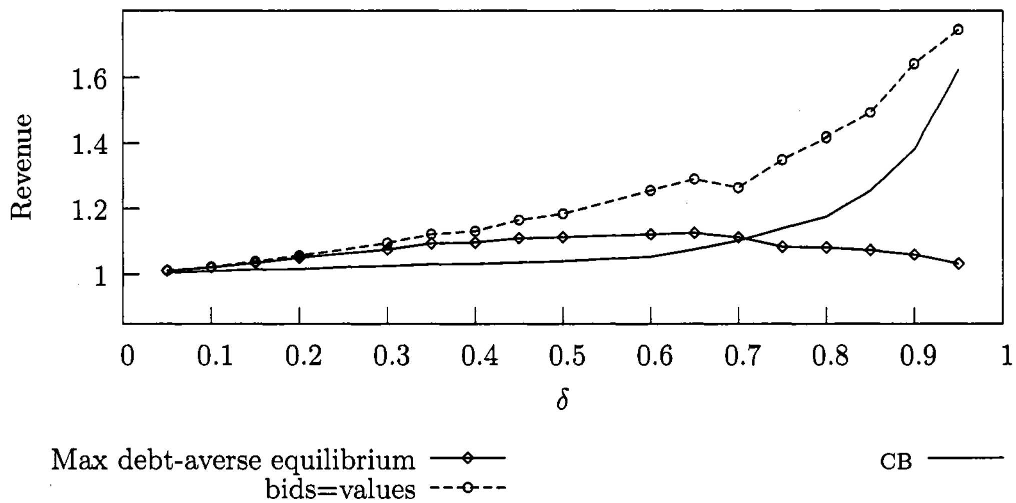

Figure 4.11: Comparing revenue from CB with the VCG revenue (VCG=1)

similar the revenue gain approaches 60%. Accordingly, player utility goes down. In fact, in this regime the revenue from CB is more than that obtained by any debt-averse equilibrium, underscoring the questionable benefits to the players of the CB strategy.

The figure was computed in the asynchronous setting; the revenue results are virtually identical under the synchronous format. This suggests that the differences between the two formats are not that large, and that the sensitivity of the BB strategy on the format might be anomalous.

Altruistic Bidding

Since competitor busting primarily benefits the auctioneer, the players might consider trying a completely different approach: altruism. A natural complement to the competitor busting strategy, where players try to hurt other bidders as much as possible, the altruistic strategy has players bidding as low as possible to win their desired slot.

Definition 4.3.3. The *altruistic* greedy strategy (AB), is the strategy for a player j that, given b_{-j}

- next targets the slot s_j^* which maximizes her utility (greedy bidding choice), that is,

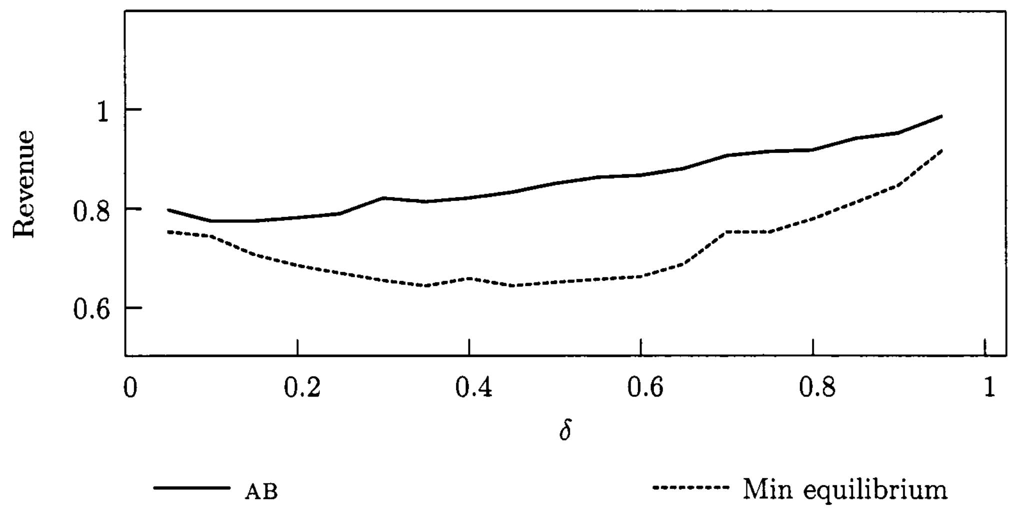

Figure 4.12: AB total revenue (VCG=1)

$s_j^* = \text{argmax}\{\theta_s(v_j - p_{-j}(s)) : s \geq s_j\}$.

- chooses her next bid as $b' = \min\{v_j, p_{-j}(s_j^*) + \varepsilon_{\text{price}}\}$. If there is no slot giving positive utility, bid v_j.

Convergence issues for AB are even more serious than for CB: the AB strategy has *no* fixed points when the θ_i's are separated.

The auctioneer might worry about how much revenue he would obtain if every player followed the AB strategy. As can be seen in Figure 4.12, this fear is only somewhat justified: AB *produces low revenue, but more than the minimum Nash equilibrium of GSP.* As with the CB revenue, these results from the asynchronous setting are very similar to those in the synchronous setting. It would be interesting to see if there are any auctions in practice where players bid in this way, or if a bidding strategy converging to the minimum Nash revenue exists.

4.4 Conclusions

In this chapter, we have shown that the natural balanced bidding strategy converges to a unique equilibrium. Using both theoretical analysis and simulations we compared the revenue from this equilibrium to various benchmarks: various extreme points from different flavors of Nash equilibria, as well as the revenue gained when using the VCG mechanism.

The convergence of the balanced bidding strategy along with the lower total search engine revenue than that from a competitor-busting strategy supports previous work, such as Edelman and Ostrovsky [15] and Asdemir [5], that advertisers gain in the long run when they do not focus solely on harming their competitors. A great challange will be to apply these results and intuitions to auctions that more accurately model practical situations, such as bidding for multiple keywords under budget constraints.

Chapter 5

PROCUREMENT AUCTIONS

Consider the problem a content delivery agency faces when attempting to provide broadcasting services or point-to-point communication over a network like the Internet. In order to broadcast to the entire network, the provider needs to acquire the usage rights to links forming a tree in the network. To achieve a desired level of fault-tolerance or bandwidth, the provider may need or want multiple such trees. We assume that the links in the network are owned by independent service providers (as is currently the case in the Internet) and that the content delivery agency must independently contract with each one (unlike the current Internet). We assume that the content provider does not require a specific amount of replication, but instead has a benefit that is a function of the degree of replication. This content delivery agency would like to maximize its own interest, the difference between its benefit for the links procured and any payments necessary to procure the links. Thus, generally, we consider the mechanism design problem this buyer faces when attempting to procure multiple disjoint sets from a set system to maximize her profit.

Another example is found in path auctions, where independent selfish agents each own an edge of a publicly known network. An agent e can transmit data along her link at some cost c_e known only to her, and the buyer wants to hire a team of agents who form a path between two given nodes s and t. Each agent submits a bid, and based on these bids, the buyer chooses a path and pays each selected agent e some amount p_e, according to the rules of the auction. The aim of each agent is to maximize her utility, the difference $p_e - c(e)$. The aim of the buyer is to minimize the total payment made. As with spanning trees, this problem can be generalized to a multi-unit case where the buyer desires several disjoint s–t paths.

These two problems are unified in a general framework by defining a family of *feasible sets* over a group of independent agents. The problem of the buyer is to procure any feasible set at as small a cost as possible. This procurement problem unifies the objectives of two areas of algorithmic mechanism design: frugality in procurement mechanisms, and profit maximization in multi-unit auctions. These areas share a common objective as maximizing profit is equivalent to minimizing cost; however the nature of these two optimization problems is quite different. The literature on procurement assumes that the buyer desires precisely one feasible set (e.g., one path or one spanning tree). There are close analogies between this setting and the problem of auctioning a single item. For both, it is not possible to do much better in the worst case than the VCG mechanism. This contrasts starkly with worst-case results for multi-unit auctions, which we describe below, where the flexibility of being able to choose the number of units to sell allows for a seller to obtain a profit significantly higher than the profit of VCG. Indeed, in our procurement setting, the most natural application of the VCG mechanism is not guaranteed to have good profit and there are instances where VCG runs a deficit.

In the very special case in path auctions where the network is simply a set of parallel links connecting s and t, the VCG mechanism reduces to simply choosing the cheapest edge and paying that edge the cost of the second cheapest edge. On the other hand, if paths can consist of multiple edges, as in the example of Figure 5.1, then not only the VCG mechanism but any truthful mechanism may overpay greatly to buy a single path, where this overpayment is measured relative to the *second cheapest path* [3, 17, 29]. For example, in Figure 5.1, the leftmost path (from Florida to Panama) will be chosen by the VCG mechanism and will result in payments of 2 to each of the six edges in the path, so that the total payment (of 12) is much more than the cost of the second cheapest path. It is possible that procuring multiple paths could lower the per-path cost. In our example, buying two paths has a per-path cost of only $9^1/_2$, and three paths are even cheaper, $8^2/_3$ per path. This raises the question as to whether, as in the digital good auction problem, the freedom to choose the number of paths procured can alleviate the necessarily high over-payments in the single-path procurement problem.

To formalize this setting, let $B(k)$ be a function specifying the buyer's value for procuring k paths. For example, this may reflect the resale value for k paths. Then the buyer's profit, if he purchases k paths at a total price of P_k, is $B(k)-P_k$. One class of problems we consider in

k	Benefit	VCG cost	Profit
1	14	12	2
2	28	19	9
3	42	26	16
4	56	60	-4

Figure 5.1: Multi-unit procurement

this paper is that of designing a truthful mechanism to achieve a target profit. For example, suppose the buyer of Figure 5.1 had a value of 14 per path (i.e., $B(k) = 14 \cdot k$) and a profit goal of 10. A prescient buyer could run the VCG mechanism specifically to procure three paths, and would make a profit of 16. The challenge is that absent foreknowledge, it is not clear how many paths to procure; furthermore, a truthful mechanism that determines the number k of paths to procure on-the-fly will not generally be able purchase them with the same payments as the VCG mechanism for k paths.

A mechanism that solves this decision problem ("Is it possible to get a profit of R?") is called a *profit extractor* [13]. The first contribution from this dissertation is to explore necessary and sufficient conditions on the structure of the procurement problem that ensure the existence of a truthful profit extractor. We prove that profit extraction based on VCG is truthful exactly on matroids.

Profit extraction is interesting in its own right, but it is also an important subroutine in the design of mechanisms for solving the corresponding optimization problem, the *profit maximization* problem. As in the case of the digital good auction problem, as well as classical optimization, a natural approach to solving an optimization problem is via reduction to a *decision problem.* In this setting the natural decision problem is profit extraction.

Our next contribution is to design a mechanism to maximize the buyer's profit. For path auctions on a graph with s and t connected by a set of parallel links (or equivalently, the procurement version of the digital goods auction problem), there are known truthful reductions from the approximate profit maximization problem to the profit extraction problem. One such reduction [19] first randomly samples some of the agents to come up with an estimate R of $\mathcal{OPT} = \max_k(B(k) - P_k)$, and then uses profit extraction on unsampled agents to try to extract a profit of R. We call an auction of this type a *random sampling profit extraction* auction. The success of this approach depends on the accuracy of the estimate $\mathcal{OPT}$ via random sampling, and on the existence of a truthful profit extractor.

To show the viability of random sampling as an estimate of the available profit, we use a theorem due to Karger [28]. This shows that if a matroid has k disjoint bases, and k is not too small, then a random sample of half the elements will have about $k/2$ bases. The challenge in using this result is to show that the VCG payments on a sample are about half of what they would be in the full set. This requires understanding, in some detail, the fine structure of optimal replacement for unions of disjoint independent sets in a matroid.

In this chapter we flesh out the background sketched above. We begin by describing the *digital goods* auction setting, where the decision problem of revenue extraction was first formulated. We then show the negative results on procurement over path auctions. This leads us to formalize the multi-unit procurement problem that we consider, and give basic facts and notation of procurement over matroids. Finally, we express the frugality results for matroids as revenue-equivalence results.

5.1 Digital Goods Auctions

5.1.1 The Digital Goods Setting

In the digital goods setting, an auctioneer has an unlimited supply of identical items. An example of this model is given by the problem of providing electronic content such as software or music. One option is to choose a single price for the auction. That is, the auctioneer elicits bids to maximize $k(p) \cdot p$, where $k(p)$ is the number of bidders whose bids are p or greater. A distinguishing feature of these auctions is that the efficient allocation is not interesting: it distributes the good to all players at the minimum value. If this value is zero, the auctioneer receives no revenue and so has no incentive to run the auction in the first place. Thus, the object is solely to maximize auctioneer revenue.

The n players each have a private value for receiving the item. We assume voluntary participation, so that the price offered to a player must be less than her value (although it may differ from her bid). For example, the auctioneer may simply allocate an item to all bidders, charging each bidder her bid. However, in this case the bidders will bid minimum values and the auctioneer will receive a minimum revenue. This is arbitrarily bad compared to the optimal auction, where each player is charged her value. Hence it seems appropriate to design a truthful mechanism that may charge players less than their bids in order to approximate an optimal revenue.

This problem has been extensively studied under competitive analysis [22, 21, 19]. We first develop a benchmark to be used for the competitive analysis, before discussing several truthful mechanisms with constant factor approximation to this benchmark. A generalization may be found in Fiat et al. [19], where the auctioneer has an unlimited supply of a set of heterogeneous goods, and the players are split into disjoint homogeneous sets where all players from the same set desire the same item. The results for general auctions are very similar to the digital goods case, and so are omitted in the interests of simplicity. The full generalization to unlimited supply combinatorial auctions is an important open problem.

5.1.2 Cooperative Games and Revenue Extraction

The mechanism developed for digital goods auctions requires a brief digression into another aspect of economic game theory, *cooperative games* and *cost sharing*. Because a digital good can be reproduced at no cost, one must devise a mechanism to sell the good that can convince additional buyers to purchase it even when the marginal cost of creating one more good is zero. In other words, the mechanism must divide up the revenue amongst the set of buyers while simultaneously determining what that set is.

Cost sharing for a cooperative game seeks to divide up the cost for a service over a set of players, each with a different utility for receiving the service. The cost of the service may vary with the subset of players served. Formally,[1] there is a set N of players and a cost function $c : 2^N \rightarrow \mathbb{R}$ which gives the cost of providing the service to each subset of players.

[1] For simplicity we present only the cooperative game with transferable utilities.

Given a cost-sharing scheme χ,

1. Initialize $S \leftarrow N$.

2. Repeat $S \leftarrow \{i \in S : v_i \geq \chi(i, S)\}$ until for all $i \in S$, $v_i \geq \chi(i, S)$.

3. Return the coalition S with payments $p_i = \chi(i, S)$ for all $i \in S$.

Figure 5.2: The group strategyproof mechanism CS_χ

Each player i in N has a value v_i that the player gains if she receives the service. The goal is to design a mechanism that, given a cost function, finds a subset of players, called a *coalition*, and a payment p_i for each player in the coalition, such that $p_i \leq v_i$, and the sum of payments equals the cost of the coalition. Given the final coalition, the computation of payments by the mechanism is specified by a *cost-sharing scheme* $\chi : N \times 2^N$, which takes a player and a coalition containing that player, and computes the player's payment.

A further requirement is to find a *group strategyproof* mechanism, where no coalition benefits by having any member bid other than her true value, a condition which parallels truthfulness for auctions. A central result from this area [36] says that if a cost-sharing scheme satisfies a certain property known as *cross-monotonicity*, then it can be used to define a group strategyproof mechanism.

Definition 5.1.1. A cost-sharing scheme χ is *cross-monotone* if for all $S, T \subset N$ and $i \in S$, $\chi(i, S) \geq \chi(i, S \cup T)$.

That is, a player's payment does not increase if the size of her coalition grows. The mechanism resulting from the cost sharing scheme due to Moulin and Shenker [36] is shown in Figure 5.2. Intuitively, the mechanism starts with a coalition consisting of the entire set. This coalition is refined using the cost-sharing scheme to select the players within it who are able to pay their share of the costs. Such a refinement changes the cost, so the process is repeated until it stabilizes.

