FAIRNESS, LEARNING AND EFFICIENCY IN MARKETS WITH BUDGETED AGENTS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Pooya Jalaly Khalilabadi

May 2018

FAIRNESS, LEARNING AND EFFICIENCY
IN MARKETS WITH BUDGETED AGENTS

Pooya Jalaly Khalilabadi, Ph.D.
Cornell University 2018

In almost all online markets with monetary transactions, the participants have a limited budget which restricts their ability to purchase their desired commodities. Models from mechanism design, algorithm design and auction theory which study these online markets often ignore this important constraint.

This dissertation presents a deep study of such markets with budget limited agents, using theoretical models as well as data from real world auction markets. In chapter 2, we study the problem of a budget limited buyer who wants to buy a set of commodities, each from a different seller, to maximize her value. The budget feasible mechanism design problem aims to design a mechanism which incentivizes the sellers to truthfully report their cost, and maximizes the buyer's value while guaranteeing that the total payment does not exceed her budget. Such budget feasible mechanisms can model a principal in a crowdsourcing market interested in recruiting a set of workers (sellers) to accomplish a task for her. We present simple and close to optimum mechanisms for this problem when the valuation of the buyer is a monotone submodular function.

In chapter 3, we present a deep study of the behavior of real estate agents in the new online advertising platform provided by Zillow. We analyze behavior of the agents through time using the provided data from Zillow. We use a

no-regret based algorithm to estimate the value of agents for impression opportunities. We observe that a significant proportion of bidders initially do not use the bid recommendation tool which has been provided by Zillow. This proportion gradually declines over time. We argue that the agents gradually trust the system by learning that the platform adequately optimizes bids on their behalf and the increased effort of experimenting with alternative bids is not worth the potential increase in their net utility.

In chapter 4, we show equilibria of markets with budget limited agents can be used to achieve fairness for problems of matching without money with agents who have preferences over commodities. A unit budget with artificial money is given to each agent for achieving fairness. We also provide polynomial time algorithms for finding the equilibria of these markets.

BIOGRAPHICAL SKETCH

Pooya Jalaly Khalilabadi was born on March 23rd, 1989 in Esfahan, Iran. He was raised in Shahr-e Kord, Iran. For guidance school and high school, he was admitted to the National Organization for Development of Exceptional Talenets, NODET. In second year of high school, he was admitted to the camp for Iranian National Olympiad in Informatics, INOI, where he received a gold medal.

As a recipient of a gold medal in INOI, he was allowed to choose any major in any university in Iran, without having to take the Iranian national entrance exam. He started his B.Sc. studies in software engineering at Sharif University of Technology in September 2007. His teams were selected as one of Sharif University's teams for ACM ICPC regional contests in 2007, 2008 and 2010, in which they placed 5^{th}, 7^{th} and 4^{th} respectively. His undergraduate research was focused on theory of computing, discrete mathematics and social networks analysis.

After receiving his bachelors degree from Sharif University, Pooya continue his studies at Cornell University to get a M.Sc. and Ph.D. in Computer Science under supervision of professor Éva Tardos. During his time at Cornell, he became interested in the intersection of computer science and economics. He currently works on problems related to algorithmic mechanism design, market algorithms, online ad auctions and algorithmic game theory.

To my parents, Mahmoud and Shahla, my beautiful sisters, Laleh and Yalda, and my lovely wife Negar FadaeiDehkordi.

ACKNOWLEDGEMENTS

I owe a deep debt of gratitude to my Ph.D. advisor Éva Tardos for helping me become the person I am today. Throughout my journey as a graduate student, her guidance and insight has been my most valuable resource. She thought me how to approach hard problems in research and life. Most importantly, she thought me how to deal with inevitable failures and hardships throughout my journey. Her genuine patience in every single one of our meetings was really inspiring. I could not ask for a more caring and inspiring person to be my advisor during my Ph.D. years at Cornell and I do not think I will ever be able to express how thankful I am for this.

I would also like to thank many Cornell professors: David Shmoys for his advice as my Ph.D. committee member and for his great *scheduling* course, Joseph Halpern for his advice as my Ph.D. committee member and his inspiring *reasoning about uncertainty* course, Robert Kleinberg for teaching me how to be compassionate with others during the time I was his teaching assistant and for his great *advanced algorithms* course, and David Williamson for his phenomenal teaching of the *approximation algorithms* course.

I would also like to thank my collaborators Denis Nekipelov and Saeed Alaei without whom most of the work I have presented in this thesis would not be possible. Denis thought me how to approach problems through the eyes of an economist. Saeed thought me to not give up when I faced with a hard problem.

I was very fortunate to have awesome lab-mates and friends like Radieh Abebe, Saeed Alaei, Soroush Alamdari, Hedyeh Beyhaghi, Dylan Foster, Daniel Freud, Sam Hopkins, Thodoris Lykoursis, Rad Niazadeh, Rahmtin Rotabi and

Yang Yuan in the Theory of Computing lab during my time as a graduate student at Cornell.

I would like to thank the professors and students at Sharif University of Technology who introduced me to Algorithmic Game Theory and Computer Science during my undergraduate studies. I would like to especially thank Mohammad Ghodsi who was my undergrad advisor. I would like to thank Vahab Mirrokni, Saieed Akbari, Mohammad Ali Safari and Mohammad Amin Fazli who made me become excited about research.

Last but not least, I would like to thank my parents Mahmoud and Shahla and my wife Negar who were there for me every single step of the way. I could not achieve anything without their love, help and support.

My work has been supported by NSF grants CCF-1563714 and CCF-0910940, ONR grant N00014-08-1-0031, a Google Research Grant. Parts of my work was done during my time as a visiting graduate student in Simons Institute for Theory of Computing at Berkeley University of California in Fall 2015.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

For centuries, markets have been the backbone of the world's economy. Emergence of the new online world was the beginning of a significant era in their evolutionary process. The possibility of instantly getting connected to the potential costumers around the world has been nothing short of a revolution for many businesses. The enormous surge in the demand, which was made possible by this opportunity, has helped businesses to grow and flourish. Customers are also beneficiaries of this connected world. The new competition meant more supply, and in some cases, the availability of commodities that were not available before. More competition resulted by these connections and the fewer middle-men or no middle-man meant lower prices for costumers, and the big boost in the demand meant more revenue for businesses.

For over a decade, many tech giants have been making billions of dollars in revenue each year, derived by connecting buisnesses to their potential costumers [EOS05]. The emerging economy of this connected world has carried over many challenging aspects of the previous economies and has been the root of many new challenging problems. These problems have been a never ending source of excitement for academics. These challenges which arose from the enormous size of this new and complex online world became a significant common ground for interdisciplinary work between computer scientists and economists, the place where algorithmic game theory was born. The computational challenges that came from the enormous number of actors in this online world required the knowledge of computer scientists and the economic challenges that were carried over from the traditional markets needed the attention of economists.

Many successful online platforms have attracted millions of people. In platforms like the Google search engine, users try to find the web pages related to their search queries. This has created an exceptional opportunity for advertisers to also use the platform for finding their potential customers. Similarly, real estate agents who are looking to sell or lease their houses use platforms like Zillow to find their target audiences who use these platforms to buy or rent a place to live in. Crowd-sourcing platforms such as Amazon Mechanical Turk provide a marketplace for work that requires human intelligence. Crowd-sourcing platforms allow companies and individuals to find their needed work force from the pool of available workers who have joined the platform to make money. Studying the process in which the participants of these various types of platforms learn to interact with each other and the platform itself as well as the properties of their final stable state (which is called an *equilibrium*) is valuable for designing, maintaining and improving these platforms. Needless to say, studying how these participants can find their (near) optimal strategies can help the platform to analyze the outcome and help the participants to benefit from these platforms as much as possible.

While many computer scientists and economists have studied these problems and challenges, there are still a lot of open problems left to be studied. Many participants like advertisers who use Google, real state agents who use Zillow and the companies that hire their work force through Amazon Mechanical Turk have a limited budget. Some studies have already considered the fact that these agents are budget limited [Sin10, HZ79, BCI+05]. However, in most studies this important property of these participants have either been ignored completely and has been replaced by assumptions such as quasi-linear utilities which assumes that spending money has a linear negative effect on the utility of these participants. This thesis focuses on deeply studying the behavior of these budget limited participants in these markets and addressing the fairness and efficiency challenges faced by these online marketplaces.

1.1 Budget Feasible Mechanism Design

Assume you have a knapsack with a limited volume. Furthermore, assume you want to go on a trip and have a set of objects that you want to take with you. Each of these objects fills some space in your knapsack and add some value to your trip. Your goal is to fit some of these objects in your knapsack while maximizing the total value that you have for the items in the knapsack. This problem is a classical optimization problem in algorithm design called the *knapsack problem* and it is known to be NP-Hard. Despite the fact that knapsack problem is NP-Hard, there are many polynomial time approximation algorithms known for this solving problem almost optimally.

Now, assume you are a budget limited principal and want to hire a set of workers in an online platform like Amazon Mechanical Turk to accomplish a task for you. Each of these workers requires to get payed for getting hired and if hired, add some value to quality of the outcome. You want to maximize the overall quality of the outcome. One way to deal with this problem is by using the optimization algorithms for the knapsack problem since the optimization problems are very similar. Here, the principal's budget is the knapsack's size and the costs are the size of the objects. However, there is a problem that the solutions from algorithm design do not address: each of the workers is a strategic human who can lie about his/her required payment. *Algorithmic mechanism design* tries to address the issue by designing mechanisms in which agents (workers in this example) do not have any incentive to lie (are truthful), while preserving the efficiency of the outcome.

The aforementioned hiring problem is studied in *budget feasible mechanism design*. In Chapter 2 of this thesis, we study prior-free budget feasible mechanism design problem, where a single buyer aims to buy a set of items, each from a different seller. Budget feasible mechanism design focuses on maximizing the value of the buyer, while keeping the total payments bellow the budget. This problem was introduced by Singer [Sin10], and models problems such as the problem of a crowdsourcing platform, like aforementioned principal using Amazon Mechanical Turk, who wishes to procure a set of workers to accomplish a set of tasks. Each worker has a private cost for her service. We offer truthful mechanisms with good approximation guarantee for this problem that incentivizes the workers to report their true cost and find a set of workers with close to optimal value [Svi04].

We focus on simple parameterized mechanisms, assuming the buyer's valuation is a general monotone (non-decreasing) submodular function. Monotone submodular functions are widely used, with submodularity capturing the diminishing returns property of adding items. Submodular value functions are the most general class of functions where the optimization problem (without considering incentives) can be solved within a constant factor in polynomial time using an oracle access to the buyers' value.

We introduce new simple parameterized mechanisms for this problem, as well as parameterizing and improving the analysis of some previously known mechanisms. Our main result is to show how these simple parameterized mechanisms can be combined for the case of large markets, where each item individually does not have a significant value compared to the optimum. We also show how our results can be used for the problem of a principal hiring in a crowdsourcing market to match the workers to tasks, subject to a total budget.

1.2 Query Based Auction Markets

Modern online platforms allow users to search for a broad spectrum of items from web search, to consumer products, services and housing. These platforms typically monetize consumer searches via a form of advertising by allowing participating firms or individuals to promote their products or services and make them appear along with the organic content. Auctions are most common mechanisms that are used to allocate and price the sponsored content or ads based on the bids submitted by the bidders. In contrast to the budget feasible mechanisms of previous section, the common auctions used here are not truthful.

These auctions run in real time for each user query. This feature, on one hand, offers flexibility for the advertisers allowing them to customize placement of their ads over time and over the types of users. On the other hand, this environment requires the advertisers to strategically optimize their bid. In such dynamic advertising auction markets, such as Google's sponsored search auction, bidders rely on software that suggests or places the bids in real time. However, if this software is created and provided by the auction platform itself, bidders may question the integrity of its design, i.e. whether it is created to optimize bidder's objective as opposed to the revenue of the platform. In chapter 3, we study the process of bidders adopting the platform-provided bid recommendation in an online advertising auction. Our goal is to understand how and when bidders learn to trust the integrity of the platform-generated recommendation.

The online platform we study in Chapter 3 is Zillow. Zillow.com is the largest residential real estate search platform in the United States used by 140 million people each month according to the company's statistic [ZG16]. Viewers are looking to buy or sell houses, want to see available properties, typical prices, and learn about market characteristic. The platform is monetized by showing ads of real estate agents offering their services. Historically, Zillow used negotiated contracts with real-estate agents for placing ads on the platform. In the experiment we study, several geographically isolated markets were switched from negotiated contracts to auction based pricing and allocation. The auction design used was a form of generalized second price, very similar to what is used in many other markets, except that agents were paying for impressions (and not for clicks). A unique feature of this experiment is that the bidders in this market are local real estate agents that bid in the auctions on their own behalf. This is unlike many existing online marketplaces (like Google and Bing) where many bidders use third-party intermediaries to assist with the bidding.

Along with the new auction platform, Zillow provided the bidders with a recommendation tool that suggests a bid to each bidder. Bidders were required to log into the system if they wanted to change their bid, and once they logged in, the system offered a suggested bid: the recommended bid for maximizing the obtained impression volume based on parameters such as the bidder's budget, the estimated impression volume, and budgets and competing bids of other bidders. Bidders in this market are limited by small budgets. During the auction, bidders typically changed their bids relatively frequently, while they tended to keep a closed to fixed budget with the average change in the budget over the time period only around %4. In light of this, we view only the bids as strategic. The bid recommendation tool was designed to suggest a bid maximizing impressions gained by spending the budget. Since the bidders eventually adopted the tool without significantly changing their budgets, we conclude that they appear to have agreed with the goal maximizing the number of impressions gained given the budget. It appears that bidders initially lacked trust in the recommendation and that is why they didn't use it, and both bidders and the platform would have been better off if the system didn't offer bidders the opportunity to avoid the recommended bid.

While it is common to use (Bayes) Nash equilibrium framework to empirically analyze auctions, this framework may not be the best fit for the experiment on Zillow, where auctions were in transition (with an increasing number of bidders invited to the auction platform over time) and bidders were learning how to bid. It is clear, however that the bidders knew their preferences over the ad impressions which we assume is fixed over time since all of the bidders have previously participated in the negotiated price-based market for impressions which existed for over 10 years.

We assume that the agents use a form of algorithmic learning to optimize their bid over time. Low-regret assumption is a simple assumption on the type of algorithmic learning used by the agents. We use the assumption to infer the values for agents: for a given regret error parameter ϵ, every possible fixed bid b implies an inequality which indicates that the agent must have a value that makes her regret for not bidding b at most ϵ. Since we do not know the actual value of agents for each impression, we use this method to infer their value indirectly.

We observe that a large proportion of bidders do not use the recommended bid to make bid changes immediately following the introduction to the new market. Many bidders could adjust their bid on their own to gain over the recommended bid since the impression volume is different in different weekdays and bid recommendation tool was not using this information. However, many bidders were not changing their bid frequently enough to take advantage of this and would have been better off by always using recommended bid. The number of bidders who outperformed the bid recommendation in our study is about the same as the number of those who did worse. The proportion of bidders following the recommendation slowly increases as markets mature.

Our work provides an empirical insight into possible design choices for auction-based online advertising platforms. Search advertising platforms (such as Google and Bing) allow bidders to submit bids on their own and there is an established market of third-party intermediaries that help bidders to bid over time. This market design allows for more complex bidding functions, for example, allowing agents to express added value for subsets of the impression opportunities via multiplicative bid-adjustments (e.g., based on the age of the viewer). In contrast, many display advertising platforms (such as Facebook) use a simpler bidding language, and optimize bids on bidders' behalf based solely on their budgets. This eliminates the need for the bidders to bid on their own or use intermediaries. Our empirical analysis shows that the latter approach may be preferred for markets where bidders are individuals who don't have access to third party tools, and who may question the fairness of platform-provided suggestions.

1.3 Fairness and Matching Markets

The problem of fairly allocating items to unit demand agents without money has been studied extensively in both Economics and Computer Science literature. In contrast with the problems in previous sections, this problem is mainly concerned with finding a fair allocation and assumes the agents report their preferences truthfully. This problem is motivated by many important problems in the real world, including assigning students to schools, jobs to applicants and people to communities. In each of these problems there is a set of agents who have their personal preferences over a set of items. For instance, in the school choice problem each students wants to be admitted in school and occupy exactly one seat (position). The students' preferences over schools can be ordinal, for instance a personal preference list that orders schools, or cardinal, for instance having a real number for each student and each school which corresponds to value of the students for attending the school. The latter allows the students to express their preferences over schools with more precision.

The assignment of items to agents, such as assigning schools to students, is called an allocation. An allocation can be either deterministic or randomized. An outcome of a deterministic mechanism is a deterministic allocation and shows exactly which student goes to which school. An outcome of a randomized mechanism is a randomized allocation which shows the probability that the student will end up in each school. In a randomized allocation, $x_{ij} \in [0, 1]$ represents the probability that student i ends up in school j and can be also seen as the *fractional allocation* of j to i. Note that for each unit demand agent i, $\sum_j x_{ij} = 1$.

If an allocation has the property that none of the students prefer another student's allocation to what they have been offered, the allocation is *envy-free*. If all outcomes of a mechanism have this property, then the mechanism itself is called an envy-free mechanism. Envy-freeness is an important and well studied fairness property in economics and computer science literature and has been studied in both deterministic and randomized settings.

Even though envy-freeness is a desired fairness property, if a randomized mechanism only has this property it can possibly lead to undesired outcomes. To see this, note that if a randomized mechanism assigns each student to each school with equal probability, the mechanism is envy free. Let a and b be two students and A and B be two schools. Assume a prefers A over B and b prefers B over A. In the mentioned randomized allocation, each student goes to each school with probability 0.5, however, it is clear that a more *efficient* allocation is to assign a to A and b to B. Both students clearly prefer the latter allocation while both allocations are envy-free. *Pareto-efficiency* generalizes this notion of efficiency. A mechanism is Pareto efficient if for any outcome of the mechanism there is no other allocation in which no one prefers the outcome of the mechanism over the other allocation and there is at least one agent who is strictly better off by using the other allocation.

A pareto-efficient and envy-free solution to this matching problem was proposed by [HZ79] using matching markets. If we give one unit of fake money to each agent and allow them to buy exactly one unit of items then the equilibria of the resulting market offer an intuitive, fair (envy-free), and Pareto-efficient solution for the problem of allocations of resources to the unit demand agents. This was first proposed and discussed by [HZ79] in the context of matching markets, and then by [DFH$^+$12] (see also [GN12] and [WM15]) in the context of allocation of resources in systems. The idea is to endow each agent with equal budget: a unit of (artificial) money. In a market with unit demand agents each of whom have one unit of budget, a set of prices $\mathbf{p}$ for the items is market clearing and are in equilibrium, if there is a fractional allocation $\mathbf{x}$ of items to agents such that the following conditions hold (i) each item is allocated at most once, (ii) each agents is allocated her favorite set of items subject to the budget constraint[1] that $\sum_j x_{ij} p_j \leq 1$, and (iii) the market clears, meaning that all items not fully allocated have price 0.

In Chapter 4 of this thesis, we view the resulting fractional (randomized) allocation $\mathbf{x}$ as a fair solution to the allocation problem without money, which is also Pareto-efficient and envy-free.[2]Aside from the fairness and efficiency criteria, we are also concerned with the running time of our mechanisms. We propose two algorithms which run in polynomial time when either the number of agents or items is fixed. Our algorithm for the case where the number items is fixed is a significantly simpler solution compared to that of [DK08] who have also offered a solution for this case. No such algorithm was previously known for the case where the number of agents is fixed.

1.4 Organization

In Chapter 2, we study and introduce many simple and efficient budget feasible mechanisms. The mechanisms try to solve problems such as the problem of a budget limited principal working on Amazon's Mechanical Turk, who wishes to procure a set of strategic workers to accomplish a set of tasks.

[1]Note that agents have no use for the (artificial) money and are simply optimizing their allocated item, subject to their budget.

[2]An alternate way to arrive to the same solution concept is to assign each agent an equal share of each resource, and then look for an equilibrium of the resulting exchange market. To see that this results in an identical outcome, we can think of each agents trade, as a two step-process, where he first sells all his allocated share on the market prices, and then uses the resulting money to buy his optimal allocation.

In Chapter 3, we study how participants (real estate agents) in Zillow's query based ad auctions learn how to interact with and trust the bid recommendation tool provided by the platform. We also study the welfare of the agents with their actual bid and recommended bids provided by the system by using no-regret learning techniques to estimate their value.

In Chapter 4, we introduce two efficient algorithms for finding the equilibria of matching markets. Matching markets model problems such as assigning students to schools, jobs to applicants and people to communities, in a fair and efficient way.

We include the formal definitions and the literature review related to the topic of each chapter in the chapter. All results presented in this thesis are based on joint work with Éva Tardos. Results presented in Chapter 3 (and Appendix B) and Chapter 4 (and Appendix C) are also based on joint work with Denis Nekipelov and Saeed Alaei respectively. Results in Chapter 2 (and Appendix A), Chapter 3 (and Appendix B) and Chapter 4 (and appendix C) are based on [JNT17], [JT17] and [AJKT17] respectively.

CHAPTER 2

BUDGET FEASIBLE MECHANISM DESIGN. FOR MONOTONE SUBMODULAR VALUATIONS

In this chapter, we study the problem of a budget limited buyer who wants to buy a set of items, each from a different seller, to maximize her value. The budget feasible mechanism design problem aims to design a mechanism which incentivizes the sellers to truthfully report their cost, and maximizes the buyer's value while guaranteeing that the total payment does not exceed her budget. Such budget feasible mechanisms can model a buyer in a crowdsourcing market interested in recruiting a set of workers (sellers) to accomplish a task for her.

Budget feasible mechanism design was introduced by Singer in 2010. There have been many improvements on the approximation guarantee of such mechanisms since then. We consider the case where the buyer's valuation is a monotone submodular function. We offer two general frameworks for simple mechanisms, and by combining these frameworks, we significantly improve on the best known results for this problem, while also simplifying the analysis. For example, we improve the approximation guarantee for the general monotone submodular case from 7.91 to 5 (7.91 was provided by [CGL11]); and for the case of large markets (where each individual item has negligible value) from 3 to 2.58 (3 was provided by [AGN14]). More generally, given an r approximation algorithm for the optimization problem (ignoring incentives), our mechanism is a $r + 1$ approximation mechanism for large markets, an improvement from $2r^2$. We also provide a similar parameterized mechanism without the large market assumption with a $4r + 1$ approximation guarantee.

2.1 Introduction

In this chapter, we study *prior-free budget feasible mechanism design* problem, where a single buyer aims to buy a set of items, each from a different seller. Budget feasible mechanism design aims to maximize the value of the buyer, while keeping the total payments bellow the budget. We offer simple and universally truthful mechanisms for this problem, significantly improving previous bounds. This problem was introduced by [Sin10], and models the problem of crowdsourcing platforms, such as Amazon's Mechanical Turk, where a requester, with a set of tasks at hand, wishes to procure a set of workers to accomplish her tasks. Each worker has a private cost for his service. We offer universally truthful mechanisms with good approximation guarantee for this problem that incentivizes the workers to report their true cost.

We give two very simple parametrized mechanisms, assuming the buyer's valuation is a general monotone (non-decreasing) submodular function. Monotone submodular functions are widely used and general, submodularity capturing the diminishing returns property of adding items. Submodular value functions are the most general class of functions where the optimization problem (without considering incentives) can be solved approximately in polynomial time using a value oracle. We consider this problem in the general case, as well as the special case where each item individually does not have a significant value compared to the optimum.

2.1.1 Our Model

We consider the problem of a single buyer with a budget B facing a set of multiple sellers A. We assume that each seller $i \in A$ has a single indivisible item, and has a private cost c_i for this item, and the buyer has no prior knowledge of the private costs. The utility of a seller for selling her item and receiving payment p_i is $p_i - c_i$. We only study universally truthful mechanisms, i.e. the mechanisms in which sellers truthfully report their costs, and do not have incentive to misreport. Since each seller $i \in A$ only has a single item, we interchangeably use i to denote the seller or her item. We assume that $v(S)$, the value of the buyer for a subset of items $S \subseteq A$, is a monotone (non-decreasing) submodular function.

The Seller has a limited budget B. The *budget feasibility constraint* requires that the total payments to the sellers may not exceed the budget. The goal of this chapter is to design simple, universally truthful and budget feasible mechanisms that approximately maximize the value of the buyer. We compare the performance of our mechanism with the true optimum, without computational or incentive limitation: maximizing the value subject to keeping the total cost bellow the budget. With this comparison in mind, incentive compatible mechanisms that do not run in polynomial time are also of some interest.

We also consider a variant of the problem modeled by a bipartite graph, where one side of the graph are agents with private costs and the other side are tasks, each with a value for the principal. An edge (a, t) represents that agent a can do task t. In this model, which was introduced by [GNS14] motivated by Crowdsourcing Markets, each agent (represented by a node) has a fixed private cost, can do a subset of the tasks, and each task has a fixed value for the principal.

2.1.2 Our Contribution

We offer two classes of parameterized mechanisms. The main result of this chapter in Section 2.3 combines these two mechanisms in a surprising way, offering a new and improved mechanisms for large markets. In Section 2.2.1, we study the class of parameterized *threshold mechanisms* that decide on adding items based on a threshold of the marginal contribution of each item over its cost (bang per buck), using a parameter γ. In section 2.2.2, we consider another parameterized class, called the *oracle mechanisms*, which also adds items in decreasing order of bang per buck, till reaching an α fraction of the true optimum, without considering the budget. In section 2.2 we analyze these two parameterized mechanisms for general monotone submodular valuations. In section 2.3 we combine the two mechanisms to get an improved result for large markets. See Table 2.1.2 for a summary of our results for the general problem. In section 2.4 we focus on the application to a problem of markets with heterogeneous tasks [CGL11, GNS14].

- In section 2.2.1 we consider threshold mechanism GREEDY-TM, and RANDOM-TM, that chooses randomly between the single item of highest value, and the output of GREEDY-TM. This framework is a direct generalization of the mechanisms presented in [Sin10], [CGL11], [SK13], and some

	Rand	Rand*	Det*	Det, LM	Rand, Oracle	Det, Oracle, LM
Previous work	7.91	7.91	8.34	3	$-$	$2r^2$
Our results	5	4	4.56	2.58	$4r$ or $4r+1$	$1+r$

Table 2.1: The top numbers are the previously known best guarantees, $r \geq 1$ is the approximation ratio of the oracle used by the mechanism, * indicates that the mechanism has exponential running time. Rand and Det stand for randomized and deterministic mechanisms, and LM indicates the large market assumption. The $4r$ guarantee requires an additional assumption for the oracle, without the assumption the bound is $4r+1$.

of the mechanisms of [AGN14], who used $\gamma = 0.5$. We show that for monotone submodular valuations, with the same choice of the parameter γ, the randomized threshold mechanism is universally truthful, budget feasible and can achieve a 5 approximation of the optimum. This improves on the best previous bound of 7.91 due to [CGL11].

- In section 2.2.2, we introduce another class of parameterized mechanisms, RANDOM-OM, and their exponential counterparts RANDOM-EOM and DETERMINISTIC-EOM discussed in Section 2.2.2, called *oracle mechanisms*, which add items in the bang per bunk order until an α fraction of the optimum value is obtained, for a parameter α. The mechanisms RANDOM-EOM and DETERMINISTIC-EOM use the true optimum value (and hence run in exponential time), while RANDOM-OM uses a polynomial time approximation instead. We show that keeping the total value of the winning set at most a fraction of the optimum guarantees that the mechanism is budget feasible. RANDOM-EOM and DETERMINISTIC-EOM use a parameterized version of the exponential time oracle mechanism of [AGN14], which we call GREEDY-EOM, as a subroutine.

  For the case when the mechanism has access to an oracle computing the true optimum value, we show that with the right choice of α, our oracle mechanism RANDOM-EOM is universally truthful, budget feasible and achieves a 4 approximation of the optimum for monotone submodular values, improving the bound of 7.91 of [CGL11]. For DETERMINISTIC-EOM, we use a derandomization idea, which is similar to that of [CGL11] and show it achieves 4.56 approximation of the optimum, improving the 8.34 bound of [CGL11].

The mechanism RANDOM-OM runs in polynomial time by using an r-approximation oracle as a subroutine instead of the optimum. We note that using GREEDY-EOM with a sub-optimal oracles breaks monotonicity; our RANDOM-OM mechanism guarantees monotonicity with any oracle. We show that with the right choice of α, RANDOM-OM is universally truthful, budget feasible, and achieves a $4r + 1$ approximation of the optimum (which improves to $4r$ when the oracle used is a greedy algorithm).

- We give the main result of this chapter in section 2.3, where we combine our two parameterized mechanisms by running both and declaring the sellers in the intersection of the two sets as winners. Taking the intersection allows our CAUTIOUS-BUYER mechanism to use larger values of the parameters γ and α and keep the mechanism budget feasible. We show that for the right choice of α and γ, our mechanism is deterministic, truthful, budget feasible and has an approximation guarantee of $1 + r$, improving the bound $2r^2$ claimed by [AGN14][1] (where r is the approximation guarantee of the oracle used). Using the greedy algorithm of [Svi04] (which was also analyzed in [KMN99] for linear valuations), the approximation guarantee of our mechanism is $1 + \frac{e}{e-1} \simeq 2.58$. In Figure 2.1.2 we

[1][AGN14] achieve the claimed $2r^2$ by replacing the subroutine computing the optimal solution in their exponential mechanisms by an r-approximation algorithm. Unfortunately, this appears to break the truthfulness of the mechanism, as we point out in Section 2.2.2

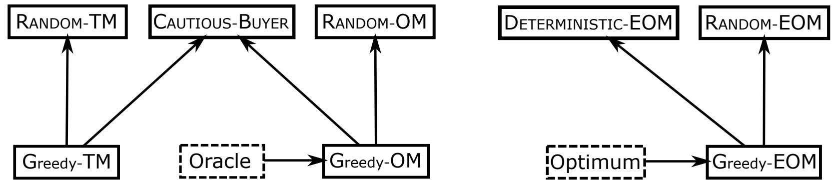

Figure 2.1: Edges show mechanisms that are used as a subroutine of others. The mechanisms on left and right run polynomial and exponential time respectively. Oracle is a polynomial time algorithm that approximately solves the budgeted optimization problem for monotone submodular valuations.

show our mechanisms and their subroutines.

- In section 2.4, we show how our results for submodular valuations can be used for the problem of Crowdsourcing Markets with Heterogeneous Tasks of [GNS14]. This implies that our large market mechanism in section 2.3 is a deterministic, truthful and budget feasible mechanism with $1 + \frac{e}{e-1} \approx 2.58$ approximation guarantee for this problem. The resulting deterministic mechanism matches the approximation guarantee of the randomized truthful (in expectation) mechanism of [GNS14] for this problem.

2.1.3 Related Work

Prior free budget feasible mechanism design for buying a set of items, each from a different seller, has been introduced by [Sin10]. For monotone submodular valuations, which is the focus of this chapter, [CGL11] improved the mechanism of [Sin10] and its analysis to achieve a 7.91 approximation guarantee, and also derandomized the mechanism to get a deterministic (but exponential time) mechanism with an approximation guarantee of 8.34. [ABM16] showed this mechanism can be derandomized given a linear programming (LP) relaxation of an integer program (IP) for this problem in polynomial time. Their approximation guarantee depends on some parameters of the input IP and its LP relaxation, but does not improve the 8.34 guarantee[2].

[SK13] considered the problem for an application in community sensing and gave a mechanism with a 4.75 approximation guarantee for large markets. [AGN14] improved the result of [SK13] achieving a 3 approximation guarantee for large markets with a polynomial time mechanism and a 2 approximation guarantee with an exponential time mechanism. [AGN14] also proposed a mechanism that given an r approximation oracle for the budgeted value maximization problem for monotone submodular functions, has a $2r^2$ approximation guarantee, however, their mechanism uses an optimization algorithm as an oracle, and loses would loose monotonicity (and hence truthfulness), when using the greedy algorithm (see Section 2.2.2 for more details). We overcome this difficulty (in addition to improving the bound) by allowing winning sets of items that are no longer contiguous in the order of their marginal bang per buck.

Budget feasible mechanism design has also been considered with special valuation functions, where better bounds are known. For example, for additive valuations the best known mechanism achieves an approximation bound of $2 + \sqrt{2} \approx 3.41$ and 3 with a deterministic and randomized mechanisms respectively due to [CGL11], who also gave a $1 + \sqrt{2} \approx 2.41$ lower bound for approximation ratio of any truthful budget feasible mechanism in this setting.

[2]The same authors somewhat improved this bound in [ABM17] using our analysis and after seeing a preliminary version of our paper. The improved bound is at least 5.45

In large markets with additive valuations, [AGN14] improved these results and gave a budget feasible mechanism with an approximation guarantee of $\frac{e}{e-1}$ with a matching lower bound.

[Sin10] also introduced the feasible mechanism design problem for matchings on bipartite graphs: the principal is required to select a matching of a bipartite graph, where each individual edge is an agent with a private cost and a public value. [CGL11] consider the knapsack problem with heterogeneous items, which is the special case of this problem where the bipartite graph is a set of disjoint stars. Their approximation bound mentioned above for additive item values, of $2+\sqrt{2}$ and 3 with a deterministic and randomized mechanisms, also extend to this case. [GNS14] considered a variant of the problem motivated by Crowdsourcing Markets, as defined above, where one side of the graph are agents with private costs, and the other side are tasks, each with a value for the principal. They give a randomized truthful (in expectation) mechanism with a $1+\frac{e}{e-1}$ approximation guarantee for this problem under the large market assumption. We consider this version of the matching problem in section 2.4. Even though adding the hard constraint that the winning set should be a matching breaks submodularity of the valuation function (see [CGL11, GNS14, ABM16]), we show that our mechanisms for the case of monotone submodular valuations still can be used for this case, matching the approximation guarantee of [GNS14] with a deterministic mechanism.

Prior free budget feasible mechanisms has also been studied for more general valuation functions. Monotone submodular valuations are the most general class of valuation functions for which a constant factor approximation guarantee with a polynomial time (with a value oracle), truthful and budget feasible mechanism is known. For subadditive valuations [DPS11] introduced a mechanism using a demand oracle (more powerful than the value oracle we use). The current best bound is an $O(\frac{\log n}{\log \log n})$ approximation guarantee due to [BCGL12]. [BCGL12] also gave randomized mechanism that achieves a constant (768) approximation guarantee for fractionally subadditive (XOS) valuations, also using a demand oracle.