The cost-sharing scheme χ_{FP} corresponding to a fixed-price digital goods auction divides the total cost evenly over all members of the set: if R is the auctioneer's revenue whose cost is to be shared over the players, then $\chi_{\mathrm{FP}}(i, S) = R/|S|$. In this case the cost-sharing mechanism of Figure 5.2 is easy to describe: it begins with all players, and raises the price until it finds a set that generates revenue R. Let the player values be listed in order $v_1 \geq v_2 \geq \cdots \geq v_n$, then $\mathrm{CS}_{\chi_{\mathrm{FP}}} = \{v_1, \ldots, v_{i^*}\}$ where i^* is the maximum i such that $i * v_i \geq R$. We will denote this digital goods cost sharing mechanism $\mathrm{CS}_{\chi_{\mathrm{FP}}}(\cdot)$.

5.1.3 Truthful Mechanisms for Digital Goods

For digital goods auctions, truthful mechanisms have a useful characterization, the intuition for which is more general. A *bid-independent auction* is one where the price p_i offered to each bidder as a function of other bids only: $p_i = f(\mathbf{b}_{-i})$. More precisely, $x_i = 1$ and $f(\mathbf{b}_{-i}) = p_i$ if $f(\mathbf{b}_{-i}) \leq b_i$, otherwise $x_i = 0$ (and p_i is immaterial). This characterizes a bid-independent deterministic mechanism completely. A randomized bid-independent mechanism works in the same way with a random function $f(\cdot)$, noting that the evaluations for different players may be correlated. One can show that bid-independence is equivalent to truthfulness.

Theorem 5.1.2 ([19]). *A deterministic auction is truthful if and only if it is bid-independent.*

5.1.4 Revenue Benchmarks

In a *fixed-price* auction for digital goods, after soliciting bids the auctioneer computes a single fixed price that is offered to all bidders, who accept only if the price is less than their value. The optimal fixed-price auction can be computed from the bidders' values as using a price p that maximizes $k(p) \cdot p$, where $k(p)$ is the number of players whose values are at least p.

Definition 5.1.3. The optimal fixed-price revenue on $\mathbf{v}$ is

$$\mathcal{F}(\mathbf{v}) = \max i \cdot v_i,$$

where the values are ordered from largest to smallest.

As noted by Fiat et al. [19], $\mathcal{F} = \Omega\big((\sum v_i)/\log n\big)$, and so is a logarithmic factor away from the optimal omniscient multi-price auction which extracts all player value. This bound is tight, as can be seen by taking $v_i = 1/i$. Thus, approximating the fixed-price optimum gives an approximation to the absolute optimal auction at the cost of a logarithmic factor. However, even such an approximation is too much to hope for. For example, if there is one bidder with value h and all other bidder have value ≤ 1, then there is no truthful mechanism to extract revenue $\Omega(h)$ [22]. Surprisingly, by relaxing this benchmark to assuming at least two buyers, a constant factor approximation can be achieved.

Definition 5.1.4. In the context of Definition 5.1.3,

$$\mathcal{F}^{(2)}(\mathbf{v}) = \max_{i \geq 2} i \cdot v_i.$$

We use this notation beyond two winners:

$$\mathcal{F}^{(m)}(\mathbf{v}) = \max_{i \geq m} i \cdot v_i.$$

Early analyses got around this problem by assuming a bound on the ratio between the highest and lowest value or by assuming that $\mathcal{F}(\mathbf{v}) \gg v_1$ [22]. Using $\mathcal{F}^{(2)}$ seems much more elegant, and generalizes very nicely to the context of matroid procurement, where the sampling algorithms require that a logarithmic number of spanning trees need to be purchased.

The constant-factor approximations to $\mathcal{F}^{(2)}$ currently known are asymptotically the best possible because of the following theorem.

Theorem 5.1.5 ([22]). *For any truthful, randomized mechanism $\mathcal{M}$ achieving revenue $\mathcal{R}(\mathbf{b})$ on bids $\mathbf{b}$,*

$$\mathrm{E}[\mathcal{R}(\mathbf{b})] \leq \frac{\mathcal{F}^{(2)}(\mathbf{b})}{2.42}.$$

Furthermore, using a randomized mechanism is essential.

Theorem 5.1.6 ([22, 19]). *If $\mathcal{M}$ is a deterministic bid-independent digital goods mechanism, then for any $1 \leq m \leq n$ there exists a set of player values $\mathbf{b}$ such that the revenue of $\mathcal{M}$ is at most $\frac{m}{n} \cdot \mathcal{F}^{(m)}(\mathbf{b})$.*

In particular, there is no deterministic mechanism that comes within a linear factor of $\mathcal{F}^{(2)}$.

Given a bid vector $\mathbf{b}$,

1. Partition bids into two sets, $\mathbf{b}'$ and $\mathbf{b}''$, by flipping an unbiased independent coin for each player.

2. Compute $\mathcal{F}' = \mathcal{F}(\mathbf{b}')$ and $\mathcal{F}'' = \mathcal{F}(\mathbf{b}'')$, the optimal fixed price revenues for the two sets of the partition.

3. Run the revenue-extractor $\mathrm{CS}_{\chi_{\mathrm{FP}}}(\mathcal{F}')$ from Section 5.1.2 on $\mathbf{b}''$ and $\mathrm{CS}_{\chi_{\mathrm{FP}}}(\mathcal{F}'')$ on $\mathbf{b}'$

Figure 5.3: The partition and cost-share mechanism PCS

5.1.5 The Partition and Cost-Share Mechanism

Figure 5.3 presents a constant-competitive mechanism PCS, for *partition and cost-share*, found in Fiat et al. [19]. The intuition behind it is to first estimate the possible revenue, then use cost sharing to extract it. In order to do this truthfully, a bidder involved in cost sharing cannot have her bid used in the revenue estimate used. Hence, the mechanism randomly partitions the bidders into two groups, and uses the revenue on one group as an estimate on that of the other group. The two-part construction of this mechanism is used in our matroid procurement mechanism.

Theorem 5.1.7 ([22, 19]). PCS *is a truthful mechanism such that the expectation of the revenue* $\mathrm{E}[R_{\mathrm{PCS}}]$ *satisfies*

$$\mathrm{E}[R_{\mathrm{PCS}}] \geq \frac{\mathcal{F}^{(2)}}{4},$$

and this bound is tight.

Proof sketch. The truthfulness of the mechanism follows from Theorem 5.1.2 and the truthfulness of the cost-sharing mechanism. To prove the bound on revenue, as the mechanism will extract at least the smaller of $\mathcal{F}'$ and $\mathcal{F}''$ from one of the two sampled halves, we need to show that in expectation the revenue $R = \min(\mathcal{F}', \mathcal{F}'') \geq \mathcal{F}^{(2)}(\mathbf{b})/4$.

Let $\mathcal{F}^{(2)}(\mathbf{b})$ sell to k bidders at price p, so that the total revenue is kp. These k bidders all bid at least p. Define k' and k'' so that k' of these bidders appear in $\mathbf{b}'$ and k'' appear in $\mathbf{b}''$. Thus, $\mathcal{F}(\mathbf{b}') \geq pk'$, and a similar statement is true for k''. Then,

$$\frac{\min(\mathcal{F}(\mathbf{b}'), \mathcal{F}(\mathbf{b}''))}{\mathcal{F}^{(2)}(\mathbf{b})} \leq \frac{\min(pk', pk'')}{pk} = \frac{\min(k', k'')}{k}.$$

Now note that $k' = i$ with probability $\binom{k}{i}2^{-k}$, so that we may compute $\mathrm{E}[R]$ by summing:

$$\frac{\mathrm{E}[R]}{\mathcal{F}^{(2)}} = \frac{1}{k} \sum_{0 \leq i \leq k} \min(i, k-i) \binom{k}{i} 2^{-k} = \frac{1}{2} - \binom{k-1}{\lfloor \frac{k}{2} \rfloor} 2^{-k},$$

which has a minimum of $1/4$ for $k = 2$ and $k = 3$. As k increases, the ratio approaches $1/2$. □

Note that because the revenue extracted from one half of the random partition formed by PCS will be more than the other half, we will lose as much as half the potential revenue. This seems to be an essential feature of this technique. Also note that the algorithm performs optimally in this sense when the optimal revenue comes from a large number of players. This confirms intuition that an auctioneer benefits from large markets.

5.1.6 *Consensus Revenue Estimate Auctions*

The factor 4 approximation achieved by the PCS mechanism has been improved to 3.39 using a new technique known as *consensus revenue estimation* (CORE) [21] that avoids the essential loss of half the revenue as with PCS. The CORE technique works best when the amount that a single bidder can change the revenue is bounded. For this reason, it does not seem to apply to the matroid domains where we are interested in profit extraction. However, the issue is not settled and it may be a promising area for future research.

5.2 Path Procurement

In this section we study *path procurement*, which has emerged as a touchstone for algorithmic mechanism design. The research on this problem has been mostly negative, showing that truthful mechanisms do not approximate the true costs of paths, while first-price auctions have equilibria with very high cost [4, 43, 17, 29].

In this problem, the buyer must buy an s–t path in a graph, where each edge is controlled by a different player. The players have a private cost of providing the edge to the buyer, who must come up with a mechanism to solicit bids from the players and then buy a selected path, minimizing the total payments. We first consider basic path procurement, where only one path is bought, before moving to the problem of buying multiple edge-disjoint paths.

Consider a graph $\mathcal{G}$ with two edge-disjoint s–t paths, one path P containing L edges each of cost 1, and the other Q a single edge with cost $(1+\varepsilon)L$. The VCG mechanism will select P, and pay each edge therein the marginal difference $1+\varepsilon L$. The total cost is thus a $L(1+\varepsilon L)$, which is an $\Omega(L)$ factor greater than the true cost, that is, the overpayment from VCG is linear in the number of edges in a path. In Archer and Tardos [4] and Elkind et al. [17], this example is extended to show that any truthful mechanism, and hence any dominant-strategy mechanism, must overpay in the same way. This is very bad news for the possibility of a mechanism that extracts a constant factor of the possible revenue.

Other mechanisms can do even worse. For example, a first-price auction has no dominant strategy, as it has a multitude of Nash equilibria. Augment the graph $\mathcal{G}$ used above with a third s–t path T consisting of a single edge of cost $(L-1)x$, for some $x > 2$. All edges bidding truthfully is an ε-Nash equilibrium, for P will win the auction, neither Q nor T will be able to lower her bid to win the auction without going under her cost, and no edge of P will be able to raise her bid by more than ε without losing the auction to P.

There is another ε-Nash equilibrium, however, which shows that the first-price auction has no dominant strategy. Let each player in P bid x, the value determining T's cost, so that the total cost for P is Lx. Let Q bid $(L-1)x-\varepsilon$, and let T again bid her value. Then Q will win the auction, but cannot increase her bid without conceding the auction to T. As before, T cannot lower her bid and so is a passive observer. Finally, the most any single edge of P can lower the cost of her path to is $(L-1)x+1$, so that there is also no possibility of P winning the auction through unilateral action. Hence the bids are in ε-Nash equilibrium, and this time the ratio between the cost of the procured path and the cheapest path is $\Omega(x)$. This shows bad behavior on the first price mechanism as well, although in a somewhat incomparable way to the example for truthful mechanisms.

The failure of these two basic mechanisms to compete with the optimal cost of a path

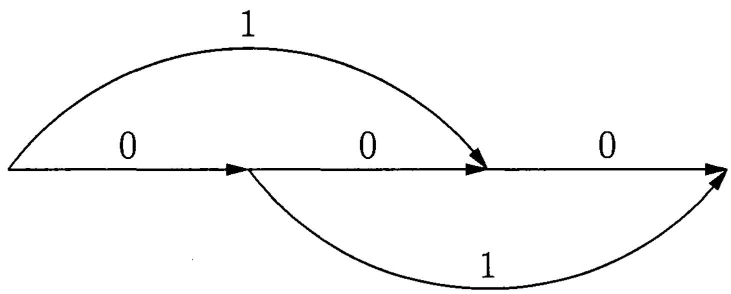

Figure 5.4: Monopoly-free path procurement where no path is edge-disjoint from the cheapest path

suggests that optimal cost may not be the right benchmark for procurement over set systems. Talwar [43] suggests an alternative benchmark using the cost of the second-cheapest disjoint feasible set, and shows that VCG is an optimal mechanism for a class of set systems that is a generalization of matroids which he calls *frugoids.* However, as was pointed out by Karlin et al. [29], there are many natural systems where using a disjoint set is too restrictive. That is, there are many natural examples where no element has a monopoly on all feasible sets, and yet there are no two disjoint feasible sets. An example for path procurement is shown in Figure 5.4. To address this objection, Karlin et al. [29] show that the VCG mechanism is revenue-equivalent with the first-price auction if and only the procurement takes place over a matroid set system. We will discuss this link in more detail in Section 5.5 as it provides the justification for our use of VCG in our procurement mechanism.

A possible way out of this dilemma is to take a cue from the success of mechanisms for digital auctions which suggest that when the buyer is free to choose the number of items sold in addition to the price, her revenue can be much more competitive than when the amount to buy is fixed in advance. Consider a graph with $k+1$ edge-disjoint paths, the first path P consisting of n edges of cost 1 each, and the remaining paths $Q_1, \ldots, Q_k$ consisting of one edge each of cost $(1+\varepsilon)n$, so that as described above the VCG overpayment for a single path is $\Omega(n)$. However, if the buyer instead procures k edge-disjoint paths, the overpayment is only $\Omega(n/k)$, so if k is $\Omega(n)$, the overpayment is only a constant factor. Is there a general way for procurement auctions to take advantage of this flexibility? One of the results of this dissertation is to show that this can be done for procurement over a matroid, but that general profit extraction cannot be truthful over non-matroids, so that frugal procurement over non-matroids requires substantially different technology.

5.3 Procurement Auctions

We now define general procurement auctions formally. We distinguish two types of procurement. In the first, which we call *simple procurement*, the buyer is trying to purchase a single set of items. In the other, *multi-unit procurement*, the buyer chooses how many sets of items to buy, and receives revenue as a function of the number of items procured

Let $\mathcal{A}$ be a set system, that is a family of *feasible* sets over a base set X. We parametrize the set system by defining $\mathcal{A}_k$ to be the set system whose feasible sets are the union of any k disjoint feasible sets from $\mathcal{A}$. We will associate *cost* or *bid* vectors with X; given such a cost vector $\mathbf{b}$ and $S \subset X$, we let $\mathbf{b}(S) = \sum_{x \in S} b_x$. We then define an optimization problem over $\{\mathcal{A}_k\}$ by letting $\mathrm{OPT}^{\mathbf{b}}(\mathcal{A}_k) = \min_{A \in \mathcal{A}_k} \sum_{x \in A} b_x$, and $S_k(\mathbf{b})$ be a cheapest optimal set. We will drop the bid vector $\mathbf{b}$ from the notation when it is clear from context.

We make the following assumptions.

Assumption 5.3.1. $\mathcal{A}$ is finite. Then there is a maximum k such that $\mathcal{A}_k$ is non-empty, which we denote K.

Assumption 5.3.2. Item costs are positive. This implies that we may assume that feasible sets are non-nesting, because if $S \subset T$ then $\mathbf{b}(S) \leq \mathbf{b}(T)$ so that there is not a benefit to choosing T over S.

A *procurement auction* over a parametrized set system is an auction between agents representing each point of X, and a centralized buyer who obtains value for each feasible set procured. In a *simple* procurement auction, the buyer is restricted to purchasing a single set, and hence the problem is simply to minimize the cost of that set. The *multi-unit* procurement auction removes this restriction, and lets the buyer choose how many sets to procure, up to the natural bound K determined by the set system. In this case it is necessary to specify a function $B(\cdot)$ which specifies the value $B(k)$ the buyer receives by buying k sets. We assume our buyer has *decreasing marginal benefit* per disjoint set from $\mathcal{A}$

procured, meaning $B(\cdot)$ satisfies $B(k+1) - B(k) \leq B(k) - B(k-1)$. An interesting special case of decreasing marginal benefit is the case where the marginal benefit is constant, i.e., $B(k) = B(1) \times k$.

To run the auction, the agents submit bids to form a vector $\mathbf{b}$. Based on these bids, the buyer selects a k and then purchases the cheapest k feasible sets according to $\mathbf{b}$, giving each agent u a payment $p_u(k)$. This payment will depend on the total number of sets procured. We will denote the total payment P_k. The buyer selects the k with the goal of maximizing her total revenue $kB - P_k$. Each agent u has a private value v_u, and if selected for a feasible set derives utility $p_u(k) - v_u$. An agent not selected has zero utility. The goal of the agents is to maximize their utilities. Note that the VCG mechanism does not make sense for this case, for maximizing social welfare will procure from as many agents as possible, that is, choose the largest k possible. In most cases this will produce much less revenue than the choice of k maximizing buyer revenue.

5.3.1 The VCG Mechanism for a Procurement Auction

Our definition of buyer revenue requires a payment mechanism for each possible value of k. A natural attempt at a mechanism will, given k, select the cheapest set in $\mathcal{A}_k$; in other words, choose an efficient allocation. Hence the VCG mechanism is an obvious first candidate.

To describe the VCG mechanism, let $\mathcal{A}_k \setminus e$ denote the set system $\{A \in \mathcal{A}_k : e \notin A\}$. Then for a given set of bids $\mathbf{b}$, the VCG payment for an $e \in S_k$ is $p_e = \mathrm{OPT}^{\mathbf{b}}(\mathcal{A}_k \setminus e) - \mathrm{OPT}^{\mathbf{b}}(\mathcal{A}_k) + b_e$, and $\mathrm{VCG}_k(\mathbf{b}) = \sum_{e \in S_k} p_e$. The payment for $e \notin S_k^{\mathbf{b}}$ is 0, as required by the VCG mechanism. We will drop the $\mathbf{b}$ from the notation when the set of bids is clear. Note this definition of p_e immediately implies that $p_e \geq b_e$, as $p_e < b_e$ implies $\mathrm{OPT}^{\mathbf{b}}(\mathcal{A}_k \setminus e) < \mathrm{OPT}^{\mathbf{b}}(\mathcal{A}_k)$, contradicting the definition of optimality.

We can describe this VCG payment as a threshold bid, below which an agent is chosen for the optimal set. In the following we will consider some element e, and let $\mathbf{b}_{-e}$ be a fixed set of bids for all elements except for e. We then define $\mathbf{b}^{(e)}(x)$ to be the bid vector that is consistent with $\mathbf{b}_{-e}$ and assigns cost x to e. That is, $b_e^{(e)}(x) = x$, and $b_u^{(e)}(x)$ is $\mathbf{b}_{-e}$ at u for all $u \neq e$.

Theorem 5.3.3. *Given any bid vector $\mathbf{b}_{-e}$ fixed outside of an element e, there is a threshold $t_k(\mathbf{b}_{-e})$ such that if $x = b_e > t_k(\mathbf{b}_{-e}))$, then e is not in the cheapest feasible set $S_k\big(\mathbf{b}^{(e)}(x)\big)$, and if $x = b_e \leq t_k(\mathbf{b}_{-e})$ then $e \in S_k\big(\mathbf{b}^{(e)}(x)\big)$ and the VCG payment p_e is equal to $t_k(\mathbf{b}_{-e})$.*

We first show the following useful lemma which is a sort of continuity result: as long as an element stays in the optimal set, the rest of that set does not change as the bid of an element increases.

Lemma 5.3.4. *Given an element e, a bid vector $\mathbf{b}_{-e}$ fixed outside of e, and two bids x and y, if $e \in S_k\big(\mathbf{b}^{(e)}(x)\big)$ and $e \in S_k\big(\mathbf{b}^{(e)}(y)\big)$ then $S_k\big(\mathbf{b}^{(e)}(x)\big) = S_k\big(\mathbf{b}^{(e)}(y)\big)$.*

Proof. For simplicity of notation we write $\mathbf{b}$ for $\mathbf{b}^{(e)}$. Assume without loss of generality that $x < y$. Suppose for contradiction that $S_k\big(\mathbf{b}(x)\big) \neq S_k\big(\mathbf{b}(y)\big)$, so that $\sum_{u \in S_k(\mathbf{b}(y))} b_u(y) < \sum_{u \in S_k(\mathbf{b}(x))} b_u(y)$ as $S_k\big(\mathbf{b}(x)\big)$ is not an optimal set for $\mathbf{b}(y)$. Then as $x < y$ and e is in both $S_k\big(\mathbf{b}(x)\big)$ and $S_k\big(\mathbf{b}(y)\big)$, and $\mathbf{b}(x)$ equals $\mathbf{b}(y)$ on all other elements, this implies $\sum_{u \in S_k(\mathbf{b}(y))} b_u(x) < \sum_{u \in S_k(\mathbf{b}(x))} b_u(x)$, contradicting the optimality of $S_k\big(\mathbf{b}(x)\big)$. $\square$

Proof of Theorem 5.3.3. Again we write $\mathbf{b}$ for $\mathbf{b}^{(e)}$. Let $Y_e = \{y : e \in S_k\left(\mathbf{b}(y)\right)\}$ be the set of bids y such that e is in an optimal set of S_k. If Y_e is nonempty, for all $y \in Y_e$, set $p(y)$ equal to the VCG payment of e for $\mathbf{b}(y)$, that is, $p_e\big(\mathbf{b}(y)\big)$ as defined above. By the previous lemma, each $y \in Y_e$ induces the same optimal set S_e. By definition we have that $\mathrm{OPT}^{\mathbf{b}_{-e}}(\mathcal{A}_k \setminus e) = p_y + \sum_{u \in S_e, u \neq e} b_u$. Hence $\sum_{u \in S_e} b_u(x) \leq \mathrm{OPT}^{\mathbf{b}_{-e}}(\mathcal{A}_k \setminus e)$ for all $x \leq p_y$, which shows that $e \in S_k\big(\mathbf{b}(x)\big) = S_e$ for all $x \leq p_y$, and so the VCG payment is p_y for all such x. If $x > p_y$, then $\mathrm{OPT}^{\mathbf{b}_{-e}}(\mathcal{A}_k \setminus e) < \sum_{u \in S_e} b_u(x)$, so $e \notin S_k\big(\mathbf{b}(x)\big)$.

Thus we may define the threshold $t_k(\mathbf{b}_{-e})$ which is equal to the uniquely defined VCG payment of e for any x such that $e \in S_k\big(\mathbf{b}(x)\big)$. Furthermore, $e \in S_k\big(\mathbf{b}(x)\big)$, that is, e is chosen as a winner, if and only if $x \leq t_k(\mathbf{b}_{-e})$. $\square$

This threshold is the pivotal value for e, as defined in Section 2.3.1. It also serves to define a bid-independent function in the sense of Section 5.1.3. This means that the intuition of VCG mechanism as thresholding function seen with digital auctions holds in the procurement setting as well.

5.4 *Matroids*

Matroids are a class of set systems with a rich history in both mathematics and computer science. In computer science they are known for describing a large class for which the greedy algorithm always finds an optimal feasible set. The minimum spanning tree problem can be expressed as a matroid optimization, and much of the terminology for matroids comes from this setting. A large chunk of the remaining terminology is related to sets of independent points in a vector space; this setting is associated with the mathematics literature and the classification of matroids. Oxley [39] is a readable and complete introduction to matroid theory.

Definition 5.4.1. A *matroid* M is a family of subsets of a base set X satisfying the following axioms. We write $I \in M$ to denote a set I in this family, and call such sets *independent.*

1. (Subset independence) If $I \in M$, then for all $J \subset I$, $J \in M$.

2. (Set augmentation) If $I, J \in M$ with $|I| > |J|$, then there exists an $x \in I \setminus J$ such that $J \cup x \in M$.

A *base* of M is an independent set of maximal size. The set of bases of a matroid gives another characterization of a matroid.

Fact 5.4.2. *$\mathcal{B}$ is a set of bases of a matroid if and only if*

1. *$\mathcal{B}$ is non-empty, and*

2. *(Basis exchange) If B_1 and B_2 are members of $\mathcal{B}$, then for all $x \in B_1 \setminus B_2$, there is a $y \in B_2 \setminus B_1$ such that $(B_1 \setminus x) \cup y \in \mathcal{B}$.*

It is easy to show using either of these axiom systems that all bases of a matroid are of the same size. In fact, set augmentation can be strengthened to the following.

Fact 5.4.3. *If S and T are two independent sets of equal size in some matroid M, then there is a bijection $\pi : S \setminus T \to T \setminus S$ such that for any $e \in S \setminus T$, $(S \setminus e) \cup \pi(e)$ is an independent set of M.*

The *rank* of a set of points in X is the size (number of elements) of the maximum independent set it contains. The rank of a matroid is the size of a maximal independent set. We will use ρ to denote the rank function for M on X. It can be shown that $\rho(A) = |A|$ if and only if A is independent. A *circuit* is a dependent set of edges, the removal of any one of which results in an independent set. Thus if C is a circuit in M, $\rho(C) = |C| - 1$, and for any $e \in C$, $\rho(C \setminus e) = |C| - 1$, so that $C \setminus e$ is independent. The following is a consequence of the matroid axioms.

Fact 5.4.4. *If C and D are two circuits of a matroid M with $C \cap D \neq \emptyset$, then for any $e \in C \cap D$, there is a circuit contained in $(C \cup D) \setminus e$.*

5.4.1 Simple Examples and Transversal Matroids

A trivial matroid on any set is the *free matroid*, where all subsets are independent. Slightly more interesting is the *uniform matroid of rank n*, where all subsets of size at most n are independent.

The *cycle matroid* of a graph has as its independent collection all sets of edges with no cycles. A circuit in this graph is thus a circuit in the cycle matroid, as removing any edge leaves an acyclic set of edges. The bases of the matroid are spanning forests; if the graph is connected the rank of the cycle matroid is one less than the number of vertices in the graph.

Another example of a matroid is to be found by taking a partition of some universe U into n disjoint sets $T_1, \ldots, T_n$. Then the family of all sets with at most one element from each T_i is a collection of independent sets for a matroid; the rank of this matroid is n and it is known as the *partition matroid.*

A function mapping subsets of a set X into $\mathbb{R}$ is said to be *increasing* if $f(A) \leq f(B)$ whenever $A \subseteq B \subseteq X$, and is called *submodular* if $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq X$. Submodularity is the combinatorial analog of convexity. It is easy to show the following.

Fact 5.4.5. *The rank function of a matroid is increasing and submodular.*

One can show that this fact has a converse. Given an increasing submodular function f, define $M(f)$ to be the family of sets $\{I \subseteq X : |J| \leq f(J)$ for all nonempty $J \subseteq I\}$. Then the following is true.

Fact 5.4.6. *$M(f)$ is a matroid. If f is integer-valued and $f(\emptyset) = 0$, then the rank function $\rho_f(\cdot)$ of $M(f)$ is given by*

$$\rho_f(A) = \min\{f(B) + |A \setminus B| : B \subseteq A\}.$$

This fact, particularly the characterization of $\rho_f(\cdot)$, is nontrivial. An exposition and proof may be found in Chapter 12 of Oxley [39].

Take a bipartite graph with vertex sets L and R. The *transversal matroid* of the graph has independent sets formed by subsets of R that have a matching into L. That is, $X \subseteq R$ is independent if there is a set of edges T where no two u and v in T have an endpoint in common, and every vertex in X is incident to a member of T. The rank of this matroid is equal to the size of a maximum matching in the associated bipartite graph. Suppose a company is looking to fill several vacant positions, each of which requires a particular skill. There is a pool of workers, each of whom possesses some but perhaps not all of the skills required. Let the vacant positions and the workers each form a side of a bipartite graph, and connect a worker to a position if she possesses the skill required for that job. An independent set in the transversal matroid of this graph represents a selection of potential hires and a set of positions they can exactly fill.

Transversal matroids arise when studying the *marriage problem* [25, 41]. Suppose we have a group S of n eligible bachelors and n unmarried women. For each women i, let A_i be the set of bachelors who she considers acceptable as husbands. The marriage problem is to determine necessary and sufficient conditions for all the women simultaneously to find acceptable husbands. Such a set of marriages is known as a *transversal* or *system of distinct representatives*: a subset $\{u_1, \ldots, u_m\}$ of the bachelors such that $u_i \in A_i$ for $i = 1, \ldots, n$. Form a bipartite graph with L equal to the set of bachelors S, $R = \{1, \ldots, n\}$, and edges $\{(u, j) : u \in A_j\}$. Then a transversal exists if and only if the rank of this transversal matroid is n. Let $K \subseteq \{1, \ldots, n\}$ and define $A(K) = \cup_{j \in K} A_j$. For a transversal to exist, certainly $|A(K)| \geq |K|$ for all $K \subseteq \{1, \ldots, n\}$. Hall [25] proves the converse.

Theorem 5.4.7 (Hall's Marriage Theorem). *Let $A_1, \ldots, A_n$ be a family of subsets of a set S and $A(K) = \cup_{j \in K} A_j$ for all $K \subseteq \{1, \ldots, n\}$. Then $A_1, \ldots, A_n$ has a transversal if and only if*

$$|A(K)| \geq |K|$$

for all $K \subseteq \{1, \ldots, n\}$.

Hall's Marriage Theorem follows as a special case of a more general theorem of matroids. Suppose that we have a matroid M defined over S, and we want to find a transversal that is also independent in M. Then the marriage problem is simply the special case when M is the free matroid on S. The following generalization of Hall's Marriage Theorem is due to Rado [41].

Theorem 5.4.8 (Rado's Theorem). *Let $A_1, \ldots, A_n$ be a family of subsets of a set S and $A(K) = \cup_{j \in K} A_j$ for all $K \subseteq \{1, \ldots, n\}$. Let M be a matroid over S with rank function $\rho(\cdot)$. Then $A_1, \ldots, A_n$ has a transversal that is independent in M if and only if*

$$\rho\big(A(K)\big)| \geq |K|$$

for all $K \subseteq \{1, \ldots, n\}$.

A natural next step is to consider the structure of all transversals satisfying Rado's Theorem. Namely, given a bipartite graph $\mathcal{B}$ with edges sets L and R, and a matroid M over L, let $\mathcal{B}(M)$ be the family of subsets of R that are matched into an independent set of M. That is, $X = \{x_1, \ldots, x_k\} \in \mathcal{B}(M)$ if there is set $Y = \{y_1, \ldots, y_k\} \subseteq L$ that is independent in M, and (x_i, y_i) are edges of $\mathcal{B}$ for each $i = 1, \ldots, k$.

Theorem 5.4.9. *$\mathcal{B}(M)$ is the set of independent sets of a matroid over L. If $\rho(\cdot)$ is the rank function of M, the rank of $X \subseteq L$ in $\mathcal{B}(M)$ is given by* $\min\{\rho\big(\Delta(Y)\big) + |X \setminus Y| : Y \subseteq X\}$, *where $\Delta(Y)$ is the set of neighbors of Y in $\mathcal{B}$.*

Proof. Define the function f over subsets of L by $f(X) = \rho\big(\Delta(X)\big)$. Clearly $f(\cdot)$ is integer-valued and $f(\emptyset) = 0$. It is straightforward to check that $f(\cdot)$ is also submodular. Hence $M(f)$ as defined for Fact 5.4.6 is a matroid over L, and a set $X \subseteq L$ is independent in $M(f)$

if for all $Y \subseteq X$, $\rho_f(Y) = \rho\big(\Delta(Y)\big) \geq |Y|$. By Rado's Theorem, this holds if and only if $\Delta(X)$ has a transversal that is independent in M, that is, there is a matching of X in $\mathcal{B}$ to an independent set of M. Hence $\mathcal{B}(M)$ is exactly $M(f)$, and the theorem follows from Fact 5.4.6. □

We call $\mathcal{B}(M)$ the *generalized transversal matroid of* $\mathcal{B}$ *and* M.

5.4.2 Basic Facts

We first describe some simple ways to form new matroids from existing matroids.

Definition 5.4.10. Given a matroid M, we let $M|_i$, the *restriction to* i *of* M, denote the family of independent sets of M of size at most i. It can be shown that this family is a matroid as well, with rank equal to the minimum of i and the rank of M.

Definition 5.4.11. Given matroids M and N over distinct base sets X and Y, the *direct sum* $M \oplus N$ is the family of sets $\{A \subseteq X \cup Y : A \cap X \in M \text{ and } A \cap Y \in M\}$. $M \oplus N$ is a set of independent sets for a matroid.

The next definition captures the structure of multi-unit procurement auctions over matroids.