Some papers consider the Bayesian setting, where cost of each agent comes from known independent distributions. [BCGL12] gave a constant-competitive mechanism for subadditive valuations (with a very large constant). [BH16] gave a $(\frac{e}{e-1})^2$-competitive posted pricing mechanism for monotone submodular valuations for large markets, using a cost version for defining the largeness of the market. The benchmark (optimum) used in [BH16] is the outcome of optimal Bayesian incentive compatible mechanism, while others (including us) have used the significantly higher, optimum with respect to the budgeted pure optimization problem as their benchmark. It is interesting to compare our results for large markets to the approximation guarantee of $(\frac{e}{e-1})^2 \approx 2.5$ of the mechanism in [BH16]. While this bound is ≈ 0.08 better than our bound, their benchmark, the optimal Bayesian incentive compatible mechanism, can be a factor of $\frac{e}{e-1}$ lower than our benchmark of the optimum ignoring incentives [AGN14]. Even when the cost of sellers come from a uniform distributions, and the value of each item is 1, the ratio between the two benchmarks is $\sqrt{2}$.

2.1.4 Preliminaries

We consider the problem of a single buyer with a budget B facing a set of multiple sellers A, each selling a single item. We let n denote the number of sellers and we assume $A = [n]$. We assume that the value $v(S)$ of the buyer for a set of items S, is a (non-decreasing) submodular function, that is, it satisfies $v(S) \leq v(T)$ for every $S \subseteq T$, and $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$, for every set $S, T \subseteq A$. For every $i \in A$ and $S \subseteq A$, we define $m_i(S) = v(S \cup \{i\}) - v(S)$, i.e. the marginal value of i with respect to subset S. Note that $v(.)$ is monotone submodular if and only if for every $S, T \subseteq A$ we have:

$$v(T) \leq v(S) + \sum_{i \in T \setminus S} m_i(S).$$

Large market assumption. For some of the results in Chapter 2, we consider large markets, assuming that the value of each agent is small compared to the optimum, i.e. $v(i) \ll opt(A)$ for all $i \in [n]$. For simplicity, we state our approximation bounds for large markets in the limit[3], assuming $\theta = \max_{i \in [n]} \frac{v(i)}{opt(A)} \to 0$.

The mechanism design problem of selecting sellers maximizing the buyer's value subject to his budget constraint, is a single parameter mechanism design problem, in which each bidder (seller) has one private value (the cost of her item). We design truthful, deterministic and individually rational mechanisms, as well as universally truthful and individually rational randomized mechanisms. A randomized mechanism is *universally truthful* if it is a randomization

[3]By having a θ-large market assumption instead, the approximation guarantees for our large market mechanisms increases by a factor of $(1 - c\theta)^{-1}$, where $c \in (0, 4)$ is a constant which is different for each mechanism. We omit stating the exact value of c for each mechanism separately.

among deterministic mechanisms, each of which are truthful. We use Myerson's characterization for truthful mechanisms:

Theorem 1. *[Mye81] In single parameter domains, a normalized mechanism $M = (f, p)$ is truthful if and only if*

- f ***is monotone****: $\forall i \in [n]$, if $c'_i \leq c_i$ then $i \in f(c_i, c_{-i})$ implies $i \in f(c'_i, c_{-i})$, or equivalently, $c'_i \notin f(c'_i, c_{-i})$ implies $c_i \notin f(c_i, c_{-i})$.*
- ***Winners are paid threshold payments****: if $i \in [A]$ is a winner and receives payment p_i, then $p_i = \inf\{c_i : i \notin f(c_i, c_{-i})\}$.*

In order to show a mechanism is universally truthful and budget feasible, it suffices to show that the allocation is monotone and by using the threshold payments, the total payments are not more than the budget. Similar to [DPS11, CGL11, BCGL12, SK13], we assume that the payments are threshold payments and only specify the allocation rule. At the end of each section, we briefly explain how the payment rule of the mechanisms in that section can be computed. In all our mechanisms if a seller bids a cost more than B, he will not be selected in the winning set, hence will have utility 0. This combined with the fact that all our mechanisms are truthful, implies *individual rationality*, i.e. in all of our mechanisms utility of sellers are non-negative.

2.2 Parameterized Mechanisms for Submodular Valuations

In this section we present two simple parameterized mechanisms. We show that these parameterized mechanisms provide good approximation guarantees, and are monotone and hence can be turned into truthful mechanisms with payments defined appropriately. We analyze the approximation guarantee of these mechanisms with and without the large market assumption and give conditions that make these mechanisms budget feasible.

Let $S_0 = \emptyset$, and for each $i \in [n]$, recursively define $S_i = S_{i-1} \cup \{\arg\max_{j \in A \setminus S_{i-1}} (\frac{m_j(S_{i-1})}{c_j})\}$, adding the item with maximum marginal value to cost ratio, to S_{i-1}. To simplify notation, we will assume without loss of generality that $\{i\} = S_i \setminus S_{i-1}$. All of our mechanisms sort the items in descending order of marginal bang for buck at the beginning and consider items in this order.

2.2.1 The Threshold Mechanisms

Our threshold mechanism generalizes the mechanisms of Singer [Sin10] and Chen et al [CGL11]. We consider items in increasing cost-to-marginal value order, as defined above. Our *greedy threshold mechanism*, GREEDY-TM, sets a threshold for the cost to marginal value ratio of the items, compared to the ratio of the budget to the total value of the set selected. Using a parameter γ, the mechanism adds items while they are relatively cheap compared to the total so far.

The GREEDY-TM mechanism works well for large markets where each individual item has small value compared to the optimum. In the general case, we will randomly choose between just selecting the item with maximum individual value and cost below the budget, or running GREEDY-TM. We call the resulting randomized mechanism RANDOM-TM(γ, A, B).

GREEDY-TM(γ, A, B)

(Greedy Threshold Mechanism)

Let $k = 1$

while $k \leq |A|$ and $\frac{c_k}{m_k(S_{k-1})} \leq \gamma \frac{B}{v(S_k)}$

do

$k = k + 1$

end

return S_{k-1}

RANDOM-TM(γ, A, B)

(Random Threshold Mechanism)

Let $A = \{i : c_i \leq B\}$

Let $i^* = argmax_{i \in [n]}(v(i))$

With probability $\frac{\gamma+1}{\gamma+2}$ **do**

return GREEDY-TM(γ, A, B)

halt

return i^*

The randomized mechanisms for submodular functions in Singer [Sin10] is similar to RANDOM-TM with parameter $\gamma = \frac{e-1}{12e-4}$ and the improved mechanism of Chen et al [CGL11] is equivalent to RANDOM-TM with $\gamma = 0.5$. In this section we offer a sketch of an improved analysis with details deferred to Appendix A.1.

Monotonicity of the mechanisms is easy to see: if someone is not chosen, he cannot be selected by increasing his cost (decreasing his marginal bang per buck).

Lemma 2. *For every fixed $\gamma \in (0, 1]$, the mechanism* GREEDY-TM(γ, A, B) *is monotone.*

We show that for every fixed $\gamma \in (0,1]$, RANDOM-TM(γ, A, B) achieves a $1+\frac{2}{\gamma}$ approximation of the optimum, improving the bound of [CGL11]. The key difference is that we compare the output of GREEDY-TM directly with the true optimum, rather than a fractional greedy solution. Doing this not only improves the approximation factor, but also simplifies the analysis. We use this idea in the proof of the following technical lemma, which is the main ingredient for proving the approximation guarantee of our mechanism. The detailed proof is deferred to Appendix A.1.

Lemma 3. *For every fixed $\gamma \in (0,1]$, if $S_{k-1} =$ GREEDY-TM(γ, A, B) then*

$$(1+\frac{1}{\gamma})v(S_{k-1}) + \frac{1}{\gamma}v(i^*) \geq opt(A). \tag{2.2.1}$$

By using the above lemma for the performance of RANDOM-TM, we can get the following approximation bound for RANDOM-TM(γ, A, B).

Theorem 4. *For every fixed $\gamma \in (0,1]$, RANDOM-TM(γ, A, B) is universally truthful, and has approximation ratio of $1+\frac{2}{\gamma}$.*

Proof Sketch. Monotonicity of the mechanism follows from Lemma 2.

The main idea for proving the approximation ratio is calculating expected value of the outcome of the mechanism and using inequality 2.2.1 from Lemma 3. ■

The mechanisms GREEDY-TM(γ, A, B) and RANDOM-TM(γ, A, B) are not necessarily budget feasible for an arbitrary choice of γ. However, [CGL11] shows that RANDOM-TM$(0.5, A, B)$ (which they call RANDOM-SM) is budget feasible. We include a simplified proof in Section A.1.

Combining the budget feasibility proof of [CGL11] and Theorem 4 for the general case, and using inequality (2.2.1) directly, instead of Theorem 4, for the case of large market, where $v(i^*) \ll opt(A)$, we get the following theorem. The bound for large markets is matching the best approximation guarantee of Anari et al [AGN14] for submodular functions with computational constraint. In Section 2.3 we improve this bound, while in Section A.1 we show that the analysis in this section for RANDOM-TM$(0.5, A, B)$ is tight.

Corollary 5. RANDOM-TM$(0.5, A, B)$ *is truthful, budget feasible and has approximation ratio of* 5*. For the case of large market case, where* $v(i^*) \ll opt(A)$, GREEDY-TM(γ, A, B) *is is truthful, budget feasible and has approximation ratio of* 3*.*

Note that we need $\gamma = 0.5$ to get the above corollary, however in Section 2.3 we use this mechanism with a larger value of γ. The threshold payment of each agent i in the winning set for the threshold mechanisms in this section can be computed by increasing i's cost until he reaches the threshold that makes him not eligible to be in the winning set, while keeping the cost of other agents fixed. In order to compute this number in polynomial time, it is enough to fix other agents' costs and see where in the sorted list of marginal bang-per-bucks this agent can be appear such that the inequality of GREEDY-TM(γ, A, B) still holds for her. The more detailed characterization of these threshold payments is similar to that of [Sin10].

2.2.2 The Oracle Mechanisms

Here, we provide a different class of parameterized mechanisms. This class of mechanisms requires an oracle $Oracle(A, B)$, which considers the optimization problem of maximizing the value of a subset of A, subject to the total cost not exceeding the budget B, and returns a value which is close to optimum. Let $\mathit{opt}(A, B)$ denote the optimum value of this optimization problem. We assume that $\mathit{opt}(A, B) \geq Oracle(A, B)$. The oracle is an r approximation, if we also have $r \cdot Oracle(A, B) \geq \mathit{opt}(A, B)$. For instance, the greedy algorithm of Sviridenko [Svi04] can be used as an oracle with $r = \frac{e}{e-1} \approx 1.58$. Since calculating the $\mathit{opt}(A, B)$ is not possible in polynomial time, we call the mechanisms that use optimum value *the exponential time oracle mechanisms*. For some of our analysis in this section we offer a proof sketch here and defer the detailed proof to Appendix A.2.

Exponential Time Oracle Mechanisms We start with a simple *exponential time oracle mechanism*, GREEDY-EOM, using the optimal solution value $\mathit{opt}(A, B)$ as an oracle. The optimum value is simpler to use, as it is monotone in the cost of the agents[4]. Later, we show how to use a polynomial time approximation oracle instead of $\mathit{opt}(A, B)$, with a small sacrifice in the approximation ratio while keeping the mechanism truthful and budget feasible. Our mechanisms in this section also sort the items in decreasing order of bang-per-buck, as explained in the beginning of the chapter.

[4]For large markets, Anari et al [AGN14] use an exponential time mechanism which is equivalent to GREEDY-EOM$(0.5, A, B)$

GREEDY-EOM(α, A, B)

(exp. time oracle mechanism)

Let $v^* = opt(A, B)$

Let $k = 1$

while $v(S_k) \leq \alpha v^*$ **do**

  $k = k + 1$

end

return S_{k-1}

RANDOM-EOM(α, A, B)

(exp. time oracle mechanism)

Let $A = \{i : c_i \leq B\}$

Let $i^* = argmax_{i \in [n]}(v(i))$

With probability $\frac{1}{2}$ **do**

  return GREEDY-EOM(α, A, B)

  halt

return i^*

DETERMINISTIC-EOM(α, A, B)

(det. exp. time oracle mechanism)

Let $A = \{i : c_i \leq B\}$

Let $i^* = argmax_{i \in [n]}(v(i))$

Let $v = opt(A \setminus \{i^*\}, B)$

if $v(i^*) \geq \frac{\sqrt{17}-3}{4} v$ **then**

  return i^*

else

  return GREEDY-EOM($0.5, A, B$)

end

Lemma 6. *For every fixed* $\alpha \in (0, 1]$, GREEDY-EOM(α, A, B) *is monotone, and if* $S_{k-1} =$ GREEDY-EOM(α, A, B) *then* $\frac{1}{\alpha} v(S_{k-1}) + \frac{1}{\alpha} v(i^*) \geq opt(A, B)$.

Proof. We first argue that the mechanism is monotone. Assume $i \in A$ and $i \notin S_{k-1}$. If i increases her cost, it cannot increase the value of $v^* = opt(A, B)$. Furthermore, by increasing i's cost, her marginal bang per buck decreases, which cannot help her get selected, so the mechanism is monotone.

To see the approximation bound simply note that by the definition of the mechanism we have $v(S_k) > \alpha v^*$. So $v(S_{k-1}) + v(i^*) > v(S_{k-1}) + v(k) > \alpha v^*$. ∎

By using Lemma 6, it is easy to prove the following theorem.

Theorem 7. *For every fixed $\alpha \in (0,1]$,* RANDOM-EOM(α, A, B) *is universally truthful, and if $S =$* RANDOM-EOM(α, A, B) *then $\frac{2}{\alpha}E[v(S)] \geq opt(A)$.*

Proof. By using Lemma 6 and similar argument to proof of truthfulness in Theorem A.1, it is easy to see that the mechanism is universally truthful.

By definition of RANDOM-EOM(α, A, B) we have

$$E[v(S)] = \frac{1}{2}v(S) + \frac{1}{2}v(i^*)$$
$$\Rightarrow \frac{2}{\alpha}E[v(s)] = \frac{1}{\alpha}v(S) + \frac{1}{\alpha}v(i^*)$$

By using Lemma 6 the proof is complete. ∎

Now we show that for the choice of $\alpha = 0.5$, GREEDY-EOM$(0.5, A, B)$ is budget feasible, so RANDOM-EOM$(0.5, A, B)$ is universally truthful, budget feasible, and a 4 approximation to the optimum. Slightly more complex analog of this Lemma for the mechanism using an approximation algorithm in place of the true optimum will be Lemma 12.

Lemma 8. *By using threshold payments,* GREEDY-EOM$(0.5, A, B)$ *is budget feasible.*

Proof. Let p_i be the threshold payment for agent i. Let S_{k-1} =GREEDY-EOM$(0.5, A, B)$. For every $i \in S_{k-1}$, we show that if i deviates to a bid of $b_i > m_i(S_{i-1})\frac{B}{v(S_{k-1})}$, he cannot be selected, implying that the threshold payment $p_i \leq m_i(S_{i-1})\frac{B}{v(S_{k-1})}$. By proving this we get that $\sum_{i \in S_{k-1}} p_i \leq$

$\sum_{i \in S_{k-1}} m_i(S_{i-1})\frac{B}{v(S_{k-1})} = B$, where the last inequality hold by recalling that $m_i(S_{i-1}) = v(S_i) - v(S_{i-1})$, so the sum telescopes.

So the mechanism is budget feasible.

We prove the inequality claimed above by contradiction: assume that i deviates to $b_i > m_i(S_{i-1})\frac{B}{v(S_{k-1})}$ and is still in the winning set. Let b be the new vector of costs with i bidding b_i and all other agents bidding their true cost. Note that the order in which items are considered after item $i-1$ is also effected by the change in i's claimed cost. Now let j be the step in the mechanism in which i is added to the winning set after he deviates to b_i and S'_j be the wining set after that step, where S'_z for $z \in [n]$ is defined similar to S_z but with cost vector b instead of c. Let S^* be the optimum solution with $v(S^*) = v^*$. Let $S^* \setminus S'_j = \{t_1, t_2, \ldots, t_q\}$, $T_0 = \emptyset$, and $T_z = \{t_l : l \in [z]\}$. Since i is the only item that has increased his cost and $i \in S'_j$, we have

$$c(S^*) \geq b(S^* \cup S'_j) - b(S'_j) = \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1})\frac{b_{t_z}}{m_{t_z}(S'_j \cup T_{z-1})}$$

By submodularity and by the fact that the mechanism choose the ordering having item i in position j (with costs b) we have that $m_{t_z}(S'_j \cup T_{z-1}) \leq m_{t_z}(S'_{j-1})$ and $b_{t_z}/m_{t_z}(S'_{j-1}) \geq b_i/m_i(S'_{j-1})$. Using these two inequalities we get

$$\begin{aligned}\sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1})\frac{b_{t_z}}{m_{t_z}(S'_j \cup T_{z-1})} &\geq \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1})\frac{b_{t_z}}{m_{t_z}(S'_{j-1})} \geq \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1})\frac{b_i}{m_i(S'_{j-1})} \\ &= \frac{b_i}{m_i(S'_{j-1})}(v(S^* \cup S'_j) - v(S'_j)) \geq \frac{b_i}{m_i(S'_{j-1})}(v(S^*) - v(S'_j))\end{aligned}$$

Now by using the contradiction assumption, and the fact that $S_{i-1} \subseteq S'_{j-1}$, we get

$$\frac{b_i}{m_i(S'_{j-1})}(v(S^*) - v(S'_j)) > B\frac{v(S^*) - v(S'_j)}{v(S_{k-1})}$$

By combining the above inequalities and using the fact that $v(S'_j), v(S_{k-1}) < \alpha v^*$ and $\alpha = 0.5$, we have $c(S^*) > B$ which is a contradiction, so the mechanism is budget feasible. ∎

Corollary 9. *By using threshold payments,* RANDOM-EOM$(0.5, A, B)$ *is universally truthful, budget feasible and a 4 approximation of the optimum.*

Next we offer a simple deterministic version of this mechanism, with a significantly better approximation factor, improving the previously known 8.34 approximation exponential time mechanism of [CGL11] to a guarantee of 4.56. In order to derandomize RANDOM-EOM, we would like to check if the optimum is large enough compared to the best valued item. To keep the mechanisms monotone, we will compare the value of the highest valued item i^* to the optimum after removing item i^*.

Theorem 10. *By using threshold payments,* DETERMINISTIC-EOM *is truthful, budget feasible and has an approximation ratio of 4.56.*

Proof. i^* cannot change $opt(A \setminus \{i^*\}, B)$, so since GREEDY-EOM$(0.5, A, B)$ is monotone and budget feasible, DETERMINISTIC-EOM is truthful and budget feasible.

In order to prove the approximation ratio, we consider two cases:

Case 1: $v(i^*) \geq \frac{\sqrt{17}-3}{4} v = \frac{\sqrt{17}-3}{4} opt(A \setminus \{i^*\}, B)$: in this case the algorithm returns i^*, and we have

$$(\frac{4}{\sqrt{17}-3} + 1) v(i^*) \geq opt(A \setminus \{i^*\}, B) + v(i^*) \geq opt(A, B)$$

Case 2: $v(i^*) < \frac{\sqrt{17}-3}{4}v \leq \frac{\sqrt{17}-3}{4}opt(A,B)$: let $S_{k-1} = \textsc{Greedy-EOM}(0.5, A, B)$. In this case, by Lemma 6, we have

$$
\begin{aligned}
&2v(S_{k-1}) + 2v(i^*) \geq opt(A) \\
\Rightarrow &2v(S_{k-1}) + 2(\frac{\sqrt{17}-3}{4})opt(A) \geq opt(A) \\
\Rightarrow &\frac{2}{1-\frac{\sqrt{17}-3}{2}}v(S_{k-1}) \geq opt(A)
\end{aligned}
$$

So in any case the approximation ratio is $\frac{4}{\sqrt{17}-3} + 1 = \frac{2}{1-\frac{\sqrt{17}-3}{2}} \simeq 4.56$ ■

Here we offer an oracle mechanism that uses a oracle in place of the optimum value $opt(A,B)$, as finding the optimum for monotone submodular maximization with a knapsack constraint cannot be done in polynomial time. Naively using the outcome of a sub-optimal oracle instead of optimum in our exponential time oracle mechanisms in appendix 2.2.2 can break monotonicity. To see this, note that if an item increases his cost, she cannot increase the value of $opt(A,B)$. However, if we replace $opt(A,B)$ with the outcome of a sub-optimal oracle (for instance a greedy algorithm), this is no longer true: if one increases the cost of all the items that are not in the optimum set to be more than the budget, any reasonable approximation algorithm for submodular maximization (for instance the greedy algorithm in [Svi04]) can detect and choose all the items that are in optimum set. Now if an item i increases her cost and this increase the value of the outcome of the approximation algorithm compared to $opt(A,B)$, this gets the mechanism to take more items, possibly including item i.

To make our mechanism monotone, we remove i before calling the oracle to decide if we should add i to the set S, making the items selected no longer contiguous in the order we consider them.

GREEDY-OM(α, A, B)

(Greedy Oracle Mechanism)

Let $S = \emptyset$

for $i = 1$ to n **do**
  if $v(S_i) \leq \alpha Oracle(A \setminus \{i\}, B)$
  then
    $S = S \cup \{i\}$
  end
end
return S

RANDOM-OM(α, A, B)

(Random Oracle Mechanism)

Let $A = \{i : c_i \leq B\}$

Let $i^* = argmax_{i \in [n]}(v(i))$

With probability $\frac{r}{\alpha+2r}$ **do**
  return GREEDY-OM(α, A, B)
  halt
return i^*

Next we show that GREEDY-OM(α, A, B) is monotone and provide its approximation ratio.

Lemma 11. *For every fixed $\alpha \in (0, 1]$, GREEDY-OM(α, A, B) is monotone. If $S =$ GREEDY-EOM(α, A, B), $k \in [n]$ is the biggest integer such that $S_{k-1} \subseteq S$, i^* is the item with maximum individual value, and assuming ORACLE is an r approximation of the optimum, then*

$$\frac{r}{\alpha} v(S_{k-1}) + (1 + \frac{r}{\alpha}) v(i^*) \geq opt(A)$$

Proof. Monotonicity of the mechanism follows from the usual argument, increasing c_i does not effect ORACLE$(A \setminus \{i\}, B)$ and decreases the item's bang per buck in any step so it can only increase the value of $v(S_i)$. To show the approximation factor, recall that $\{k\} = S_k \setminus S_{k-1}$. Since k was not chosen by the mechanism we have

$$v(S_{k-1}) + v(k) \geq v(S_k) > \alpha Oracle(A \setminus \{k\}) \geq \frac{\alpha}{r} opt(A \setminus \{k\})$$
$$\geq \frac{\alpha}{r}(opt(A) - v(k)) \geq \frac{\alpha}{r}(opt(A) - v(i^*)) \quad \blacksquare$$

Next Lemma shows that GREEDY-OM$(0.5, A, B)$ is budget feasible. We include a sketch of the proof here with details deferred to Appendix A.2.

Lemma 12. *By using threshold payments,* GREEDY-OM$(0.5, A, B)$ *is budget feasible.*

Proof Sketch. We prove that if each agent i bids higher than $m_i(S_{i-1})\frac{B}{v(S)}$, where $S =$GREEDY-OM$(0.5, A, B)$, then she will not be in the winning set. We use contradiction to prove this showing that, if an agent bids higher than this amount and is still in the winning set, then cost of optimum has to be higher than the budget which is a contradiction. ■

Note that in large markets $v(i^*) \ll opt(A)$, so by combining Lemma 11 and Lemma 12, we get the following corollary.

Corollary 13. *In large markets,* GREEDY-OM$(0.5, A, B)$ *is truthful, budget feasible and given an r-approximation oracle, achieves $2r$ approximation of the optimum.*

The previously known best oracle mechanism for large markets is due to Anari et al [AGN14] achieves $2r^2$. We will improve this bound for the case of large markets to $r + 1$ in section 2.3.

By using Lemma 11, we get the following theorem, whose proof is deferred to Appendix A.2.

Theorem 14. RANDOM-OM(α, A, B) *is truthful and in expectation achieves* $1 + \frac{2r}{\alpha}$ *of the optimum, assuming the oracle used is an* r*-approximation.*

By combining Lemma 12 and Theorem 14 we have the following theorem.

Theorem 15. RANDOM-OM$(0.5, A, B)$ *is truthful, budget feasible and in expectation achieves* $1 + 4r$ *of the optimum.*

By using the greedy algorithm of [Svi04] as an oracle, we can improve the approximation ratio to $\frac{2r}{\alpha}$. To achieve this, we change GREEDY-OM(α, A, B), so that instead of using $Oracle(A \setminus \{k\}, B)$, it uses $\max_{c'_i \geq c_i} Oracle(A, (c'_i, c_{-i}))$. We also change the probability of choosing the greedy mechanism's outcome in RANDOM-OM(α, A, B) to $\frac{1}{2}$. By doing so, RANDOM-OM(α, A, B) can achieve $\frac{2r}{\alpha}$ instead of $1 + \frac{2r}{\alpha}$. By using the greedy algorithm of [Svi04], as an oracle, finding $\max_{c'_i \geq c_i} Oracle(A, (c'_i, c_{-i}))$ can be done in polynomial time, since we only have to check polynomial number of cases for c'_i. Furthermore, if i increases his cost, he cannot increase the value of $\max_{c'_i \geq c_i} Oracle(A, (c'_i, c_{-i}))$. We omit the proof of the following theorem, as it is analogous to our previous proofs.

Theorem 16. *The above modification of the* RANDOM-OM$(0.5, A, B)$ *mechanism is truthful, budget feasible, get expected value of a* $4r$ *fraction of the optimum. With the greedy algorithm as the oracle, it can be implemented in polynomial time, and is a* $4e/(e-1)$*-approximation mechanism.*

For calculating the agents' threshold payments of our oracle mechanisms in this section, it is enough to check what is the maximum cost that each agent i can declare such that she is still in the winning set. Similar to section 2.2.1, for each agent i, this number can simply be computed by checking where in the sorted list of agents by their marginal bang-per-buck this agent can appear such that the inequality of GREEDY-EOM(α, A, B) (for the exponential time mechanisms) and the inequality of GREEDY-OM(α, A, B) (for polynomial time mechanisms) still hold. The characterization of these threshold payments is similar to the payment characterization of the oracle mechanisms in [AGN14].

2.3 A Simple $1+\frac{e}{e-1}$ Approximation Mechanism for Large Markets

In this section we combine the two greedy parameterized mechanisms of Section 2.2, GREEDY-OM(α, A, B) and GREEDY-TM(γ, A, B) to improve the approximation guarantee for large markets. Given a polynomial time r approximation oracle, our simple, deterministic, truthful, and budget feasible mechanism in this section has an approximation ratio of $1 + r$ and runs in polynomial time.

CAUTIOUS-BUYER(α, γ, A, B)

Let $A = \{i : c_i \leq B\}$

Let $S_\alpha =$ GREEDY-OM(α, A, B)

Let $S_\gamma =$ GREEDY-TM(γ, A, B)

return $S_\alpha \cap S_\gamma$

At first glance, CAUTIOUS-BUYER seems worse than both of GREEDY-OM and GREEDY-TM, since its winning set is the intersection of the wining sets of these mechanisms. However, taking the intersection of these mechanisms will allow us to choose the value of α and γ to be higher than 0.5 while keeping the mechanism budget feasible.

It is easy to see that the intersection of two monotone mechanisms is monotone.

Proposition 17. *For two monotone mechanisms M_1 and M_2, the mechanism M that outputs the intersection of the winning set of M_1 and the winning set of M_2 is monotone.*

Next we give a parameterized approximation guarantee for CAUTIOUS-BUYER(α, γ, A, B).

Lemma 18. *Assuming the large market assumption, for every fixed value of $\alpha, \gamma \in (0, 1]$ CAUTIOUS-BUYER(α, γ, A, B) is monotone. Furthermore, with an r approximation oracle, it has a worst case approximation ratio of $\max(1 + \frac{1}{\gamma}, \frac{r}{\alpha})$.*

Proof. From Proposition 17, Lemmas 3, 11 it follows that CAUTIOUS-BUYER is monotone.

Let S be the outcome of the mechanism. Let k be the biggest integer such that $S_{k-1} \subseteq S$, i.e., $S_{k-1} \subseteq S_\alpha$ and $S_{k-1} \subseteq S_\gamma$. By definition $k \notin S$, so there are two cases

- $k \notin S_\alpha$: By lemma 11, the large market assumption and monotonicity of $v(.)$, we have $\frac{r}{\alpha}v(S) \approx \frac{r}{\alpha}v(S) + (1 + \frac{r}{\alpha})v(i^*) \geq \frac{r}{\alpha}v(S_{k-1}) + (1 + \frac{r}{\alpha})v(i^*) \geq opt(A, B)$.

- $k \notin S_\gamma$: By lemma 3, the large market assumption and monotonicity of $v(.)$, we have $(1+\frac{1}{\gamma})v(S_{k-1}) \approx (1+\frac{1}{\gamma})v(S_{k-1}) + \frac{1}{\gamma}v(i^*) \geq opt(A,B)$.

In both cases we have, $\max(1+\frac{1}{\gamma}, \frac{r}{\alpha})v(S) \geq opt(A,B)$ assuming that $v(i^*)$ is negligible. ■

Now we provide a simple condition for the budget feasibility of CAUTIOUS-BUYER(α,γ,A,B).

Lemma 19. *If $\alpha \leq \frac{1}{1+\gamma}$ for any $\alpha,\gamma \geq 0$, then by using threshold payments, CAUTIOUS-BUYER(α,γ,A,B) is budget feasible.*

Proof. Let p_i be the threshold payment for agent i. Let S =CAUTIOUS-BUYER(α,γ,A,B). For every $i \in S$, we show that if i deviates to bidding a cost $b_i > m_i(S_{i-1})\frac{B}{v(S)}$, he cannot be in the winning set. By proving this and by using the definition of threshold payments we get $\sum_{i\in S} p_i \leq \sum_{i\in S} m_i(S_{i-1})\frac{B}{v(S)} \leq \sum_{i\in S} m_i(S_{i-1}\cap S)\frac{B}{v(S)} = B$, so the mechanism is budget feasible.

We prove above claim by contradiction: assume that i deviates to $b_i > m_i(S_{i-1})\frac{B}{v(S)}$ and is in the winning set. Let b be the new cost vector and j be position of i in the new order of items. Let S'_z for $z \in [n]$ be defined similar to S_z but with cost vector b instead of c. So S'_j is the set of items that are in the winning set of GREEDY-TM(γ,A,B) at the end of step j once i is added. Note that S'_j is also equal to the set of all the items that has been considered by GREEDY-OM(α,A,B) at the end of its j-th step. So by using the same argument as proof of lemma 12 we get

$$c(S^*) > B\frac{v(S^*) - v(S'_j)}{v(S)} \qquad \text{and} \qquad v(S'_j), v(S) \leq \alpha v(S^*)$$

By defining $x = \frac{v(S'_j)}{v(S)}$, we have $v(S'_j) = xv(S) \leq \alpha x v(S^*)$. So we get

$$c(S^*) > B\frac{v(S^*) - v(S'_j))}{v(S)} > B\frac{(1-\alpha)v(S^*)}{\alpha x v(S^*)} = B\frac{1-\alpha}{\alpha x}$$

so if $1 - \alpha \geq \alpha x$, or equivalently, $\alpha \leq \frac{1}{1+x}$ then we get $c(S^*) > B$ which is the desired contradiction.

Since $i \in S_\gamma$, we also have

$$\frac{b_i}{m_i(S'_j)} \leq \gamma \frac{B}{v(S'_j)} = \frac{\gamma}{x}\frac{B}{v(S)}$$

So since $m_i(S'_j) \leq m_i(S_{i-1})$, if $\gamma \leq x$, we get to a contradiction with the assumption about b_i.

The only remaining case is when $\gamma > x$ and $\alpha > \frac{1}{1+x}$. This means that $\alpha > \frac{1}{1+\gamma}$ which is a contradiction with the property in the statement of lemma, so the mechanism is budget feasible. ■

By using Lemmas 18 and 19, the main theorem of this section follows. We include the detailed proof in Appendix A.3.

Theorem 20. *By using threshold payments,* CAUTIOUS-BUYER $(\frac{r}{r+1}, \frac{1}{r})$ *is truthful, budget feasible, and* $1+r$ *approximation of the optimum. By using the greedy algorithm with* $r = e/(e-1)$ *we get a mechanism with approximation guarantee of* ≈ 2.58.

The threshold payment of an agent in this mechanism is the minimum of the threshold payment of two mechanisms we intersected to get CAUTIOUS-BUYER(α, γ, A, B).

2.4 Application to Hiring in Crowdsourcing Markets

In this section, we consider an application of our mechanisms for the problem of a principal hiring in a Crowdsourcing Market. We consider the model where there is a set of agents A that can be hired and a set of tasks T that the principal would like to get done. Each agent $i \in A$ has a private cost c_i. We represent the abilities of the agents by a bipartite graph $G(A, T)$, where edge $e = (a, t)$ in the G indicates that agent a can be used for task t, where each agent hired can be used for at most one of the tasks she can do. The value of buyer for each edge e is v_e, which can be different for each edge. The principal would like to hire agents to maximize the total value of tasks done, while keeping the total payment under her budget B. The optimal solution for this problem is a maximum value matching, subject to the budget constraint on the cost of the hired agents.

This model is also known as knapsack with heterogeneous items. Knapsack with heterogeneous items and buyer with a matching constraint for budget feasible mechanism design was also studied by [Sin10, CGL11, GNS14, ABM16]. There are many ways for modeling the heterogeneity of items. In [CGL11], this heterogeneity has been defined by having types for items where at most one item can be chosen from each type, corresponding to a bipartite graph where agents have degree 1. [GNS14] consider our model of agents and tasks with a bipartite graph, but assume that the principal has a fixed value for each task completed, independent of the agent that took care of the task, so values of the edges entering a task node t are all equal.

In this section, we apply our technique from section 2.3 for this problem. For the model used by [GNS14], where the principal has a value for each task independent of who completes the task, we show that a small change in our mechanism (stopping it when the marginal increase in value is 0) results in the same approximation guarantee. The small modification is needed as the value for the buyer is not a submodular function of the set of agents hired (see [ABM16]).