Definition 5.4.12. If M and N are two matroids over the same element set X, we define their *join*, denoted $M \vee N$, as the family of all the unions of one independent set from each of M and N:

$$M \vee N = \{I \cup J : I \in M, J \in N\}.$$

We let M_i denote the self-join a matroid M with itself i times.

It is not hard to show that $M \vee N$ is a matroid. Note that by subset independence, if $I \in M \vee N$, we may take $I = A \cup B$, $A \in M$ and $B \in N$ with $A \cap B = \emptyset$ without loss of generality. $S \in M_i$ iff S can be partitioned into disjoint $S_1, \ldots, S_i$ such that $S_j \in M$ for all j. We use ρ_i to denote the rank function of M_i. We can use Theorem 5.4.9 to give a decomposition of the rank of a set in M_i in terms of the base set. Create a bipartite graph $\mathcal{B}$ by forming L from two copies of X, the base set of M, and let M' be the matroid $M \oplus M$

over L. Let R be a third copy of X, and connect each $x \in R$ to its two corresponding copies in L with an edge. Then the generalized transversal matroid $\mathcal{B}(M')$ as defined in Section 5.4.1 describes $M \vee M$. Extending this to larger self-joins, the rank characterization from Theorem 5.4.9 implies the following lemma.

Lemma 5.4.13.

$$\rho_i(A) = \min_{Y \subseteq A} i \cdot \rho(Y) + |A \setminus Y|$$

From this we can prove the following about a nice witness to the rank of a set.

Lemma 5.4.14. *Any A contains a set Y with $\rho_i(Y) = i \cdot \rho(Y)$ and $\rho_i(A) = i \cdot \rho(Y) + |A \setminus Y|$. If $A \notin M_i$, then $Y \neq \emptyset$.*

Proof. Take Y a minimizing set in Lemma 5.4.13. If $Y = \emptyset$, then $\rho_i(A) = |A|$ and $A \in M_i$. Otherwise, if $\rho_i(Y) < i\rho(Y)$, then take T the minimizing set for Y itself in Lemma 5.4.13, so that $i\rho(T) + |Y \setminus T| < i\rho(Y)$. As $|A \setminus T| = |A \setminus Y| + |Y \setminus T|$, we have that $i\rho(T) + |A \setminus T| < i\rho(Y) + |A \setminus Y|$, contradicting the minimality of Y. If $\rho_i(Y) > i\rho(Y)$, then Y contains an independent set $U \in M_i$ of size $> i\rho(Y)$. By definition, this means U can be decomposed into i disjoint independent sets of M whose average size is strictly greater than $\rho(Y)$, meaning there exists a $Z \subset Y$ with $\rho(Z) > \rho(Y)$, a contradiction. □

We will be interested in looking at decompositions of sets into independent sets of self-joins of various orders. For example, we may have $S \in M$ and $S \cup T \in M_2$. This does not imply that $T \in M$, as shown in Figure 5.5. We can, however decompose $S \cup T$ into two independent sets, one of which has the same size as S.

Lemma 5.4.15. *Let M and N be matroids over X. Let $S \in M$ and $T \cup U \in M' = M \vee N$, with $T \in M$, $U \in N$, and $T \cap U = \emptyset$. If $S \subset T \cup U$ with $|S| > |T|$, then there exists $T' \cup U' \in M'$ such that $|T'| = |S|$, $T' \cap U' = \emptyset$, $T \subset T'$ and $U' \subset U$.*

Proof. We augment T with elements from S to produce T' with $|T'| = |S|$ and $T' \in M$. Then as the augmenting elements $T' \setminus T \subset U$, we set $U' = U \setminus T'$, and $U' \in N$ as subsets of independent sets are independent. $T' \cup U'$ then gives the desired decomposition. □

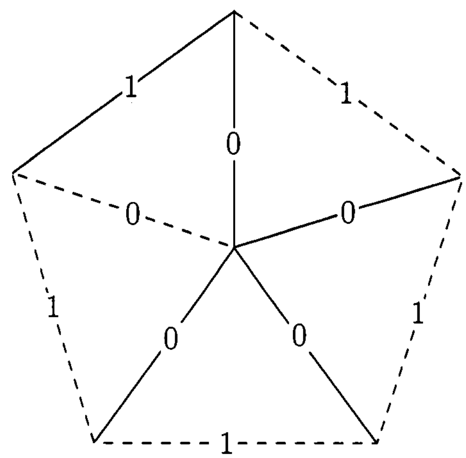

The spokes of the wheel (edges labeled 0) form a spanning tree, and when augmented with the rim will form two disjoint spanning trees as illustrated by the solid and dashed lines. However, the rim by itself (edges labeled 1) is not a spanning tree.

Figure 5.5: The wheel

Note that the lemma depends on $|S| > |T|$. This will not be a problem for our applications, as we apply this when S is a maximal independent set of M in $T \cup U$. The last introductory lemma is due to Karger, and discusses how the rank of a base of high order changes when it is sampled.

Lemma 5.4.16 (Karger [28], Theorem A.7). *Let $M(p)$ be the matroid obtained by sampling the elements of a matroid M independently with probability p. If k is the maximum number of disjoint bases contained in M, then the maximum number of disjoint bases in $M(p)$ is at least $pk(1-\varepsilon)$ with probability at least $1 - \rho(M) \cdot e^{-\varepsilon^2 pk/2}$.*

The number k used above is known as the *packing number* of a matroid.

5.4.3 Optimization over Matroids

In computer science, matroids are linked to the greedy optimization algorithm. If the elements of a matroid base set are weighted, the natural optimization problem is to find a minimum cost base of the matroid. Let M be a matroid over a weighted base set X. For

$Y \subseteq X$, OPT(Y) denotes a minimum weight base of M in X. We will assume throughout that no single element of X is dependent in M (for a graphic matroid, this means there are no self-loops).

The greedy algorithm starts with the empty set, and considers the elements of X in order of increasing weight, adding an element to the current set if it is independent. The resulting set is always a minimal base. A proof of this may be found in Oxley [39]; a disguised proof is found in most algorithms textbooks (*e.g.*, Cormen, Leiserson, Rivest, and Stein [12]) in the guise of Kruskal's algorithm for minimum spanning trees.

More precisely, let $n = \rho(M)$ and let $X = \{x_1, x_2, \ldots\}$ be the base set ordered by weight. $S(t)$ will be the state of the algorithm's set after x_t has been considered, with $S(0) = \emptyset$ being the initial state. We think of the algorithm as operating temporally in order of the elements, so that we say that x_t is *considered at time t*. The update rule is $S(t) = S(t-1) \cup x_t$ if $S(t-1) \cup x_t$ is independent, and $S(t) = S(t-1)$ otherwise. The algorithm is finished when $|S(u)| = n$ and will denote this set by S.

Lemma 5.4.17. *$S(t)$ is an independent set of maximum size in $\{x_1, \ldots, x_t\}$.*

Proof. Let $|S(t)| = m$, and let T be an independent set in $\{x_1, \ldots, x_t\}$ of maximum size. Suppose the lemma is false, so that $|T| > m$. Consider the matroid M' over $X' = \{x_1, \ldots, x_t\}$ whose independent sets are $\{I \cap X' : I \in M\}$. As T is a maximal independent set in M', $\rho(M') = |T| > m$. The greedy algorithm run on M' will produce $S(t)$ as the final set. But this is not a base as $|S(t)| < \rho(M')$, contradicting the optimality of the greedy algorithm. □

Referring to the Figure 5.6, we see in our notation $|S_i(t)| = |R| + h - 1$. But this is impossible, as the greedy algorithm guarantees a minimal cost basis, for any choice of inputs. Let y be the next element to be added to S_i (that is, after time t). No matter the cost of $c(U + T_j)$, we can set the cost of y to be large enough that $c(U \cup T_j) < c(S_i \cup y)$, which contradicts the optimality of the greedy algorithm.

Now consider the self-join matroids $M_1, M_2, \ldots M_K$ of M. As M_i is itself a matroid, the greedy algorithm can be run on M_i to produce a minimal set of i disjoint bases of M. We now prove the following lemmas, which will enable us to run the greedy algorithm simultaneously on all M_i. For any time t, extend the $S(t)$ notation to the M_k by defining $S_k(t)$ to be the independent set of M_k of maximum size in $\{x_1, \dots, x_t\}$

Lemma 5.4.18. *For all t, $S_k(t) \subseteq S_{k+1}(t)$.*

Proof. We proceed by induction over t. The base case at time $t = 0$ is trivial. Suppose e appears at time $t + 1$, and is added to $S_k(t)$. As the inductive hypothesis assumes that $S_k(t) \subseteq S_{k+1}(t)$, by Lemma 5.4.15, $S_{k+1}(t)$ can be decomposed into the disjoint sets A and B, such that A is independent in M_k, B is independent in M, and $|A| = |S_k(t)|$. Now as $S_k(t) \cup e$ is an independent set in M_k that is larger than A, there is an element $x \in S_k(t) \cup e$ such that $A \cup x$ is independent in M_k as well. If $x \neq e$, then $A \cup x$ is contained in $S_{k+1}(t)$ but is larger than $S_k(t)$, contradicting the optimality of $S_k(t)$. Hence $x = e$, and the decomposition $A \cup e$ and B witness that $S_{k+1}(t) \cup e$ is independent in M_{k+1}. $\square$

To simultaneously compute optimal bases for all M_i, we consider the elements of X in sorted order as before. Maintain sets $\Delta_k(t)$, and add x_t to $\Delta_k(t-1)$ where k is the smallest i such that $\left(\bigcup_{j \leq i} \Delta_j(t-1)\right) \cup x_t$ is independent in M_i. The current independent set for a particular order k is given by $S_k(t) = \bigcup_{1 \leq i \leq k} \Delta_i(t)$. Finally, we let S_k or Δ_k refer to the final set, after all input elements have been considered.

We also have the following extension to Lemma 5.4.15.

Lemma 5.4.19. *There exists a decomposition of $S_k(t)$ into k independent sets $T_1, \dots, T_k$ of M such that $|T_i| = |\Delta_i(t)|$.*

Proof. We proceed by induction down from k. Let $R_k = S_k(t)$. Then by Lemma 5.4.15, R_k can be decomposed into $R_{k-1} \in M_{k-1}$ and $T_k \in M$, with $|R_{k-1}| = |S_{k-1}|$. Suppose now that we have $R_i \subseteq R_{i+1} \subseteq \cdots \subseteq R_k$ and $R_j \cup T_{j+1} \cup \cdots \cup T_k = S_k(t)$ for $j = i, \dots, k-1$, with $T_h \cap T_\ell = \emptyset$ and $R_j \cap T_h = \emptyset$ for all $\ell, h > j$ with $\ell \neq h$. Let $U = T_{i+1} \cup \cdots \cup T_k$, and let R_i decompose into R^1 and R^2 with $R^1 \in M_{i-1}$ and $R^2 \in M$. Apply Lemma 5.4.15 on $R^1 \cup (R^2 \cup U)$ and $S_{i-1}(t)$ to get $R_{i-1} \cup U' = S_k(t)$ with $R_{i-1} \in M_{i-1}$, $U' \in M_{k-i+1}$ and $|R_{i-1}| = |S_{i-1}(t)|$. Then $R_{i-1} = R^1 \cup R' \cup U'$ for some $R' \subseteq R^2$ and $U' \subseteq U$. If $R_{i-1} \cap U' \neq \emptyset$, then $|R_{i-1} \cup R^2| > |R^1 \cup R^2| = |S_i(t)|$ but is independent in M_i, contradicting the maximality

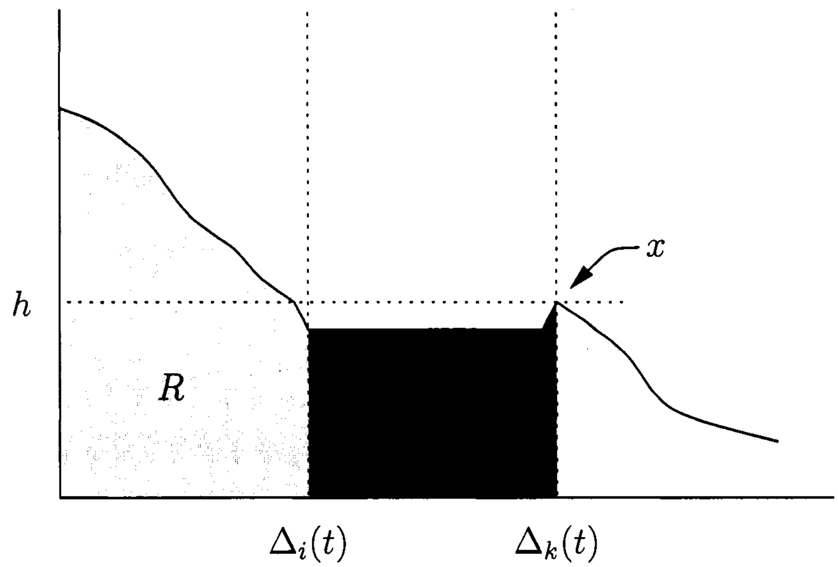

Figure 5.6: The impossible case of Lemma 5.4.20

of $S_i(t)$ guaranteed by Lemma 5.4.17. Hence $R_{i-1} \subseteq R_i$, $T_i = R^2 \setminus R'$, and T_j is unchanged for $j > i$. □

We will also need the following structural lemma on the Δ_is.

Lemma 5.4.20. $|\Delta_i(t)| \geq |\Delta_{i+1}(t)|$ *for all* t *and* $1 \leq i < K$.

Proof. Suppose the lemma is false. Choose t to be the time immediately after the first violation occurs, so that the situation can be described as in the Figure 5.6, with $|\Delta_j(t)| \geq |\Delta_{j'}(t)|$ for all $k > j' > j$, $h = |\Delta_k(t)| = |\Delta_{k-1}(t)| + 1$, $|\Delta_j(t)| = |\Delta_i(t)|$ for all $i < j < k$ (note that k may equal $i + 1$), and the element x was just added to Δ_k.

Decompose $S_k(t)$ into $R \cup U$ with $R \in M_{i-1}$, $U \in M_{k-i+1}$ and $R \cap U = \emptyset$; by Lemma 5.4.15 we can assume that $|R| = |S_{i-1}(t)|$. Note that $|S_i(t)| = |R|+h-1$. Let $U_1 \cup \cdots \cup U_{k-i+1}$ be a partition of U with each $U_j \in M$. As $|U| = (k-i)(h-1)+h$, there must exist a j with $|T_j| \geq h$. Then $R \cup U_j$ is an element of M_i of rank $|R| + h > |S_i(t)| = |R| + h - 1$, contradicting Lemma 5.4.17. □

Corollary 5.4.21. $\mathrm{c}(\Delta_i) \leq \mathrm{c}(\Delta_{i+1})$ *for all* $1 \leq i < K$.

Proof. As $i < K$, the maximum number of disjoint bases in M, both S_i and S_{i+1} are bases in M_i and M_{i+1}, respectively, so that $|\Delta_i| = |\Delta_{i+1}| = \rho(M) = n$. Let $\Delta_i = \{x_{h_1}, \ldots, x_{h_n}\}$ and $\Delta_{i+1} = \{x_{k_1}, \ldots, x_{k_n}\}$, where the x_j are indexed in order of weight as before. By the previous lemma, $|\Delta_i(t)| \geq |\Delta_{i+1}(t)|$, implying $h_j < k_j$ for each $j = 1, \ldots, n$. Hence $\mathrm{c}(x_{h_j}) \leq \mathrm{c}(x_{k_j})$ for each such j, and the corollary follows. □

5.4.4 VCG Payments & Matroids

We close this section on matroids by examining the structure of VCG payments. For any set system, given the bids associated with other elements, each element has a threshold bid that determines its payment. In the case of matroids, we will now show that this threshold bid for a element e corresponds to bid of another player, which is a replacement for e in the optimal set when e is removed from the auction. Given a base B containing e, we define r_e^B to be the cheapest element x such that $(B \setminus e) \cup x$ is a base. When the base is clear from context, we will simply write r_e. This element will be called the *cheapest replacement* for e.

Lemma 5.4.22. *If $S = \mathrm{OPT}(X)$, and if $e \in S$ has cheapest replacement r_e, then* $\mathrm{OPT}(X \setminus e) = \mathrm{c}\big((S \setminus e) \cup r_e\big)$.

Proof. Suppose not. Let $S' = \mathrm{OPT}(X \setminus e)$, and let x be an element not equal to r_e in S', that is not in S. As $x \notin S$, it must have been rejected by the greedy algorithm while constructing S, which means it forms a cycle C_x with elements in S. As it does not form a cycle in S', and e is the only element in S not in S', $e \in C_x \subset S \cup \{x, r_e\}$. As r_e is a replacement for e, there is a cycle $C_r \subset S \cup r_e$. As $e \in C_x \cap C_r$, by Fact 5.4.4, there is a cycle $B \subset (C_x \cup C_r) \setminus e \subset (S \setminus e) \cup \{x, r_e\} \subset S'$. But S' is a base and so cannot contain a cycle. □

Corollary 5.4.23. *The VCG payment to an element is the cost of its cheapest replacement.*

The following lemma says that we can find replacements as we build up the sets in the greedy algorithm, without needing to wait until the sets are completed.

Lemma 5.4.24. *Suppose e is added to S at time t, so that $e \in S(t) \setminus S(t-1)$. Then if u is the smallest time $> t$ such that $x_u \notin S(u)$ but $(S(u) \setminus e) \cup x_u$ is independent, then $x_u = r_e$.*

Proof. As x_u is dependent on $S(u)$ but independent of $S(u) \setminus e$, $e \in C$ where C is the canonical cycle of x_u in $S(u)$. Then $C \subset S$ is still a cycle, so that $S \setminus e \cup x_u$ is independent and hence x_u is a replacement for e. Suppose there were a cheaper replacement r for e. Then and $S \setminus e \cup r$ is independent, $S(t) \setminus e \cup r$ is independent as a subset. However, x_u is defined to be the cheapest such element. □

The existence of a single element replacement characterizes matroids.

Lemma 5.4.25. *If $\mathcal{A}$ is not the collection of bases of a matroid, then there is a cost vector c, feasible sets A and B in $\mathcal{A}$ and distinct elements e and u in $A \setminus B$ such that* $\mathrm{OPT}(\mathcal{A}, c) = A$ *is the unique optimal set for c, and* $\mathrm{OPT}(\mathcal{A} \setminus u, c) = B$*, for some integer cost vector c, with $e \notin B$. In other words, the cheapest replacement set for u replaces e as well.*

Proof. As $\mathcal{A}$ is not a matroid, then as the basis exchange axiom of Fact 5.4.2 is not satisfied, there are feasible sets S and T with an element $x \in S \setminus T$ such that for all $y \in T$, $(S \setminus x) \cup y$ is not feasible. Let C be a minimum cardinality feasible set in $S \cup T$ such that $x \notin C$.

Case 1: $S \setminus C$ contains at least one element other than x. In this case, let $A := S$, $B := C$, $u := x$ and let e be any other element in $S \setminus C$, with the cost vector

$$c(z) = \begin{cases} 0 & \text{if } z \in A, \\ 1 & \text{if } z \in B \setminus A, \\ |B \setminus A| + 1 & \text{otherwise.} \end{cases}$$

Then $A = \mathrm{OPT}(\mathcal{A}, c)$, and by the minimality of C, $B = \mathrm{OPT}(\mathcal{A} \setminus u, c)$.

Case 2: $S \setminus C$ contains only x. Let $X = S \cap C = S \setminus x$ and $Y = C \setminus S$. Note $|Y| \geq 2$ by our initial (non-matroid derived) assumption on S and T. Then as feasible sets do not nest, the only possible feasible sets in $S \cup C$ aside from S and C are of the form $x \cup W \cup Z$ where W and Z are strict subsets of X and Y, respectively, and Z is nonempty. Define the cost vector c to be

$$c(z) = \begin{cases} 0 & \text{if } z \in X, \\ 1 & \text{if } z \in Y, \\ |Y| + 1 & \text{if } z = x, \\ 2|Y| + 3 & \text{otherwise.} \end{cases}$$

Then $C = \text{OPT}(\mathcal{A}, c)$, and for any $y \in Y$, $S = \text{OPT}(\mathcal{A} \setminus y, c)$. To see the latter, observe that $c(x \cup W \cup Z) \geq |Y| + 2$ where W and Z are as above. Thus with $A := C$, $B := S$ and u and e two elements from Y, we have what we need. □

5.5 *Mechanism Equivalence*

While digital goods auctions have mechanisms that are competitive with an omniscient auctioneer with access to private player values, the path procurement problem shows that in general this is probably too strong a result to achieve. Instead of using an omniscient auction as a lower bound for an optimal mechanism, we must bound the optimal revenue more directly.

The theory of auction mechanisms in the Bayesian setting has developed a very strong class of results known as *revenue equivalence*. They assert that a market of identical bidders will in equilibrium pay the same prices, and produce the same revenue, in all auction mechanisms that produce the same allocation. In particular, any mechanism that produces an efficient allocation will produce the same revenue, so that any such mechanism is optimal. This may be interpreted as an expression of the efficiency of markets, and the limitations of mechanism design.

Klemperer [33] gives a survey of revenue equivalence, and a very readable development is found in Milgrom [35]. As a formal development of this theory would take us too far off track, we only sketch the result and its implications. The equivalence results follow simply from the so-called envelope theorem, which gives conditions on when the utilities from optimal strategies will depend on the allocation function alone. One of the conditions is that the player utilities be quasi-linear, which is a common assumption and one that we use here. The result is that in settings where it applies, mechanisms which produce the same allocation have optimal strategies that lead to the same revenue and player payoffs. Results from this include the fact that *ex post*, any efficient mechanism produces the same revenue as VCG. In particular, the revenue from the first and second price single-item auctions is the same. This gives a powerful way to compute optimal strategies: one can first compute payoffs under some mechanism, like VCG, where payoffs are easy to compute. Then, using payoff equivalence, one finds the strategy that results in the same payoff.

Revenue equivalence gives limits of mechanism design, by showing that the market will arrive at the same outcome regardless of the mechanism forced upon it. Equivalence between a truthful mechanism and non-truthful mechanisms means that a designer is free to use a truthful mechanism, which may be much easier to analyze than an alternative, knowing that she is not sacrificing revenue in the process.

5.5.1 Revenue Equivalence in Matroid Procurement Auctions

The following definition is taken from Karlin et al. [29], where it was used to characterize the revenue from both the first-price and VCG mechanisms for matroid procurement auction.

Definition 5.5.1. Let M be a matroid, and let S be an optimal base under a cost vector $\mathbf{c}$. Then $\nu(\mathbf{c})$ is defined to be the solution to the following optimization problem.

Minimize $\sum_{e \in S} b_e$ subject to

1. $b_e \geq c_e$ for all e
2. $\sum_{e \in S \setminus T} b_e \leq \sum_{e \in T \setminus S} c_e$ for all $T \in \mathcal{B}(M)$
3. For every $e \in S$, there is a $T_e \in \mathcal{B}(M)$ such that $e \notin T_e$ and $\sum_{u \in S \setminus T_e} b_u = \sum_{u \in T_e \setminus S} c_u$.

Thus $\nu(\mathbf{c})$ is the total cost of the minimum efficient Nash equilibrium (if one exists). The first constraint says that all players bid at least their values. The second says that the efficient allocation is indeed the winning set in the equilibrium. The final constraint asserts that this is an equilibrium, that is, for every player in S, if that player raises her bid she will no longer be in the winning set.

Lemma 5.5.2. *Let M be a matroid and S an optimal base under cost vector* $\mathbf{c}$ *where all costs are distinct. If S is the unique optimal base and a Nash equilibrium exists for* $\mathbf{c}$*, then all Nash equilibria for* $\mathbf{c}$ *have S as the winning set.*

Proof. Let $\mathbf{b}$ be the bid vector for a Nash equilibrium of $\mathbf{c}$, and suppose T is the winning set under $\mathbf{b}$. If $T \neq S$, then there is an $e \in S \backslash T$ as feasible sets do not nest by Assumption 5.3.2. As $e \in S$, e improves T under $\mathbf{c}$, and as S is unique there is a u such that $T' = (T \backslash u) \cup e \in M$ and $\mathbf{c}(T') < \mathbf{c}(T)$, also implying that $c_e < c_u$ (here we use the distinctness of costs to get the strict inequalities). As $e \notin T$, it must be that $b_e > b_u$. We know that $b_u \geq c_u$, otherwise u would change her bid to b_u and make zero rather than negative utility. Hence e can bid b' with $c_e < b' < b_u < b_e$ and be included in the winning, making positive utility. As with her current bid e has zero utility, this contradicts $\mathbf{b}$ being a Nash equilibrium. □

Lemma 5.5.3. *If M is a matroid and* $\mathbf{c}$ *a monopoly-free cost vector for M, then for any $\varepsilon > 0$,* $\mathbf{c}$ *has an ε-Nash equilibrium.*

Proof. Adjust $\mathbf{c}$ within ε to make $\mathbf{c}'$ with all costs are distinct. Then there is a unique optimal base S under $\mathbf{c}'$. For all $e \in S$, let b_e be the cost of the cheapest replacement for e under $\mathbf{c}'$, and for all $e \in M \setminus S$, let $b_e = c'_e$. Then S is the winning set under $\mathbf{b}$, by the definition of cheapest replacement. Any $e \in S$ has no incentive to raise her bid, as that will cause her replacement to be chosen instead. No $e \notin S$ can lower her bid and make positive utility. Hence, $\mathbf{b}$ is a Nash equilibrium for $\mathbf{c}'$, and so is an ε-Nash equilibrium for $\mathbf{c}$. □

Thus $\nu(\mathbf{c})$ characterizes equilibrium bids in first-price auctions for monopoly-free matroids. The following result from Karlin et al. [29] shows that the VCG makes the same revenue. It furthermore states that VCG makes the revenue $\nu(\mathbf{c})$ if and only if the set system is a matroid. In this case, $\nu(\mathbf{c})$ may no longer be the minimum Nash equilibrium revenue, as if the set system is not a matroid, Lemma 5.5.2 does not hold, so that $\nu(\mathbf{c})$ may not be the minimum Nash equilibrium revenue. However, $\nu(\mathbf{c})$ does describe some Nash equilibrium, namely the one corresponding to the VCG mechanism, so that the existence of a lower revenue equilibrium only strengthens the separation between matroids and general set systems.

Lemma 5.5.4 ([29]). *In a procurement auction, the VCG mechanism generates at least the revenue from any first-price auction in Nash equilibrium if and only if the set system is a matroid.*

This theorem may be interpreted as a revenue equivalence result for procurement auctions. Instead of comparing revenue of the VCG mechanism against that achieved by some omniscient auctioneer as we did for digital goods auctions, we more directly address the question of the optimality of VCG by showing it always produces the revenue of a first-price auction. While we have not excluded the possibility of some mechanism creating much more revenue, the first-price auction is such a well-known benchmark that this result does show that VCG is a reasonable mechanism to use for matroid procurement.

Such a result is not possible for path auctions. The following lemma shows how the first-price auction can achieve smaller cost for an optimal path at equilibrium, suggesting that a non-truthful mechanism is needed for frugal procurement of a single path.

Lemma 5.5.5. *There is a path auction instance where a minimum cost first-price auction with players in Nash equilibrium is achieved on the optimal path, and the cost is less than the VCG cost.*

Proof. Consider the graph with two disjoint paths: P, with n edges each of cost 1, and Q, with 1 edge of cost $n + \varepsilon$. Then a minimum ε-Nash equilibrium has P winning the auction, with the total payments to all edges of P equal to $n+\varepsilon$. In contrast, as described previously, the VCG cost is larger, $(1 + \varepsilon)n$. □

Chapter 6

NEAR-OPTIMAL MATROID PROCUREMENT

In this chapter we describe our mechanism for multi-unit matroid procurement. Like the partition and cost-share mechanism for digital goods, our mechanism consists of two parts. We first discuss profit extraction, showing that cost-sharing is truthful for matroids. Furthermore, we show that profit extraction is generally not truthful for non-matroids, and hence any procurement mechanism that is effective over non-matroids will require new technology. We then examine the VCG cost of a random sample from a matroid and show that it scales proportionally to the size of the sample, analogous to proof of Theorem 5.1.7.

Recall from Section 5.3 that the profit of an auctioneer is determined by the benefit function $B(\cdot)$, in that procuring k disjoint feasible sets at cost P_k generates profit $B(k) - P_k$. We would like the mechanism to obtain a profit close to a natural benchmark of optimality, $\mathcal{OPT} = \max_k B(k) - P_k$, when the payments P_k are given by VCG_k, the VCG mechanism for procuring k disjoint sets. We review how the benchmark $\mathcal{OPT}$ can be approximated through a reduction to the decision problem. The mechanism-design decision problem for objective $\mathcal{OPT}$ is to give a truthful mechanism parametrized by a target profit R that gives an outcome and payment with profit at least R when $R \leq \mathcal{OPT}$. A solution to the mechanism design decision problem is called a *profit extractor.*

Definition 6.0.6 (RSPE). The *Random Sampling Profit Extraction* auction (RSPE) on a set system $\mathcal{A}$ over a base set X.

1. Randomly partition the agents X into two parts X' and X'' by flipping an unbiased coin for each agent.

2. Compute the optimal benchmark on each part: $R' = \mathcal{OPT}(X')$ and $R'' = \mathcal{OPT}(X'')$.

3. Run the profit extractor with R'' on X' and likewise with R' on X''.

Clearly, RSPE is truthful for bidders in X' (likewise for X'') as no bidder in X' can affect the value of $R'' = \mathcal{OPT}(X'')$ and because the profit extractor with R'' on X' is truthful. The profit of this auction is at least $\min(R', R'')$. Thus if the expected minimum of $\mathcal{OPT}(X')$ and $\mathcal{OPT}(X'')$ is a good approximation to $\mathcal{OPT}(X)$ then this reduction approach gives a good approximation. This approximate reduction is from Fiat et al. [19].

To apply this reduction to our problem domain, we need

1. to design a truthful profit extractor, and

2. to prove that the benchmark, $\mathcal{OPT}$, on a random sample is large enough.

We will consider the following algorithm as a candidate solution to the decision problem. This algorithm is a generalization of one given in Deshmukh et al. [13] for the double auction problem which is based on the cost-sharing mechanism discussed in Section 5.1.5

Definition 6.0.7 ($\mathcal{OPT}$-profit Extraction). The *$\mathcal{OPT}$-profit extraction* algorithm with target R:

1. Find the largest k such that the the payments P_k of VCG$_k$ satisfy $B(k) - P_k \geq R$.

2. If such a k exists, output output $S = S_k$ and the VCG$_k$ payments.

3. Otherwise, output $S = \emptyset$ and zero payments.

If all agents bid truthfully, this algorithm gives a profit of at least R if and only if $R \leq \mathcal{OPT}$. We show that if $\mathcal{A}$ are the bases of a matroid then this algorithm gives a truthful mechanism.

6.1 Profit Extraction and Matroids

Theorem 6.1.1. *The $\mathcal{OPT}$-profit extractor is truthful for matroid set systems.*

Proof. We first fix b_{-e} and show that the threshold bid $p_k(e)$ of agent e in VCG$_k$ is nondecreasing with k. Suppose not. Then for some set system $(E, \mathcal{F})$, $S_k \subseteq S_{k+1}$ for all k, but for some k and some bid vector b_{-e}, $p_e(k) > p_e(k+1)$. Let $p_e(k+1) < b_e < p_e(k)$. Then $e \in S_k$, but not in S_{k+1}, contradicting Lemma 5.4.18.

We now show that if we fix the winners S_k of VCG$_k$ and the bids of all losers in $T = E \backslash S_k$ then the payment P_k of VCG$_k$ is also fixed. Corollary 5.4.23 states that if S_k is the cheapest feasible set in $\mathcal{A}_k$ then there is some element $y \in T$ that we would replace agent $e \in S_k$ with if we were to choose the cheapest feasible set from $E \setminus \{e\}$. Since the bids of the agents in T are fixed, this replacement y of minimal cost is fixed. By the definition of the VCG payment rule, the payment to edge e is

$$p_e(k) = \min_{S' \in \mathcal{A}_k \,:\, e \notin S'} \sum_{e \in S'} c_e - \sum_{e \in S_k \backslash \{e\}} c_e = c_y,$$

which is fixed. As P_k is the sum over all of these fixed payments it is also fixed. This implies that no winner can change the total payment P_k from VCG$_k$ without losing.