Theorem 21. *The truthful budget feasible threshold and oracle mechanisms, as well as the large markets mechanism for submodular valuations of this chapter without any loss in the approximation ratio can be also used for the case of heterogeneous tasks, with the constraint that each agent in the winning set should be assigned to a unique task (matching constraint).*

Before proving this theorem, we consider the general problem defied above, but similar to [Sin10], we relax the assumption that the allocation should always assigned each agent in the winning set to a unique task, and allow instead that multiple agents get assigned to a given task, with only one contributing to the value. We define the value of the buyer for the winning set S to be the value of the maximum matching on the induced subgraph $G[S,T]$. We'll see that allowing the principal to hire extra agents makes her valuation a monotone submodular function. This can model the case where the principal asks more than one worker to accomplish a task, but only keep the best results.

General Crowdsourcing Markets It is well-known and not hard to see that the function defined by value of maximum matching adjacent to a subset of agents S is a monotone and submodular function of S. The following proposition formalizes this statement.

Proposition 22. *For $S \subseteq A$, let $f(S)$ be the maximum value of a matching of the induced subgraph $G[S,T]$ of the bipartite graph $G(A,T)$, then $f(S)$ is a monotone submodular function.*

This proposition implies that all our truthful budget feasible mechanisms for submodular valuations can be used for this model.

Corollary 23. *Without the strict matching constraint, budgeted valuations with heterogeneous tasks (items) are a special case of monotone submodular valuations, so all the mechanisms from the previous sections can be applied to this problem.*

Hiring with Strict Matching Constraint. Consider the case where the buyer's value is defined by summation of her value for each task, i.e. for all the edges that are directly connected to the same task, the value of the buyer for those edge is the same. We argue that in this model, if we add the hard constraint that each agent in the winning set should be assigned to a unique task (similar to [CGL11, GNS14, ABM16]), then with a small change in our mechanisms, all our results still hold. This problem was considered by [GNS14] for large markets, who gave a randomized truthful (in expectation) and budget feasible mechanism with a $1 + \frac{e}{e-1}$ approximation guarantee for large markets (the main result of [GNS14]). Next lemma shows how one can use our truthful budget feasible mechanism for large markets to get the same approximation guarantee with a deterministic mechanism.

Lemma 24. *For $S \subseteq A$, let $f(S)$ be the maximum value of matching of the induced subgraph $G[S,T]$ of the bipartite graph $G(A,T)$ in which all the edges that connect to the same node of T has the same value. If a maximum value matching induced by $S \subseteq A$ connects all vertices in S to a vertex in T, and for $a \in A \setminus S$, $f(S \cup \{a\}) - f(S) > 0$, then there is a maximum value matching induced by $S \cup \{a\}$ which is also assigning each agent to a unique task.*

Proof. We use contradiction. Assume that there is a subset of agents $S \subseteq A$ such that there is a maximum matching M in the subgraph of G, induced by vertices of S and T that connects each agent in S to a task in T. Let $a \in A$ be an agent such that $f(S \cup \{a\}) - f(S) > 0$ and there is no maximum matching in the subgraph induced by $S' = S \cup \{a\}$ and T that connects each agent in S' to a task in T. Let M' be a maximum matching of this induced subgraph. Let G' be the union of edges in M and M' and let C and P be the set of cycles and paths that contain all the edges of G'. Since M and M' are both maximum matchings, we have $W(M \cap c) = W(M' \cap c)$ for all $c \in C$. Since the only difference between S and S' is having a, there can only be one path $p \in P$ such that $W(p \cap M') > W(p \cap M)$. Furthermore, one of the end points of p should be a and for all other paths $p' \in P$ that $p' \neq p$, $W(p' \cap M) = W(p' \cap M')$. For p there are two cases

- The edge that is connected to the other endpoint of p is in M: in this case, since the value of matching is defined by tasks, $W(p \cap M) = W(p \cap M')$. Therefore, $W(M) = W(M')$ and $F(S') - F(S) = 0$ which is a contradiction.

- The edge that is connected to the other endpoint of p is in M': In this case if we define a matching $M^* = (M \setminus p) \cup (M' \cap p)$, then M^* will connect each agent in S' to a unique task, which is a contradiction.

This means that we reach contradiction in both cases, and the proof is complete.
■

Proof of Theorem 21 We use our mechanism this problem, using Corollary 23, but stopping to consider items in the sorted list of marginal bang per bucks whenever the marginal bang-per-buck of the item is 0. Note that since the items are listed in decreasing order of marginal bang per buck and we know that the valuation is submodular, doing this will not have any effects on the approximation ratio (since the marginal bang-per-buck of the next items is also 0) and truthfulness (since the threshold payment of an agent whose item has 0 marginal value is 0) of our mechanisms.

By Lemma 24 in the resulting subset of agents, there is a maximum value matching assigning each agent to a unique task. ■

For the case of large markets, by this corollary and Theorem 20, CAUTIOUS-BUYER is a deterministic truthful and budget feasible mechanism for this problem, matching the $1 + \frac{e}{e-1}$ guarantee of the randomized truthful (in expectation) mechanism of [GNS14].

CHAPTER 3

LEARNING AND TRUST IN AUCTION MARKETS

In this chapter, we study the behavior of bidders in an experimental launch of a new advertising auction platform by Zillow, as Zillow switched from negotiated contracts to using auctions in several geographically isolated markets. To help bidders, Zillow has provided a recommendation tool that suggests a bid to each bidder. Our main focus in this chapter is on the decisions of bidders whether or not to adopt the platform-provided bid recommendation. We observe that a significant proportion of bidders initially do not use the recommended bid. That proportion gradually declines over time. For bidders not following the recommendation, we use their bids to infer their likely value for the advertising opportunity, assuming that the agents are learning to bid in the auction, which a weaker assumption than assuming they know how to instantaneous best respond. Using this inferred value, we find that for half of the agents not following the recommendation, the increased effort of experimenting with alternate bids results in decreased net cumulative utility out of the system. Over time agents discover that the platform adequately optimizes bids on their behalf and *learn to trust* the platform-provided tool.

3.1 Introduction

Auctions are most common mechanisms that are used to allocate and price the sponsored content or ads based on the bids submitted by the bidders. These auctions are run in real time for each user query. While offering the advertisers flexibility in scheduling and changing their campaigns, auction platforms require them to use complex dynamic bidding strategies to compete in the market.

To help humans make decisions in these complex environments there are a variety of automated tools that provide recommendations or make choices on humans' behalf. In online advertising auctions, such as Google's sponsored search advertising auctions, advertisers have access to automated platform-provided or third party bidding tools that automatically adjust their bids based on the declared goals of the advertising campaign. However, to advertisers these tools often appear as black boxes and they may question the integrity of their design, maybe in part as the technology behind these tools is typically proprietary. In this case the advertisers lack understanding why specific recommendations are made and thus may question the accuracy or integrity of these recommendations. This issue is further amplified when the recommendation tool is provided by the platform itself and, thus, advertisers may not trust that the recommendations are designed to benefit them and not the platform.

In this chapter we study the process of bidders adopting the platform-provided bid recommendation in an online advertising auction. Our goal is to understand how and when bidders learn to trust the integrity of the platform-generated recommendation.

Our data comes from an experimental launch of a new advertising auction platform by Zillow. Zillow.com is the largest residential real estate search platform in the United States used by 140 million people each month according to the company's statistic [ZG16]. Viewers are looking to buy or sell houses, want to see available properties, typical prices, and learn about market characteristic. The platform is monetized by showing ads of real estate agents offering their services. Historically, Zillow used negotiated contracts with real-estate agents for placing ads on the platform. In the experiment we study, several geographically isolated markets were switched from negotiated contracts to auction based pricing and allocation. The auction design used was a form of generalized second price, very similar to what is used in many other markets, except that agents were paying for impressions (and not for clicks). A unique feature of this experiment is that the bidders in this market are local real estate agents that bid in the auctions on their own behalf. This is unlike many existing online marketplaces where many bidders use third-party intermediaries to assist with the bidding.

Along with the new auction platform, Zillow provided the bidders with a recommendation tool that suggests a bid to each bidder. Bidders were required to log into the system if they wanted to change their bid, and once they logged in, the system offered a suggested bid: the recommended bid for maximizing the obtained impression volume based on parameters such as the bidder's budget, the estimated impression volume, and budgets and competing bids of other bidders. Bidders in this market are limited by small budgets. During the auction, bidders typically changed their bids relatively frequently, while they tended to keep a closed to fixed budget with the average change in the budget over the time period only around %4. In light of this, we view only the bids as strategic. The bid recommendation tool was designed to suggest a bid maximizing impressions gained by spending the budget. Since the bidders eventually adopted the tool without significantly changing their budgets, we conclude that they appear to have agreed with the goal maximizing the number of impressions gained given the budget. It appears that bidders initially lacked trust in the recommendation and that is why they didn't use it, and both bidders and the platform would have been better off if the system didn't offer bidders the opportunity to avoid the recommended bid.

While it is common to use (Bayes) Nash equilibrium framework to empirically analyze auctions, this framework may not be the best fit for the experiment on Zillow, where auctions were in transition (with an increasing number of bidders invited to the auction platform over time) and bidders were learning how to bid. It is clear, however that the bidders knew their preferences over the ad impressions since all of the bidders have previously participated in the negotiated price-based market for impressions which existed for over 10 years.

We assume that the agents use a form of algorithmic learning to optimize their bid over time. Low-regret assumption is a simple assumption on the type of algorithmic learning used by the agents. We use the assumption to infer the values for agents: for a given regret error parameter ϵ, every possible fixed bid b implies an inequality which indicates that the agent must have a value that makes her regret for not bidding b at most ϵ. Since we do not know the actual value of agents for each impression, we use this method to infer their value indirectly.

We observe that a large proportion of bidders do not use the recommended bid to make bid changes immediately following the introduction to the new market. Many bidders could adjust their bid on their own to gain over the recommended bid since the impression volume is different in different weekdays and bid recommendation tool was not using this information. However, many bidders were not changing their bid frequently enough to take advantage of this and would have been better off by always using recommended bid. The number of bidders who outperformed the bid recommendation in our study is about the same as the number of those who did worse. The proportion of bidders following the recommendation slowly increases as markets mature.

Our work provides an empirical insight into possible design choices for auction-based online advertising platforms. Search advertising platforms (such as Google and Bing) allow bidders to submit bids on their own and there is an established market of third-party intermediaries that help bidders to bid over time. This market design allows for more complex bidding functions, for example, allowing agents to express added value for subsets of the impression opportunities via multiplicative bid-adjustments (e.g., based on the age of the viewer). In contrast, many display advertising platforms (such as Facebook) use a simpler bidding language, and optimize bids on bidders' behalf based solely on their budgets. This eliminates the need for the bidders to bid on their own or use intermediaries. Our empirical analysis shows that the latter approach may be preferred for markets where bidders are individuals who don't have access to third party tools, and who may question the fairness of platform-provided suggestions.

Related Work Number of papers in recent years focus on estimating bidder's value in online advertising auctions. One the earliest papers in this area is

[AN10], who study bidder values in Bing's GSP auction for search ads. They use the equilibrium characterization of GSP to recover values of players.

In dynamic or new markets where interaction is repeated, the value of each individual interaction is small, and bidders are not (yet) knowledgeable about the system, it is better to model players as learners. [NST15] suggest this assumption for studying bidders in Bing's market for search ads, and show how to infer values based on bidding behaviour under this weaker assumption on the outcome. To evaluate the effectiveness of the bid-recommendation tool for the bidders, we need to estimate their value for impressions. We do this by using the methodology developed in [NST15], making the assumption that agents are low-regret learners, rather than relying on the much stronger assumption that their bid is best response to the environment.

In a recent paper, [NN17] report on a human subject experiment on the reliability of regret based inference. In their experiment, human subjects participated in bidding games (including the GSP format). The paper asks the question if human behavior can be modeled as no-regret learning, and to what extent the inference based on the low regret assumption can be used to recover the bidders' values from their bidding behavior. Their finding are mixed. They find the players whose value is high behave rationally, experiment with the best bidding behaviour, achieve very low regret, and inference based on this assumption accurately recovers their value. The finding for players with low types is less positive. Some participants in the experiments were given values so low, that rational behavior would have them drop out of the auction (or bid so low they are guaranteed to lose). Such low value players were frustrated by the game, and behaved rather irrationally at times. It is interesting to think about the contrast between the participants in the Nisan-Noti laboratory experiment and the agents in the Zillow field experiment. The players in the Nisan-Noti experiment were paid to participate (even if frustrated), while in contrast participation in Zillow's ad-auctions is optional, and for typical real estate agents Zillow may not be the main channel through which they get "client leads". Frustrated agents can drop out. In fact, there were many short lived agents in our data. We focus our analysis on agents that stay in the system for an extended period of time. In addition, we note that [NN17] as well as [NST15] identify the value with smallest regret error relative to the value (smallest error to value ratio). This method favors larger values, that make the relative error smaller. Using the value with smallest absolute error would have made the identification more successful even for bidders with relatively smaller values. This is the method we will use in this chapter.

A distinctive feature of Zillow's field experiment is that the bidders were provided a bid recommendation tool. Such tools are not unique to Zillow and are routine in search advertising on Google and Bing such as in [goo]. [ATCO16] report experiments with adding bid-recommendations at LinkedIn, where they find that the advertisers and the publisher both benefit having recommendations. On those platforms, there is also a set of third-party tools (not provided by platforms) that facilitate bidding. However, on Zillow the bidders were faced with the choice between trusting the recommendation provided by Zillow's tool or learning on their own. Our work thus bridges the gap betwen the literature on empirical analysis of algorithmic learning in games and the literature on recommender systems without trust (e.g. see [RRS11] for a survey of the latter).

3.2 Data Description

In the period between 2014 and 2016[1], Zillow has run a series of large-scale experiments where the mechanism for selling ad impressions was switched from negotiated contracts to auctions. During this period, Zillow defined markets as zip codes. The real estate agents were not allowed to use targeting within the zip code (i.e. by advertising only on the pages of specific real estate listings) or buy "packages" of impressions across multiple zip codes. In fact, a vast majority of real estate agents that we observe in the data only compete for a single zip code.

The experiments were rolled out in a large number of clearly isolated markets with zip codes coming from either separate states or sufficiently far from each other within a state. In order to facilitate this experimental mechanism rollout, Zillow has engaged in a significant marketing and training effort to ensure that real estate agents in the experimental markets understand the structure of the auction and to help agents learn how to bid well in the auction, akin the set of tutorials provided by Google for its advertisers.

Our data comes from 57 experimental markets from Zillow. These markets are close to the entirety of markets that were switched to auction-based prices and allocations. Zillow's experiments were designed not just to evaluate the performance of auctions in selected markets per se, but also to compare key characteristics of monetization and impression sales in the incumbent mechanism with negotiated contracts and the new auction mechanism. To provide

[1]We withhold the exact start and end date of the experiments for confidentiality purposes

data for credible comparison of these variables (which we do not analyze in this chapter) Zillow did not convert entire markets to auctions. Instead, a fixed proportion of impressions was reserved for fixed price contracts and the remaining inventory was released to the auction-based platform. In each market Zillow selected several agents that were brought to the auction platform (and who were not allowed to buy impressions from fixed contracts in the same markets). We dropped a few markets from the data that either did not have reliable data due to possible malfunction of the implementation of the auction mechanism, or the data span was too short to produce reliable results. For example, towards the end of our period of observation, more agents were getting enrolled in the auction markets. For most of those new agents the period of observation is too short for statistically valid inference. As a result, we chose to drop such short-living agents. For data confidentiality purposes all dollar-valued variables, such as prices and budgets, in our data were re-scaled and do not reflect the actual amounts.

Our structural analysis in this chapter is concentrated on the much smaller set of 6 very active markets. Our goal in selecting these markets was to (i) ensure that these markets are sufficiently geographically separated, yet have the typical statistical properties of all markets, such as impression prices, in all characteristics except the activity of agents; (ii) have sufficient number of observations of bid changes for different bidders. To understand the behavior of bidders in these auctions, we need to infer their values. Agents that are not active on the platform, do not provide enough data for us to reasonably estimate their value. As it will become clear in Section 4, the second part is crucial for us to be able to produce reliable evaluation of payoffs and bidding strategies of the agents. To select the 6 markets, we first filter the markets where the number of participating agents is 15 or less, which gets down the number of regions from 57 to 12. As it is typical in display advertising, the impression bids are expressed per mille (1000 impressions). Agents have to enter a monthly budget and a per mille bid (for their maximum willingness to pay for 1000 ad impressions). Most agents changed their monthly budgets extremely infrequently, and changed their bids more regularly. The average frequency of bid changes per day in these regions was 0.43. The 6 markets we use for our structural estimation, are the markets with above average number of bid changes.

Variable	Selected Regions				All Regions			
	Mean	STD	25%	75%	Mean	STD	25%	75%
Number of agents	19.33	2.29	18.0	20.75	10.74	5.32	6.0	15.0
Bids	23.94	14.14	17.3	19.31	18.79	9.71	14.06	23.84
Recommended Bids	29.2	21.84	18.96	20.25	20.89	12.05	15.03	24.75
Budgets (daily)	8.92	3.0	6.31	11.71	9.22	4.96	5.9	12.44
Bid change per month	2.9	0.61	2.56	3.26	1.42	0.99	0.59	2.16
Budget change per month	0.28	0.1	0.23	0.29	0.3	0.24	0.18	0.38
Active duration	85.97	10.38	78.03	91.5	96.04	20.74	86.53	107.33
Reserve price	11.65	7.03	7.99	10.74	13.39	9.55	6.0	16.93
Bid changes	0.73	0.26	0.54	0.85	0.22	0.28	0.03	0.32
Impression Volume	5.52	1.72	4.25	5.89	5.29	3.19	2.73	6.89

Table 3.1: Basic information for all regions and the selected regions. The impression volume's unit is 1000 impressions per day. Bids, recommended bids, reserve prices and budgets are also per 1000 impressions. Active Duration is in days. Bid changes is the average number of agents that change their bid per day in a region. The average of bids, budgets and active duration has been calculated for each agent first and then their averages has been taken over all agents of each region.

In Table 3.1 we display basic statistics from our data. The table contrast the statistics for our selected 6 markets with the statistics of the entire set of 57 markets that we analyzed. Presented statistics correspond to the number of participating bidders, their bids and budgets, period of time when the bidder is active in an auction (i.e. she is bidding above the reserve price, has a positive budget), daily frequency of bid changes and the market reserve prices. The Table indicates that our selected 6 markets have similar values of monetary variables (e.g. average bid of 23.9 in selected markets vs. 18.8 in the entire set of markets and average daily budget of 8.9 in selected markets vs 9.2 in the entire set of markets). However, there are two key statistics that are clearly different in our selected set of markets: the time-average number of participating bidders (19.3 in selected markets vs. 10.7 in the entire set of markets) and the average frequency of bid changes (.7 per day in the selected markets vs. .3 per day in the entire set of markets).

This means that while the per impression values in our selected markets should be similar with those in the entire set of experimental markets, our selected markets have more intense competition and, therefore, we would expect smaller markups of the bidders and faster convergence of bidder learning towards the optimal bids.

Our data also contains the predicted monthly impression volume for the month ahead (from the start date of each bid). This estimated impression volume is an input in Zillow's bid recommendation tool whose goal is to compute the bid that will guarantee that the bidders win impressions uniformly over time, and wins the maximum expected number of impressions for the future month for the given budget. To guarantee uniform service, Zillow implemented budget smoothing explained in the next section.

An important takeaway from Table 3.1 is the magnitude of the relative scale of bids and budgets of bidders across the markets. Recall that the impression bids are expressed per mille (1000 impressions), and note that the Table also reports daily impression volumes in mille. The striking fact is the small scale of budgets relative to the bids, far lower than needed to pay for all available impressions. We note that we computed daily budgets for bidders using the period they were active, which is often only a subset of time. The limited size of the budgets is in contrast with the evidence for large advertisers from sponsored search advertising (on Google and Bing) where budgets declared to the advertising platform are often not binding. This means that the issue of smooth supply of impressions to each agent becomes one of the central issues of the platform design. The platform needs to engage in active management of eligibility of bidders for impression auctions to ensure that each bidder participates in the auctions at uniform rate over time.

Also note that bidders change their budget an order of magnitude less frequently than changing their bid. They appear not to change their budget to gain more impressions. In our analysis, we will treat the bidder's budgets as their true monthly budgets, i.e., assume that budgets are not strategic. This is in contrast to bid values that were clearly strategic.

Due to limitations of the data collection, we do not have the data for eligible user impressions for the entire duration of our auction dataset. For most of the period we only have Zillow's estimate for user impressions, and only have the actual impression volume for three months. For this period we noticed that impression volume fluctuates with the days of the week, as shows on Figure

3.1, while the estimated impression volume doesn't show such fluctuation.

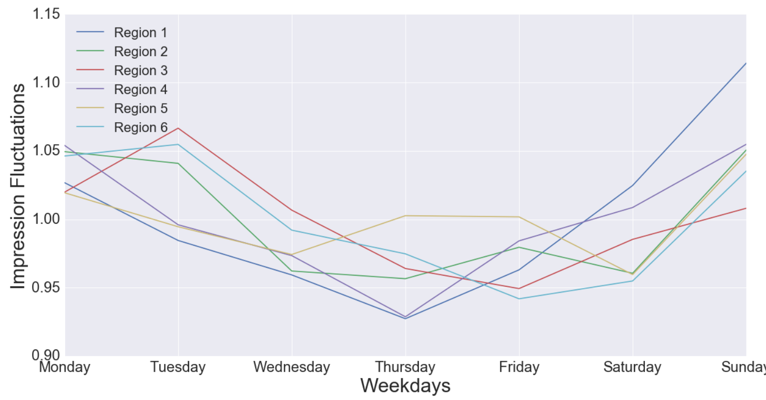

Figure 3.1: Impression volume fluctuations in weekdays shown for 6 regions with the most number bid changes. The impression volume of each region is normalized by the average daily impression volume of that region.

We also do not have data for each bidder on the impressions for which that bidder was *eligible* including those that she wasn't served. To address this data deficiency, we take the predicted volume of eligible impressions (which is the most reliable proxi for the total number of eligible ad impressions). Using the seasonal modulation (mostly reflecting the intra-week changes of the impression volume) that we observe in the detailed impression records, we augment the impression volume predictions to produce a more reliable proxi for daily user impressions. This generates a realistic pattern for daily impressions for the entire time period that we observe.

3.3 The Auction Format

3.3.1 The Mechanism

Zillow's mechanism in the large scale experiments can be characterized as

- real-time: with ads placed in real time as opportunities arise
- weighted: higher bids have a higher chance of being shown,
- agents are paying per impression, unlike the per-click payment used for search ads
- generalized second price auction
- with reserve prices and budget smoothing.

We now discuss in detail these components of the implemented mechanism.

Allocation

The mechanism is designed to allocate impressions to bidders on the basis of their bids. The feature of the mechanism is that all impressions are made *ex ante* equivalent since each ad is placed randomly into one of the three available slots. However, smaller bids will have a smaller change of being shown. Each bidder i submits her bid b_i. The auction platform identifies a subset of eligible bids, based on the budget smoothing discussed below, and selects the four highest eligible bids. The auction platform takes fixed position "weights" γ_j. These weights are used by the platform to induce the dependence of the impression allocation on the rank of each bidder's bid: weights sum to 1, and the agent with jth highest bid is shown with probability $3\gamma_j$ on a page, so 3 of the top 4 agents are shown on each page. The weights used in the system are $0.33, 0.28, 0.22, 0.17$, so only the highest 4 bids have a chance of being shown, with the top bidder shown almost surely, the next one with probability $3 \cdot 0.28$, that is, 84% probability, etc.

Pricing

The mechanism uses the generalized second price. Once the eligibility decision is made, the eligible bidders are ranked according to their bids. Three of the top four bidders are allocated an impression (based on the probability weights γ_i), and placed on the page in a randomly chosen advertising slot. Then the price that the bidder i ranked j pays for that impression is the bid of the bidder ranked $j + 1$. If fewer than 4 bidders are eligible, the per impression price of the last bidder in the ranking is equal to the reserve price. To avoid having to deal with ties, Zillow effectively implemented a priority order via assigning each agent "quality score" very close to 1, to determine the order of agents with identical bids. These quality scores were fixed during the entire period of our observation. In our simulations, we take them into account to determine the ranking of bidders. However, given that these scores simply operated as fixed scaling, we do not include them in the further discussion for simplicity and clarity of our exposition.

Budget smoothing (pacing)

In addition to bids, bidders also submit the advertising budgets B_i. These budgets are typically set for a month, but the bidders also have an option to set their budgets for shorter periods, and some do. Real-estate agents typically have relatively small budgets, so the system needs to implement a form of "budget-smoothing" or pacing to have the agents participate in auctions evenly across the time interval. For each bidder j, the system determines a budget-smoothing probability π_j, that in expectation will ensure that the agents don't run out of budget before the end of the period.

We describe the details of budget smoothing, and analyze the properties of the auction mechanism from the *expected impression* perspective. Although, this gives a simplified view of the system (e.g. avoiding dynamics and the fluctuation of the impression volume), that allows to discuss the incentives in the auction mechanism in the most crystallized way. If N_{pv} is the total number of page views in thousands over the time period of the bidder i and $\bar{B}_i$ is the total budget of the bidder, then the *per thousand impression opportunity budget* can be expressed as $B_i = 3\,\bar{B}_i \,\big/\, N_{pv}$. This reflects the fact that on each page there are 3 ad impression opportunities for 3 available slots and each bidder can only appear in one of the three slots, so if the agent i has bid $b_i \leq B_i$, then she is guaranteed to not run out of budget, even if participating in all opportunities. However, no bidders had budgets this high in our data.

Budget smoothing process can be characterized by a probability π_i that determines eligibility of bidder i to participate in the auction. *Conditional on not* *being budget smoothed*, bidder i participates in a generalized second price auction. Since all bidders may be budget smoothed, the group of auction participants becomes random. As a result, the auction outcomes for each bidder i are random. Let I denote the set of bidders with bids above the reserve price. The participation of each bidder i in a given auction can be represented by a binary variable where 1 indicates that the bidder is made eligible for the auction, and 0 means that the bidder was made ineligible due to budget smoothing. Then the set of all possible participating bidders can be represented by 2^I binary I-digit numbers $N = n_1^N\, n_2^N\, \ldots\, n_I^N$, where $n_i^N \in \{0,\, 1\}$. Denote this set $\mathcal{N}$. Also denote $\text{rank}_i(N)$ the rank of bidder i in the subset of bidders represented by N. Then the expected cost and the expected impression share for bidder i conditional on her not being budget-smoothed are computed as

$$\text{eCPM}_i(b_i) = \sum_{N \in \mathcal{N}, n_i^N = 1} 3\, \gamma_{\text{rank}_i(N)} \prod_{j \neq i} \pi_j^{n_j^N} \, (1 - \pi_j)^{n_j^N} \; \text{PRICE}^N_{\text{rank}_i(N)} \tag{3.3.1}$$

where $\gamma_j = 0$ for $j > 4$, and we use PRICE_j^N to denote price paid by bidder ranked jth, which is either the $(j+1)$th bid in N, or the reserve price, if there only j bids in N. Further,

$$eQ_i(b_i) = \sum_{N \in \mathcal{N}, n_i^N = 1} 3\, \gamma_{\text{rank}_i(N)} \prod_{j \neq i} \pi_j^{n_j^N} \, (1 - \pi_j)^{n_j^N}$$

Here we use notation $\text{eCPM}_i(\cdot)$ to reflect that impression costs are typically expressed per thousand impressions (a.k.a *per mille*).

The expected cost and expected price per impression are also functions of bids and budget smoothing probabilities of competing bidders, a dependence that we will make explicit in notation when useful. We note that both these objects are defined via the conditional expectation, i.e. they determine the spend and the impression probability conditional on bidder i not being budget smoothed. The impression eligibility probability is determined by the balanced budget condition:

$$\pi_i = \min\left\{1,\ \frac{B_i}{\mathrm{eCPM}_i(b_i)}\right\}. \tag{3.3.2}$$

Note that if the expected per impression spent does not exceed the per impression budget, then such a bidder should not ever be budget smoothed.

Budget smoothing is one of the most technically challenging components of the implemented experimental mechanism. For large advertisers on platforms like Google, budgets typically play a minor role essentially working as "insurance" from surges in spent generated by idiosyncratic events. In contrast, advertisers on Zillow are real-estate agents, and typically have small monthly budgets relative to per impression cost. In particular, our data shows that virtually all bidders are budget smoothed over certain periods.

In these settings a carefully constructed system for budget smoothing is essential. Take the vector of current eligible bids and budgets for bidders $i = 1, \ldots, I$. The idea for recovering the filtering probabilities will be to solve for a set of probabilities $\pi_1, \ldots, \pi_I$ such that (3.3.2) is satisfied for each i. The main ingredient in computing π_i is the expected cost per opportunity $\mathrm{eCPM}_i(b_i)$ with expectation taken with respect to the distribution of $\pi_1, \ldots, \pi_I$ of other bidders, and $eQ_i(b_i)$ is the probability of being shown conditional on not being filtered.

3.3.2 Our Simulations

We calculate the outcomes of the empirical markets by simulating Zillow's budget smoothing algorithm. We start by stating the main theorem and intuitive description of the algorithm.

Overview

We compute the total payments and the total impressions share of all agents inside a market by simulating the outcome for each day then summing the daily payments and impression shares of each agent. For each day, we initialize the budget smoothing probabilities with 1 and adjust them until none of the agents overspend their budget or have leftover money while they still can get more impressions by increasing their budget smoothing probabilities. In order to make these adjustments, at each iteration, we find the expected spent of each agent ($\text{eCPM}_i(b_i, \pi_i, \ldots, \pi_I)$) given the budget smoothing probabilities from the previous iteration. Let π_i^t be the budget smoothing probability of agent i at time t. The budget smoothing probability adjustments are as follows

$$\pi_i^{t+1} = \min\left\{1,\ \frac{B_i}{\text{eCPM}_i(b_i; \pi_1^t, \ldots, \pi_I^t)}\right\},\quad i = 1, \ldots, I.$$

The main difficulty in each iteration is calculating the expected payment of each agent. Note that directly using equation 3.3.1, the running time of the algorithm in each iteration would be $O(I2^I)$. We significantly reduce this running time by using dynamic programming and the fact that budget smoothing probabilities are independent random variables. The following theorem summarizes the main contribution of our algorithm for calculating the expected payments and expected impressions share given budget smoothing probabilities.

Theorem 25. *Given budget smoothing probabilities of agents (π), the expected payments (eCPM) and expected impressions shares (eQ) of all agents can be computed in $O(I)$ time.*

Description

To characterize the empirical market outcomes we numerically simulate the budget smoothing algorithm. Zillow's budget smoothing algorithm has the following features:

1. Sort the bidders i by their bid b_i and assume bidders are numbered in this order.
2. Construct an array of 2^I binary I-digit numbers from $\{0, 0, \ldots, 0\}$ to $\{1, 1, \ldots, 1\}$ where the number $N = n_1^N n_2^N \ldots n_i^N \ldots n_I^N$ will correspond to bidders j with $n_j^N = 1$ not being filtered out. Call the set of elements in this array $\mathcal{N}$
3. Take a subset of elements of $\mathcal{N}$ where i-th digit is equal to 1. Call this set $\mathcal{N}_i$
4. Let $N = n_1^N n_2^N \ldots n_i^N \ldots n_I^N$ with $n_i^N = 1$ and $n_j^N \in \{0, 1\}$ be a specific row in $\mathcal{N}_i$, corresponding to the outcome of filtering when the agents with $n_j^N = 1$ remained.

5. Include the bid of each bidder j for whom $n_j^N = 1$, and determine the price of bidder i, calling it PRICE_i^N which is the maximum of the reserve price, and the bid b_j first agent $j > i$ with $n_j^N = 1$. Let j_i^N be the position of agent i after filtering, and let γ_j^N the corresponding probability $\gamma_{j_i^N}$.
6. Compute the expected spent as

$$\text{eCPM}_i(b_i; \pi_1, \ldots, \pi_I) = \sum_{N \in \mathcal{N}_i} \gamma_i^N \prod_{j \neq i} \pi_j^{n_j^N} (1 - \pi_j)^{(1 - n_j^N)} \, \text{PRICE}_i^N.$$

7. Solve for $\pi_1, \ldots, \pi_I$ by solving a system of nonlinear equations

$$\pi_i = \min\left\{1, \; \frac{B_i}{\text{eCPM}_i(b_i; \pi_1, \ldots, \pi_I)}\right\}, \quad i = 1, \ldots, I.$$

For instance, we can find an approximate solution by minimizing the sum of squares using gradient decent or Newton's method.

$$\sum_{i=1}^{I} \left(\pi_i - \min\left\{1, \; \frac{B_i}{\text{eCPM}_i(b_i; \pi_1, \ldots, \pi_I)}\right\}\right)^2$$

with respect to $\pi_1, \ldots, \pi_I$.

The main part of the above iterative algorithm for finding a fixed point is to compute the expected cost (eCPM) and expected impressions share (eQ) for all agents for a given set of probabilities $\pi_1, \ldots, \pi_I$. Considering all subsets of agents, this can take exponential time in number of agents. We need to run the simulations for every day and each region separately and compute these many times to get the filtering probabilities. So these calculations can significantly increase the running time of our simulations. Furthermore, for auctions where the number of agents is big, using an exponential time algorithm to calculate the outcome of each iteration is not feasible. In order to get around this issue, we use the fact that in each underlying GSP, only impressions of the first four agents are eligible to be shown, as well as the fact that the filtering probabilities of agents are independent. Our algorithm to find expected cost (eCPM) and expected impressions share (eQ) for given filtering probabilities runs in linear time (in number of agents).

For each agent $i \in [I]$, our algorithm first computes the expected impression share of the agent eQ_i (assuming she has not been filtered). Note that eQ_i only depends on the number of unfiltered agents (r) that have a higher bid than i and the probability (γ_{r+1}) associated with i's rank. We first find the probability that there are exactly $r \in [0, 3]$ number of agents with higher bid than i, called $p_{i,r}(\pi)$. By multiplying $p_{i,r}(\pi)$ by the impressions of the $(r+1)$-th position (γ_{r+1}) we can find the expected impression share of agent i.