These two facts imply the theorem as follows. Fix the bids b_{-e} of all agents except e. Let k^* be the maximum k such that $B(k) - P_k \geq R$, when e bids truthfully and the other agents bid b_{-e}. Suppose $e \in S_{k^*}$. Then e's actual value c_e is at most $p_e(k^*)$. e cannot bid below c_e without risking a net loss. Bidding between c_e and $p_e(k^*)$ will not change the payment to e, and risks losing. If e bids above $p_e(k^*)$, she will lose for sure, since for any $k > k^*$ and fixed b_{-e}, for all $b_e \leq p_e(k)$, P_k is constant, and therefore, $B(k) - P_k < R$ by definition of k^*. Now suppose that $e \notin S_{k^*}$ when e bids truthfully. Then $c_e > p_e(k^*)$. Raising her bid won't change the outcome, since for all $k > k^*$, the profit extracted is less than R, while lowering her bid could cause her to win, but only at a net loss. □

6.2 Limits of Cost-Sharing Profit Extraction

The goal of this section is to show that cost-sharing profit extraction is truthful exactly for matroids.

Theorem 6.2.1. *The $\mathcal{OPT}$-profit extraction algorithm does not give a truthful mechanism for non-matroid set systems. In other words, for any set system that is not a matroid and any marginally decreasing $B(\cdot)$, there is a set of private values c and a choice of R for which the profit extractor is not truthful.*

Proof. Let $\mathcal{A}$ be a set system that is not a matroid. We begin by applying Lemma 5.4.25 to find a cost vector c, $A, B \in \mathcal{A}$, and distinct elements $e, u \in A \setminus B$ such that $\mathrm{OPT}_1(\mathcal{A}, c) = A$ and $\mathrm{OPT}_1(\mathcal{A} \setminus u, c) = B$. We now modify the cost vector c. Let S be a union of two disjoint feasible sets that minimizes $|S \setminus (A \cup B)|$. Raising the costs of elements outside of $A \cup B \cup S$ does not change the properties of A and B, so we may change c so that the cost of any such element is very large; $M = 1 + K \cdot \big(|A| \cdot c(S) + c(A \cup B \cup S)\big)$ will suffice, where K is the maximum number of disjoint feasible sets in $\mathcal{A}$.. Then the following holds.

Claim 6.2.2. *Under the cost vector c, $\mathrm{VCG}_k > k \cdot \mathrm{VCG}_1$ for all $k > 1$. In addition, $A = \mathrm{OPT}_1(\mathcal{A}, c)$ and $S = \mathrm{OPT}_2(\mathcal{A}, c)$.*

To see this, as it is clear that $A = \mathrm{OPT}_1(\mathcal{A}, c)$, we first show that for any $k > 1$, $\mathrm{OPT}_k(\mathcal{A}, c)$ is not strictly contained in $A \cup B$. Let $\mathrm{OPT}_k(\mathcal{A}, c) = S_1 + S_2 + \cdots + S_k$, where S_i, $1 \leq i \leq k$ are disjoint feasible sets, and labeled so that $u \notin S_1$. Then $c(S_1) \geq c(B)$, as B is an optimal feasible set among those not containing u, and $c(S_2) \geq c(A)$, as A is an overall optimal feasible set. Hence, $c\big(\mathrm{OPT}_k(\mathcal{A}, c)\big) \geq c(S_1) + c(S_2) \geq c(A) + c(B) \geq c(A \cup B)$. As costs are all positive, either $\mathrm{OPT}_k(\mathcal{A}, c) = A \cup B$, or $\mathrm{OPT}_k(\mathcal{A}, c) \setminus (A \cup B) \neq \emptyset$. In particular, this shows that $S = \mathrm{OPT}_2(\mathcal{A}, c)$.

We next upper bound VCG_1. Let S_1 and S_2 be a decomposition of S into two disjoint feasible sets. Given any $x \in A$, at least one of S_1 or S_2 does not contain x and serves as a replacement set, so that $c\big(\mathrm{OPT}_1(\mathcal{A} \setminus x), c\big) \leq c(S)$. As the payment to x is $c\big(\mathrm{OPT}_1(\mathcal{A} \setminus x), c\big) - c\big(\mathrm{OPT}_1(\mathcal{A}, c)\big) + c(x)$, $c(S)$ also (loosely) bounds the payment to x as well. Hence $\mathrm{VCG}_1 \leq |A| \cdot c(S)$.

We can now finish the claim. Let $Z = \mathrm{OPT}_k(\mathcal{A}, c)$ for any $k > 1$. If Z contains an element outside of $A \cup B \cup S$, then as payments are greater than costs, we have that $\mathrm{VCG}_k \geq M > K \cdot \mathrm{VCG}_1 \geq k \cdot \mathrm{VCG}_1$, and we are done. The second possibility is that $Z \subseteq S \cup A \cup B$ and contains an element $x \in S \setminus (A \cup B)$. Let $r = |S \setminus (A \cup B)|$, so that $\mathcal{A} \setminus x$ contains only $r - 1$ elements outside $A \cup B$ of cost less than M. As $\mathrm{OPT}_k(\mathcal{A} \setminus x, c)$ contains at least two disjoint feasible sets, and S is such a pair of sets that minimizes the number of elements outside $A \cup B$, $\mathrm{OPT}_k(\mathcal{A} \setminus x, c)$ must also contain at least r elements outside of $A \cup B$, and so must contain an element of cost M. Thus we have that the payment to any such element is at least $M - c(Z) + c(x) > K \cdot \mathrm{VCG}_1 \geq k \cdot \mathrm{VCG}_1$. The final possiblity, as Z is not strictly contained in $A \cup B$, is that $Z = A \cup B$. Then for any $x \in Z$, $\mathrm{OPT}_k(\mathcal{A} \setminus x, c) \setminus (A \cup B) \neq \emptyset$, otherwise strict containment would be violated. Hence we again have that the payment to x is at least $M - c(Z) + c(x) > K \cdot \mathrm{VCG}_1 \geq k \cdot \mathrm{VCG}_1$.

We can now contradict truthfulness by showing that e can raise its bid and cause VCG_1 to go down. We will choose the revenue goal and benefit function so that if all elements bid truthfully, the buyer can nearly but not quite meet the goal. In particular, e will receive no utility from the canceled auction. However, e's overbidding and subsequent reduction in VCG_1 will cause the buyer to meet the goal at $k = 1$, and provide e with positive utility, and thus incentive to bid non-truthfully.

Note that the chosen cost vector is integral, so that as $\mathrm{OPT}(\mathcal{A}, c)$ is unique, $c(e)$ is at least one less than the threshold for e to be included in the optimal set. Hence if e bids $c(e) + 1/2$, e will remain in the optimal set. Let c' be this cost vector, that is, $c'_{-e} = c_{-e}$ and $c'(e) = c(e) + 1/2$. Furthermore, as the bid of e is increased, e will still be excluded from $\mathrm{OPT}(\mathcal{A} \setminus u, c')$. Define $p_1(v, c)$ to be the VCG_1 payment to v under c. Then $p_1(u, c) - p_1(u, c') = \big(c(B) - c(A) + c(u)\big) - \big(c'(B) - c'(A) + c(u)\big) = -1/2$, so that the payment to u decreases. For any other element $v \in A$, the threshold for e to be in $\mathrm{OPT}(\mathcal{A} \setminus v, c)$ is an integer, so $e \in \mathrm{OPT}(\mathcal{A} \setminus v, c)$ if and only if $e \in \mathrm{OPT}(\mathcal{A} \setminus v, c')$, and so $p_1(v, c) - p_1(v, c')$ is either 0 or $-1/2$. Summing these payment differences over all $v \in A$ thus gives $\mathrm{VCG}_1(c') \leq \mathrm{VCG}_1(c) - 1/2$, where $\mathrm{VCG}_1(v)$ is the VCG payment under cost vector v.

Now choose $L_0 > 1$ such that $\mathrm{VCG}_1(c') < L_0$ and $L_0 \cdot k < \mathrm{VCG}_k(c)$ for all $k \geq 1$; such an L_0 exists as c is an integer cost vector and $\mathrm{VCG}_k(c) > K \cdot \mathrm{VCG}_1(c)$ for all $k > 1$. If the benefit function were $B_0(k) = L_0 \cdot k$, then choosing a revenue target R as $R = B_0(k) - \mathrm{VCG}_1(c) + 1/4$ would suffice. If all elements bid their values, revenue R cannot be extracted, as $B(k) - \mathrm{VCG}_k(c) < 0$. On the other hand, if e raises her bid to $1/2$, then revenue at least $R + 1/4$ can be extracted at $k = 1$. Hence there is incentive for e to overbid. Given $B(\cdot)$ from the lemma of the statement, choose L so that $B(k) \leq L \cdot k$; such an L exists as $B(\cdot)$ is marginally decreasing. Then scaling the costs of all elements by L/L_0 gives the required cost vector. $\square$

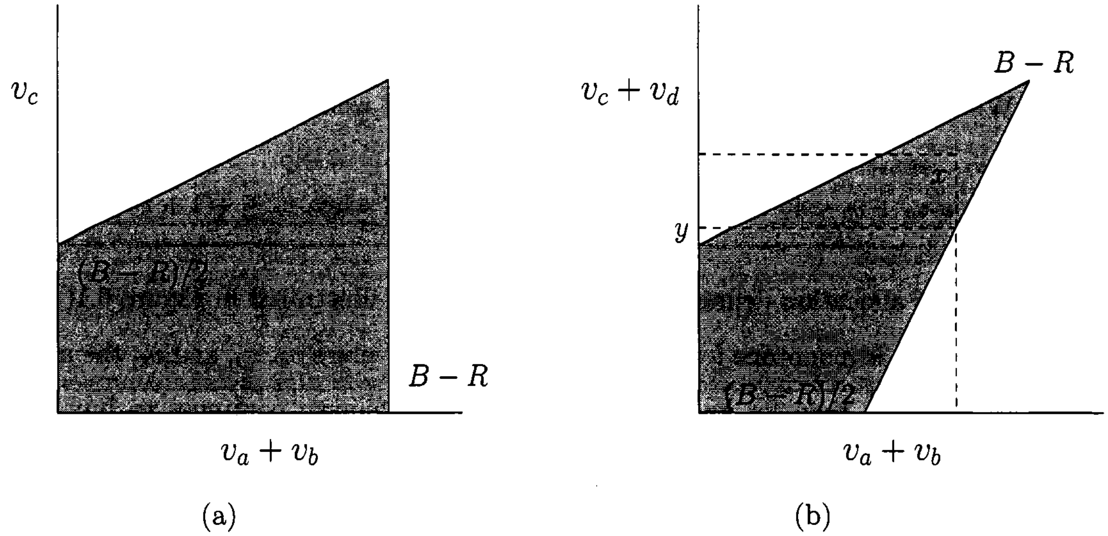

Figure 6.1: Allocation regions used in Lemmas 6.2.3 and 6.2.4

6.2.1 General Extraction

Theorem 6.2.1 shows that cost-sharing profit extraction, based on VCG, is not truthful. We now show two examples, which consider general profit extraction, which may not use the VCG mechanism to compute prices. In particular, we show one non-matroid set system which has a truthful profit extractor, and one for which no truthful profit extractor exists. The first example is rather artificial, and is not obviously extended to a general class of non-matroids. It is our conjecture that there is no natural profit extractor for non-matroids; the challenge is to determine what "natural" means in this context.

Lemma 6.2.3. *There exists non-matroid set system for which there is a truthful profit extractor.*

Proof. Define a set system $\mathcal{A} = \big\{\{a, b\}, \{c\}\big\}$ with agent values given by v_a, v_b and v_c. Note that it is possible to procure at most one set from $\mathcal{A}$ with the VCG mechanism, as there is no alternative pair of sets when any agent is removed. Hence the benefit function is expressed simply as a single number B. Given target profit R, define a mechanism for profit extraction buys the set $\{c\}$ if an only if $\mathcal{OPT}(\{a, b, c\}) \geq R$. The mechanism never buys the set $\{a, b\}$. This gives two constraints on allocation:

1. When $v_a + v_b > v_c$ then $\mathcal{OPT}$ meets the target R when $v_a + v_b < B - R$.

2. When $v_a + v_b < v_c$ then $\mathcal{OPT}$ meets the target R when $v_c \leq \frac{1}{2}(1 + v_b + B - R)$.

The region of allocation defined by these constraints is illustrated in Figure 6.1(a). Notice that this region is monotone for $\{c\}$, that is, if for some values v_a, v_b, and v_c, the mechanism buys the set $\{c\}$, then for lower values of v_c the mechanism continues to buy $\{c\}$. The threshold bid and payment for c is given by $(v_a + v_b + B - R)/2$. This implies that the mechanism is truthful: if c overbids, she risks going over her threshold and missing the allocation, while underbidding only changes the allocation when the payment is less than v_c.

By construction this mechanism allocates whenever $\mathcal{OPT}(\{a, b, c\}) \geq R$. It remains to show that the our profit is at least R whenever we allocate. The payment to c when we allocate is $p_c = (v_a + v_b + B - R)/2$. We show that $B - p_c \geq R$ whenever we are in the region of allocation.

$$\begin{aligned} B - p_c &= B - \tfrac{1}{2}(v_a + v_b + B - R) \\ &= \tfrac{1}{2}(B - v_a - v_b + R) \end{aligned}$$

However, in the region of allocation $v_a + v_b \leq B - R$, implying that $R \leq B - v_a - v_b$, which gives $B - p_c \geq R$. □

Lemma 6.2.4. *There exists a non-matroid set system and benefit such that there is no truthful profit extractor for* $\mathcal{OPT}$.

Proof. Let $\mathcal{A} = \{\{a, b\}, \{c, d\}\}$. As before, it is only possible for the VCG mechanism to procure one set. If $v_a + v_b < v_c + v_d$, then $\mathrm{VCG}_1 = 2(v_c + v_d) - (v_a + v_b)$, and conversely for $v_c + v_d < v_a + v_b$. Thus to compete with the omniscient VCG extractor, an extractor must produce revenue over the allocation region defined by $B - 2(v_c + v_d) + (v_a + v_b) > R$ and $B - 2(v_a + v_b) + (v_c + v_d) > R$, as shown in Figure 6.1(b).

Consider the point x at $\left(\frac{5}{6}(B-R), \frac{5}{6}(B-R)\right)$. It cannot be allocated to $\{c, d\}$ by a truthful mechanism. To see this, note that if it were allocated to $\{c, d\}$, then the allocation would not be monotone as $v_c + v_d$ varies. Consider the value y shown in Figure 6.1(b). If $v_c + v_d < y$, then c has incentive to raise her bid until $v_c + v_d$ is large enough to force allocation, showing that the mechanism would not be truthful. For similar reasons, it cannot be allocated to $\{v_a, v_b\}$. However, as x is in the region of allocation, this means no truthful extractor is possible. □

6.3 Random Sampling of Matroids and VCG payments

As discussed at the beginning of this chapter, the random sampling reduction to the decision problem requires that the value of $\mathcal{OPT}$ on a random sample of the elements in the ground set be close to that of the full set. In this section we prove that with high probability $\mathcal{OPT}$ of a random sample is a constant fraction of $\mathcal{OPT}$ on the full set. This shows that the Random Sampling Profit Extraction auction is a constant approximation. This section is analogous to the proof of Theorem 5.1.7, where the revenue from looking at one part of a random partition of bidders is shown to be about half of the optimum when there is a lower bound on the number of bidders participating in the optimal solution.

With the following two technical lemmas we prove our main theorem.

Lemma 6.3.1. *Let M be a matroid over X. Let $m = \lfloor (1-\varepsilon)k/2 \rfloor$ for some constant $\varepsilon > 0$ and $k \geq \frac{8}{\varepsilon^2} \log n$, where $n = \rho(M)$. With probability $1 - 1/n$, P'_m, the* VCG_m *payments for m disjoint bases in the sample X' satisfies:*

$$P'_m \leq m \cdot \mathrm{c}(\Delta_k).$$

Lemma 6.3.2. *Let M be a matroid over X. The cost P_k paid by* VCG_k *satisfies*

$$P_k \geq k \cdot \mathrm{c}(\Delta_k).$$

These enable our main theorem.

Theorem 6.3.3. *Let M be a matroid over X. Let $k^* = \mathrm{argmax}_k\, B(k) - P_k$ and $R = B(k^*) - P_{k^*} = \mathcal{OPT}$. For any $\varepsilon > 0$, the RSPE procurement mechanism obtains a profit* *that is at least* $\alpha = (1-\epsilon)/2$ *of* R *with probability* $1 - 2/n$*, where* $n = \rho(M)$*, provided* $k^* \geq \frac{8}{\epsilon^2} \log n$.

Proof. Sample to get X' and X'', and compute the optimal revenues R' and R''. We claim that both R' and R'' are at least αR. To see this, observe that by Lemma 6.3.1, with probability at least $1 - 1/n$, $P'_{\alpha k^*} \leq \alpha k^* \mathrm{c}(\Delta_{k^*})$ and by Lemma 6.3.2, $k^* \mathrm{c}(\Delta_{k^*}) \leq P_{k^*}$. Thus $P'_{\alpha k^*} \leq \alpha P_{k^*}$. It is easy to show that if $B(\cdot)$ is marginally decreasing, then $B(\alpha k) \geq \alpha B(k)$, so that with probability at least $1-1/n$, $B(\alpha k^*) - P'_{\alpha k^*} \geq \alpha B(k^*) - \alpha P_{k^*} = \alpha(B(k^*) - P_{k^*}) = \alpha R$, and so $R' \geq \alpha R$. Similarly, with probability at least $1 - 1/n$, $R'' \geq \alpha R$. Thus, by a union bound we see that the RSPE mechanism will obtain a profit of $\min\{R', R''\} \geq \alpha R$ with probability at least $1 - 2/n$. $\square$

6.4 An Upper Bound for the Sampled Cost: Lemma 6.3.1

Lemma 5.4.16 shows that if we sample each element of S_k with probability 1/2, then the sampled set will contain at least $m = \lfloor (1-\epsilon)k/2 \rfloor$ disjoint bases with high probability in the rank of the matroid. The main challenge we face is to show that the VCG replacement costs for a base of M_m of this size in the sampled set is not too large. Our starting point is the following lemma.

Lemma 6.4.1. *Let* $S_{k+1} \in M_{k+1}$ *be partitioned into* $\Delta_1, \ldots, \Delta_{k+1}$ *such that* $S_j = \bigcup_{1 \leq \ell \leq j} \Delta_\ell$ *is a maximal subset in* M_j *of* S_{k+1}*, for each* $j = 1, \ldots, k+1$*. If* $|\Delta_j| \geq |\Delta_{j+1}|$ *for all* j*, then there are* $k \cdot |\Delta_{k+1}|$ *elements in* S_k *whose total replacement cost is at most* $k \cdot \mathrm{c}(\Delta_{k+1})$.

Proof. Consider any decomposition of S_{k+1} into $k+1$ disjoint independent sets $T_1, \ldots, T_{k+1}$ of M. Define $U_j = S_{k+1} \setminus T_j$. By the decomposition of S_{k+1}, $U_j \in M_k$ for all j. Therefore, by the maximality of S_k, $|U_j| \leq |S_k|$. We will now perform $k+1$ rounds of element exchange, one between each U_j and S_k. First augment U_j with elements of S_k to create U'_j with $|U'_j| = |S_k|$. Then use Fact 5.4.3 to associate each element of $U'_j \setminus S_k$ with an element of $S_k \setminus U'_j$. Note that $U'_j \setminus S_k \subseteq \Delta_{k+1}$, so that if $R_j = S_k \setminus U'_j$, each element of R_j has been replaced with an element of Δ_{k+1}. Furthermore, each element $y \in \Delta_{k+1}$ is used for replacement exactly k times, once for each j such that $y \notin T_j$.

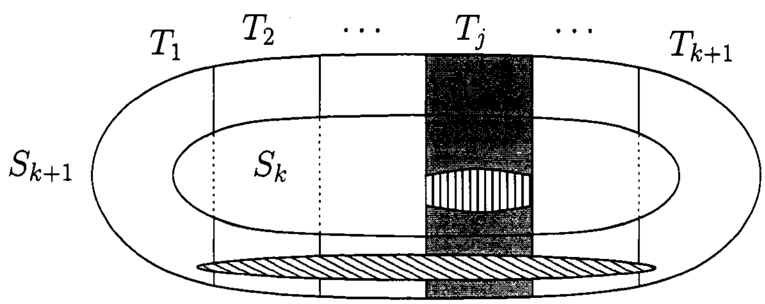

Above we show $S_k \subset S_{k+1}$. $T_1, \ldots, T_{k+1}$ partition S_{k+1}. The diagonally hatched area is Δ_{k+1}. T_j is darkened, and $U_j = S_{k+1} \backslash T_j$ is the remaining area of S_{k+1}. R_j is the vertically hatched area in $S_k \cap T_j$, which will be associated with $U_j \cap \Delta_{k+1}$.

Figure 6.2: The construction of Lemma 6.4.1

Let $R = \bigcup R_j$. We first show that an element of S_k is in R iff it is in $T_j \setminus U'_j$ for some j. Note that $S_k \setminus U_j = S_k \cap (S_{k+1} \setminus U_j) = S_k \cap T_j$. Therefore $R_j = (S_k \setminus U'_j) = (S_k \setminus U_j) \setminus (U'_j) = (S_k \cap T_j) \setminus U'_j$. This implies that R_j is contained in $T_j \setminus U'_j$. On the other hand, if an element is both in S_k and some $T_j \setminus U'_j$ then it is in $(S_k \cap T_j) \setminus U'_j = R_j$ and thus in R.

Let $d = |\Delta_{k+1}|$ and let $|T_j| = d + t_j$; we have that $t_j \geq 0$ as $|U_j| \leq |S_k| = |S_{k+1}| - d$ implying $|T_j| \geq d$. Exactly t_j elements in S_k are added to U_j to form U'_j, so at most $\sum_{j=1}^{k+1} t_j = |S_{k+1}| - (k+1)d$ elements of S_k are disqualified from being in R. Thus, $|R| \geq |S_k| - \sum_{j=1}^{k+1} t_j = (|S_{k+1}| - d) - (|S_{k+1}| - (k+1)d) = kd$. Thus as there are kd exchanges, each element $x \in R$ is involved in exactly one exchange, with U_j when $x \in T_j$. Hence the d elements of Δ_{k+1} each replace k different elements in R with total cost $k \cdot \mathrm{c}(\Delta_{k+1})$, as required. $\square$

The intuition behind our proof is that the VCG cost of S_k can be estimated by $k \cdot c(\Delta_k)$. In order to bound the VCG cost of S'_m in the sampled set, we must show that the sampled remainder of Δ_k bounds the replacement cost of S'_m. For example, suppose all the cheap elements in Δ_k that are sampled become part of S'_i for some $i \ll m$, and only expensive elements of Δ_k are left for Δ'_m, of cost like $\max_{v \in \Delta_k} c(v)$. Using this to estimate the VCG

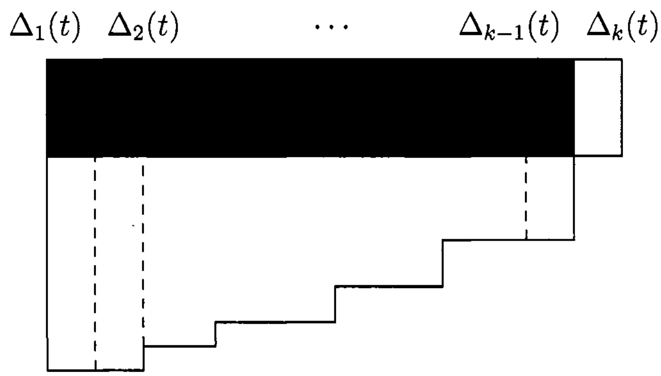

Figure 6.3: A greedy profile

cost of S'_m gives $m \cdot n \cdot \max_{v \in \Delta_k} c(v)$ rather than the desired $m \cdot c(\Delta_k)$. Let v_i be the i-th element of Δ_k, and let t be the time at which it is considered. Then it will be enough to show for each such i, in the sampled set $|\Delta'_m(t)| \geq i$. If this is the case, since all the elements of $S'_m(t)$ have cost at most $c(v_i)$, we can apply Lemma 6.4.1 and show that at least $i(m-1)$ elements of $S'_{m-1}(t)$ can be replaced at cost no more than $c(v_i)$.

We show this in two steps. Consider some time t. First, there is an *incompressible* set R whose rank in M is the same as $|\Delta_k(t)|$. As sampling a set does not increase its rank, the sampled version R' of this set also has rank at most $|\Delta_k(t)|$, and hence has at least $|\Delta_k(t)|$ elements D, analogous to $\Delta'_m(t)$, that are not in a maximal subset of R' independent in $M_{\alpha k-1}$. D is analogous, but not the same as $\Delta'_m(t)$, because we run the greedy algorithm *after* sampling. Hence elements of R may be replaced by other elements not in R. Therefore our second step is to show that in the actual greedy decomposition of S'_k, $\Delta'_m(t)$ has the same size as the subset D of our incompressible set.

Consider $S_k(t)$ as pictured in Figure 6.3, where the Δ_is are arranged in columns. We know that the lower margin is monotone nondecreasing by Lemma 5.4.20. Generally, a set like the shaded one in Figure 6.3 will have rank in M larger than its height, because elements in it are dependent only on elements in the unshaded region. An incompressible set is one where this is not the case, and the dependencies that support the height of $\Delta_k(t)$ are all internal. Note that it is not the case that the incompressible set is a subset of S_{k-1}:

it may contain elements in $\Delta_k(t)$ even though the set is independent in M_{k-1}.

Lemma 6.4.2. *If an independent set S_k in M_k has $|\Delta_k| = i$, then there exists an $R \subset S_k$ with $R \in M_{k-1}$, $|R| = (k-1)i$ and $\rho(R) = i$. (Such an R is called* incompressible.*)*

Proof. Let $S_{k-1} \subset S_k$ be independent and maximal in M_{k-1}. Then, by Lemma 5.4.14, $|S_{k-1}| = \rho_{k-1}(S_k) = (k-1)\rho(X) + |S_k \setminus X|$ for some $X \subset S_k$ with $\rho_{k-1}(X) = (k-1)\rho(X)$. Let $r = \rho(X)$, so that X contains a set R of $k-1$ disjoint independent sets of size r. Since X is a subset of S_k, it can be decomposed into k disjoint independent sets in M and each of these sets has size at most r by the rank of X, so we have that $|X| \leq kr$. Now suppose $r < i$. As S_{k-1} is independent in M_{k-1} and $\rho_{k-1}(X) = (k-1)r$, $|S_{k-1} \cap X| \leq (k-1)r$. Thus $|S_{k-1}| = (k-1)\rho(X) + |S_k \setminus X| \geq (k-1)r + |S_{k-1}| + |\Delta_k| - |X| \geq |S_{k-1}| + |\Delta_k| - r$. But since $|\Delta_k| - r > 0$, the right hand side of this inequality is strictly larger than $|S_{k-1}|$ which is a contradiction. Thus, $\rho(X) = i$, and R as defined above satisfies the lemma. $\square$

We can now get the main lemma we need to bound the replacement costs.

Lemma 6.4.3. *Let $m = (1-\varepsilon)k/2$ for some constant $\varepsilon > 0$. Let t be the time that the i-th cheapest element v_i of Δ_k is added to Δ_k. Then with probability $1 - n \cdot \exp(-\varepsilon^2 k/2)$, $S'_{m-1}(t)$ has at least $(m-1) \cdot i$ elements that can each be replaced by an element of cost at most* $\mathrm{c}(v_i)$.

Proof. We first use Lemma 6.4.2 to find an incompressible set $R \subset S_k(t)$ in M_{k-1} such that $\rho(R) = |\Delta_k(t)| = i$. Let R' be the sampled portion of R. By applying Karger's theorem (Lemma 5.4.16) to the matroid $M|_i$, the matroid whose bases are all independent sets of M of cardinality at most i, we find that with the desired probability there is a subset S' of R' of cardinality $m \cdot i$ that is independent in M_m. Moreover, since $\rho(S') \leq \rho(R') \leq \rho(R) = i$, it must be that $\rho(S') = i$, otherwise it could not contain as many as $m \cdot i$ elements that are independent in M_m.

Notice however that the cheapest base of M_m in X' may not include all the elements of S'. We now show that in the process of improving S' with elements of X' to form the cheapest base of M_m in the sample, $|\Delta'_m(t)|$ does not decrease. Precisely, to determine the cheapest base of X', we begin with S' and consider the elements of $X' \setminus S'$ in order, replacing elements of S' when they can be improved. We begin with $S^{(0)} = S'$. Given $S^{(\ell-1)}$ and the next element x_ℓ from $X' \setminus S'$, we either add x_ℓ to $S^{(\ell-1)}$, if it is independent, use it to replace a point in $S^{(\ell-1)}$ with higher cost, or, if it cannot replace a more expensive element, do nothing. The resulting set is $S^{(\ell)}$. At each step, we have a decomposition $\Delta_1^{(\ell)}, \ldots, \Delta_m^{(\ell)}$ where $\bigcup_{1\le i\le j} \Delta_i^{(\ell)}$ is a maximal cheapest set of M_j in $S^{(\ell)}$. Using this notation, because $\rho(S') = i$ while $|S'| = m \cdot i$, we have that $|\Delta_m^{(0)}| = i$. The following claim shows that the size of this set does not decrease as we improve it.

Claim 6.4.4. *$|\Delta_m^{(\ell)}| \ge |\Delta_m^{(\ell-1)}|$ for all $1 \le \ell \le |X' \setminus S'|$.*

For ease of notation in proving this claim, let $x = x_\ell$, $\Delta_j = \Delta_j^{(\ell-1)}$ be an old set, and $\Delta'_j = \Delta_j^{(\ell)}$ a new set. First suppose x is added without replacing a point. Let j be the first j so that $|\Delta'_j| > |\Delta_j|$. If $j = m$, we are done. Otherwise, note $|\Delta'_h| = |\Delta_h|$ for $h < j$, as the size of a maximal set of some order will only increase after adding x. Note that

$$|S_j^{(\ell)}| = 1 + |S_j^{(\ell-1)}|, \tag{6.1}$$

as by adding one element we can only increase the size of a maximal set by one, as the subset not including x is an independent subset of $S^{(\ell-1)}$. Hence $\sum_{s>j} |\Delta'_s| = \sum_{s>j} |\Delta_s|$. Then if $|\Delta'_m| < |\Delta_m|$, $\sum_{j<s<m} |\Delta'_s| > \sum_{j<s<m} |\Delta_s|$, which taken together with (6.1) gives $|S_{m-1}^{(\ell)}| > |S_{m-1}^{(\ell-1)}|$. But then $|S_{m-1}^{(\ell)} \setminus x| > |S_{m-1}^{(\ell-1)}|$, a contradiction.

Now suppose x cannot be added, but replaces a point y. Let j be so that $x \in \Delta'_j$. $|\Delta'_h| = |\Delta_h|$ for each $h < j$, because $\bigcup_{s\le h} \Delta'_s$ is a maximal independent set of M_h in $S' \cup \{x_1, \ldots, x_{\ell-1}\}$, so it must be the same size as $\bigcup_{s\le h} \Delta_h$, which is also maximal. By similar reasoning, $|\Delta'_j| \ge |\Delta_j|$. Suppose $|\Delta'_j| > |\Delta_j|$. Decompose $S_m^{(\ell-1)}$ into $T_j \cup T_r$, and $S^{(\ell)}$ into $T'_j \cup T'_r$, where $r = m - j$, $T_j \in M_j$, $T_r \in M_r$, T_j and T_r are disjoint, and similarly for T'. By Lemma 5.4.15, we can assume without loss of generality that $|T_j| = \sum_{h\le j} |\Delta_h| < |T'_j|$, so that by basis augmentation there is a $z \in T'_j$ that can be added to T_j. If $z = x$, then $T_j \cup x$ with T_r is a decomposition showing that $S_m^{(\ell-1)} \cup x \in M_m$, contradicting the fact that x could not be added. Otherwise, $z \in S_m^{(\ell-1)}$, so that $T_j \cup z$ and $T_r \setminus z$ is a decomposition for M_m with $T_j \cup z$ in M_i, contradicting the maximality of T_j in $S_m^{(\ell-1)}$ for M_j. Hence $|\Delta'_j| = |\Delta_j|$, and we can continue by induction to show that $|\Delta'_h| = |\Delta_h|$ for $h \ge j$. In particular $|\Delta'_m| = |\Delta_m|$ as required to prove the claim.