After finding the expected impression share of i, we calculate $eCPM_i$ by using eQ_i. In order to do this, we use the fact that the filtering probabilities of agents are independent. Furthermore, the expected cost per impression for agent i is only a function of bids of agents who are bidding lower than i and their filtering probabilities. So by calculating the expected cost per impression and multiplying it by the expected impression share of agent i, we can calculate her expected payment conditioned on i being in the auction ($eCPM_i$). Recall that for each agent i, eQ_i and $eCPM_i$ are calculated conditioned i not getting filtered, so in order to calculate the total expected number impressions that she wins in the auction and her expected spent, it is enough to multiply eQ_i and $eCPM_i$ by π_i respectively. In algorithm 1 we have marked these steps for each agents.

Input: b: bids (sorted) , π: filtering probabilities, γ: rewards, $reserve$: the reserve price
Output: $eCPM, eQ$
Let $\{1, 2, \ldots, I\}$ be the list of all agents such that $b_1 \geq b_2 \geq \ldots \geq b_I \geq reserve$
Let $p_{1,0}(\pi) = 1$ and $p_{1,r} = 0$ for $0 < r \leq 3$
for $i \in [I]$ **do**
    if $i > 1$ **then**
        Let $p_{i,0}(\pi) = (1 - \pi_{i-1})p_{i-1,0}(\pi)$
        for $r \in [1, 3]$ **do**
            Let $p_{i,r}(\pi) = (1 - \pi_{i-1})p_{i-1,r}(\pi) + \pi_{i-1}p_{i-1,r-1}(\pi)$
        end
    end
    Calculating $p_{i,r}(\pi)$ from $p_{i-1,r}(\pi)$
    Let $eQ_i = 0$
    for $r \in [0, 3]$ **do**
        $eQ_i = eQ_i + \gamma_{r+1}.p_{i,r}(\pi)$
    end
    Calculating expected impression share $eQ_i$from $p_{i,r}(\pi)$
    Let $j = i + 1$
    Let $CPM_i = 0$
    if $i = 1$ or $\pi_i = 1$ **then**
        while $j \in [I]$ and $\pi_j < 1$ **do**
            if $j = i + 1$ **then**
                $q_{i,j}(\pi) = \pi_j$
            else
                $q_{i,j}(\pi) = \frac{\pi_j(1-\pi_{j-1})}{\pi_{j-1}} q_{i,j-1}(\pi)$
            end
            $CPM_i = CPM_i + b_j.q_{i,j}(\pi)$
            $j = j + 1$
        end
    else
        $CPM_i = \frac{CPM_{i-1} - \pi_i b_i}{1-\pi_i}$
    end
    if $j = I + 1$ **then**
        if $i = I$ **then**
            $q_i(\pi) = 1$
        else
            $q_i(\pi) = \frac{1-\pi_{j-1}}{\pi_{j-1}} q_{i,j-1}(\pi)$
        end
        $CPM_i = CPM_i + reserve \cdot q_i(\pi)$
    end
    Calculating cost per impression
    $eCPM_i = eQ_i \cdot CPM_i$
end

return $eCPM, eQ$

Algorithm 1: Calculating eCPM and eQ of agents in linear time.

If this algorithm is implemented naively, the most expensive computation is computing for all the agents $p_{i,r}(\pi)$, the probability that there are exactly r agents unfiltered with bid above b_i. For each agent, it takes $O(I^3)$ to find all the configurations where there are $r \in [0, 3]$ agents who are not filtered and have higher bid than i. For each configuration, it takes $O(I)$ to compute this probability. By using dynamic programming, we can compute $p_{i,r}(\pi)$ from $p_{i-1,r}(\pi)$ by considering the cases where agent i is getting filtered and is not getting filtered separately. For initialization, we set $p_{1,0}(\pi) = 1$ and $p_{1,r}(\pi) = 0$ for $0 < r \leq 3$. We use the following update rule to computer $p_{i,r}(\pi)$ for $i > 0$:

$$p_{i,r}(\pi) = \begin{cases} (1 - \pi_{i-1})p_{i-1,r}(\pi) & r = 0 \\ (1 - \pi_{i-1})p_{i-1,r}(\pi) + \pi_{i-1}p_{i-1,r-1}(\pi) & 0 < r \leq 3 \end{cases}$$

This reduces the running time of computing $p_{i,r}$ for each agent from $O(I^4)$ to $O(1)$. The calculations for finding CPM and $eCPM$ can be done in a similar way: instead of computing eCPM for each agent from scratch, we can compute the expected price per impression from the previous calculations. When $i = 1$ (she is the highest bidder), or $\pi_i = 1$ (she is never getting filtered), first, we set the expected cost per impression to 0 and find the smallest $j > i$ such that $\pi_j = 0$. Then, for each $k \in (i, j]$, we find the probability that all the agents $z \in (i, j)$ are getting filtered and k is not getting filtered ($q_{i,k}(\pi)$), multiply it by bid of k (b_k) and add the result number to the expected cost per impression. If for all $j > i$, $\pi_j < 1$ then we set $j = I + 1$ and we assume that $I + 1$ is an agent who is never getting filtered and has a bid equal to the reserve price. Note that we also compute $q_{i,k}(\pi)$ from $q_{i,k-1}(\pi)$ in $O(1)$, instead of computing it for each k from scratch.

When $i > 1$ and $\pi_i < 1$, we use the previous expected cost per impression that we had from the previous agent (lets call it CPM_{i-1}) to calculate the cost per impression of agent i (CPM_i) by doing the following calculation

$$CPM_i = \frac{CPM_{i-1} - \pi_i b_i}{1 - \pi_i}$$

This operation nullifies the effect of agent i in the cost per impression of the previous agent ($i-1$) and calculates the new expected cost per impression. Finally, we set $eCPM_i = eQ_i \cdot CPM_i$. Note that even though the running time of this algorithm may be $O(I)$ for some agents, the total (amortized) running time of these calculations for all the agents combined is still $O(I)$. So overall the algorithm requires amortized $O(1)$ number of calculations for each agent and it takes linear time ($O(I)$) to calculate eCPM and eQ for all the agents, given the sorted list of agents based on their bids. Since we need to sort the agents by their bids at the beginning, the total running time of each iteration in computing the filtering probabilities is $O(I\log(I))$. This improvement in the running time (from $O(I2^I)$ to $O(I)$) is crucial for simulating the outcome of the auction, specially in the auctions where the number of agents is big.

3.3.3 Strategic Actions and Preferences of Bidders

Agent i is able to choose her per impression bids b_i. We will assume that the bid space is the bounded segment $\mathcal{B}$. While agents were also allowed to choose and their budgets B_i (which we normalize to express it per impression), those budgets were typically fixed. Our summary Table 3.1 demonstrates that budgets were changed on average 10 times less frequently than bids. Based on this observation, we treat budgets as fixed parameters of the bidders rather than strategic variables.

We assume that utilities of bidders are quasilinear as long as they observe the budget constraint, that is, if bidder i is allocated an impression and she is charged the price p_i for it, then she receives the surplus of $v_i - p_i$. Then the expected utility conditional on not being budget smoothed for bidder i is

$$u_i(v_i, B_i, b_i) = eQ_i(b_i)\, v_i - \text{eCPM}_i(b_i). \tag{3.3.3}$$

Classical analysis of such an economic system would assume that the outcome is a Nash equilibrium of the game, where each bidder maximizes her utility by setting the bid. We will skip here the details of the resulting equilibrium analysis. We note that identifying the right bid can be challenging for the bidders, who are real-estate agents, and often don't have the data or the analytic tools to do a good job optimizing their bid. To help the advertisers, the platform provided a bid recommendation, suggesting the bid that maximizes the expected number of impression the agent can achieve based on her budget.

3.3.4 Bid recommendation tool

Overview

To help the advertisers, the platform also provides a bid-recommendation. The agents in this market are real-estate agents, who often don't have the data or the analytic tools to do a good job optimizing their bid, and for whom Zillow may not be the main channel through which they get the "client leads". As a result, some agents may be reluctant to engage in active exploration of optimal bidding. In order to facilitate the bidding for those agents, the platform has developed a tool that recommends a bid for a given bidder based on this bidder's monthly budget. The tool was designed to set the bid that maximizes the expected number of impressions that the bidder gets given her budget.

For low bids an increase in the agent's bid leads to the increased number of delivered impressions, as when the agent's bid is sufficiently low she spends less than her budget, and thus she is not budget smoothed. Her utility is then characterized by equation (3.3.3), which increases by increasing $eQ_i(\cdot)$, as long as the marginal increase in cost is not exceeding v_i. Once the expected spent starts exceeding the budget, the pacing mechanism starts operating to ensure that balanced budget condition (3.3.2) holds. In this case, a further increase in the bid does not change the expected spent (which is equal to the budget). However, it increases the expected per impression rank, given that bidders with higher bids are ranked higher. In combination with budget smoothing, the bidder pays more per impression while getting fewer impressions in expectation. The bid recommendation tool suggests the smallest bid b_i^* where the agent spends her budget, that is

$$b_i^* = \arg\min_{b_i} \text{eCPM}(b_i) \geq B_i.$$

Bidding above this value will only decrease the number of impressions gained without decreasing the cost of the agent.

Note that with this bid, if the $eCPM(\cdot)$ function were continuous, then the agent would not be budget smoothed, i.e. $\pi_i = 1$. In auction with a small set of opponents, the $eCPM(\cdot)$ function is not continuous, so this is only approximately true. This bid also corresponds to the best response of the budget-constrained bidder whose value per impression exceeds marginal per impression price that she can get while spending her budget.

Description

Here we describe the details of the bid-recommendation tool that was implemented in Zillow.

For each actual realization of the group of competing bidders, we define the cost function $CPM_j(b_j)$ as a mapping from the bid of bidder j to the price she pays per impression in an auction. Recall that we defined $Q_j(b_j)$ as a probability of being allocated an impression as an outcome of an auction conditioned on that the agent wasn't filtered. Note that without the effect of filtering both functions are step functions: whenever bidder j outbids bidder ranked i but ranks below bidder ranked $i-1$, then this bidder j pays the price determined by i b_j between b_i and b_{i-1}. In figures 3.2,3.3 we illustrate the concept of price and the impression probability using the bidders in one of our 6 markets and take the bidder ranked 4 in the first week in the market and bid of 30 (recall that the price units were scaled not to reflect the actual market prices). The figure shows the price and the impression share that this bidder gets if all other bidders are made eligible for this impression.

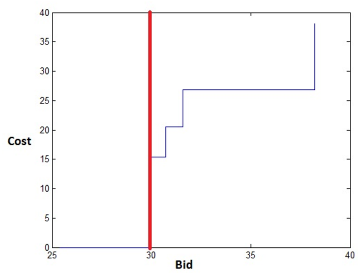

Figure 3.2: Spent (agent's bid in red)

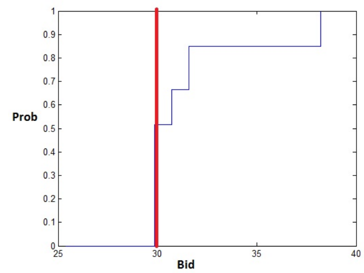

Figure 3.3: Impression share (agent's bid in red)

We note that the bidder under consideration has a per impression budget significantly below the cost of an impression (as much as a factor of 10 below). If the eligibility status for impressions was recorded by the system, then we could compute the empirical fraction of impressions where this given bidder was made eligible for an auction: divide the number of impressions where a given bidder was made eligible for an auction by the total number of arrived impressions. Depending on the impression volume, this can be computed using all impressions from the beginning of the month or some smaller window of time (e.g. the week before). This would be our estimated probability π_j of actually getting displayed for an impression.

Expected spent and expected impression allocations If bidders are affected by the budget smoothing then the participation of bidders in an auction is random (where randomness is activated by the budget smoothing mechanism). Then for each bidder, instead of the actual spent and the impression share, we will have the expected spent and the expected impression share, where the expectation is taken with respect to the randomness of participation of competing bidders. Given the filtering probabilities, the expected spent and expected impression share are represented by the expectation of spent and impression shares in all possible bidder configurations which are then weighted by the probabilities of those configurations. The participation of each bidder i in a given auction can be represented by a binary variable where 1 indicates that the bidder is made eligible for the auction and 0 means that the bidder was made ineligible due to budget smoothing. Then the set of all possible participating bids can be represented by an array of 2^I binary I-digit numbers from $\{0, 0, \ldots, 0\}$ to $\{1, 1, \ldots, 1\}$. Call the set of elements in this array $\mathcal{N}$. Then a subset of elements of $\mathcal{N}$ where i-th digit is equal to 1 corresponds to the configurations where bidder i is made eligible for an auction. Call this set $\mathcal{N}_i$. Then the expected cost and the expected impression share are computed as

$$eCPM_i(b_i) = \sum_{N \in \mathcal{N}_i} \gamma_i^N \prod_{j \neq i} \pi_j^{n_j^N} (1 - \pi_j)^{n_j^N} \, \mathrm{PRICE}_i^N(b_i), \tag{3.3.4}$$

and

$$eQ_i(b_i) = \sum_{N \in \mathcal{N}_i} \prod_{j \neq i} \pi_j^{n_j^N} (1 - \pi_j)^{n_j^N} \, \gamma_i^N,$$

where N corresponds to the index of the set of eligible of bidders.

Computation of filtering probabilities If the impression allocations are not available, then the probabilities of participation of bidders in impressions have to be computed. We can consider the actual budget smoothing as an iterative process: we continuously evaluate the actual spent for each bidder and when the spent exceeds the allocated budget, then the bidder is made ineligible for some impressions. This iterative process reaches the steady state when the expected spent in an impression for a given bidder becomes equal to the budget:

$$\pi_i \times eCPM_i(b_i) = \mathrm{Budget}_i.$$

We can simulate this iterative process for the bidder in one of our selected markets. Note that that bidder has a very low per impression budget of 3.87. We start the process assuming that all the bidders are eligible for an impression. Then using the spent in that impression, we compute the filtering probabilities for all by dividing the budget by the spent and then iterate the process to set

$$\pi_i = \frac{\text{Budget}_i}{eCPM_i(b_i)}$$

using the previous iteration values of the eligibility probabilities. The algorithm for computing the probabilities of being displayed on the page is the following.

Iteration 0: Initialize probabilities of being eligible for an impression at $\pi_i^{(0)} = 1$.

Iteration k: Take the probabilities of being eligible for an impression $\pi_i^{(k-1)}$ computed from the previous iteration. Compute eCPM from (3.3.1) for each bidder $i = 1, \ldots, I$. If $eCPM_i(b_i) = 0$, then set the probability $\pi_i = 1$ (bidder is always displayed, this bidder never gets any impressions as an outcome of the auction). If $eCPM_i(b_i) > 0$, then set

$$\pi_i^{(k)} = \min\left\{1, \frac{\text{Budget}_i}{eCPM_i(b_i)}\right\}.$$

Stopping criterion: Stop when the probabilities become close across the iterations: $\max_i |\pi_i^{(k)} - \pi_i^{(k-1)}| < \epsilon$, for a given tolerance criterion.

We illustrate the trajectory across the iterations the bidder of interest on the following figure.

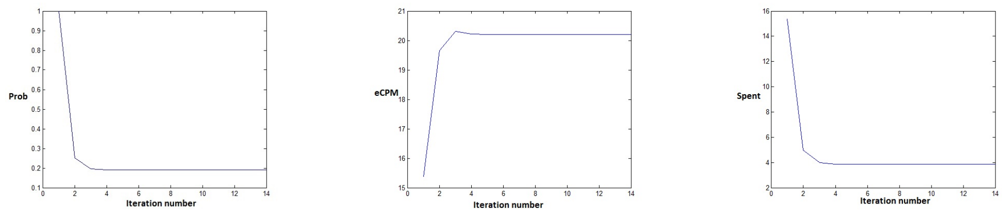

Figure 3.4: Iteration path for probability of eligibility, eCPM and expected spent

At the end of the iterative process the expected spent approaches the budget due to the increase in the filtering probability. Note that the expected CPM significantly increases in response to the change in the filtering probabilities for all other bidders.

The randomness increases the eCPM and the probability of allocation into an impression as compared to the fully deterministic case (i.e. when bidders are not randomly removed from impressions due to budget smoothing). The figure below demonstrates the expected CPM and expected fraction of impressions (after budget smoothing) for bidder with bid of 30 and per impression budget of 3.87.

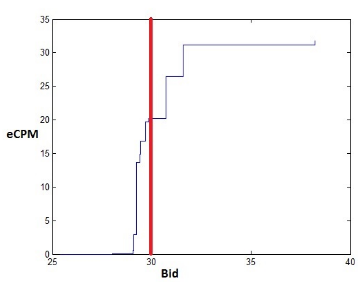

Figure 3.5: Spent (actual bid in red)

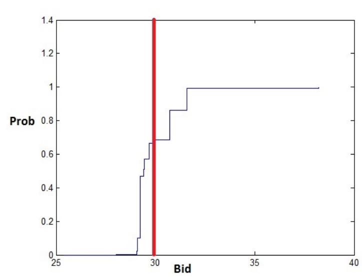

Figure 3.6: Impression fraction (actual bid in red)

Note that an increase in the auction outcomes (probability of being allocated an impression and the eCPM is compensated with a decrease in the probability of being eligible for an auction).

Computation of the optimal bid for impression ROI optimizers We note that the eCPM and allocation probabilities are monotone functions of the bid. As a result, if a given bidder maximizes the probability of appearing in an impression as a function of the bid, the optimal bid will be set such that (a) the expected spent does not exceed the per impression budget; (b) an increase in the bid will result in an increase in the spent exceeding the budget. Note that the spent is (non-strictly) monotone increasing until it reaches the level of the per impression budget and then it stays at the level equal to the budget due to budget smoothing. Assuming that the bidders do not have the "values of residual budget", this means that the bidder whose budget per impression exceeds any other budgets should set the bid at the level equal to the budget. Note that the deviation from this strategy will not be optimal: a decrease in the bid for such a bidder results in "budget savings" that have no value for this bidder, but at the same time it will result in a (weak) decrease in the number of impressions.

The tool that computes the optimal best response for such a bidder proceeds in the following way.

Construction of the grid of bids We construct the grid of bids of opponent bidders. These are the points where the spent function exhibits jumps.

Construction of the eCPM curve We construct the eCPM curve. By choosing a small ϵ (smaller than the minimum distance between the closest score-weighted bids), we evaluate the changes in the eCPM after a given bidder outbids and under-bids the opponent by ϵ.

Computation of the optimal bid Set the bid to the level where the eCPM curve intersects the horizontal line corresponding to the budget.

Adjustment of the bid for top/bottom bidders The top bidder sets the bid at the level equal to the per impression budget, the bottom bidder sets the bid to the maximum level that makes the spent positive.

Note that if there are I bidders, this approach amounts to $2 \times I$ evaluations of the eCPM function. The picture below demonstrates the shift to the optimal bid for the bidder under consideration by equating this bidder's eCPM with the per impression budget.

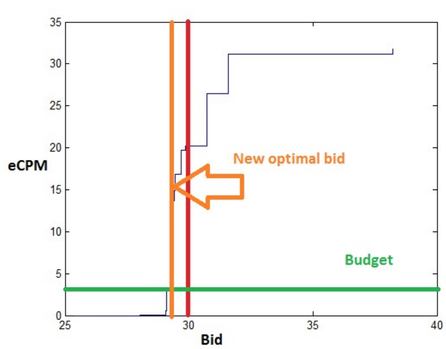

Figure 3.7: Optimization of impression ROI

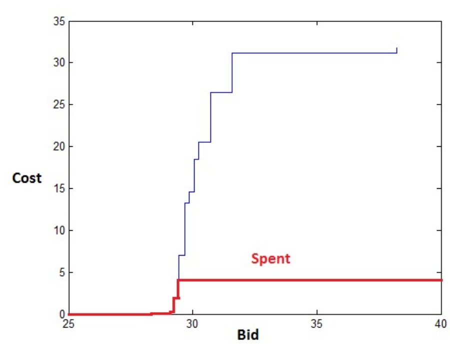

Figure 3.8: The impact of budget smoothing on expected spent

We note what happens to the actual spent of the bidder whenever the bid exceeds the recommended level. Given that at the recommended level the bidder's spent is at or below the per impression budget, if the bid increases then the budget smoothing gets initiated. The overall spent per impression is equal to the product of the probability of being eligible for an auction (π_i) and the expected outcome of an auction ($eCPM_i(b_i)$)

$$\pi_i \, eCPM_i(b_i).$$

Thus, whenever the budget smoothing is initiated ($\pi_i < 1$) then the spent is exactly equal to the budget. Thus the spent as a function of the bid will become flat once the optimal bid has been exceeded.

The probability of being allocated an impression is equal to the product of the probability of being eligible for an auction (π_i) and the probability of being allocated an impression as an outcome of an auction ($eQ_i(b_i)$). We note that since the GSP is monotone, the probability of being allocated an impression as an outcome of an auction increases in the bid: the higher the bid, the higher the probability of being displayed. When the budget smoothing is not initiated, then $\pi_i = 1$ and the probability of appearing on the page is simply $eQ_i(b_i)$ (increasing in the bid). When the bid exceeds the optimal level, then the probability of being eligible for an impression is $\pi_i = \text{Budget}_i / eCPM_i(b_i)$ leading to the probability of being allocated an impression of

$$\text{Budget}_i \times \frac{eQ_i(b_i)}{eCPM_i(b_i)}.$$

This function decreases as a function of bid. This means that if the per impression budget does not warrant a given bidder the top position without filtering, then the probability of getting an impression increases up to the optimal bid level and then decreases whenever the bid starts exceeding the optimal level.

Budget and bid recommendations based on the impression targets We can use the "expected impression" model to make the recommendations for the choice of the monthly budget and the corresponding bid that meet a given impression target. Note that due to the budget smoothing, the expected spent in a given impression

$$\text{Spent}_i(b_i) = \pi_i \, eCPM_i(b_i) \leq \text{Budget}_i.$$

The inequality may not be binding due to the possible jumps in the $eCPM$ curve. We note that the expected probability of appearing in the impression is

$$\text{Prob}_i \left(b_i, \, \text{Budget}_i\right) = \pi_i \, eQ_i(b_i).$$

Consider this probability as a function of the bid and the budget, taking into account our model of filtering due to budget smoothing.

$$\text{Prob}_i \left(b_i, \, \text{Budget}_i\right) = \begin{cases} eQ_i(b_i), & \text{if } eCPM_i(b_i) \leq \text{Budget}_i, \\ \text{Budget}_i \frac{eQ_i(b_i)}{eCPM_i(b_i)}, & \text{if } eCPM_i(b_i) > \text{Budget}_i. \end{cases}$$

The expected impression count is obtained by multiplying the probability of appearing on the page $\text{Prob}_i \left(b_i, \, \text{Budget}_i\right)$ by the total projected impression inventory. Note that function $\text{Prob}_i \left(b_i, \, \text{Budget}_i\right)$ is increasing in the bid b_i up to the bid $b_i^*(\text{Budget}_i)$ such that

$$eCPM_i \left(b_i^*(\text{Budget}_i)\right) = \text{Budget}_i$$

and decreases when the bid is greater than $b_i^*(\text{Budget}_i)$. Therefore, the expected number of impressions is maximized for a given budget at $b_i = b_i^*(\text{Budget}_i)$.

Let Inventory be the total impression inventory and Goal_i be the impression target for bidder i. Then the optimum bid for a given budget is set as

$$\text{Prob}_i \left(b_i, \, \text{Budget}_i\right) \leq \frac{\text{Goal}_i}{\text{Inventory}}.$$

The minimum budget per impression for which the impression goal is met is

$$\text{Budget}_i = eCPM_i(b_i),$$

leading to

$$eQ_i(b_i) = \frac{\text{Goal}_i}{\text{Inventory}}.$$

To determine the recommendations of the bid and the budget based on the expressions above, we formulate the following problem: find the profile of the probabilities of eligibility for an auction $\pi_1, \pi_2, \ldots, \pi_I$ and the optimal bid of bidder i b_i^* such that

1. $\pi_i = 1$ for bidder of interest i;
2. $eQ_i(b_i^*) = \frac{\text{Goal}_i}{\text{Inventory}}$.

Note that this is equivalent to solving a system of equations

$$\begin{aligned}
&\pi_j = \min\left\{\frac{\text{Budget}_j}{eCPM_j(b_j)}, 1\right\}, \; j \neq i, \\
&\pi_i = 1, \\
&eQ_i(b_i^*) = \frac{\text{Goal}_i}{\text{Inventory}}.
\end{aligned} \tag{3.3.5}$$

with unknowns π_j, $j \neq i$ and b_i^*. The recommended bid is the solution b_i^* and the budget recommendation is given by multiplying the per impression budget

$$\text{Budget}_i^* = eCPM_i(b_i^*)$$

by the projected impression inventory.

We also need to account for possible corner solutions that do not allow the equality $eQ_i(b_i^*) = \frac{\text{Goal}_i}{\text{Inventory}}$ to be satisfied. First, suppose that for the grid of bids $\{b_k^g = \frac{s_k\, b_k}{s_i}\}_{k=1}^I$ we observe

$$\max_k \, eQ_i(b_k^g) < \frac{\text{Goal}_i}{\text{Inventory}}.$$

Then the optimal bid $b_i^* = \max_k \, b_k^g$. In that case we set the budget $\text{Budget}_i^* = b_i^* > eCPM_i(b_i^*)$. The rationale for this is that the top bidder does not have an incentive to set the bid below the budget as that would lead to a weak decrease in the expected number of impressions while not decreasing the expected spent.

Second, suppose that

$$\min_k \, eQ_i(b_k^g) > \frac{\text{Goal}_i}{\text{Inventory}}.$$

In that case for any bid level the bidder will be subject to budget smoothing. Thus the optimal bid will correspond to

$$b_i^* = \min_k \, b_k^g.$$

The recommended budget will correspond to

$$\text{Budget}_i^* = \frac{\text{Goal}_i}{\text{Inventory}} \frac{eCPM_i(b_i^*)}{eQ_i(b_i^*)}.$$

Simultaneous optimization for multiple bidders The automated optimization for multiple bidders is based on a simple generalization of the single bidder problem. We note that for all bidders whose bids and budgets are optimized to meet the impression goals, we need to solve a system of equations equivalent to (3.3.5) to find bids b_j^*. Let $\mathcal{J}$ be the subset of bidders who use the bid and budget recommendation. Then we find the set of recommended bids $\{b_j^*,\, j \in \mathcal{J}\}$ by solving the system of equations

$$
\begin{aligned}
&\pi_k = \min\left\{\frac{eCPM_k(b_k)}{\text{Budget}_k},\, 1\right\},\; k \notin \mathcal{J},\\
&\pi_k = 1,\; k \in \mathcal{J} \\
&eQ_k(b_k^*) = \frac{\text{Goal}_k}{\text{Inventory}},\; k \in \mathcal{J}.
\end{aligned} \tag{3.3.6}
$$

We also take into account the "corner solutions" corresponding to very low and very high impression goals relative to the available inventory.

Formal integrity tests for tool performance The structure of the rank-based auction leads to the set of properties that have to be satisfied by the optimal solutions for bids and budgets. We can use these properties to construct the tests for the performance of the recommendation tool.

1. For any $\tau > 0$, if $b_k,\ k = 1, \ldots, I$ is the solution of (3.3.6), then if the filtering probabilities are fixed, then the replacement of b_k with $\tau\, b_k$ does not change the predicted impression counts $eQ_k(\tau\, b_k) \times$ Inventory.
2. For any $\tau > 0$, if $b_k,\ k = 1, \ldots, I$ is the solution of (3.3.6), then if the filtering probabilities are fixed, then the replacement of b_k with $\tau\, b_k$ leads to the proportional increase in the predicted total spent $eCPM_k(\tau\, b_k) \times$ Inventory $=$ $\tau\, eCPM_k(b_k) \times$ Inventory.
3. For any $\tau > 0$, if the inventory changes to τ Inventory and all impression goals change to $\tau\,$Goal$_i$, then the optimal bids and per impression recommended budgets remain the same.
4. The ratio $\frac{eQ_i(b_i)}{eCPM_i(b_i)}$ is a (weakly) monotone decreasing function of the bid. In other words, for grid points $\{b_k^g\}_{k=1}^I$, $b_m^g > b_n^g$ should lead to $\frac{eQ_i(b_m^g)}{eCPM_i(b_m^g)} <$

$\frac{eQ_i(b_n^g)}{eCPM_i(b_n^g)}$.

3.4 Inference using No-Regret Learning

To evaluate the agent's ability to adjust their bids in the Zillow's auction markets, we need to infer the player's value for an impression in the auction. We assume each bidder i has value per impression v_i which is her fixed characteristic, and that the bidder is aware of her value. One can think of the value as an expected value, based on the likelihood that the ad-impression will lead to a new real-estate client. The assumption that agents know their value is justified, as Zillow has been selling impressions to real estate agents in these areas for 10 years before the auction market was launched. Therefore, the agents should have a good idea of how user impressions convert to potential client leads.

Based on the above assumption, we view the agents experimenting with different bids, not as a reflection that their value is changing, but rather as a sign that they are experimenting with the strategic aspect of bids in the new auction system. In a dynamically changing market bids will vary over time and agents would not necessarily maximize utility at each instant. For understanding the first (few) months of experimenting with auctions, we believe that it is best to model them as off-equilibrium. A characteristic feature of this market that we analyze in this chapter is the relatively small stakes of a single action (measured in terms of per impression prices relative to the budgets). In such markets, exploration is a good way to learn the best response. We use the methodology developed in [NST15] to infer the values for the agents as they are learning to bid in the new auction system. For agents who change their bid relatively frequently, we model their behavior as *no-regret learning* which then allows us to infer their value using the notion of the *rationalizable set* from [NST15].

Recall that we assumed that agents have quasi-linear utility: utility is expressed as (3.3.3), assuming the agent is not subject to budget smoothing. To simplify notation in this section, for an agent i let $\mathrm{exCPM_{it}}(b_i)$ and $\mathrm{exQ_{it}}(b_i)$ denote the expected cost and expected impression share for bidder i at time t if she bids b_i. Recall that we used $eCPM_i(b_i)$ and $eQ_i(b_i)$ to denote these functions conditional of not being budget smoothed, but considering multiple steps, we now need to at the time t also to the notation. Assuming bidder i uses bid b_{it} and then get dropped with probability π_{it} by the budget smoothing process, then we get that $\mathrm{exQ_{it}}(b_{it}) = \pi_{it} eQ_{it}(b_{it})$ and $\mathrm{exCPM_{it}}(b_{it}) = \pi_{it} eCPM_{it}(b_{it})$. Using these notations, the utility of bidder i at instance t is expressed as as

$$u_{it}(b_{it}, v_i) = v_i \, \mathrm{exQ_{it}}(b_{it}) - \mathrm{exCPM_{it}}(b_{it}).$$

Note that the impression share $\mathrm{exQ_{it}}(\cdot)$, as well as the expected cost $\mathrm{exCPM_{it}}(\cdot)$, depends on the bids of all other bidders $\vec{b}_{-i,t}$, as well as on the budgets, and on other environment variables, such as the impression volume available at time t. In our discussion, we will leave these variables implicit (where it does not affect mathematical clarity), as we have done in the sections so far.

The idea behind the no-regret learning is the assumption that agents vary their bids and succeed to get utility close to as high as any single bid with hindsight. More formally, we say that agent i has average regret at most ϵ_i if the following holds.

Definition 1 (Average Regret)**.** *A sequence of play that we observe has average regret at most ϵ_i for bidder i if:*

$$\forall b' \in \mathcal{B} : \frac{1}{T}\sum_{t=1}^{T} u_{it}(b_{it}, v_i) \geq \frac{1}{T}\sum_{t=1}^{T} u_{it}(b', v_i) - \epsilon_i \tag{3.4.7}$$

The concept of average regret is comparing the utility of the agent to a baseline bidding strategy, which is the best fixed bid over the whole time. The overall regret of the agent is measured by the difference in the cumulative utility that agent's bidding strategy yields and that of the best fixed bid. We note that an agent who successfully takes advantage of fluctuations of impression opportunities, could do better than any single bid with hindsight, and hence achieve negative regret. Agents achieving negative regret satisfy the definition with $\epsilon_i = 0$.

Our inference of agent values will be based on the assumption that each agent will experiment with bids to try to find a consistent good bid. We use this weaker assumption, in place of the classical assumption that the agents best respond to the environment in each time step.

In light of the above way of thinking of agents as experimenting with bids, and learning what may be the best bid, each bidder will then be characterized by two parameters: her value v_i and the average regret that evaluates the success of the dynamic bid adjustment. Average regret of a player reflects the properties of bidder's learning algorithm used. This leads to the following definition of a *rationalizable set under no-regret learning* (or more precisely, small average regret learning). From the Econometric perspective, the rationalizable set corresponds to the identified set of values and average regret consistent with small average regret learning by the agents.

Definition 2 (Rationalizable Set). *A pair (ϵ_i, v_i) of a value v_i and average regretr ϵ_i is a rationalizable pair for player i if it satisfies Equation* (3.4.7)*. We refer to the set of such pairs as the* rationalizable set *and denote it with $\mathcal{N}R_i$.*

The rationality assumption of the inequality (3.4.7) models players who may be learning from experience while participating in the game. We assume that the strategies b_{it} and functions $\mathrm{exQ}_{\mathrm{it}}(\cdot)$ and $\mathrm{exCPM}_{\mathrm{it}}(\cdot)$ are input simultaneously, so agent i cannot pick her strategy anticipating the realization of uncertainty regarding the impression opportunities or the strategies of other agents $b_{-i,t}$. This makes the standard of a single best strategy b_i natural, as chosen strategies cannot depend on the uncertainty in impression volume or the bids of opponents. Beyond this, we do not make any assumption on what information is available for the agents, and how they choose their strategies.

We can specialize the definition of the rationalizable set in (3.4.7) to auctions for randomly arriving impressions by introducing functions

$$\Delta\,\mathrm{exCPM}_{\mathrm{i}}(b') = \frac{1}{T}\sum_{t=1}^{T}\left(\mathrm{exCPM}_{\mathrm{it}}(b') - \mathrm{exCPM}_{\mathrm{it}}(b_{it})\right), \quad \text{and} \tag{3.4.8}$$

$$\Delta\,\mathrm{exQ}_{\mathrm{i}}(b') = \frac{1}{T}\sum_{t=1}^{T}\left(\mathrm{exQ}_{\mathrm{it}}(b') - \mathrm{exQ}_{\mathrm{it}}(b_{it})\right), \tag{3.4.9}$$

corresponding to an aggregate outcome in T time periods from switching to a fixed bid b' from the actually applied bid sequence $\{b_{it}\}_{t=1}^{T}$. The ϵ-regret condition reduces to:

$$\forall b' \in \mathbb{R}_+ : v_i \cdot \Delta\,\mathrm{exQ}_{\mathrm{i}}(b') - \Delta\,\mathrm{exCPM}_{\mathrm{i}}(b') \leq \epsilon_i \tag{3.4.10}$$

for each bidder i. Hence, the rationalizable set $\mathcal{N}R_i$ is an envelope of the family of half planes obtained by varying $b' \in \mathbb{R}_+$ in Equation (3.4.10).

Assuming that values of agents are contained in $\mathcal{B}$, the rationalizable set is closed, convex and bounded. Since closed convex bounded sets are fully characterized by their boundaries, we can use the notion of the *support* function to represent the boundary of the set $\mathcal{N}R_i$. The support function of a closed convex set X is $h(X, u) = \sup_{x \in X} \langle x, u \rangle$, where in our case $X = \mathcal{N}R_i$ is a subset of $\mathbb{R}^2$ or value and error pairs (v_i, ϵ_i), and then u is also an element of $\mathbb{R}^2$.

An important property of the support function is the way it characterizes closed convex bounded sets. Denote by $d_H(A, B)$ the Hausdorf distance between convex compact sets A and B. Recall that the Hausdorf norm for subsets A and B of the metric space E with metric $\rho(\cdot, \cdot)$ is defined as

$$d_H(A, B) = \max\{\sup_{a \in A} \inf_{b \in B} \rho(a, b), \sup_{b \in B} \inf_{a \in A} \rho(a, b)\}.$$

It turns out that $d_H(A, B) = \sup_u |h(A, u) - h(B, u)|$. Therefore, if we find $h(\mathcal{N}R_i, u)$, this will be equivalent to characterizing $\mathcal{N}R_i$ itself.

The characterization of the support function is based on the following idea. Let (u_1, u_2) be a unit vector. Suppose that there exists a bid b such that $\Delta\,\mathrm{exQ_i}(b) = -\frac{u_1}{u_2}$. In that case the boundary of the half-space $v_i \cdot \Delta\,\mathrm{exQ_i}(b) - \Delta\,\mathrm{exCPM_i}(b) \leq \epsilon_i$ transforms to $u_1\, v_i + u_2\, \epsilon_i = |u_2|\, \Delta\,\mathrm{exCPM_i}(b)$. In other words, the support function in this case corresponds to the right-hand side $u_2\, \Delta\,\mathrm{exCPM_i}(b)$. While previous work in [NST15] relied on the continuity of the functions $\Delta\,\mathrm{exQ_i}(\cdot)$ and $\Delta\,\mathrm{exCPM_i}(\cdot)$, for our data they are discrete (and the rationalizable set is a polyhedron). To modify the construction of the support function for this case, for each $\mathbf{u} = (u_1, u_2)$ we find $b^u_{\mathbf{u}}$ and $b^l_{\mathbf{u}}$ such that $\Delta\,\mathrm{exQ_i}\left(b^l_{\mathbf{u}}\right) < -\frac{u_1}{u_2} < \Delta\,\mathrm{exQ_i}\left(b^u_{\mathbf{u}}\right)$. Function $\Delta\,\mathrm{exQ_i}\left(\cdot\right)$ is constant in the interval $[b^l_{\mathbf{u}},\, b^u_{\mathbf{u}})$. Thus, all normal vectors $\mathbf{u}$ that satisfy the inequality above correspond to the hyperplanes passing through the vertex of the rationalizable set. As a result, their support function will be determined by $|u_2|\ \Delta\,\mathrm{exCPM_i}\left(b^l_{\mathbf{u}}\right)$.

The following result fully characterizes the support function of the set $\mathcal{N}R_i$ based on the aggregate auction ouctomes $\Delta\,\mathrm{exQ_i}(\cdot)$ and $\Delta\,\mathrm{exCPM_i}(\cdot)$:

Theorem 26. *Under monotonicity of $\Delta\,exCPM_{\mathrm{i}}(\cdot)$ the support function of $\mathcal{N}R_i$ is function $h\,:\,\{(u_1,u_2)\,:\,u_1,u_2\in\mathbb{R},\ u_1^2+u_2^2=1\}\mapsto\mathbb{R}_+$ such that*

$$h(\mathcal{N}R_i,u)=\begin{cases}|u_2| \inf\limits_{u_1+u_2\,\Delta\,exQ_{\mathrm{i}}(b)>0} \Delta\,exCPM_{\mathrm{i}}(b), & \textit{if}\ \ u_2<0, \\ \qquad\qquad\qquad\qquad \textit{and}\ \frac{u_1}{|u_2|}\in\left[\inf_b \Delta\,exQ_{\mathrm{i}}(b),\,\sup_b \Delta\,exQ_{\mathrm{i}}(b)\right], & \\ +\infty, & \textit{otherwise.}\end{cases}$$

This theorem is the identification result for valuations and algorithm parameters for ϵ-regret learning algorithms. Unlike equilibrium settings that we discussed above, we cannot pin-point the values of players. At the same time, the characterization of the set $\mathcal{N}R_i$ reduces to evaluation of two one-dimensional functions. We can use efficient numerical approximation for such an evaluation. The shape of the set $\mathcal{N}R_i$ will generally depend on the parameters of a concrete algorithm used for learning. Thus the analysis of the geometry of $\mathcal{N}R_i$ can help us not only to estimate valuations of players but the learning algorithm as well.

The inference for the set $\mathcal{N}R_i$ reduces to the characterization of its support functions which only requires to evaluate the function $\inf_{u_1+u_2 \Delta \text{exQ}_{\text{i}}(b)>0} \Delta \text{exCPM}_{\text{i}}(b)$. It is a one-dimensional function that can be computed directly from the data.

One potential challenge in our case is the evaluation of the "sampling uncertainty." We have a dataset that contains the entire universe of auction outcomes and, therefore, is the superset of the information that the agents used to make decisions. We can potentially consider a different version of our model where we place agent's decisions in the Bayesian framework in which they make decisions trying to anticipate future uncertainty and thus their actions are based on the expected auction outcomes with the expectation taken over the Bayesian posterior of agents over uncertainty of the impression volume and opponent actions. In that case the uncertainty characterizes the difference between agent's expectations and averages that we compute from the data. We can model this uncertainty, for instance, by subsampling the data (assuming that the sampling noise is distributed identically over time). We do not focus on this source of uncertainty in this chapter and adhere to our model.

Another important source of uncertainty is the sampling error in our estimated intra-week impression volume. This corresponds to the usual error of fitting a linear regression (that is further used to predict daily impression volume). When we recover rationalizable sets of bidders from the data we account for this uncertainty using subsampling. We will discuss this in more detail in the next section.

3.5 Empirical analysis of learning dynamics

3.5.1 Characterization of agent heterogeneity

In our structural inference we study the success of agents' bid adjustment over time. We find that a key characteristic of agents is the frequency by which they update bids. On Figure 3.9 we plot the histogram of the distribution of daily frequency of bid changes for all agents across the 6 markets that we study. The histogram shows fairly spread distribution of frequencies, close to uniform between the once-a-month update to once a week update with some agents updating the bids more frequently.

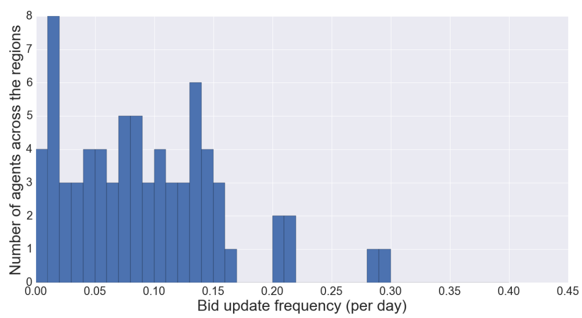

Figure 3.9: Histogram of bid change frequencies across the 6 selected regions.

However, the analysis of frequency of bid changes within individual markets shows much less concordance in the bid update frequency across bidders. In fact the frequency of bid changes turns out to be the key variable that allows us to cluster the bidders into distinct types. We run k-mean clustering algorithm using the bid change frequency variable to partition the agents inside each region into 3 clusters, agents in cluster 1 and cluster 3 have the lowest and the highest number of bid changes per day respectively. This allows us to give a straightforward interpretation to the types identified by the cluster. If the bid changes are triggered by the benefit of the bid change out-weighting the cost of the bid change, then the three clusters can be interpreted as identifying the the bidder with high, medium and low cost of bid changes.

The results of the clustering is demonstrated in Table 3.2. The results show fairly balanced cluster sizes across markets with the high and medium cost clusters containing the largest number of bidders and the lowest cost cluster containing the smallest number of bidders (less than a quarter of bidders).

Region #	1	2	3	4	5	6
Number of agents	20	23	18	21	16	18
Average bid changes per day	1.23	0.91	0.65	0.54	0.54	0.51
Selected agents in clusters 1,2,3	5,5,4	5,6,4	6,5,1	3,2,2	8,3,3	3,6,1

Table 3.2: Basic information for the 6 selected regions. We removed the agents that are in auction for less than 7 days or do not change their bid at all.

3.5.2 Learning to trust the recommendation

Bid recommendation tool was designed by the platform to help the bidders to transition from the fixed price contracts to the auction-based system for impression pricing and delivery. The bidding tool provided a simple interface that allowed the bidders to submit their monthly budget and their bid for a specific market. The bid recommendation tool provided the bid that maximizes the number of impressions that could be purchased with the given budget. The tool would adjust the recommendation with any change that occurs in the system, such as the arrival of the new bidder, changes of bids by existing bidders, or the change in the predicted number of market impressions. However, bidders would need to log into the system to see (and possibly adapt) the updated recommendation.

During the market rollout Zillow made the agents aware of the tool's existence and explained the principles that were used to design the tool. However, despite the outreach and marketing work when auction platform was introduced to the set of experimental markets, the actual utilization rate of the tool was initially low.

On Figure 3.10 we display the percentage of time when recommended bid was used for the bid change as a function of the agent's time in the auction platform, averaged over all agents in the experimental markets. The figure shows that when agents were introduced to the platform they tend to use the recommendation tool for less than 50% of their bid changes. This number tends to grow to almost 100% as the agent is exposed to the auction platform for more than 5 months.

Figure 3.11 presents the same trend of utilization of the bid recommendation tool but decomposed by clusters. Recall that we cluster bidders based on their bidding frequency with cluster 1 being the cluster with the lowest frequency of bid changes and cluster 3 being the cluster with the highest frequency of bid changes. Figure 3.11 shows the persistence of the trend with the relatively low utilization of the bidding tool when the agents are just introduced to the

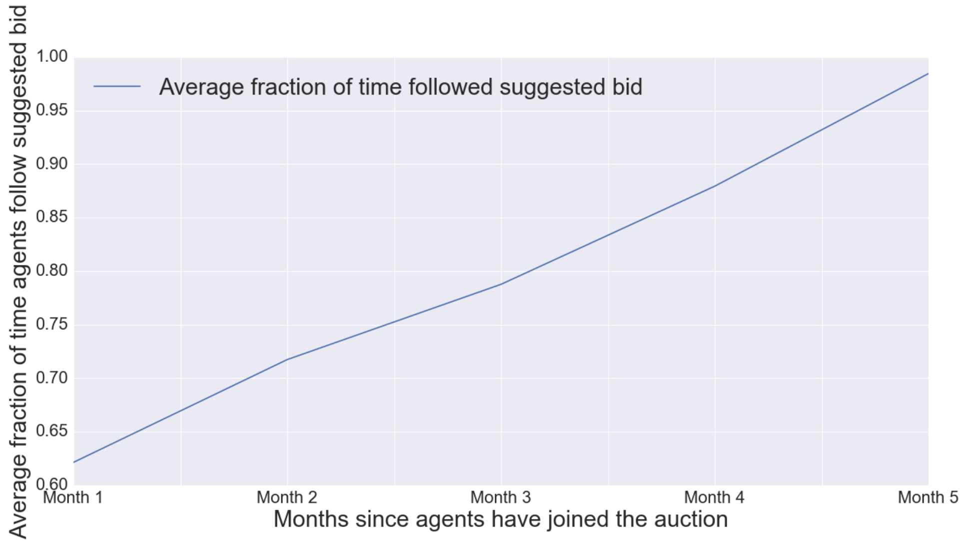

Figure 3.10: Average fraction of time agents follow the recommended bid across the 6 selected regions.

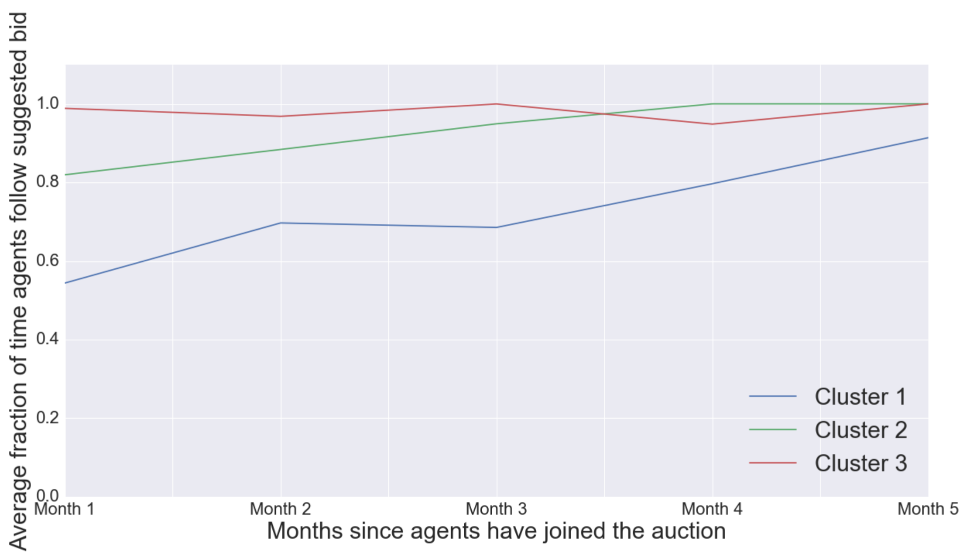

Figure 3.11: Average fraction of time agents follow the recommended bid separated by clusters.

auction platform with an increased degree of this utilization as the agent spends more time bidding in the auctions. We note that this trend is most rapid for the cluster of bidders with the highest frequency of bid changes. These bidders start using the bid recommendation tool for almost all of their bid changes after 3 months of exposure to the auction platform. At the same time the bidders that are least frequently changing their bids do not get to the point of fully using the recommendation tool even after 5 months of experience with the auction platform.

We further illustrate the persistence of this trend across the 6 markets that we study in the Appendix. Figure B.1 in the Appendix confirms the consistency of the aggregate trend of the utilization of bid recommendation tool with those trends in individual markets. Moreover, for some markets the percentage of utilization of the bid recommendation tool is even smaller than that on average, especially for the bidders that change their bids the least frequently.

This leads us to two important observations. First, the bidders in all observed markets were willing to "experiment" with their bids by deviating from the recommended bid. The proportion of bids devoted to experimentation is large, especially when the agents are newly introduced to auction markets, and remains relatively large among the bidders that do not frequently change their bids even after 5 months of them bidding on the auction platform. Second, even though the bid recommendation tool was designed to optimize bids on behalf of the bidders, assuming they want to spend their budget, and a big fraction of agents (overall a 39.4% of all agents) essentially spend their budgets, the bidders did not have the full faith that the recommendations benefit them (as opposed to the auction platform). The increasing adherence to the recommended bid is observed only after the bidders experiment with alternative bids for a sufficiently long time.

Our observation that the agents used the recommended bid at the initial stage of their tenure in the auction market less than after they spent a sufficient time bidding in the new auctions clearly demonstrates that the agents *learned to adopt the recommended bid*. In our further analysis we will study how the dynamics of agents' payoffs has helped them make adoption decisions.

3.5.3 Identified sets for agent's parameters

From the econometric perspective the rationalizable set $\mathcal{N}R$ is the *identified set* of structural parameters for each agent: agent's per impression value v_i and the bound of the average regret of agent's learning strategy ϵ_i. Since our theory characterizes the rationalizable set $\mathcal{N}R$ as a collection of half-spaces generated by inequalities (3.4.10), we can estimate the rationalizable set of pairs (v_i, ϵ_i) by constructing functions $\Delta\,\mathrm{exCPM}_i(\cdot)$ and $\Delta\,\mathrm{exQ}_i(\cdot)$ from the data for each player i. We use the grid over bid space to approximate these functions. On Figures 3.12 we show the structure of the rationalizable sets for 3 of the bidders most frequently changing their bids in region 1. The structure of the rationalizable set is similar in all the 6 markets we analyzed, see the corresponding figures B.2-B.4 in the appendix. The vertical axis on these plots is the per impression value of the bidder (expressed in monetary units) while the horizontal axis is the additive average regret.

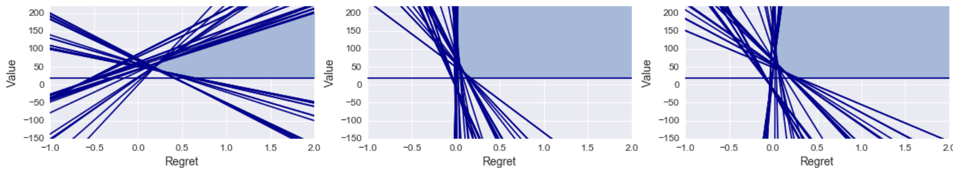

Figure 3.12: Rationalizable set for 3 of the agents most frequently changing bids in region 1

We note a dramatic difference in the shape of these sets with the rationalizable sets for the bidders in advertising auctions on Bing estimated in [NST15]. While the rationalizable sets in [NST15] have smooth convex shape, the rationalizable sets for the agents on Zillow are polyhedra. This is due to a much higher degree of uncertainty for the bidders on Bing (induced by variation of estimated clickthrough rates and user targeting across user queries) that smoothes out the boundary of the rationalizable set.

Another important observation is the highly concentrated set of hyperplanes that pass close to the origin for many the bidders in the observed markets. In fact, as we mentioned before, agents' budget constraints on Zillow are binding for most bidders. For a bidder that is exactly spending their budget, there is a set of bids that correspond to them completely spending the budgets (i.e. that all have identical expected spent). Possible alternate bids that resulting in the same spent, i.e., have $\Delta\,\text{exCPM}_i(b) \approx 0$, give a boundary of the rationalizable set passing through (or near) the origin. For agents where such a line is close to vertical, they can only gain a very small amount of impression volume by the best bid with hindsight.

Discussion: It is insightful to discuss the economic meaning of the rationality assumption for different lines. Recall the small regret condition of (3.4.10). It is useful do distinguish cases depending on the sign of $\Delta exQ_i(b')$ and $\Delta\text{exCPM}_i(b')$.

- When $\Delta exQ_i(b')$ and $\Delta\text{exCPM}_i(b')$ are both positive, the agent had the possibility for getting more impressions by paying more using a fix bid b'. If switching to this fixed bid doesn't increase the agent's utility, this implies an upper bound on the agent's value: the extra expense is not worth the extra utility. For example, on the left image of Figure 3.12, the rationalizable region of this agent has such a bid, so this agent must not be spending his budget.
- When $\Delta exQ_i(b')$ and $\Delta\text{exCPM}_i(b')$ are both negative, the agent had the possibility for getting spending less money by getting less impressions using a fix bid b'. If switching to this fixed bid doesn't increase the agent's utility, this implies a lower bound on the agent's value: the saved expense is outweighed by the lost value of impressions. This is a possible option for all agents, for example, on the same left image on Figure 3.12 has this kind of line also.
- Finally, when $\Delta exQ_i(b')$ is positive and $\Delta\text{exCPM}_i(b')$ is negative (or zero), there is a fixed bid b' that would simultaneously get more impressions, and decrease the cost. This fixed bid would be an improvement for any agent, so the agent definitely has regret, if there is such a bid. For example in region 2 on Figure B.2 , the image in the middle row on the right has such a boundary. Now the magnitude of the regret depends on the agent's value v_i, and the small regret rationality assumption asserts that if the agent has

large value, this regret would be of large magnitude, hence we assume the agent would have noticed it and acted on it. So this also gives an upper bound on the agent's value.

For bidders who spend their budgets, the small regret constraint of the rationalizable set may not give any upper bound on their valuation: these bidders are constrained by their budget, and not by their value of each impression. The rationalizable set of an example of such an agent is the one right on Figure 3.12, or even more clearly, the right of the top row of region 1 on Figure B.2.

One specific point of the rationalizable set that we consider is some of our further analysis corresponds to the pair of value and regret, where the observed bid sequence has the smallest possible average regret, or if the left boundary of the rationalizable region is vertical, the smallest such value. Essentially, this point corresponds to the value of the bidder for which the observed bid sequence corresponds to the best possible learning strategy, breaking ties for small values. Since the average additive regret and the value are expressed in the same monetary units, we can directly compare them. A simple visual analysis of plots of rationalizable sets on Figures B.2-B.4 indicates that while for some bidders the smallest rationalizable average regret is small relative to the corresponding value, there is a large number of bidders with high relative regret. This is particularly pronounced for the bidders with narrow cone-shaped rationalizable sets, such as the two first regions on Figure B.2 in Region 2. From the economic perspective, this shape indicates that few (frequently two) bids dominate all other bids in terms of total cumulative utility for large ranges of possible values of bidders.

Accuracy of estimated rationalizable sets: In the previous section, we noted that an important source of sampling uncertainty in our inference procedure is the error in evaluation of the daily impression volume. We had the data on the daily impression volume for all markets only for half of the time period. For the first half the volume was only available as a monthly aggregate. To account for the intra-week variation of the number of impressions, we fit a linear model with day-of-the-week dummy variables and estimate its coefficients. The error in the estimation of the day of the week effects impacts the estimation of the rationalizable sets. To measure the impact of this error, we implemented a block bootstrap procedure where we repeatedly drew weeks of data (keeping the days of the week in the correct sequence) such that each week was drawn with an equal probability .5. Then we estimated the day-of-the-week effects and use them co construct rationalizable sets. On Figure 3.13 we demonstrate the 90% confidence sets for the boundary of rationalizable sets for the same bidders whose rationalizable sets we presented on Figure 3.12. We obtain these by taking the highest and lowest regret on the boundary of the rationalizable set constructed from a given subsample for a given value of the bidder. We note that these errors are small. At the point of the lowest rationalizable regret, 90% confidence interval for the corresponding value is less than 1% of that value. These are typical for the bidders in the markets that we analyze. For this reason, we do not show confidence sets on any further figures.

Robustness of inference: While the rationality assumption behind the concept of low regret learning is much weaker than the Nash equilibrium assumption, it still requires the players to change the bid in response to any changes in the expected auction outcomes that lead to the violation of the low regret con-

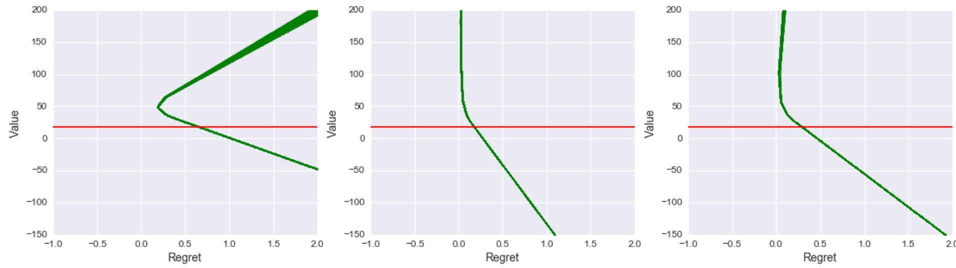

Figure 3.13: 90% Confidence sets for boundaries of rationalizable set for 3 of the agents most frequently changing bids in region 1, with the reserve price is shown in red

straint for the current bid value. That means that the bidders potentially need to adjust their bids if there is an alternative bid that results in large change in the per impression cost even if it does not result in large absolute change in cost or impressions. It is clear that this may not be true in reality: if the bid change does not result in the expected utility change that exceeds the opportunity cost of searching for the new bids, the bidders may choose not to engage in this search. To robustify the empirical rationalizable sets we would like to eliminate the rationality constraints that do not result in "detectable" changes in the overall payment or impressions from an alternative bid. We use a heuristic approach to perform such a robustness analysis.

On Figures B.11 and B.11 each point is a pair of changes in the payment and impression in the rationality constraint that determines the upper and lower bound for the value with which the bidder has the smallest rationalizable regret. These plots show that there are only few bidders whose inferred values maybe determined by the alternative bids that generate a very small change in the overall impressions or payment. As a result, we use a threshold of 2% of the change in either impression volume or the payment in the rationality constraint to consider it "robust." Then we only consider constraints that lead to larger changes in payment or impressions.

On Figures B.5-B.7 we compare our original rationalizable sets and the rationalizable sets that we changed by removing the constraints that imply small changes. We observe that each market has a small number of bidders whose rationalizable set was changed. At the same time, for all bidders who were affected by the changes the shape of the rationalizable did not change. This indicates the robustness of our original rationalizable sets to the constraints that require the bidders to pay attention to small changes in payments and impressions.

3.5.4 Comparing outcomes to recommended bid outcomes

We want to understand why agents may not be following the platform provided recommendation: do they use a different bidding strategy as that improves their obtained utility, or is it simply a question of lack of trust? We note that the bid-recommendation tool didn't take into account the weekly impression volume fluctuation shown on the figure in Section 2.2, so with bidding differently on weekdays and weekends, the agents could have done better than the recommendation, and in fact would be able to achieve negative regret. However, for the agents we study achieve negative regret: none of the regions reach into the $\epsilon < 0$ area, as shown on Figures Figures B.2-B.4. This is not surprising, as they do not change their bids frequently enough to take advantage of the opportunity provided by the fluctuation. Their behavior appears to be well modeled by searching for a fixed good bid.

Our approach to evaluate the advantages and disadvantages of following the bid recommendation would be by evaluating the *counterfactual regret*. In other words, we want to explicitly compare the cumulative utility that bidder's applied bid sequence delivered, the cumulative utility that would have been achieved if the bidder simply used the recommended bid every time they logged in, and the cumulative utility that could have been achieved if the bidder chose the constant bid that optimized the cumulative utility in the hindsight. We display this for a range of possible bidder values. Note that the optimal fixed bid with hindsight does depend on the player's value.

To compute agent's utility we need to have her value per impression. While in Section 3.4 we suggested that the value with smallest regret may be most likely the player's value, here we display the utility as a function of the value for a range of values, as the rationalizable set does not provide a conclusive single value that is compatible with agent's observed behavior. We construct the identified sets that characterize three performance characteristics for all possible values of agents:

(i) the cumulative utility of agent's applied bid;
(ii) the cumulative utility of recommended bid;
(iii) the cumulative utility of the optimal bid for the given value.

For each possible value of the agent the difference between the cumulative utility of the optimal bid and the cumulative utility of the applied bid is the agent's actual regret for that value. The gap between the cumulative utility of the optimal bid and the cumulative utility of the recommended bid is the *counterfactual regret* that the agent would have had if she had adhered to the recommended bid.

For a given generic bid sequence $\{b_{it}\}_{t=1}^{T}$ we can construct the set

$$\mathcal{U}\left(\{b_{it}\}_{t=1}^{T}\right) = \left\{(U_i, v_i)\ :\ U_i = v_i \sum_{t=1}^{T} \text{exQ}_{\text{it}}(b_{it}) - \sum_{t=1}^{T} \text{exCPM}_{\text{it}}(b_{it})\right\}.$$

which characterizes cumulative utility achieved with this bid sequence for all possible values of the agent. Due to linearity of the cumulativre utility in the value, this set is a line with positive slope. Denote by $\{b_{it}^{a}\}_{t=1}^{T}$ the sequence of applied bids of agent i and $\{b_{it}^{r}\}_{t=1}^{T}$ the sequence of recommended bids. Then the identified set of cumulative utility and value pairs (U_i, v_i) for the applied bids is the line $\mathcal{U}\left(\{b_{it}^{a}\}_{t=1}^{T}\right)$ and the identified set of cumulative utility and value pairs (U_i, v_i) for the recommended bids is the line $\mathcal{U}\left(\{b_{it}^{r}\}_{t=1}^{T}\right)$. The utility from the optimal bid is a nonlinear convex function of the value, given that the optimal bid in general will be different for each possible value. It can be expressed as $\mathcal{U}^{*} = \left\{(U_i, v_i)\ :\ U_i = \max\limits_{b \in [0,\, B]} \left(v_i \sum_{t=1}^{T} \text{exQ}_{\text{it}}(b) - \sum_{t=1}^{T} \text{exCPM}_{\text{it}}(b)\right)\right\}$. We note that the curve corresponding to the best constant bid can lie below of the line corresponding to the agent's own bid, since we noticed that the bidding strategy that depends on the day of the week can potentially deliver more impressions to bidders than the fixed bid strategy, however, this tends not to be the case.

On Figure 3.14 we illustrate the structure of the constructed identified sets for selected bidders. The first panel of the figure shows the case where for the values with which the agent can achieve positive utility both the applied and the recommended bid deliver cumulative utility close to that of the optimal bid on high enough values. Clearly, both the applied bid as well as the recommended bid can bring non-negative cumulative utility only if the actual value of the agent is high enough. For this agent, the recommended bid brings fewer expected user impressions than the agent's own bid as indicated by the lower slope of the cumulative utility of the recommended bid. The second panel shows the case where recommended bid is clearly better than the applied bid for any possible value. Finally, the third panel describes the case where the agent does better than the recommended bid while there is significant gap between the optimal utility and the utility that both applied and the recommended bid could generate.

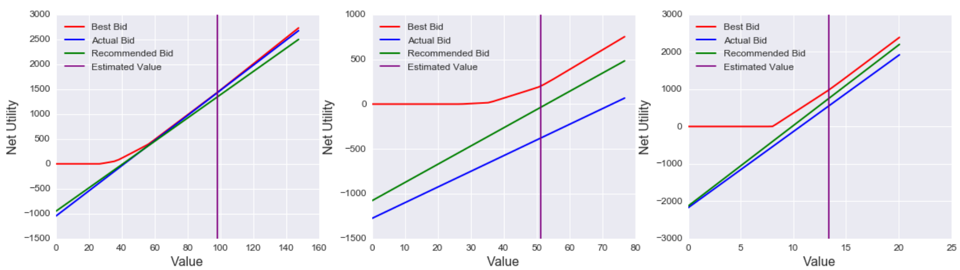

Figure 3.14: Identified sets for cumulative utility - value pairs with actual (blue), recommended (green), and optimal (red) bids

We present the identified sets for utility of the applied bid and utilities of counterfactual optimal and recommended bids for bidders most frequently changing bids for 6 markets in our data in Figures B.8-B.10. We note that for several bidders the applied bid has higher cumulative utility for the entire range of per impression values meaning that those bidders take advantage in the weekly variation of the impression volume.

While the identified sets for the cumulative utility and value pairs of agents

Region	reg 1	reg 2	reg 3	reg 4	reg 5	reg 6	cl 1	cl 2	cl 3	all
worse	21.4%	20%	8.3%	0%	21.4%	50%	23.3%	25.9%	6.7%	20.8%
better	42.9%	20%	41.7%	28.6%	21.4%	20%	46.7%	22.2%	6.7%	29.2%
equal	35.7 %	60%	50%	71%	57.1%	30%	30%	51.8%	86.7%	50%

Table 3.3: The percentage of agents that do worse (or better) with their bids than following the recommendation. The three columns on the right of the table offer aggregate statistics across the 6 markets segmenting agents by the frequency they update.

make it possible to provide universal comparison of regret of the recommended bid and the regret of the applied bid, it is hard to make market-level and cross-market comparisons with those sets. To be able to do such an aggregate analysis, we select the value that corresponds to a specific point in the rationalizable set of each agent. We use the value with smallest rationalizable additive regret as our selected value, but to display the values in context, we also want to account for two features.

In Table 3.3 we show summary statistics on whether agents would decrease or increase their regret by not using a platform provided recommendation in each of the 6 markets. All regrets are computed using the value we inferred using the agent's own bid. Whenever we say that the regret of a bidder's learning strategy is the same as the regret of the recommended bid either her own learning strategy is as good as the recommended bid or that she simply adhered to the recommended bid. The overall distribution of the regret difference indicates that sizeable fractions of bidders both have the regret that exceeds the regret of the recommended bid and the regret that is smaller than that if the recommended bid was used. We also show the same statistics by regions and clusters, where cluster 3 are the agents who update their bids most frequently.

We further illustrate this point on Figure 3.15. To consider the size of regrets, it's useful to measure regret relative to bidder's value (i.e. the bidders may be prone to evaluate the "loss" associated with their learning strategies in the increments of the total "gain"). Second, it is also convenient to normalize the regret by the number of impressions the agent won, so we measure "per impression" regret, viewing the value of the regret compared to the total value obtained from the impressions won. This normalization allows is the compare the regret of agents with rather different budgets. We consider the distribution of differences between the regret of agents' own bidding strategy and the recommended bid. Figure 3.15 shows the distribution of these differences grouped by what percentage of time were the agents following the recommendation. The figures suggest that agents rarely following the recommendation were doing very badly compared to the recommended bid, while the agents following recommendation 20-80% of the time, tended to do better than the recommendation. Note that these differences are calculated using the value for the agent that makes the regret of his own bid choices as small as possible, so strongly favors this conclusion. However, the utility tradeoff between the two options, depicted on Figures B.8-B.10 for all possible values of the agents, suggests the same conclusion: when the blue and green lines differ significantly, the blue (corresponding to the agents own bid) tends to be higher. However, even though the bidders chose to experiment with their bids, the experimentation did not necessarily lead to an improvement of regret over the recommendation. Further, this increase in utility (or decreased regret) appears not to be worth the effort of experimentation, as most agents consistently follow the recommendation by the end of the 5 months period.

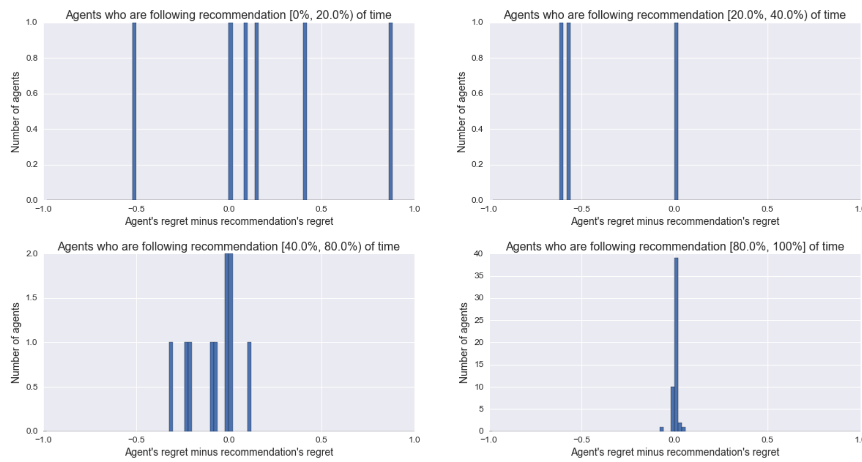

Figure 3.15: Distribution of difference between the regret of own bidding strategy and recommended bid across agents in selected markets separated by the percentage of time the agent follows the recommendation.