Thus after considering all elements of $X' \setminus S'$ we have that $|\Delta'_m(t)| \geq i$. Now we apply Lemma 6.4.1 to show that $(m-1) \cdot i$ elements of S'_{m-1} can be replaced at cost at most $(m-1) \cdot i \cdot \mathrm{c}(v_i)$ to finish the lemma. □

Finally, we can put it all together:

Proof of Lemma 6.3.1. Using a union bound, if $k \geq \frac{8}{\varepsilon^2} \log n$, we have that Lemma 6.4.3 holds with probability at least $1 - 1/n$ for all $1 \leq i \leq n$. Let $\Delta_k = \{v_1, \ldots, v_n\}$ ordered by cost. Taking $i = n$, we have that $P'_{m-1} \leq n(m-1)\mathrm{c}(v_n)$. Now considering $i = n-1$, we have at least $(n-1)(m-1)$ elements in S'_{m-1} that can be replaced with cost at most $\mathrm{c}(v_{n-1})$, so that $P'_{m-1} \leq (n-1)(m-1)\mathrm{c}(v_{n-1}) + (m-1)\mathrm{c}(v_n)$. Induction shows that $P'_{m-1} \leq \sum_{i=1}^{n}(m-1)\mathrm{c}(v_i) = (m-1)\mathrm{c}(\Delta_k)$, as required. □

6.5 A Lower Bound for the Total Cost: Lemma 6.3.2

Now we prove our lower bound on the VCG payments of $\mathcal{OPT}$.

Recall that a *circuit* in a matroid is a dependent set that is independent after removing any element, and that if an element d is dependent on an independent set A, there is a unique circuit in $\{d\} \cup A$.

Lemma 6.5.1. *Let S be an independent set in M_k partitioned into U and T such that U is a maximal independent set of S in M_{k-1}, and T is independent in M. Let D be a set of elements each of which is dependent on S in the matroid M_k. If R is the set of all elements in S that may be replaced by an element of D, then $|R| \leq k \cdot |T|$.*

Proof. Any $d \in D$ forms a unique M_k-circuit $C(d)$ with S. As d is dependent on S in M_k, it must be dependent on U in M_{k-1}, with a unique M_{k-1}-circuit C' in $d \cup U$. $C' \subseteq C(d)$, as otherwise $d \cup (C(d) \cap U)$ would be independent in M_{k-1}, and the decomposition of $C(d)$ into $d \cup (C(d) \cap U)$ and $C(d) \cap T$ would show that $C(d)$ is independent in M_k. We now show that in fact $C' = d \cup (C(d) \cap U)$. Suppose not, so that there exists $x \in (C(d) \cap U) \setminus C'$. We first show the existence of an M_{k-1} circuit $B_0 \subset (U \cup T) \cap C(d)$. As $C(d)$ is a circuit, $C(d) \setminus x$ is independent in M_k, and so has a decomposition into X and Y, where X is independent in M_{k-1}, Y is independent in M, and $X \cup x$ is dependent in M_{k-1}. Let Q be the unique M_{k-1}-circuit in $X \cup x$. If $d \notin Q$, then we may take B_0 to be Q as $Q \subset (U \cup T) \cap C(d)$. Otherwise, if $d \in Q$, then applying circuit reduction (Fact 5.4.4) to Q and C' on d gives B_0 as required.

If B_0 is wholly contained in U or T, we have a contradiction, since each of these is independent in M_{k-1}. Thus, we can assume that B_0 intersects both U and T. From B_0 we now construct a sequence of nonempty circuits $B_1, \ldots, B_\ell$ such that for all $i < \ell$, B_i intersects both U and T, but B_ℓ intersects exactly one of U or T. If B_i intersects both U and T, take some $e \in B_i \cap T$. As U is maximal, $e \cup U$ is a dependent set in M_{k-1}, so there is an M_{k-1}-circuit $C_e \subset e \cup U$. Using circuit reduction again, we have a nonempty $B_{i+1} \subset (B_i \cup C_e) \setminus e$, and $|B_{i+1} \cap T| < |B_i \cap T|$. Thus this sequence will end at some ℓ with $B_\ell \cap T = \emptyset$ or $B_\ell \cap U = \emptyset$ as required. But then B_ℓ is an M_{k-1}-circuit that is contained in an independent set in M_{k-1}, a contradiction.

Hence we can describe the M_k-circuit of each $d \in D$ with S as $C(d) = C_1(d) \cup C_2(d)$, where $C_1(d)$ is an M_{k-1}-circuit in $U \cup d$ and $C_2(d) \subset T$. Let $R = \bigcup_D C_1(d)$ be the set of all elements in U that can be replaced by an element in D. Let $U_1, \ldots, U_{k-1}$ be a partition of U into independent sets of M. If $|R| > (k-1)|T|$, then there is an $i < k$ with $|R \cap U_i| > |T|$, so that by set augmentation there is an $r \in R \cap U_i$ such that $T' = r \cup T$ is independent in the matroid M. Now as $r \in C_1(d)$ for some $d \in D$, $U' = (U \setminus r) \cup d$ is independent in M_{k-1}. But then $U' \cup T'$ gives a decomposition showing that $S \cup d$ is independent in M_k, which is impossible by the definition of d. Therefore $|R| \leq (k-1)|T|$, and so after including the elements of T, at most $k \cdot |T|$ elements of $S_k(t)$ can be replaced by elements from D, as claimed. □

Corollary 6.5.2. *Let v_j be the j-th least expensive element in Δ_k. Then the total number of elements in S_k that can be replaced by elements cheaper than v_j is at most $(j-1)k$.*

Proof. Let v_j be added to S_k just after time t. Let $E(t)$ be all elements at time t, and let $D(t) = E(t) \setminus S_k(t)$. $D(t)$ is the set of possible replacement elements at time t, and each element of $D(t)$ is dependent on $S_k(t)$ in M_k. In addition, as elements are ordered in increasing cost by time, $D(t)$ is also the set of all possible replacements of cost at most that of v_j for the final S_k. Note that no element of $D(t)$ can ever replace an element of S_k added at time greater than t, as that would contradict the fact that S_k is optimal.

By Lemma 5.4.19 and the definition of $\Delta_k(t)$, $S_k(t)$ can be decomposed into independent sets $T_1, \ldots, T_k$ of M, with $|T_k| = |\Delta_k(t)|$ and $(T_1 \cup \cdots \cup T_{k-1})$ maximal as an independent set in M_{k-1}. By our choice of t, $|\Delta_k(t)| = j - 1$. By Lemma 6.5.1, $D(t)$ replaces at most $(j-1)k$ elements of $S_k(t)$, and as noted above, no element in $D(t)$ can possibly replace any other element in S_k. As $D(t)$ is the set of all possible replacements with cost at most that of v_j, this proves the corollary. □

The main lemma of the section now follows quickly.

Proof of Lemma 6.3.2. Let $v_1, \ldots, v_n$ be the elements of Δ_k. For any $1 \leq j \leq n$, by Corollary 6.5.2, at least $(n - j + 1)k$ elements in S_k have replacement cost $\geq \mathrm{c}(v_j)$. Hence we can partition S_k into $L_1, \ldots, L_n$, where $|L_j| = k$ and the replacement cost for any $x \in L_j$ is at least $\mathrm{c}(v_j)$. Summing over all L_j proves the lemma. □

Chapter 7

CONCLUSION AND OPEN RESEARCH DIRECTIONS

In this dissertation we have explored auctions both from the perspective of the players and that of the auctioneer. For sponsored search auctions, we describe a framework of natural bidding strategies, and describe the convergence properties of some of them. We show that a certain balanced strategy converges to a well-known equilibrium in several variations of a standard auction model. We examine the revenue implications of various equilibria and our strategies in our class.

We study mechanism design from the point of view of the auctioneer by demonstrating a mechanism that chooses and procures a number of feasible sets when the auctioneer can resell the sets at a decreasing marginal benefit. We then demonstrate that the auctioneer can achieve near-optimal revenue when the system of feasible sets is a matroid, and present evidence that matroids are the only class of set systems with such a mechanism, given current technology. The mechanism adapts techniques from the design of auctions for digital goods.

We are far from having the last word on these topics. The area of sponsored search auctions is still poorly understood. The most important direction of research is to better understand advertiser strategies for bidding simultaneously in multiple keyword auctions under budget constraints. While some of the equilibria properties are known, their relationship to efficient allocations is not clear. Efficiency, social welfare, and the VCG mechanism have not been extensively studied or characterized in this context. While strategies based on return-on-investment have been considered both theoretically and practically, it is not clear if they are the only natural strategy for this problem.

Even the model for sponsored search auctions is not completely settled. While we have assumed an open-cry format, where all bids are public, the biggest sponsored search market, run by Google, operates with sealed-bids. Can the repeated nature of the auction be used to estimate actual bids? It is not clear how effective estimation is in practice, and we are not aware of a comprehensive theoretical analysis of the problem. The auction format is also changing. For example, Google will regulate an advertiser's bidding to adapt to changing user traffic over the course of a day, as well as to target particular slots. It is not clear exactly how Google incorporates these features into its mechanism, and what the impact on bidding strategies as well as Google's revenue might be. Is it possible, for example, to regulate bidding over the course of the day in such a way as to increase Google's revenue as well as advertiser click traffic? Might naive methods of regulation result in catastrophic revenue decline, or open the system to strategic advertiser exploitation?

For procurement auctions, one dangling thread left by this dissertation is to what extent cost-efficient mechanisms exist only for matroids. While we have nearly shown that only matroids have truthful profit extractors, settling this question would only leave open the larger question of the extent to which the profit estimation and extraction framework is necessary for procurement. While techniques such as CORE auctions do not seem to directly apply to procurement over set systems, there may be others. In particular, can we relax the requirement of truthfulness to one where the mechanism is only truthful with high-probability? This has been successfully used to create *envy-free* auctions in contexts where being envy-free and truthful simultaneously is impossible [20].

Finally, the structural results we have found for matroids should apply in other contexts, for example forward auctions. A holy grail for mechanism design is *combinatorial*, or *package*, auctions, where players bid for subsets from a set of items. The value of a subset for a player can be a complex function of its constituents, taking advantage of substitutability or complementarity between items. Perhaps our techniques could be used in situations where bidder valuations have a matroid structure. This is the most speculative research direction mentioned so far, as to our knowledge combinatorial auctions have not been considered in such a model before.

BIBLIOGRAPHY

[1] A. Animesh, V. Ramachandran, and S. Vaswanathan. Online advertisers' bidding strategies for search, experience, and credence goods: An empirical investigation. In *Second Workshop on Sponsored Search Auctions.* ACM Electronic Commerce, 2005.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science (FOCS)*, 2004.

[3] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[4] A. Archer and E. Tardos. Frugal path mechanisms. *ACM Trans. Algorithms*, 3(1):3, 2007. ISSN 1549-6325. An early version appeared in SODA 2002.

[5] K. Asdemir. Bidding patterns in search engine auctions. In *Second Workshop on Sponsored Search Auctions.* ACM Electronic Commerce, 2006.

[6] Atlas Search. URL `http://www.atlasonepoint.com`. Operated under Atlas DMT, LLC.

[7] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, J. Jain, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *16h International World Wide Web Conference*, 2007.

[8] F. Brandt and G. Weiss. Antisocial agents and Vickrey auctions. In *Proceedings Of The 8th Workshop on Agent Theories Architectures*, 2001.

[9] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. In *EC '07: Proceedings of the 8th ACM Conference on Electronic Commerce*, 2007.

[10] M. Cary, A. Flaxman, J. Hartline, and A. Karlin. Matroid procurement auctions, 2007. Unpublished manuscript.

[11] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[12] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms.* The MIT Press, 2nd edition, 2001.

[13] K. Deshmukh, A. Goldberg, J. Hartline, and A. Karlin. Truthful and competitive double auctions. In *Proc. 10th European Symposium on Algorithms*. Springer–Verlag, 2002.

[14] eBay. URL `http://www.ebay.com`.

[15] B. Edelman and M. Ostrovsky. Strategic bidder behavior in sponsored search auctions. In *First Workshop on Sponsored Search Auctions*. ACM Electronic Commerce, 2005.

[16] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. To appear in *American Economic Review*, 2006.

[17] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proc. 15th Symp. on Discrete Alg.* ACM/SIAM, 2004.

[18] J. Feng, H. K. Bhargava, and P. D.M. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 2006. Forthcoming.

[19] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.

[20] A. Goldberg and J. Hartline. Envy-free auction for digital goods. In *Proc. of 4th ACM Conference on Electronic Commerce*. ACM Press, New York, 2003.

[21] A. V. Goldberg and J. D. Hartline. Competitiveness via consensus. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.

[22] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 2006.

[23] Google. Adwords. URL `http://adwords.google.com`.

[24] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[25] P. Hall. On representatives of subsets. *J. London Math. Soc*, 10:26–30, 1935.

[26] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *EC '05: Proceedings of the 6th ACM Conference on Electronic Commerce*. ACM Press, 2005.

[27] iProspect. URL `http://www.iprospect.com`. Search Engine Marketing Firm.

[28] D. R. Karger. Random sampling and greedy sparsification for matroid optimization problems. *Math. Programming*, 82(1-2, Ser. B), 1998.

[29] A. Karlin, D. Kempe, and T. Tamir. Beyond VCG: Frugality in truthful mechanisms. In *Proc. of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.

[30] R. Karp, E. Koutsoupias, C. Papadimitriou, and S. Shenker. Optimization problems in congestion control. In *FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, page 66. IEEE Computer Society, 2000.

[31] B. Kitts and B. Leblanc. Optimal bidding on keyword auctions. *Electronic Markets*, pages 186–201, 2004.

[32] B. Kitts, P. Laxminarayan, B. LeBlanc, and R. Meech. A formal analysis of search auctions including predictions on click fraud and bidding tactics. In *First Workshop on Sponsored Search Auctions*. ACM Electronic Commerce, 2005.

[33] P. Klemperer. Auction theory: A guide to the literature. In *Journal of Economic Surveys*, pages 227–286. 13(3), 1999.

[34] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM Conference on Electronic Commerce*. ACM Press, 2006.

[35] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[36] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, 18:511–533, 2001.

[37] N. Nisan. Algorithms for Selfish Agents: Mechanism Design for Distributed Computation. In *16th Annual Symp. on Theoretical Aspects of Comp. Sci.*, 1999.

[38] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. of 31st Symposium on Theory of Computing*, pages 129–140. ACM Press, New York, 1999.

[39] J. G. Oxley. *Matroid Theory*. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1992.

[40] C. Papadimitriou. Algorithms, games and the internet. In *Proc. of 33rd Symposium Theory of Computing*, pages 749–752. ACM Press, 2001.

[41] R. Rado. A theorem on independence relations. *Quart. J. Math. Oxford*, 13:83–89, 1942.

[42] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 2002.

[43] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proc. of 20th International Symposium on Theoretical Aspects of Computer Science*, 2003.

[44] E. Tardos. Network games. In *Proc. 20th Annual Symp. on Theoretical Aspects of Computer Science*, 2003.

[45] H. Varian. Position auctions. To appear in *International Journal of Industrial Organization*, 2006.

[46] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[47] Web-Cite. Shooting the moon and budget busting with Overture's auto-bidding, April 2003. URL `http://www.web-cite.com/search_marketing/000112.html`.

[48] Yahoo! Sponsored search. URL `http://searchmarketing.yahoo.com`.

[49] X. Zhang. Finding edgeworth cycles in online advertising auctions. MIT Sloan School of Management, Working Paper, 2005.

[50] Y. Zhou and R. Lukose. Vindictive bidding in keyword auctions. In *Second Workshop on Sponsored Search Auctions*. ACM Electronic Commerce, 2006.

VITA

Matthew Cary was born in Seattle, Washington. He attended Brown University where he received a B.A. in Mathematics in 1996. He returned to Seattle to attend the University of Washington, where he received a M.Sc. in Computer Science in 2000. In addition to mechanism design, he has published in the areas of cryptography, software security and algorithms.