One a priori possible explanation for why agents don't initially follow the platform recommendation is that the recommended bid does not provide satisfactory outcomes for the agents and switching to an alternative bidding sequence improves their long-run performance sufficiently to be worth the added effort. Figure 3.16 shows that this hypothesis is not consistent with the data. On average, the agents who use the recommended bid least do not show any improvements over the recommended bid measured by the average regret.

Combining this information with the previous observation of an increasing trend of utilization of bid recommendation tool, we conclude that the key element that explains our results is the *trust* of the agents in the platform-provided bid recommendations. While the bid recommendation tool is optimized for the agents, upon entry to the platform the agents do not trust the tool. Instead, they experiment with alternative bids and compare the performance of those deviating bids with the performance of the recommended bids that they also

Figure 3.16: Scatter plot of difference between the regret of own bidding strategy and recommended bid across agents and the percentage of time agents use the recommended bid. The dashed line shows the best linear fit.

occasionally choose. Once the agents empirically verify that the tool indeed optimizes the bids on their behalf, they start using the tool for most of the bid changes.

While all markets in our data transitioned towards a complete reliance on the bid recommendation tool by the most frequent bidders, we can evaluate the "efficiency cost" of the adjustment period. To do so, we computed cumulative market welfare that corresponds to the actual allocations and payments by the bidders. Then we re-computed payments and allocation for the counterfactual scenario where all bidders always use recommended bids. In Table 3.5.4 we compare the actual welfare of the agents with the counterfactual welfare of the recommended bid (taking the values that correspond to the lowest rationalizable regret).

Table 3.5.4 demonstrates that in 5 out of 6 markets the welfare corresponding to the applied bid exceeds the welfare of the recommended bid. While this

Region	Welfare w. applied bid	Welfare w. rec. bid	Diff	percentage
1	46540.31	46042.84	-497.47	-1.1%
2	22222.87	21988.92	-233.95	-1.1%
3	5095.84	5004.31	-91.53	-1.8%
4	6381.98	6358.47	-23.51	-0.4%
5	51267.45	51737.09	469.64	0.9%
6	7789.30	7741.52	-47.78	-0.6%

Table 3.4: Aggregate market welfare (expressed in normalized $) with applied bids and recommended bids for 6 markets. The values of the agents correspond to the point of the rationalizable set with the lowest regret.

may indicate that own learning strategies of bidders performed better than the recommended bid, our methodology for selecting the value for the analysis "favors" bidder's own strategy. This means that it may over-estimate the actual cumulative welfare that corresponds to bidders' own learning strategy. On the other hand, the welfare difference between the applied bid and the recommended bid is small in percentage terms and is below 2% for all markets. While bidders exerted significant effort experimenting with the bid recommendation relative to their own perception of the correct bid the 2% welfare gain does not seem sufficient to warrant bidders' engagement in this learning. From this perspective, the best design of the system would simply default the bidder into the recommended bid (similar to the design of the advertising auction run by Facebook).

In Tables 3.5.4 and 3.5.4 we investigate the robustness of our conclusion. Table 3.5.4 reports the welfare figures when we re-compute the values of the bidders by dropping the rationality constraint that involve insignificant changes in the overall impressions or payment of that bidder. The table indicates that after re-computation of the values the gap between the welfare from the applied bid

Region	Welfare before				Welfare after			
	applied bid	recommend. bid	Δ	% Δ	applied bid	recommend. bid	Δ	% Δ
1	46540.31	46042.84	-497.47	-1.07	47188.34	46703.43	-484.91	-1.03
2	22222.87	21988.92	-233.95	-1.05	27689.09	27448.44	-240.64	-0.87
3	5095.84	5004.31	-91.53	-1.80	6137.43	6040.30	-97.13	-1.58
4	6381.98	6358.47	-23.51	-0.37	6381.98	6358.47	-23.51	-0.37
5	51267.45	51737.09	469.64	0.92	27844.60	28314.23	469.64	1.69
6	7789.30	7741.52	-47.78	-0.61	7835.94	7790.47	-45.48	-0.58

Table 3.5: [

caption]Comparison of the estimated aggregate market welfare (expressed in normalized $) with applied bids and recommended bids for 6 markets between the case when we use the original values of bidders ("Before") and re-computed values with robust rationalizable sets dropping non-robust constraints ("After"). New bidder values were constructed by dropping constraints that lead to less than 2% change in payment or impressions.

and the welfare of the recommended bid diminishes. In Table 3.5.4 we report the welfare dopping the bidders whose inferred value is based on the rationality constraint that is involved with an insignificant change in impressions or payment or whose inferred value changed after this correction by more than a factor of 2. Similarly to Table 3.5.4 we observe that the gap between the welfare of applied and recommended bid decreases. This confirms the robustness of our results and further supports our conclusion.

Region	Welfare before				Welfare after			
	applied bid	recommend. bid	Δ	% Δ	applied bid	recommend. bid	Δ	% Δ
1	46540.31	46042.84	-497.47	-1.07	47188.34	46703.43	-484.91	-1.03
2	22084.50	21850.55	-233.95	-1.06	21741.53	21500.89	-240.64	-1.11
3	5095.84	5004.31	-91.53	-1.80	6137.43	6040.30	-97.13	-1.58
4	6381.98	6358.47	-23.51	-0.37	6381.98	6358.47	-23.51	-0.37
5	22245.68	22715.31	469.64	2.11	22482.62	22952.25	469.64	2.09
6	7789.30	7741.52	-47.78	-0.61	7835.94	7790.47	-45.48	-0.58

Table 3.6: [

caption]Comparison of the estimated aggregate market welfare (expressed in normalized $) with applied bids and recommended bids for 6 markets between the case when we use the original sets of bidders ("Before") and dropping the bidders whose value corresponding to the lowest rationalizable regret is either based on the constraint that results in less than 2% change in payment or impressions or whose new recomputed value changes by a factor of more than 2 ("After").

CHAPTER 4

COMPUTING EQUILIBRIA IN MATCHING MARKETS

Market equilibria of matching markets offer an intuitive and fair solution for matching problems without money with agents who have preferences over the items. Such a matching market can be viewed as a variation of Fisher market, albeit with rather peculiar preferences of agents. These preferences can be described by piece-wise linear concave (PLC) functions, which however, are not separable (due to each agent only asking for one item), are not monotone, and do not satisfy the gross substitute property– increase in price of an item can result in increased demand for the item. Devanur and Kannan in FOCS 08 showed that market clearing prices can be found in polynomial time in markets with fixed number of items and general PLC preferences. They also consider Fischer markets with fixed number of agents (instead of fixed number of items), and give a polynomial time algorithm for this case if preferences are separable functions of the items, in addition to being PLC functions.

Our main result is a polynomial time algorithm for finding market clearing prices in matching markets with fixed number of different agent preferences, despite that the utility corresponding to matching markets is not separable. We also give a simpler algorithm for the case of matching markets with fixed number of different items.

4.1 Introduction

We consider the problem of matching without money with n agents who have preferences over m items. This problem models a range of situations from assigning students to schools, applicants to jobs, or people to committees. We call such an assignment problem a matching problem, if all agents are required to get a fixed number of items. An intuitive application is the school choice problem. Students have preferences over schools, and each student needs to get assigned to exactly one school. In this chapter we will consider computing the fair randomized solution to this problem proposed by [HZ79] based on market equilibria.

We model the preferences of agents with the value of agent i for being assigned to item j is v_{ij}. Using values allows agents to express the intensity of their preferences. An important property of an allocation is its efficiency. Since agents utilities are meaningless to compare (without money, there is no natural unit to express utility), the best we can hope for is a Pareto-efficient allocation. An allocation is Pareto inefficient if there is an alternate allocation where no agent is worse off, and at least one agent has improved utility. We will consider fractional or randomized allocation. The value of an agent i for the fractional allocation x_{ij} is $\sum_j v_{ij}x_{ij}$, if it obeys the matching constraint. This is the agent's expected value, if x_{ij} is the probability of assigning item j to agent i.

Market equilibria offer an intuitive, fair, and Pareto-efficient solution for problems of allocations of resources to agents who have their own (incomparable) preferences over the items in systems with no money. This was proposed by

[HZ79] in the context of matching markets, and by [DFH$^+$12] (see also [GN12] and [WM15]), in the context of allocation of resources in systems. The idea is to endow each agent with equal resource: a unit of (artificial) money. A set of prices $\mathbf{p}$ for the items is market clearing, if there is a fractional allocation $\mathbf{x}$ of items to agents such that the following conditions hold (i) each item is allocated at most once, (ii) each agents is allocated her favorite set of items subject to the budget constraint[1] that $\sum_j x_{ij} p_j \leq 1$, and (iii) the market clears, meaning that all items not fully allocated have price 0. [HZ79] showed that such market equilibria is guaranteed to exists, see also Appendix C.3. We view the resulting fractional (or randomized) allocation $\mathbf{x}$ as a fair solution to the allocation problem without money, which is also clearly Pareto-efficient and envy-free (no agent prefers the allocation of another agent).[2] We are concerned with computing this solution efficiently.

Computing Market Equilibria and the Odd Demand Structure of Matching Markets. Market equilibrium problems where demands satisfy the *gross substitute* condition are well understood [CV07], and can be computed efficiently. The demand structure of our matching problem does not satisfy the gross substitutability condition, which requires that decreasing the price of an item (while keeping all other prices fixed) should never decrease the demand for that item. We show an example in Appendix C.1 that decreased price can cause decreased demand in a matching markets. It is not hard to gain intuition for the phenom-

[1] Note that agents have no use for the (artificial) money and are simply optimizing their allocated item, subject to their budget.

[2] An alternate way to arrive to the same solution concept is to assign each agent an equal share of each resource, and then look for an equilibrium of the resulting exchange market. To see that this results in an identical outcome, we can think of each agents trade, as a two step-process, where he first sells all his allocated share on the market prices, and then uses the resulting money to buy his optimal allocation.

ena: with the decreased price the agent could get her old allocation, and would have money leftover. In other markets, money can be used to buy additional items. However, in a matching market additional item makes no sense, and instead, the agent may want to exchange his share of a cheaper and less favorable item (possibly the item whose price decreased) for share of a more valuable expensive item.

[DK08] gave an algorithm to compute market equilibria in markets with a fixed number of items, where agents have piece-wise linear concave (PLC) utility functions, despite the fact that PLC utility functions can give rise to demand not satisfying the gross substitute condition. They also gave a polynomial time algorithm to compute market equilibria in markets with a fixed number of agents, where agents have piece-wise linear concave and *separable* utility functions. They leave as an open problem to give a polynomial time algorithm to compute market equilibria in markets with a fixed number of agents and general PLC utilities that are not separable. We have shown in Section 4.1.3 that demand structure of the matching market can be modeled by a piece-wise linear concave (PLC) utility function, which however, is neither separable nor monotone. This allows us to use the algorithm of [DK08] to find a market equilibrium is the number of goods is fixed, but leaves open the question whether market equilibrium can be found in polynomial time if the number of different agents is finite instead.

4.1.1 Our Results

We give a polynomial time exact algorithm for finding market equilibria of matching markets with a fixed number of agents, extending the work of [DK08] to the case of matching markets with a fixed number of agents, despite the fact that utilities describing matching markets are not separable. Our algorithm in Section 4.2 is based on the structural Theorem 37, and explores a polynomial number of possible player utility values and allocation structures, and finds a market equilibrium in polynomial time when the number of agents is fixed. The algorithm also extend to the case when there are only a finite number of different agent utility types.

In case of large number of items and finite number of agents, when each individual item is insignificant, our allocation can be used for finding an approximately optimal integer solution. We achieve this by showing in Lemma 35 that we find an allocation in which the number of items which are shared by the agents is $O(n^2)$, which is constant when the number of agents is constant.

In Section 4.3, we consider the problem with a fixed number of goods. In this case, the algorithm of [DK08] can find market equilibria in polynomial time. We give a simpler algorithm which is tailored for matching markets. With m different goods and n agents, our algorithm enumerates a polynomial number of different set of prices and allocation structures for the equilibrium.

4.1.2 Related Work

The problem of fairly allocating items to unit demand agents without money has been studied extensively in both Economics and Computer Science literature. Perhaps the most well known solution to this problem is the *random serial dictatorship (RSD)* [AS98] — also known as *random priority (RP)* — in which agents are served sequentially according to a random permutation, and each agent in turn receives her most preferred item among the remaining ones. Clearly, serial dictatorship is Pareto efficient, and as a result, RSD is ex post Pareto-efficient, i.e., Pareto-efficient given the order used. However, the expected allocation of RSD may not be Pareto-efficient, i.e., its not interim Pareto-efficient. In Appendix C.1 we give an example where the expected allocation of RSD can be Pareto improved just using the order of player preferences, showing that RSD may be Pareto-inefficient even with ordinal preferences.

An alternative solution called probabilistic serial (PS) was proposed by [BM01] which is both envy-free and Pareto-efficient with respect to ordinal preferences. The PS mechanism is, however, not Pareto-efficient with cardinal preferences. This is possible, as ordinal preferences are not always sufficient for ranking the randomized (interim) allocations, (i.e., ranking of distributions over outcomes).

The mechanism we study in this chapter, based on market equilibrium from equal income, has been proposed in this context by [HZ79], is both envy-free and Pareto-efficient even with respect to cardinal preferences. Note that neither PS nor the market equilibrium mechanism is strategy-proof. However, [Zho90] shows that for $n \geq 3$ agents there is no mechanism that is strategyproof, Pareto-efficient, and envy-free.

[HZ79] proves that equilibrium is guaranteed to exist (see also Appendix C.3), and propose an exponential time algorithm for finding approximate equilibrium, whereas this chapter proposes an exact algorithm for computing equilibrium which runs in polynomial time when either the number of agents or the number of items is constant. Most of the recent work on the problem of assignment without money has been focused on analyzing the efficiency of RSD and PS mechanism under cardinal and/or ordinal preferences, e.g., [BCK11, ASZ14].

The main techniques used in the current chapter are based on the cell decomposition result of [BPR98] which has also been used by [DK08] to derive a polynomial time algorithm for a related market equilibrium computation problem. We show how to find equilibria of matching problems in polynomial time when the number of agents is fixed. In Appendix 4.3 we also give an algorithm to find equilibria in matching markets with a fixed number of goods. While the algorithm of [DK08] can be used for this latter case, our algorithm is simpler: we avoid some complications (for instance their primal dual technique for checking the market clearing conditions), and we use a simpler cell decomposition theorem. The case of fixed number of agents has been studied by [EW12]. However, they assume that the agents' utility functions are strictly concave and strictly monotone, which does not apply to our problem. They also approximate the Walrasian equilibrium, while our main goal is to find the exact value of equilibrium prices and allocations.

4.1.3 Preliminaries

In this section we review the matching problem with additive preferences and the market equilibrium solution we aim to compute. Then we'll discuss our main technical tool, the cell decomposition technique of Basu, Pollack, and Roy [BPR98].

The Matching Problem The problem is defined by a set of m items, and $C_j \geq 0$ amount available of item j, and a set of n agents. The matching problem requires that we allocate exactly 1 unit of these items to all agents. The amount C_j available of each item j may be very small, so the 1 unit allocated to an agent may need to be combined from small fractions of many different items. An allocation $\{x_{ij}\}$ for all agents i is a feasible solution of the matching problem, if

$$\begin{aligned}\sum_j x_{ij} &= 1 \text{ for all } i \in [n]\\ \sum_i x_{ij} &\leq C_j \text{ for all } j \in [m]\\ x_{ij} &\geq 0 \text{ for all } i \in [n] \text{ and } j \in [m]\end{aligned}$$

We assume agent i has value v_{ij} for a unit of item j. So her value for a set of $\{x_{ij}\}$ amounts of each item j is $\sum_j v_{ij}x_{ij}$, assuming $\sum_j x_{ij} = 1$.

More generally, we can require to allocate different amounts for different agents, and allow the matching constraint to be only an upper bound, that is, allocate at most 1 unit to each agent. For simplicity of presentation, in this chapter we will use equal amount required for the agents, and normalize that value to 1. Further, we will assume, also for simplicity of the presentation only, that $\sum_j C_j = n$, so a feasible solution to the matching problem will fully allocate all items.

Fair Allocation: Matching Market. We use the Fisher market proposed by [HZ79] to make the allocation fair. Fisher market is defined by giving each agent a unit of (artificial) money. A market equilibrium is defined by a set of prices $p_j \in \mathbb{R}_{\geq 0}$ for each item j. Given a set of prices, the agent i's optimization problem can be written as follows:

$$\begin{aligned}
\text{maximize} \quad & \sum_j v_{ij} x_{ij} & \\
\text{subject to} \quad & \sum_j x_{ij} = 1 & \text{(The matching constraint)} \\
& \sum_j x_{ij} p_j \leq 1 & \text{(The budget constraint)} \\
& x_{ij} \geq 0 & \forall j \in [m] \, .
\end{aligned}$$

A *market equilibrium* for this market is a set of prices $\mathbf{p}$, and a feasible allocation $\mathbf{x}$ such that (i) $\mathbf{x}_i$ is an optimal solution to agent i's optimization problem with respect to prices $\mathbf{p}$ and her budget of 1 unit of money. Note that the requirement of market clearing that all items are allocated, is automatically satisfied due to the matching constraint and our assumption that $\sum_j C_j = n$.

The matching constraint in the agents' preferences creates an odd demand structure. For some prices the agent's optimization problem is not feasible, and even on prices when all agent optimization problems are feasible, the preferences do not satisfy the gross substitute property, i.e increasing price of an item may increase the demand of that item, as explained in the introduction, and an example is shown in Appendix C.1.

A natural idea to convert the problem to one with a simpler structure is to allow agents to have free disposal, i.e., if assigned more than 1 unit of item, they value the best unit. Unfortunately, this change in the model significantly changes the structure of the problem, and can result in market equilibria that are simply not feasible for the original problem. Consider the following market with 2 agents and 2 items.

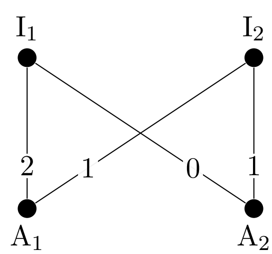

Let $v_1 = (2, 1)$ and $v_2 = (0, 1)$. Assume that the prices are $\mathbf{p} = (0.5, 1.5)$. When agents have free disposal utilities, these prices are equilibrium, since $\mathbf{x}_1 = (1, \frac{1}{3})$ and $x_2 = (0, \frac{2}{3})$ is optimal for both agents, and also clears the market. However, with matching market preferences these prices are not equilibrium, since I_1 will not get any of A_2 and I_2 cannot afford all of A_2 so the prices are not market clearing.

We note that it is possible to express this matching market problem as a classical Fisher market with agent preferences that are piece-wise linear and concave (general PLC) functions, though non-monotone and non-separable. To do this, we first relax the matching constraint in the agents' optimization problem to requiring only that $\sum_j x_{ij} \leq 1$. We will show that the market clearing condition of a Fisher market requiring that all items are allocated (or have 0 price) will help to enforce that all agents get exactly 1 unit. We can also express the agents' utility as a piece-wise linear and concave function. To do this, let $v_i^* = \max_j v_{ij} + \epsilon$ where $\epsilon > 0$. Let agent i's utility for an allocation $\mathbf{x}_i$

$$\min(\sum_j v_{ij}x_{ij}, \sum_j v_{ij}x_{ij} + v_i^*(1 - \sum_j x_{ij}).$$

When $\sum_j x_{ij} \leq 1$, the first term is smaller, but when $\sum_j x_{ij} \geq 1$ the minimum is taken by the second term, so the total value strictly decreases as the allocated amount exceeds 1. Since the agent's utility decreases with more than one unit allocated, an optimal solution to the agent's optimization problem will allocate at most 1 unit of item to each agent, and hence if the market clears, this is also a solution to the matching market problem.

[DK08] provides a polynomial time algorithm to find the equilibrium prices and allocation for the case with fixed number of goods. In Appendix 4.3 we give a simpler exact algorithm for this case, taking advantage of the matching structure. Our main result is to extend this to the case of fixed number of agents, instead of fixed number of items. [DK08] also offers an algorithm for finding market clearing prices for with fixed number of agents, but only with *separable* PLC utilities, i.e., when the utility of each agent is a separable function of the items allocated to the agents. Note that the utility function of a matching problem is necessarily not separable, as it needs to express the upper bound on the total allocation.

The Cell Decomposition Technique Our algorithms will search the space of optimal utility values of each agent, and for each possible utility value, will search through the possible structures of allocations. There are only a fixed number of agents, however, the space of possible optimal values is huge. We use a cell decomposition technique to make the search space discrete, facilitated by a characterization of equilibria. The beginning of Section 4.2 has a more detailed outline.

A main technical tool for our work will be the following theorem concerning the way polynomials divide the space in a d dimensional space. Given a set of polynomials on M variables, the sign of the polynomials define an equivalence between vectors in the M dimensional space $\mathbb{R}^M$, where two vectors y and y' are equivalent if all polynomials have the same sign on y and y'. We call the equivalence sets of this relations the *cells* of the way the polynomials divide up the space. In principle N functions can divide up the space into at many as 3^N cells (as each polynomial can be 0 positive or negative). However, Basu, Pollack, and Roy [BPR98] showed that bounded degree polynomials in small dimensional spaces define much fewer cells.

Theorem 27. *[BPR98] If we have a set of M number of variables, and N number of polynomials whose degree is at most d, then the number of non-empty cells, and the time required to enumerate them is $O(N^{M+1})d^{O(M)}$.*

We will use this decomposition to find the equilibrium for our matching problem. To illustrate the idea, and let $\mathbb{R}^n$ be the space of all possible agent utilities. Roughly speaking the idea is as follows. If we could describe whether a set of utilities $u \in \mathbb{R}^n$ arises from an equilibrium by the signs of a few bounded degree polynomials in these variables, then we could use Theorem 27 to enumerate all cells defined by these polynomials, and test which of the cells satisfies the condition required for being an equilibrium. Unfortunately, the equilibrium condition cannot be described this way, so we will need to introduce additional variables (helping us infer the prices and assignment, despite the fact that these are not in fixed dimensional space) to be able to carry out this plan.

4.2 Computing Market Equilibrium with Fixed Number of Agents

In this section, we give an exact algorithm to find an equilibrium in the case where the number of agents n is constant, and the number of different goods is an arbitrary number $m \in \mathbb{N}$, under the mild technical assumption that each agent has a unique most preferred item. More formally, for every agent $i \in [n]$, there is exactly one item j such that $v_{ij} = \max_{k \in [m]}(v_{ik})$. The goal of this section is proving the following theorem.

Theorem 28. *Exact equilibrium (prices and allocations) in a matching market with fixed number of agents, in which agents have additive values one unit of money, and a unique most preferred item, can be found in polynomial time.*

General Outline and Techniques. Our algorithm searches the space of agents' optimal utilities and item prices to find an equilibrium. We divide this space into a polynomial number of cells, where each cell contains utility and price vectors that have the same properties. We use Theorem 27 as the basis of the cell decomposition. The space of possible agent utilities is finite dimensional. However, since the number of items is not constant, we cannot use a separate variable for each item price. In section 4.2.1, we provide a bundling technique and a characterization of the equilibrium structure that allow us to define equilibria using only a finite set of variables.

Cell. Consider the vector of player utilities $\mathbf{u} = \{u_i\}$, a constant dimensional space for fixed number of agents. Now consider the linear functions $u_i - v_{ij}$ for each agent i and item j. A cell of the space of utilities $\mathbf{u}$ defined by these functions is the region of this space in which each of these functions has a fixed sign. Within each region, the items are divided for each agent i into those with value above u_i, same as u_i, and below u_i. This division also has implications on prices: if the utilities are part of an equilibrium, the price p_j of any item j with value $v_{ij} > u_i$ will have to be above 1. We will add further variables and polynomials, until each cell provides enough information for checking all the equilibrium properties.

Layered Cell Decomposition. Next we would like to add the item prices as variables. However, to keep the running time polynomial when using Theorem 27, we can only have a constant number of variables, so we cannot afford a variable for each item price. To get around this problem we will try a polynomial number of different structures for the price vector, where for any fixed structure, we can define all item prices via a fixed number of variables. To do this, for each agent we will fix a special item that is at least partially allocated to the agent. Lemma 33 will show that given prices for the fixed number of special items, we can infer prices for all items.

Finally, we also need to be able to find the assignment variables. We will show in Lemma 35 that each pair of agents only shares a few items (at most 5), and given the set of shared items, as well as the utilities and item prices, the allocation can be fully determined. Our algorithms iterates over all structures of specially assigned items and shared sets of items. For each of these structures, the algorithm iterates over all cells of the cell-decomposition given by agent utilities and prices of special items and the constraint (polynomials) described in the next subsections that ensure that these describe an equilibrium, and finds the ones which correspond to equilibria.[3]

Bundles. Rather than thinking about individual items in isolation, it is useful to think of items in pairs. In equilibria each agent spends exactly one unit of money and gets exactly a total of one unit of items. This means that in equilibrium, if an agent gets some amount an item whose unit price is less than one, she should also get some amount of an item whose unit price is more than one. As suggested by this fact, we pair items of price below and above 1. We define a *bundle* as either a single item of price 1, or fractions of two items of a total of one unit of items, where the total unit price of the bundle is exactly 1. First, we show that in a market equilibrium, the allocations of items to agents, can be rewritten as the allocation of bundles to agents.

Lemma 29. *In an equilibrium pricing* $\mathbf{p}$ *with equilibrium allocation* $\mathbf{x}$*, there exists a bundling* B *of items such that*

[3]Recall that by [HZ79] (see also Appendix C.3) an equilibrium must exists.

a. Each bundle b consists of at most two items. Two items j and k are in a bundle $b = (j, k)$ if and only if $p_j < 1 < p_k$. One item j forms a bundle $b = (j)$ if and only if $p_j = 1$.

b. Each bundle $b = (j, k)$, associated with a unique mix $0 < \alpha_b < 1$ such that $\alpha_b p_j + (1 - \alpha_b)p_k = 1$ (recall that $p_j < 1 < p_k$). For bundles with one item we use $\alpha_b = 1$.

c. The optimum allocation of items to agent i satisfying the matching constraint can be rewritten as allocation of bundles.

Proof. Let $\mathbf{p}$ be an equilibrium pricing and $\mathbf{x}$ be an arbitrary allocation associated with $\mathbf{p}$. For each agent i, we know that the total allocation of items to i and the the total cost of i's items is 1 (due to the market clearing condition), i.e.

$$\sum_{j \in [m]} x_{ij} = 1$$

and

$$\sum_{j \in [m]} p_j x_{ij} = 1$$

We start rewriting i's allocation with bundles. Let y_{ib} be the amount of bundle b that i uses in the new allocation (y_{ib} is zero at the beginning). At each step we consider the following cases

1. For every j that $x_{ij} > 0$, we have $p_j = 1$: in this case all such items should be in a bundle $b = (j)$ ($y_{ib} = x_{ij}$), so our claim is correct.
2. There exists j that $x_{ij} > 0$ and $p_j > 1$: This means that there exists another item j' such that $p_{j'} < 1$ and $x_{ij} > 0$, otherwise if i gets t unit of items, her total cost cannot be t (note that $t = 1$ here). So by the definition, j and j' should be in a bundle $b = (j, j')$.

Let $z = \min(\frac{x_{ij}}{\alpha_b}, \frac{x_{ij'}}{1-\alpha_b})$. In the second case, we increase y_{ib} by z, and reduce x_{ij} by $z\alpha_b$ and $x_{ij'}$ by $z(1-\alpha_b)$. This means that both the total cost of remaining allocation of $\mathbf{x}$ and the total remaining items in allocation of $\mathbf{x}$ decrease by z. Note that by doing this the total allocation of i (counting her allocation in $\mathbf{x}_i$ and y_i) does not change. Furthermore, either x_{ij} or $x_{ij'}$ becomes zero.

If we repeat this process, y_i gives us a way to rewrite the allocation of items to i as an allocation of some bundles (B_i) to i. If we repeat this for all the agents, we get what we want. ∎

For each agent i, the value of a bundle $b = (j, k)$ is $v_{ib} = \alpha_b v_{ij} + (1 - \alpha_b)v_{ik}$, while the value of a single item bundle $b = (j)$ is just the value $v_{ib} = v_{ij}$. Let $B_i \subseteq B$ the set of bundles of maximum value for agent i (called i's optimum bundles).

Corollary 30. *The optimum allocation of items to agent i is any allocation of bundles using only bundles in B_i. Furthermore, utility of i in the equilibrium is $u_i = v_{ib}$ for the bundles $b \in B_i$.*

Proof. We prove this by contradiction. Assume that $\mathbf{y}$ is the allocation of bundles to agents. Also assume that there exists an agent i and a bundle b for which $y_{ib} > 0$ but there exists another bundle b' such that $v_{ib} < v_{ib'}$. Since the unit price of the bundles is 1, and there is one unit of items in them, i can trade her share of b for the same amount of b' and increase her value. This is a contradiction.

Now since the unit values of bundles that i uses are the same, and i gets exactly one unit of bundles, her utility in equilibrium is equal to the unit value of these bundles. So the second claim is also true. ■

A key observation for using the bundles to define prices in Lemma 33 is the following exchange property of optimal bundles.

Lemma 31. *If $b = (j, k)$ and $b' = (j', k')$ are optimum bundles of agent i (are in B_i), such that $p_j, p_{j'} < 1 < p_k, p_{k'}$, then $b'' = (j, k')$ and $b''' = (j', k)$ are also in B_i.*

Proof. We prove this lemma by using contradiction. W.l.o.g assume that $b_3 \notin B_i$. If in equilibrium i trade her bundles to get $y_{ib_1} > 0$ unit of b_1 and $y_{ib_2} > 0$ unit of b_2, then by lemma 30, her value will not change. Now, since $x_{ij} > 0$ and $x_{ik'} > 0$, similar to proof of lemma 29, we can rewrite her allocation so that it includes some of b_3, and similarly rewrite the remaining allocation of i with other bundles. Since the unit value of i for b_3 is less than value of i in equilibrium, this is a contradiction with corollary 30 and we are done.

■

4.2.1 Characterizing the Prices and Optimum Bundles with Polynomials

In this subsection we define a set of variables and polynomials that help us determine agent utilities and prices of all items at equilibrium. We consider assignments in the next subsection.

For each agent i, we define a variable u_i which is i's utility in the equilibrium. By Lemma 29, we know that any equilibrium allocates bundles to agents, where each agent only gets one unit of her optimum bundles B_i. Since we did not define variables for the prices yet, we cannot use prices to define bundles, so we start by defining a set of item bundles for each agent just based on the fact from Corollary 30 that optimum bundles must give value u_i.

Candidate Bundles. For each agent i, we define a *candidate bundle* to be the items whose value is u_i, or the pair of items j and k such that $v_{ij} < u_i < v_{ik}$, so there exists a unique $0 < \alpha^i_{jk} < 1$ such that $\alpha^i_{jk} v_{ij} + (1 - \alpha^i_{jk}) v_{ik} = u_i$. Note that the optimum bundles of i also satisfy this constraint (by Corollary 30). This means that the optimum bundles of an agent is a subset of her candidate bundles. In addition, the price of optimum bundles is exactly 1.

In order to find the set of candidate bundles of agent i, we define a polynomial $v_{ij} - u_i$ for each agent i and item j. This way, each cell tells us for each item j, whether $v_{ij} < u_i$, $v_{ij} = u_i$ or $v_{ij} > u_i$. For any two items $j, k \in [m]$, j and k form a candidate bundle if $v_{ij} < u_i$ and $v_{ik} > u_i$. Similarly if for an item j, if $u_i = v_{ij}$ then j alone forms a candidate bundle. By the information provided by each cell, for each agent $i \in [n]$ and item $j, k \in [m]$ that form a candidate bundle, we define the ratio for the candidate bundle to be $\alpha^i_{jk} = \frac{u_i - u_k}{u_j - u_k}$.

Not all the candidate bundles of agent i are in the set of her optimum bundles, since the price of optimum bundles should be exactly 1. We first observe that the unit price of a candidate bundle cannot be less than 1. This property of candidate bundles is useful in proof of lemma 33, allowing us to infer all prices from prices on only a few items.

Lemma 32. *In an equilibrium with prices* $\mathbf{p}$*, all candidate bundles have price at least* 1.

Proof. We prove this by contradiction. Assume that there exists an agent i such that the price of one her candidate bundles b is less than one. Assume that i's most preferred item is j. We have the following cases

- $b = (z) \neq (j)$: In this case, since i's most preferred item is unique, there exists an $\epsilon > 0$ such that if i gets $1 - \epsilon$ unit of item z and ϵ unit of item j, her total price is still less than 1, but her utility is more than v_{iz}. This is a contradiction since by the definition of candidate bundles, i's utility in equilibrium cannot be more than v_{iz}.
- $b = (j)$: In this case, we claim that $\mathbf{x}$ does not allocate any item to i other than j. Assume this is not true. If i trades whatever she gets in the equilibrium with j, her price will be less than 1, and since the maximum value for i is unique, her utility will increase. So it should be the case that i does not spend all her money and market will not be cleared. This is a contradiction with the assumption that the prices are market clearing.
- $b = (k, z)$: Since the maximum value item of i is unique, her value b should be less than v_{ij}. Now since the unit price of b is less than one, there exists $\epsilon > 0$ such that $\epsilon p_j + (1-\epsilon)p_b < 1$. However, if i gets ϵ from item j and $1-\epsilon$

from bundle b, his utility is more than v_{ib}. This is a contradiction with the definition of candidate bundles.

So all the possible cases reach a contradiction and we are done. ■

Next we wish to find the prices for all items. We will show that if one knows for each agent her utility, the price of only one item in her optimum bundles, and we use the set of candidate bundles defined above, we can find a the price of all the items which are in one of her optimum bundles.

Lemma 33. *Consider an equilibrium where we know for each agent i, the utility u_i of the agent, and the price of a single item j which is in a bundle of B_i with two items. Using these values, and the notion of candidate bundles defined above, we can find the price of all items in polynomial time.*

Proof. The key for finding the prices is the observation that if for a bundle $b = (j, k)$, we have α_b and p_j, then we can find p_k. This fact, combined with lemma 31, imply that if for each agent i, we know the price of one item in one of her optimum bundles, then we can find the price of all the items in her optimum bundles. The only problem is that we do not know which one of her candidate bundles is also one of her optimum bundles.

Assume that for each agent i, we know a good $g_i \in [m]$ is in one of her optimum bundles and we have a variable for g_i's price. To find a formula for price of other items with the variables we have defined so far, consider the following game. Assume that i is playing a game, in which she wants to maximize the number of her optimum bundles. She can participate in this game by proposing a price for all the items in her candidate bundles, knowing p_{g_i} and her candidate bundles. W.l.o.g assume that $p_{g_i} < 0$. For each of her candidate bundles $b = (g_i, j)$, i first proposes price $p_j = \frac{1-\alpha_b g_i}{1-\alpha_b}$ for item j, then for each of her candidate bundles $b' = (k, j)$, she proposes price $p_k = \frac{1-(1-\alpha_{b'})p_j}{\alpha_{b'}}$ for item k. Finally she proposes price 1 for all the items which form a bundle alone. By doing this, if any of her proposed prices is chosen, that item will be in her optimum bundles. Now, for each item j, we, the game coordinator, choose the maximum price for j among all the prices which were proposed by the agents for j and set that to be the price of item j.

Note that in equilibrium, we have to choose the maximum proposed price, since if we choose less than that, the agents with higher proposals will have candidate bundles whose price is less than 1, this is a contradiction with lemma 32. ∎

In order to use the above lemma, we should be able to do two things: (i) For each agent i, select an item j that is in one of i's optimum bundle with two items (if such an item exists), and set its price p_j. (ii) For each item j, find the maximum proposed price among all the proposed prices for that item, with proposed price of 1 of items in single item candidate bundles for any agent (including agents with no special item). We can do the first task by checking all possible assignments of the special items with defining $O(m^n)$ separate equilibrium structures for each possible selection of one special item assigned to each agent, and checking them separately. Since the number of agents is constant, the number of different equilibrium structures is polynomial. For each case, we use $O(n)$ variables, one for agent utility, and one for the price of the proposed special items for agents. To define prices of other items, we use candidate bundles for each agent to define candidate prices, add polynomial comparing the expressions for candidate prices. The actual price of the item is the highest of all prices as shown in the proof of Lemma 33, which is now set uniquely in each cell. Note that if an agent i is proposing a price for an item j which is higher than the proposed price of another agent i' and j is the special item of i', then this cell cannot contain equilibria.

Lemma 34. *Consider the space of at most $2n$ variables including agents' utilities and price of one item in each agent's optimum bundles. Now we add $O(mn^2)$ polynomials: comparing utilities to item values, and comparing candidate item prices using candidate budges, as defined in Lemma 33, the sign of these polynomials gives us a formula for the price of each item as well as the set of optimum bundles of each agent.*

4.2.2 Characterizing the Equilibria

In this section, we add a set of new variables and polynomials to the set of variables and polynomials defined in Section 4.2.1, in order to determine whether each cell contains equilibria. The new variables will help us define assignments. We cannot directly define a variable representing the allocation of all goods/bundles to agents, since the number of these is not constant.

The key idea is to show that for every equilibrium pricing, there is a specific allocation of items to agents where the number of items which is being shared

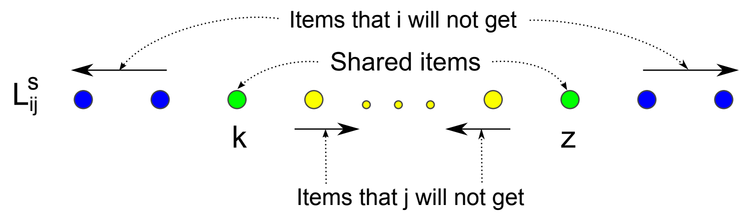

Figure 4.1: The nodes are the items in S (one side of G) that are in a optimum bundle of both i and j ($i < j$), which are sorted by their prices. The figure shows the items that i and j share (green nodes), the items that of i or j will not get in the special allocation (blue and yellow nodes).

between multiple agents is very small, and the allocation has a special structure. This helps us to significantly reduce the number of variables needed to define allocations. Consider two agents $i < j$, and all items of price $p_k < 1$ in sorted by price as shown by Figure 4.1. We will show that there is an equilibrium allocation with only two of these items shared between i and j, and the structure indicated by the figure, and the analogous structure for items of price $p_k > 1$. If we know the shared items, this structure helps us with finding the owner of the items which only have a single owner in the allocation, hence finding the allocation of each agent using allocation variables only for shared items. Before stating the properties of this allocation in Lemma 35, we have to define some notations.

Consider a bipartite graph $G = (S, T, E)$ in which vertices in S are the items with price less than 1 and vertices in T are the items with price more than 1. For simplicity, we sort items in each side of G by their price and break ties with items' indexes. Figure 4.1 is one side of this bipartite graph. For each bundle, we put an edge in the graph which connects the two items in the bundle. For every pair of agents i and j, let $B^S_{ij} = S \cap (V(B_i) \cap V(B_j))$ and $B^T_{ij} = T \cap (V(B_i) \cap$

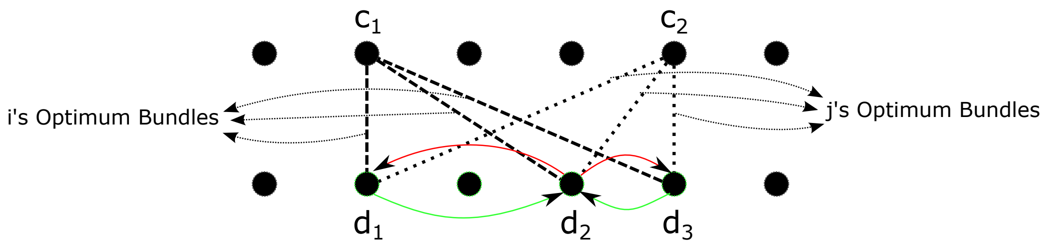

Figure 4.2: The arrows show the trades of items between agents i and j in the proof of Lemma 35.

$V(B_j))$. Furthermore, let H be the set of items whose price is 1 and let B_i^H be the optimum bundles of agent i which has exactly 1 item.

For each pair of agents i and j, let L_{ij}^S and L_{ij}^T be the list of items in B_{ij}^S and B_{ij}^T sorted by their price in increasing order (break ties with index of the items), respectively, the set of items of price above 1 (and below 1) that are part of optimum bundles for both i and j. Figure 4.2 is showing a part of the above bipartite graph for a pair of agents.

The following Lemma claims that there is an equilibrium allocation of a special structure suggested by Figure 4.1: on each side of the bipartite graph the agents share at most two items, and the agent with lower index only gets items between the two shared items, while the agent with higher index only gets items outside this interval as shown by Figure 4.1.

The main idea of the proof is that given any equilibrium allocation, we can make each pair of agents that violate the properties trade their items, as illustrated by the Figure 4.2. Note that the running time of this trading process is not important, since we only use it to show that for any equilibrium pricing an allocation with the desired structure exists.

Lemma 35. *For every equilibrium pricing* $\mathbf{p}$*, there exists an equilibrium allocation* $\mathbf{x}$ *of items to agents such that for every pair of agents* i *and* j *(*$1 \leq i < j \leq n$*)*

1. *There are at most* 2 *items that* X *is allocating to both* i *and* j *on each side of* G*. Furthermore, if* k *and* z *are two items in* S *(*T*) such that* $p_k \leq p_z$ *and* i *and* j *are sharing* k *and* z*,* i *only gets items between* k *and* z *in the order sorted by price, while* j *does not get any of the items from* B_{ij}^S *(*B_{ij}^T*) whose whose position in* L_{ij}^S *(*L_{ij}^T*) is between* k *and* z *in the order (see Figure 4.1).*
2. *There is at most one item* k *in* $B_i^H \cap B_j^H$ *that is shared by* i *and* j *in* $\mathbf{x}$*, and* i *only get items from* B_i^H *whose index is lower than* k *and* j *only gets items whose index is higher than* k*.*

Proof. Suppose that we have an allocation of bundles to agents $\mathbf{y}$ in an equilibrium. We want to reallocate these bundles so that it satisfies the conditions of the lemma.

Assume that there exist two agents $1 \leq i < j \leq n$, such that there are tree bundles $a_1 = (c_1, d_1)$, $a_2 = (c_1, d_2)$ and $a_3 = (c_1, d_3)$ such that $a_1, a_2, a_3 \in B_i$ and three bundles $b_1 = (c_2, d_1)$, $b_2 = (c_2, d_2)$ and $b_3 = (c_2, d_3)$ such that $b_1, b_2, b_3 \in B_j$. Suppose that $p_{d_1} \leq p_{d_2} \leq p_{d_3}$ and $Y\mathbf{y}$ is allocating $y_{id_1}, y_{id_3} > 0$ of a_1 and a_3 to agent i, and $y_{jd_2} > 0$ of d_2 to agent j. Furthermore, assume that $y_{ic_1}, y_{jc_2} > 0$.

Since $p_{d_1} \leq p_{d_2} \leq p_{d_3}$, there is $0 < \beta < 1$ such that $\beta p_{d_1} + (1-\beta)p_{d_3} = p_{d_2}$. So if we remove βz from y_{id_1} and $(1-\beta)z$ from y_{id_3} and add z to y_{id_2} then the total cost of agent i will be the same. Similarly, if we add βz to y_{jd_1} and $(1-\beta)z$ to y_{jd_3} and remove z from y_{jd_2}, then the total cost of j will not change. Furthermore, it

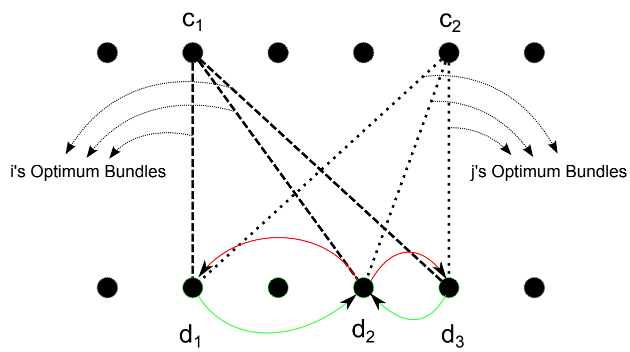

Figure 4.3: Red and green arrows shows the trades of items between agent i and j in proof of lemma 35.

is easy to see that doing this does not affect the matching constraints. Figure 4.3 demonstrates this procedure.

Now, we have to show that doing this does not change the utility of i and j. From the assumption that $Y\mathbf{y}$ is an equilibrium allocation, it follows that by doing this, the utility of i and j cannot increase. Assume that by doing this utility of agent i decreases. We have

$$\begin{aligned}u_i &= \alpha_{a_2} v_{id_2} + (1-\alpha_{a_2}) v_{ic_1} \\ &< \alpha_{a_2}(\beta v_{id_1} + (1-\beta) v_{id_3}) + (1-\alpha_{a_2}) v_{ic_1} \\ &= \beta\alpha_{a_2} v_{id_1} + (1-\beta)\alpha_{a_2} v_{id_3} + (1-\alpha_{a_2}) v_{ic_1}\end{aligned}$$

similarly we have

$$\beta\alpha_{a_2} p_{d_1} + (1-\beta)\alpha_{a_2} p_{d_3} + (1-\alpha_{a_2}) p_{c_1} = 1$$

This means that if agent i only use these three items with these ratios, he gets one unit of item, spends at most 1 unit of money and her utility will be more than u_i. This is a contradiction. The argument for agent j is similar to the argument for agent i.

For each pair of agents $1 \leq i < j \leq n$, we find the cheapest item k and the most expensive item z in S that i owns, if there is an item r in between k and z in the ordered list of items in S that j owns, we switch the allocation we switch the ownership of these items until, one of the following cases happen

- j runs out of item r.
- i runs out of item k.
- i runs out of item z

Let Φ be a potential function which is equal to $|S|$ minus the the position of highest positioned item that i owns in the ordered list, plus the position of the lowest positioned item that i owns in the ordered list, plus the number of items for which j is a shareholder and are between the the two in the ordered list. If we repeat the above procedure, each time this potential function will decrease by 1. Since the potential function is always non-negative, we cannot continue the above procedure forever. This means that at after some iterations, we will reach an equilibrium allocation which satisfies the first condition for this pair of agents.

We repeat this procedure until such pair of agents and set of items with these properties do not exist. We also do the same to the allocation of items on the other side of G (T) to agents. For the items in $B_i^H \cap B_j^H$, one can also transfer items between i and j to satisfy the second condition by finding two items that violate the condition and switch their ownership. The procedure defined for satisfying these conditions are an easier version of the previous procedure. ■

From now on, we focus on finding and characterizing the specific equilibrium allocation which is guaranteed to exist by this lemma. By using Lemma 35, for every pair of agents $i, j \in [n]$ we can use 5 item indexes f_{ij}^S, r_{ij}^S, f_{ij}^T, r_{ij}^T, and h_{ij} that tell us which items are shared by agent i and j. We can assume that we know what are these shared items by simply checking all the possible $O(m^{5n^2})$ of these combinations. At the start of the algorithm, we fix the 5 shared items for each pair of agents and an item in optimum bundle of each agent (which we discussed in Section 4.2.1). We call this set of items associated with each agent and pairs of agents the *structure* of the equilibrium. For each such a structure, we will aim to decide if there is an equilibrium with the given structure.

When considering equilibria of a given structure, we define variables for the allocation of the at most $5n^2$ shared items, but do not define variable for allocations of other items. Next we show that given the allocations of the shared items, we can (i) infer allocations of all other items using the structure of Lemma 35, and (ii) can also define polynomials whose sign will tell us if there is an equilibrium with the given structure and allocation of shared items.

Since all the items should be sold in the equilibrium, for each agent i, all the items that are in only in i's optimum bundles, should get allocated to i. Second, we have defined a variable that indicates the share of each agent for the shared items. The only thing left is to consider items that are in the set of the optimum bundles of multiple agents, but are not shared by these agents.

Lemma 35 helps us find allocation of this set of items. We start with the agent with the lowest index (agent 1) and one side of G, say S. The items that agent 1 gets, are the ones that satisfy all the constraints given to us by second part of the lemma. In order to check whether for an item j all these constraints hold, for each agent i, we can look at the position of item j in L^S_{1i}. If j's position is between position of f^S_{1i} and r^S_{1i} in the list, for all the choices of $i > 1$, then it satisfies all the constraints. By lemma 35, we know this is a necessary and sufficient condition for j to get allocated to agent 1. We do this for S and T separately, and remove the items that get allocated to agent 1. Now, we repeat this procedure for agent 2, but only check the constraints for agents $i > 2$, then remove the items that agent 2 gets. We repeat this for rest of agents based on their indexes. We can do the same procedure for items in H to find which agent is getting what item. Now, we are ready to exactly specify what are the necessary and sufficient conditions for the prices in each cell to be the equilibrium prices. This process is summarized in the following lemma.

Lemma 36. *Consider a structure of special items for agents, and shared items between agents (as defined after Lemma 35). Now consider a cell in the space of variables of agent utilities, prices of the special items, and allocation shares of the shared items, defined by the polynomials that help define prices of all items. The prices and allocation of this cell forms an equilibrium, if and only if the allocation defined above satisfies the following constraints*

1. *All the items get fully allocated to agents.*
2. *Each agent gets exactly one unit of items.*
3. *For each agent i, the total cost of buying the items allocated to i is exactly 1.*
4. *For each agent i, each of the items allocated to i is in one of her optimum bundles.*

Proof. $\Rightarrow$ The first, second and third condition follow from the market clearing conditions. The fourth condition directly follows from part c of lemma 29.

$\Leftarrow$ From the first and second condition, we know that the pricing and allocation are market clearing. From the second, third and fourth condition, and the argument in proof of part c of lemma 29 we know that the allocation of items to each agent can be rewritten as her optimum bundles to her. So, from the definition of optimum bundles we know that the allocation of items to agents is optimal. Therefore, the prices and allocation are in equilibrium state. ■

The final thing we need to do is to define a set of polynomials for checking the above conditions. The first conditions holds for an item j, if summing up the share of each agent for that item, the sum is equal to C_j. This can be handled by adding one polynomial for each item j.

The second condition holds for an agent i, if when we sum all the items (including the proportion of the shared items) that i gets, this sum is exactly 1. So we can also check this condition by adding a polynomial for each agent. Note that we can do this since we explained how to find out what is the allocation of items to agents for this cell.

The third condition holds by Lemma 36 if multiplying the share of each agent for an item by its price and summing over all items, we get 1 for all the agents. We can do this by defining a polynomial for each agent and checking its sign. The fourth condition is guaranteed to hold by definition. The following theorem summarizes how we equilibria are cells of the constraints discussed throughout this section.

Theorem 37. *Consider an equilibrium structure, and the space of the $O(n)$ variables for agents' utilities and price of one special item in each agent's optimum bundle, the $O(n^2)$ variables for allocation of shared items between the agent. Divide this space into cells by the signs of the polynomials defined in the previous section, along with the $O(m^2 + n)$ polynomials defined just above, for checking the existence of the special equilibrium allocation. The sign of these polynomials fully determines either that the vectors in this cell can be extended to form equilibria.*

Using this structure Theorem allows us to prove Theorem 28.

Proof of Theorem 28. We use theorem 27 as the base of our algorithm. We start by fixing the structure of the equilibrium, selecting a single item from the optimum bundles of each agent, and selecting fixing 5 items shared for every pair of agents. We check all the possible combinations, at most $O(m^{n+5n^2})$ options, which is polynomial in m for fixed n.

For a given equilibrium structure, we use the $O(n)$ variables, the agent utilities, and $O(mn)$ polynomials defined in section 4.2.1 to find a set of candidate bundles for each agent. Then we use an additional $O(n)$ variables, the prices of the special items for each agent, and $O(mn^2)$ polynomials in order to find a formula for the prices and the set of optimum bundles of each agent in each cell. Finally, we use the last set of $O(n^2)$ variables, the assignments of shared items, and $O(m^2 + n)$ additional polynomials, in order to check whether the set of prices in the feasible cell are equilibrium prices with the given structure. The degree of the defined polynomials is polynomial. We check all the non-empty cells of the resulting system, taking time polynomial on m for any fixed n by

Theorem 27. Since the equilibrium exists, it should be in one of the non-empty cells.

Finally, if the prices of the cell are equilibrium prices, we take any vector from that cell, and extend it to get an equilibrium pricing and allocation. After we have the equilibrium prices $\mathbf{p}$, we can also find each agent's allocation by finding a solution of the following set of inequalities.

$$\sum_{b \in B_i} x_{ib} = 1 \qquad \forall i \in [n]$$

$$\sum_{i} \Big(\sum_{b:b=(j,k)} \alpha_b x_{ib} + \sum_{b:b=(k,j)} (1-\alpha_b) x_{ib} + \sum_{b:b=(j)} x_{ib} \Big) = C_j \qquad \forall j \in [m]$$

$$x_{ib} \geq 0 \qquad \forall i \in [n], b \in B$$

In which, x_{bi} is the amount of bundle b used by agent i. Note that since we know this bundling is associated with an equilibrium, the feasible region of the above inequalities is not empty. Finally, the allocation of each agent i for item j in this equilibrium is

$$x_{ij} = \sum_{b:b=(j,k)} \alpha_b x_{ib} + \sum_{b:b=(k,j)} (1-\alpha_b) x_{ib} + \sum_{b:b=(j)} x_{ib}$$

4.2.3 Relaxing Full Budget Spent Assumption of Agents in Equilibria

So far, we have assumed each agent fully spend her budget and introduced an algorithm for finding an equilibrium with this assumption. While equilibria where all agents spend their budget do exists (as shown in Appendix C.3), there may be equilibria where not all agents spend their budgets. In this section, we describe how the same algorithm can be modified to find an equilibrium which does not require this assumption. The following lemma shows that if an agent does not fully spend her budget in an equilibrium, she will get exactly one unit of her most desired item, hence she cannot get any part of the other items. Recall that we assume that for each agent, her most preferred item is unique.

Lemma 38. *In an equilibrium, with price vector* $\mathbf{p}$ *and equilibrium allocation* $\mathbf{x}$*, for each agent* i *whose most preferred item is* j*, if* i *does not fully spend her one unit of money then* $x_{ij} = 1$*.*

Proof. We use contradiction. Assume i is not fully spend her budget in an equilibrium, but there exists an item z such that $x_{iz} > 0$ and z is not her most preferred item. Let b'_i be i's leftover money and j be her most preferred item. If $p_z \geq p_j$ then i can trade her share of item z for the same amount of item j and increase her utility, which contradicts the equilibrium condition. If $p_z < p_j$ then, since i has b'_i unit of extra money, i can afford to trade $\min(\frac{p_j - p_z}{b'_i}, x_{iz}) > 0$ unit of item z with the same amount of item j and increase her utility. This contradicts the fact that in equilibrium, utility of each agent should be optimum and we are done. ∎

Since the number of agents is fixed, we can iterate over all the 2^n settings each of which tells us whether each agent fully spend her budget or not. In each setting, before we start the algorithm, for each agent that does not fully spent her budget, we remove one unit from her most preferred item. From lemma 38 we know that if this setting leads to an equilibrium, this will be the only valid allocation of items to these agents. We also remove these agents themselves and continue the algorithm assuming all the remaining agents fully spend their budget as before. If in a setting, enough of the most preferred items of the agents who do not fully spend their budget is not available, then it simply means that the setting cannot lead to an equilibrium, so we ignore the setting and iterate to the next one.

Note that if an agent is not fully spending her budget, from lemma 38 we know that the price of her most preferred item should be at most 1. Therefore, in our algorithm after formulating the prices in the algorithm (described in section 4.2.1), for each item that was partially allocated to the agents who do not fully spent their budget, we add a constraint that its price should be at most 1. Then we continue with the algorithm to find the equilibrium given the initial setting. Because such equilibrium always exists, at least one of the initial settings will lead to finding an equilibrium price vector and allocation.

4.3 Fixed Number of Goods

In this section, we give a polynomial time algorithm which finds the exact value of equilibrium prices and allocations when the number of goods is constant m and the number of agents is an arbitrary number $n \in \mathbb{N}$. In order to find an equilibrium, our algorithm uses a cell decomposition technique which uses the algorithm proposed in [BPR98].

The goal of this section is to prove the following theorem by providing a polynomial time algorithm.

Theorem 39. *Finding an equilibrium of a market with fixed number of goods in which agents have additive values, one unit of money and matching constraints can be done in polynomial time.*

4.3.1 Characterizing the Bundles with Polynomials

In this section, we define a set of variables, which represent the prices, and polynomials so that the sign of polynomials in each cell determines the bundling and the optimum bundles of each agent for the set of prices in that cell.

First, for each item j we define a variable p_j to represent its price. We define m^2 polynomials $p_j - p_k$ for all $j, k \in [m]$, p_j for all $j \in [m]$ and $p_j - 1$ for all $j \in [m]$ that give us the order of items' prices, check whether price is non-negative and whether p_j is less than, greater than or equal to 1, respectively. Recall lemma 29.

Instead of finding equilibrium allocation of items to agents, we focus on finding the allocation of bundles to agents.

By lemma 29 we know that each bundle should have price 1. It means that if items j and k are in a bundle b, then $p_b = \alpha_b p_j + (1 - \alpha_b)p_k = 1$. We can rewrite this equation to get α_b from p_j and p_k, i.e. $\alpha_b = \frac{1-p_k}{p_j-p_k}$. So for every $j, k \in [m]$, we define $q_{jk} = \frac{1-p_k}{p_j-p_k}$, and when $0 \leq q_{jk} \leq 1$, it means j and k form a valid bundle. Note that we don't put j and k in a bundle unless $p_j \neq p_k$. Since in each cell we know the order of prices, we can check whether $p_j = p_k$, $p_j > p_k$ or $p_j < p_k$, then for the last two cases, check whether $0 \leq 1 - p_k \leq p_j - p_k$ and $0 \geq 1 - p_k \geq p_j - p_k$ respectively. The only thing left now is to check the bundles with only one item. Since an item forms a bundle if and only if its price is exactly 1, we can do this by checking sign of $p_j - 1$ for all the items. It is easy to see that we can find whether each of these conditions hold by defining $O(m^2)$ polynomials and checking their sign in each feasible cell.

Now, each output cell of the cell decomposition algorithm gives us the valid bundling associated with the signs of the polynomials in that cell. We also wish to find the optimum bundles for each agent. Note that for a valid bundle $b = (j, k)$, $q_{jk} = \alpha_b$. By corollary 30, we know that value of each agent should be maximum for all of her optimum bundles. So given that for each bundle b, we can write α_b in terms of the prices, we can also write the value v_{ib} of each agent i for each bundle b. So to find out which of the valid bundles is optimum, we can check the sign of $v_{ib} - v_{ib'}$ for all $i \in [n]$ and $b, b' \in B$, and order the value of agents for bundles. This can be done by defining $O(nm^4)$ number of polynomials. By adding these polynomials, each cell also tells us what is the order of each agent over all the bundles, hence we can find the optimal bundles of each agent.

We can sum up the above arguments in the following lemma.

Lemma 40. *For each cell, defined by the above $O(nm^4)$ polynomials and $O(m)$ variables, we can use the signs of the defined polynomials in order to find the bundles and the optimum bundles of each agent for the set of prices in that cell.*

4.3.2 Characterizing the Equilibria

In the previous section we described how to define polynomial number of polynomials with constant number of variables (prices) such that each cell specifies all the valid bundles and the optimum bundles of each agent. We also showed how we can use these variables to get formulate α_b (q_{jk}) for each bundle $b = (j, k)$. In this section we describe how we can define a set of polynomials whose sign tells us if a given bundling satisfies the equilibrium conditions.

One difficulty is that if we want our algorithm to run in polynomial time, we cannot directly define a variable x_{ij} for each agent i and each item j, since the number of variables needs to remain constant. In order to get around this problem, we define a variable for each bundle instead that shows how much this bundle is used in equilibrium, and use these variables to check equilibrium conditions.

Assume that we have an equilibrium $E(\mathbf{x}, \mathbf{p})$ and its valid bundling B. Let

$\mathbf{x}_b$ be the total amount of bundle $b \in B$, which is used in this equilibrium. For a subset of agents S, let B_S be the union of all the optimum bundles of agents in S, i.e. $B_S = \bigcup_{i \in S} B_i$. In order to characterize equilibrium bundles with variables x_b and polynomials, we use the following lemma.

Lemma 41. *A pricing and its corresponding bundling is an equilibrium if and only if there exists $\mathbf{x} \in \mathbb{R}_{\geq 0}^{|B|}$ such the following hold*

a. For every subset of agents S, we have $\sum_{b \in B_S} \mathbf{x}_b \geq |S|$.

b. For every item j, $\sum_{b:b=(j,k)} \alpha_b \mathbf{x}_b + \sum_{b:b=(k,j)} (1-\alpha_b)\mathbf{x}_b + \sum_{b:b=(j)} \mathbf{x}_b = C_j$.

Proof. The proof of this lemma is similar to proof of the Hall's theorem. On one side of the bipartite graph we have agents with capacity 1 and on the other side we have the bundles with capacity x_b (for each bundle $b \in B$). Furthermore, there is an edge between bundle b and agent i iff $b \in B_i$.

$\Rightarrow$ This case directly follows from lemma 29 and the market clearing conditions.

$\Leftarrow$ Assume that there is an $X \in R^{|B|}$ such that a and b hold. We have to show that there exists an allocation of bundles to agents such that every agent gets exactly 1 unit of her optimum bundles. We use contradiction. Assume that this allocation does not exists. Let Y be an allocation which allocates maximum total amount of bundles to agents among all the valid allocations. Let i be an agent for which Y allocates less than 1 unit of her optimum bundles. If there is an augmenting path from i to a bundle b which is not fully consumed, then we can force the agents along this path to deviate and increase the allocation of i without changing the total allocation, cost and utility of other agents. This is a contradiction. Now assume that an augmenting path from i does not exist. Let S be the set of agents and T be the set of bundles which are reachable from i. Since all the bundles in T have been fully consumed by agents in S, and each agent in S has at most 1 unit of bundles in T, we have

$$\sum_{b \in B_S} x_b < |S|$$

This is a contradiction, and we are done. ∎

The problem with lemma 41 is that in order to check whether the first set of conditions hold, for every subset of agents we have to define a polynomial and there are exponentially many of these subsets. For a subset of bundles T, let $A(T)$ be the set of agents that for every agent i, $i \in A(T)$ if and only if $B_i \subseteq T$. In order to fix the problem, we prove the following lemma.

Lemma 42. *For every subset of agents S, we have $\sum_{b \in B_S} x_b \geq |S|$ if and only if for every subset of bundles T, $\sum_{b \in T} x_b \geq |A(T)|$.*

Proof. $\Rightarrow$ Let $S = A(T)$. By definition we know that $B_S \subseteq T$. So we have

$$|A(T)| = |S| \leq \sum_{b \in B_S} \mathbf{x}_b \leq \sum_{b \in T} \mathbf{x}_b$$

$\Leftarrow$ Let $T = B_S$. By definition we know that $S \subseteq A(T)$. So we have

$$|S| \leq |A(T)| \leq \sum_{b \in T} x_b = \sum_{b \in B_S} \mathbf{x}_b$$

∎

Note that since the number of goods is constant, the number of subsets of bundles is also constant. Lemma 42 shows that when we know the bundles and optimum bundles of all agents, we can find out whether the the first condition of lemma 41 holds by defining a variable for each bundle and checking polynomial number of inequalities. It is also easy to see that we can check the second set of conditions by defining a polynomial (whose degree is a function of number of bundles) for each agent, using the formula from the previous section for each α_b.

We summarize the arguments in this section in the following lemma

Lemma 43. *For each output cell of theorem 27, defined by variables and polynomials in the previous section and above $O(m^2)$ variables and $O(2^{m^2})$ polynomials, we can use the sign of the defined polynomials to see if the prices in that cell are equilibrium prices.*

Proof. The proof of this lemma is similar to proof of the Hall's theorem. On one side of the bipartite graph we have agents with capacity 1 and on the other side we have the bundles with capacity $\mathbf{x}_b$ (for each bundle $b \in B$). Furthermore, there is an edge between bundle b and agent i iff $b \in B_i$.

$\Rightarrow$ This case directly follows from lemma 29 and the market clearing conditions.

$\Leftarrow$ Assume that there is an $\mathbf{x} \in \mathbb{R}^{|B|}_{\geq 0}$ such that a and b hold. We have to show that there exists an allocation of bundles to agents such that every agent gets exactly 1 unit of her optimum bundles. We use contradiction. Assume that this allocation does not exists. Let Y be an allocation which allocates maximum total amount of bundles to agents among all the valid allocations. Let i be an agent for which Y allocates less than 1 unit of her optimum bundles. If there is an augmenting path from i to a bundle b which is not fully consumed, then we can force the agents along this path to deviate and increase the allocation of i without changing the total allocation, cost and utility of other agents. This is a contradiction. Now assume that an augmenting path from i does not exist. Let S be the set of agents and T be the set of bundles which are reachable from i. Since all the bundles in T have been fully consumed by agents in S, and each agent in S has at most 1 unit of bundles in T, we have

$$\sum_{b \in B_S} \mathbf{x}_b < |S|$$

This is a contradiction, and we are done. ■

4.3.3 Finding an Equilibrium

So far, we have characterized the equilibria by defining constant number of variables and polynomial number of polynomials. Now we put the pieces together to prove theorem 39.

Proof of Theorem 39. We use theorem 27 with variables and polynomials defined in section 4.3.1 and 4.3.2. Theorem 39 iterates all the feasible cells of these polynomials. Each cell gives us the sign of all these polynomials. The signs of these polynomials tell us

1. What are the bundles and optimum bundles for each agent, by lemma 40.

2. Whether these bundles are associated with a market equilibrium, by lemma 43.

Finally, if a bundles associated with a cell is characterizing an equilibrium bundling, we sample a set of prices from that cell. Now that we know equilibrium prices, it is easy to find the bundles as before. We can also find each agent's allocation by finding a solution of the following set of inequalities

$$
\begin{aligned}
&\sum_{b \in B_i} x_{ib} = 1 && \forall i \in [n] \\
&\sum_{i} \Big(\sum_{b:b=(j,k)} \alpha_b x_{ib} + \sum_{b:b=(k,j)} (1-\alpha_b) x_{ib} + \sum_{b:b=(j)} x_{ib} \Big) = C_j && \forall j \in [m] \\
&x_{ib} \geq 0 && \forall i \in [n], b \in B
\end{aligned}
$$

In which, x_{bi} is the amount of bundle b used by agent i. Note that since we know this bundling is associated with an equilibrium, the feasible region of the above inequalities is not empty. Finally, the allocation of each agent i for item j in this equilibrium is

$$x_{ij} = \sum_{b:b=(j,k)} \alpha_b x_{ib} + \sum_{b:b=(k,j)} (1-\alpha_b) x_{ib} + \sum_{b:b=(j)} x_{ib}$$

Note that since theorem 27 iterates over all the possible feasible cells, if the equilibrium exists, our method finds it. Furthermore, our method can characterize all the equilibrium cells.

Since the number of variables is constant ($O(m^2)$), the number of polynomials is polynomial in the number of agents ($O(2^{m^2}+nm^4)$) and all the polynomials that we define have constant degree ($O(m^2)$), our algorithm runs in polynomial time.

APPENDIX A

MISSING PROOFS FROM CHAPTER 2

A.1 Deferred Proof of Theorem 2.2.1

Recall Lemma 3 which states that for every fixed $\gamma \in (0,1]$, if $S_{k-1} = $ GREEDY-TM(γ, A, B) then

$$(1+\frac{1}{\gamma})v(S_{k-1}) + \frac{1}{\gamma}v(i^*) \geq opt(A)$$

Proof. Let S^* be the optimum. By monotonicity and submodularity of $v(.)$, we have

$$v(S^*) - v(S_{k-1}) \leq v(S^* \cup S_{k-1}) - v(S_{k-1}) \leq \sum_{i \in S^* \setminus S_{k-1}} m_i(S_{k-1}) = \sum_{i \in S^* \setminus S_{k-1}} c_i \frac{m_i(S_{k-1})}{c_i}$$

By using the fact that $k \in \arg\max_{i \in A \setminus S_{k-1}} \frac{m_i(S_{k-1})}{c_i}$, we get $\frac{m_i(S_{k-1})}{c_i} \leq \frac{m_k(S_{k-1})}{c_k}$. Since k is not in the winning set $\frac{c_k}{m_k(S_{k-1})} > \gamma \frac{B}{v(S_k)}$. Using these we get

$$\sum_{i \in S^* \setminus S_{k-1}} c_i \frac{m_i(S_{k-1})}{c_i} \leq c(S^* \setminus S_{k-1}) \frac{m_k(S_{k-1})}{c_k} < B(\frac{1}{\gamma}\frac{v(S_k)}{B}) = \frac{1}{\gamma} v(S_k)$$

Finally, by definition of $v(S_k)$ and using submodularity, we have

$$\frac{1}{\gamma}v(S_k) \leq \frac{1}{\gamma}(v(S_{k-1}) + v(k)) \leq \frac{1}{\gamma}(v(S_{k-1}) + v(i^*))$$

By putting all the above together and rearranging the terms, we get the desired inequality. ∎

Recall Theorem which states that for every fixed $\gamma \in (0,1]$, RANDOM-TM(γ, A, B) is universally truthful, and has approximation ratio of $1 + \frac{2}{\gamma}$.

Proof. By Lemma 2 the mechanism is monotone and hence universally truthful. To prove the approximation ratio, let S be the outcome of the mechanism. The mechanism chooses $S_{k-1} = \textsc{Greedy}\text{-TM}(\gamma, A, B)$ with probability $\frac{\gamma+1}{\gamma+2}$ and i^* with probability $1 - \frac{\gamma+1}{\gamma+2} = \frac{1}{\gamma+2}$. So we have

$$\begin{aligned} E[v(S)] &= \frac{\gamma+1}{\gamma+2} v(S_{k-1}) + \frac{1}{\gamma+2} v(i^*) \\ \Rightarrow (1+\frac{2}{\gamma}) E[v(S)] &= \frac{\gamma+1}{\gamma} v(S_{k-1}) + \frac{1}{\gamma} v(i^*) \end{aligned}$$

So by using Lemma 3 we have

$$(1+\frac{2}{\gamma}) E[v(S)] \geq opt(A)$$

■

[CGL11] show that $\textsc{Random}\text{-TM}(0.5, A, B)$ and $\textsc{Greedy}\text{-TM}(0.5, A, B)$ are budget feasible. Here we include a simpler proof for completeness.

Proof. In order to prove $\textsc{Random}\text{-TM}(0.5, A, B)$ is budget feasible, we only have to show $\textsc{Greedy}\text{-TM}(0.5, A, B)$ is budget feasible, since when i^* is selected, his threshold payment is B and the mechanism is budget feasible.

Let p_i be the threshold payment for agent i. Let $S_{k-1} = \textsc{Greedy}\text{-TM}(0.5, A, B)$ the set returned by the greedy threshold mechanism. For every $i \in S_{k-1}$, we show that if i deviates to bidding a cost $b_i > m_i(S_{i-1}) \frac{B}{v(S_{k-1})}$, then he would not be selected. This will imply that the threshold payment p_i for player i is at most $p_i \leq m_i(S_{i-1}) \frac{B}{v(S_{k-1})}$, and so we get $\sum_{i \in S_{k-1}} p_i \leq \sum_{i \in S_{k-1}} m_i(S_{i-1}) \frac{B}{v(S_{k-1})} = B$, so the mechanism is budget feasible.

Consider the run of GREEDY-TM where i deviated to bidding $b_i > m_i(S_{i-1})\frac{B}{v(S_{k-1})}$ and is selected, while all other players bid truthfully. Let b denote the resulting cost-vector. Note that by bidding higher, player i would occur later in the order. Let j be the step in which i would occur, after he deviates to bidding cost b_i. Let S'_j be the items that are in the winning set at the end of this step (S'_z for $z \in [n]$ is defined similar to S_z but with cost vector b instead of c, with the change in one cost also effecting the order of items after item i). If i is in the winning set, we have

$$\frac{b_i}{m_i(S'_{j-1})} \leq 0.5\frac{B}{v(S'_{j-1})}$$

Since the items are sorted by their marginal bang per buck, for every $z \in S_{k-1}$, $c_z \leq 0.5B\frac{m_z(S_{z-1})}{v(S_{k-1})}$, so we have $c(S_{k-1}) \leq 0.5B$. Let $T = S_{k-1} \setminus S'_j = \{t_1, t_2, \ldots, t_q\}$, $T_0 = \emptyset$ and $T_z = \{t_l : l \in [z]\}$. So we have

$$v(S_{k-1}) - v(S'_j) \leq v(S_{k-1} \cup S'_j) - v(S'_j) = \sum_{z\in[q]} m_{t_z}(S'_j \cup T_{z-1}) = \sum_{z\in[q]} c_{t_z}\frac{m_{t_z}(S'_j \cup T_{z-1})}{c_{t_z}}$$

By submodularity we have

$$\sum_{z\in[q]} c_{t_z}\frac{m_{t_z}(S'_j \cup T_{z-1})}{c_{t_z}} \leq \sum_{z\in[q]} c_{t_z}\frac{m_{t_z}(S'_{j-1})}{c_{t_z}}$$

Since i is selected at step j it means it has the highest marginal bang per buck at that step. The cost vectors b and c are also only different in the cost of i. So we have

$$\sum_{z\in[q]} c_{t_z}\frac{m_{t_z}(S'_{j-1})}{c_{t_z}} \leq \sum_{z\in[q]} c_{t_z}\frac{m_i(S'_{j-1})}{b_i} \leq \frac{m_i(S'_{j-1})}{b_i}\sum_{z\in[q]} c_{t_z}$$

Since i has increased his cost, he cannot be selected before step i, so $S_i \subseteq S'_j$ and $v(S_i) \leq v(S'_j)$. By using this and contradiction assumption we have

$$\frac{m_i(S'_{j-1})}{b_i}\sum_{z\in[q]} c_{t_z} < \frac{v(S_{k-1})}{B}\frac{B}{2} = \frac{v(S_{k-1})}{2}$$

So $v(S_{k-1}) \leq 2v(S'_j)$. By adding this to the previous inequality and replacing $m_i(S'_{j-1})$ with $m_i(S_{i-1})$ (note that we can do this since $S_{i-1} \subseteq S'_{j-1}$), we get to a contradiction. ■

Our Approximation Analysis for RANDOM-TM$(0.5, A, B)$ is tight. To see this, consider the following example (with additive valuation): assume we have 5 items numbered from 1 to 5 with budget 4. Let $v_1 = 1$ and $c_1 = 0$, and let $v_i = 1 - \epsilon$ and $c_i = 1$ for $2 \leq i \leq 5$. The mechanism chooses item 1, however since $1 > \frac{4}{2}\frac{1-\epsilon}{2-\epsilon}$, none of the other items is in the winning set of GREEDY-TM$(0.5, A, 4)$. So the value of the GREEDY-TM$(0.5, A, 4)$ as well as the value of RANDOM-TM$(0.5, A, 4)$ is 1. However, optimum can select all of the items and get the value $5 - 4\epsilon$. Since ϵ can be arbitrarily small, RANDOM-TM$(0.5, A, B)$ is at most a 5-approximation.

A.2 Deferred Proofs from Section 2.2.2

Recall Lemma 12 which states that by using threshold payments, GREEDY-OM$(0.5, A, B)$ is budget feasible.

Proof. Let p_i be the threshold payment for agent i. Let S =GREEDY-OM$(0.5, A, B)$. For every $i \in S$, we show that if i deviates to a cost $b_i > m_i(S_{i-1})\frac{B}{v(S)}$, he cannot be selected. By proving this and by using the definition of threshold payments we get $\sum_{i \in S} p_i \leq \sum_{i \in S} m_i(S_{i-1})B/v(S) \leq \sum_{i \in S} m_i(S_{i-1} \cap$

$S)B/v(S) = B$, so the mechanism is budget feasible.

Assume that i deviates to $b_i > m_i(S_{i-1})\frac{B}{v(S)}$ and is still in the winning set. Let b be the resulting costs, with i's cost as b_i and for all other players $b_j = c_j$. Note that this change in cost for i changes the order in which the mechanism considers the items. Let j be the step that the mechanism considers item i with costs b. For $z \in [n]$, S'_z is defined similar to S_z but with cost vector b instead of c. [1]

$$c(S^*) \geq b(S^* \cup S'_j) - b(S'_j) = \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1}) \frac{b_{t_z}}{m_{t_z}(S'_j \cup T_{z-1})}$$

By submodularity we have $m_{t_z}(S'_j) \geq m_{t_z}(S'_j \cup T_{z-1})$, and by the fact that i was in position j in the ordering considered, we have that $b_{t_z}/m_{t_z}(S'_{j-1}) \geq b_i/m_i(S'_{j-1})$. Using these we get:

$$\begin{aligned}\sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1}) \frac{b_{t_z}}{m_{t_z}(S'_j \cup T_{z-1})} &\geq \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1}) \frac{b_{t_z}}{m_{t_z}(S'_{j-1})} \geq \sum_{z \in [q]} m_{t_z}(S'_j \cup T_{z-1}) \frac{b_i}{m_i(S'_{j-1})} \\ &= \frac{b_i}{m_i(S'_{j-1})}(v(S^* \cup S'_j) - v(S'_j)) \geq \frac{b_i}{m_i(S'_{j-1})}(v(S^*) - v(S'_j))\end{aligned}$$

Now by using fact that $S_{i-1} \subseteq S'_{j-1}$, as the first $i-1$ steps of the algorithm are not effected by i's change of bid, and by the assumption about b_i, we have

$$\frac{b_i}{m_i(S'_{j-1})}(v(S^*) - v(S'_j)) > B\frac{v(S^*) - v(S'_j)}{v(S)}$$

So by combining the above inequalities we have that the total cost $c(S^*)$ of set S^* is $c(S^*) > B(v(S^*) - v(S'_j))/v(S)$.

Since i was selected after deviating, we have $v(S'_j) \leq \alpha\textsc{Oracle}(A \setminus \{a_i\}, B) \leq \alpha v(S^*)$. Let $k' \in [n]$ be the minimum integer such that $S \subseteq S_{k'}$.

[1] Note that in GREEDY-TM, as well as GREEDY-EOM, S'_j would be exactly the winning set of the mechanism at the end of step j. However, in GREEDY-OM, S'_j may be a super set of the set of items that has been added to the winning set at the end of step j.

Since k' was chosen, we have $v(S) \leq v(S_{k'}) \leq \alpha Oracle(A \setminus \{a_{k'}\}, B) \leq \alpha v(S^*)$. Since $\alpha = 0.5$ we get $C(S^*) > B$ which is contradiction. ∎

Recall Theorem 14 which states that RANDOM-OM(α, A, B) is truthful and in expectation achieves $1 + \frac{2r}{\alpha}$ of the optimum, assuming the oracle used is an r-approximation.

Proof. Let S be the outcome of the mechanism. The mechanism chooses $S =$ GREEDY-OM(γ, A, B) with probability $\frac{r}{\alpha+2r}$ and i^* with probability $1 - \frac{r}{\alpha+2r} = \frac{\alpha+r}{\alpha+2r}$. Let $k \in [n]$ be the biggest integer such that $S_{k-1} \subseteq S$. So we have

$$\begin{aligned} E[v(S)] &\geq \frac{r}{\alpha+2r} v(S_{k-1}) + \frac{\alpha+r}{\alpha+2r} v(i^*) \\ \Rightarrow (1 + \frac{2r}{\alpha}) E[v(S)] &\geq \frac{r}{\alpha} v(S_{k-1}) + \frac{\alpha+r}{\alpha} v(i^*) \end{aligned}$$

So by using Lemma 11 we have

$$(1 + \frac{2r}{\alpha}) E[v(S)] \geq opt(A)$$

∎

A.3 Deferred Proof from Section 2.3

Recall Theorem 20 which states that by using threshold payments, CAUTIOUS-BUYER $(\frac{r}{r+1}, \frac{1}{r})$ is truthful, budget feasible, and $1 + r$ approximation of the optimum. By using the greedy algorithm with $r = e/(e-1)$ we get a mechanism with approximation guarantee of ≈ 2.58.

Proof. By Lemma 18 we know that the mechanism is monotone, so by using threshold payments, the mechanism is truthful.

Let S be the winning set of the mechanism. By Lemma 18, we have

$$\begin{aligned}
&\max(1+\frac{1}{\gamma},\frac{r}{\alpha})v(S) \geq opt(A,B)\\
\Rightarrow &\max(1+r,\frac{r}{\frac{r}{r+1}})v(S) \geq opt(A,B)\\
\Rightarrow &(1+r)v(S) \geq opt(A,B)
\end{aligned}$$

So the mechanism is $1+r$ approximation of the optimum.

We also have that

$$\alpha = \frac{r}{r+1} = \frac{\frac{1}{\gamma}}{\frac{\gamma+1}{\gamma}} = \frac{1}{1+\gamma}$$

So by Lemma 19 the mechanism is budget feasible.

By choosing the oracle to be the greedy algorithm in [Svi04], which is $\frac{e}{e-1}$ approximation of the optimum, the approximation ratio of CAUTIOUS-BUYER is $1+\frac{e}{e-1} \approx 2.58$. ■

APPENDIX B

ADDITIONAL FIGURES FOR CHAPTER 3

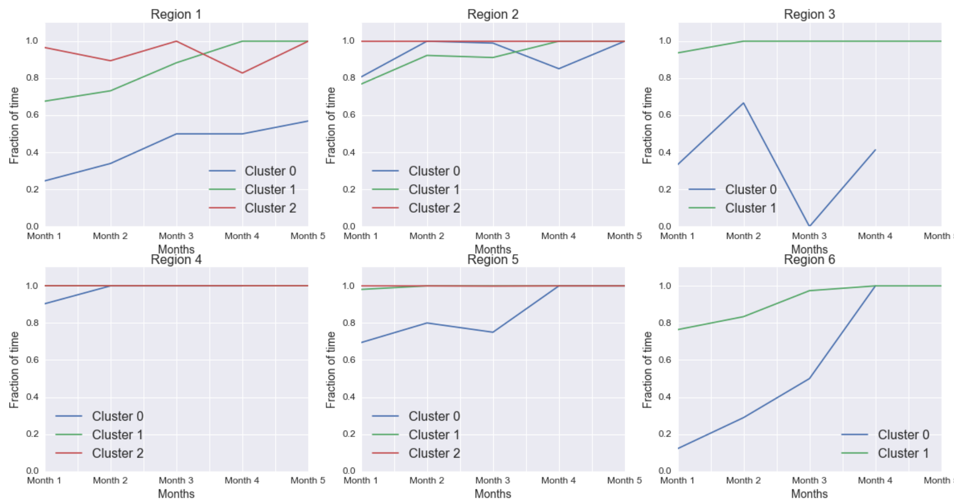

Figure B.1: Average fraction of time agents follow the recommended bid separated by clusters and regions.

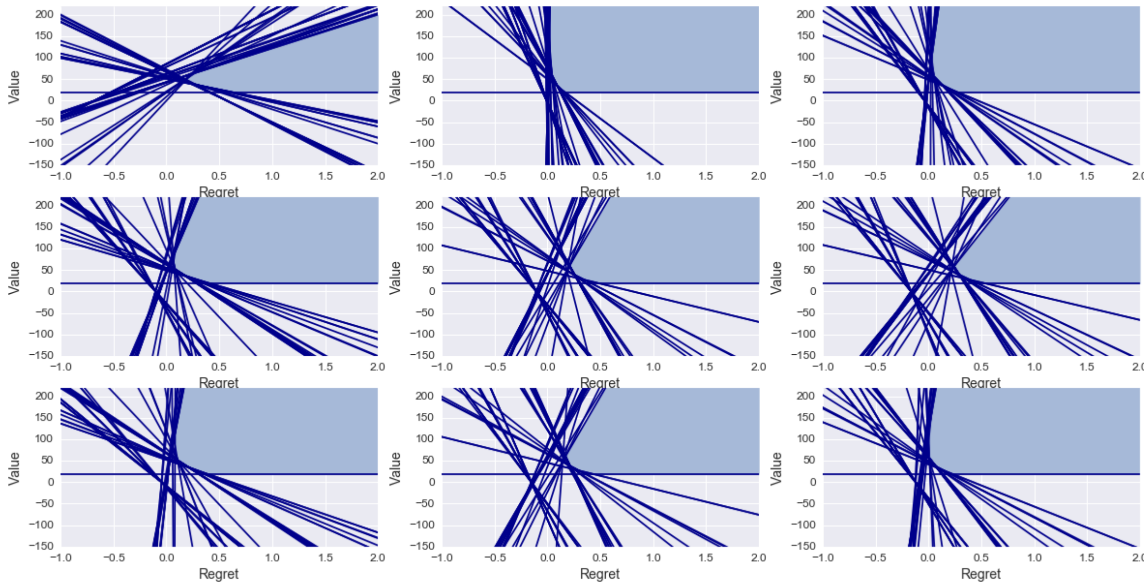

Region 1

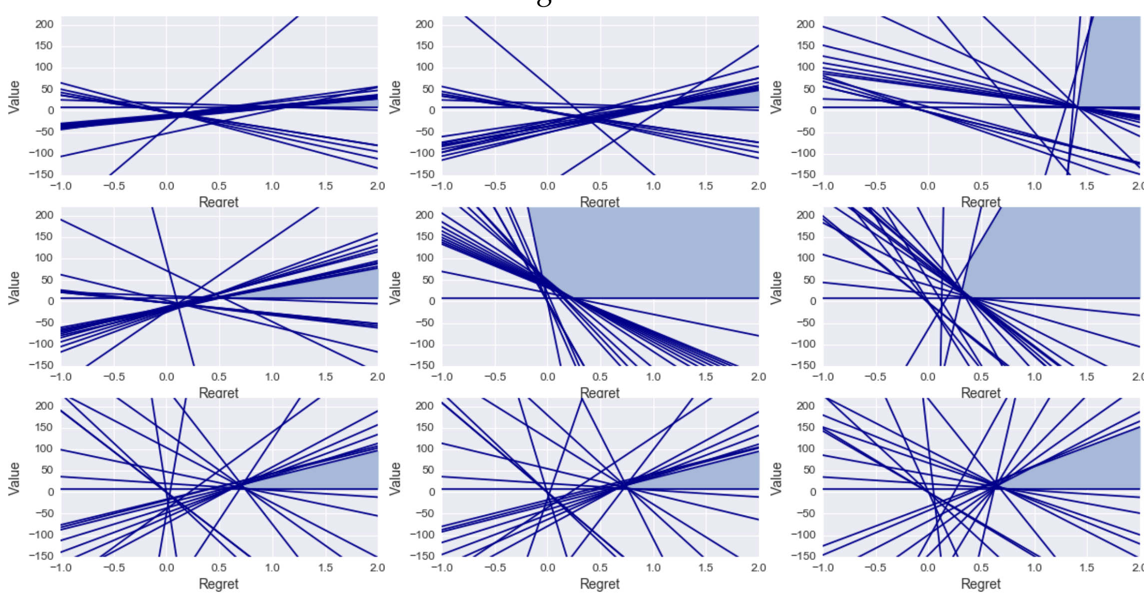

Region 2

Figure B.2: Rationalizable set for 9 agents most frequently changing bids

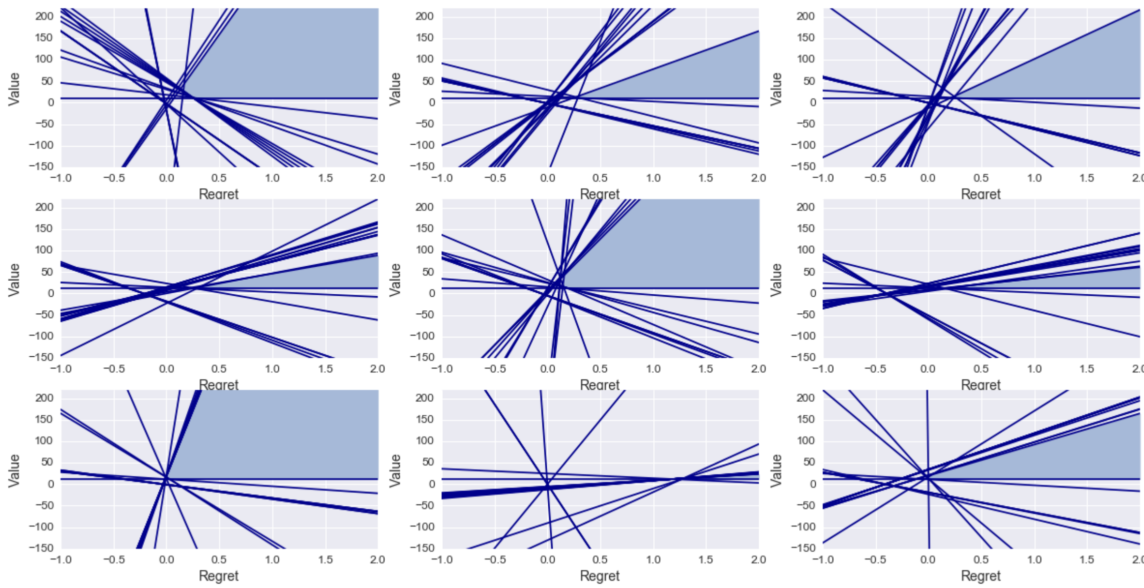

Region 3

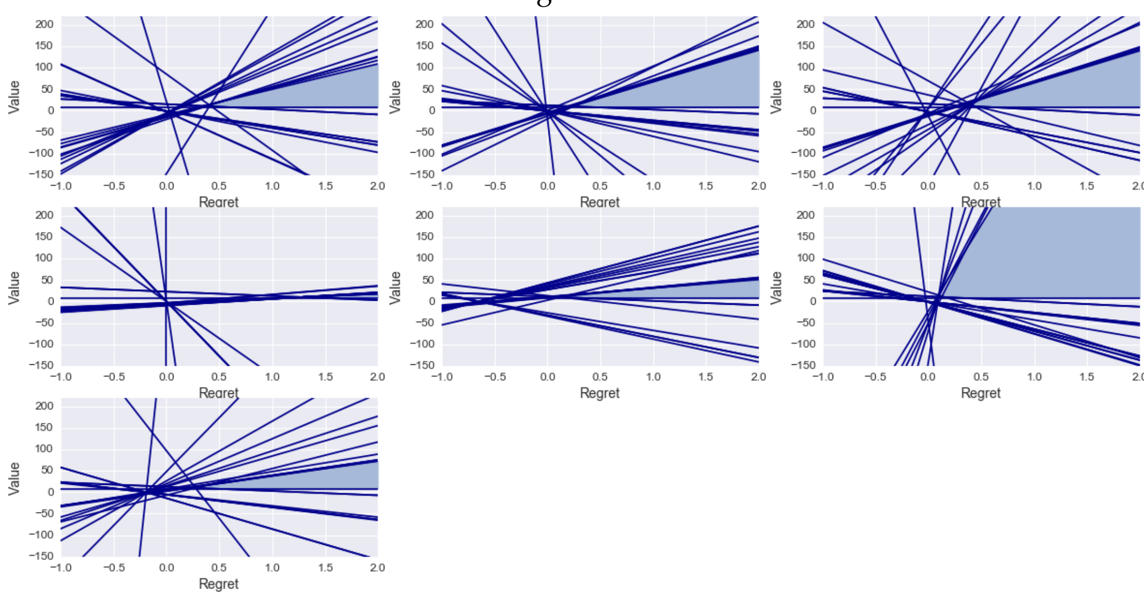

Region 4

Figure B.3: Rationalizable set for 9 agents most frequently changing bids

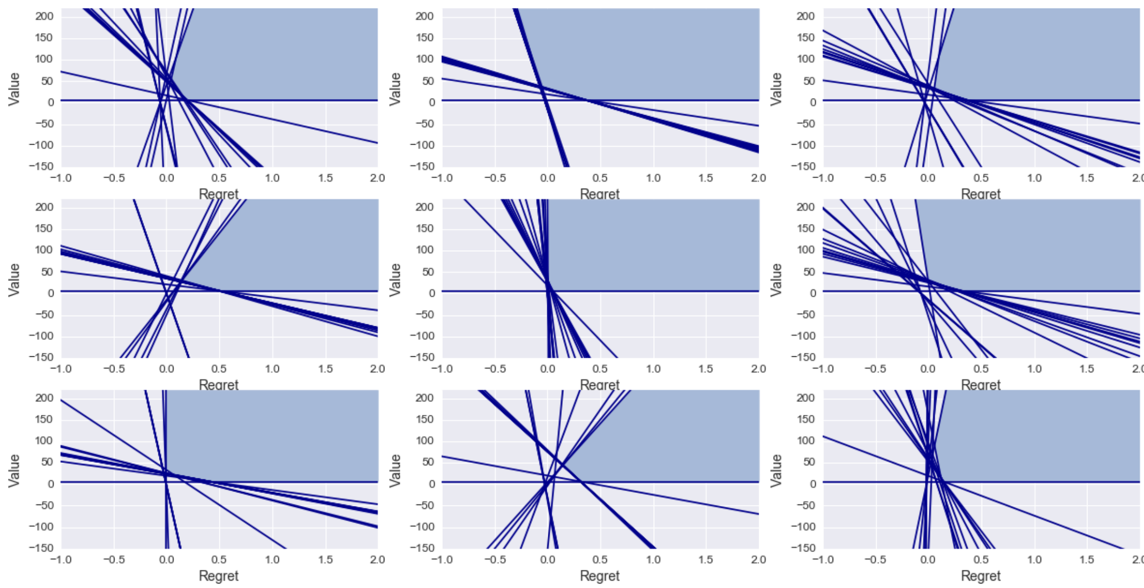

Region 5

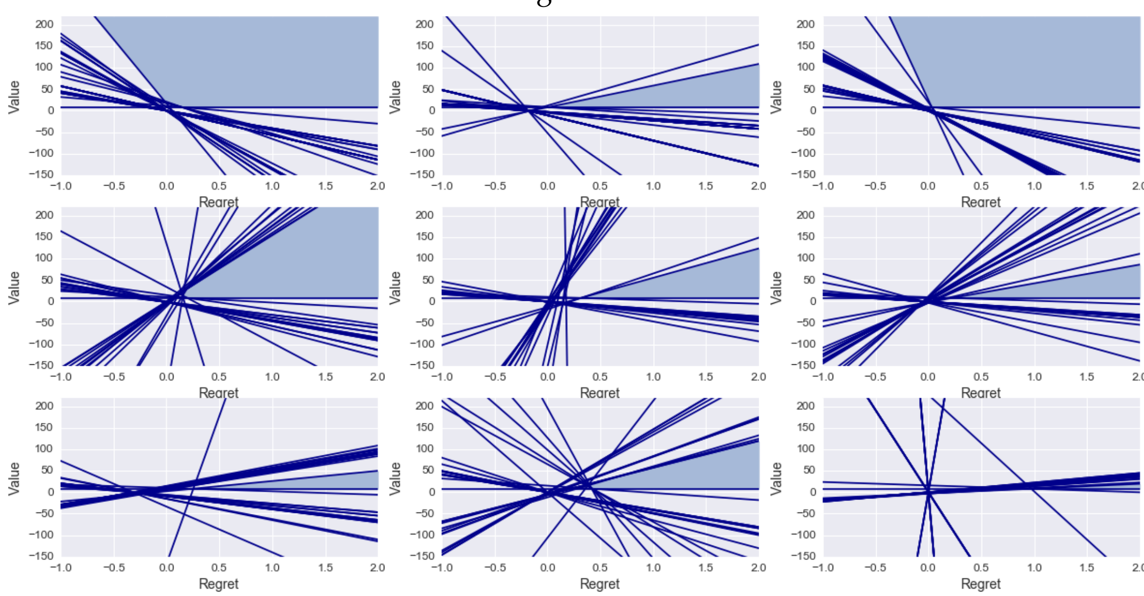

Region 6

Figure B.4: Rationalizable set for 9 agents most frequently changing bids

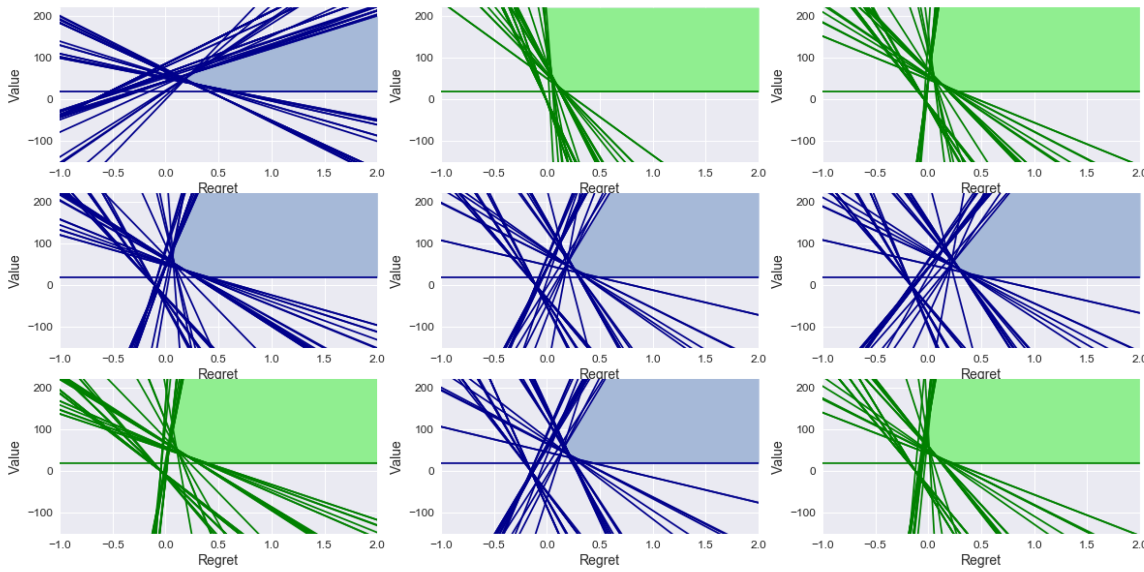

Region 1

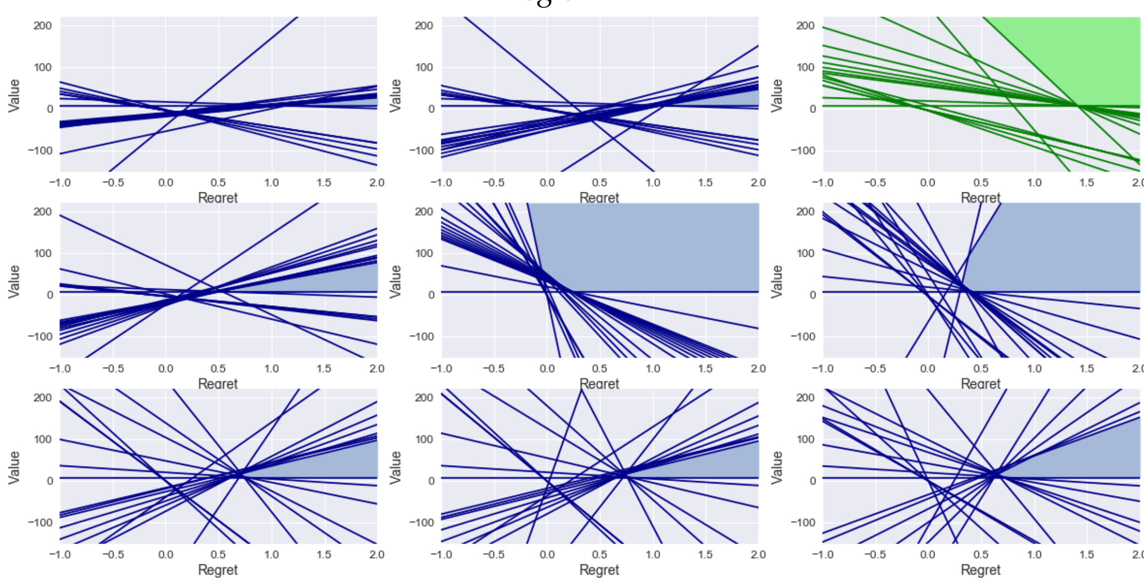

Region 2

Figure B.5: Rationalizable set for 9 agents most frequently changing bids after removing the deviations with at most 2% payment per impression or impression gain. The changed rationalizable sets are marked by green.

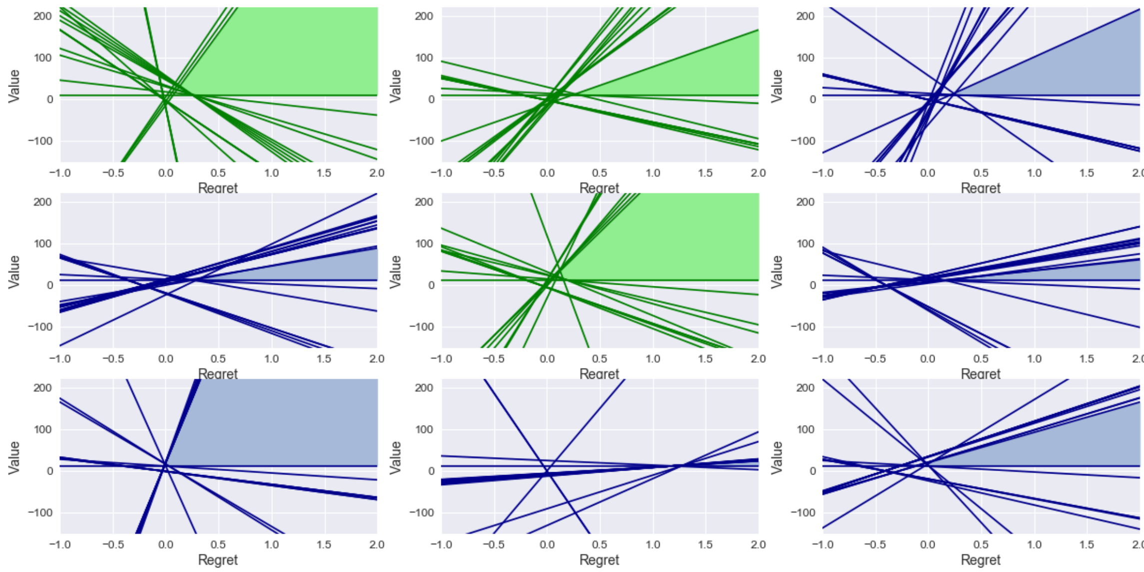

Region 3

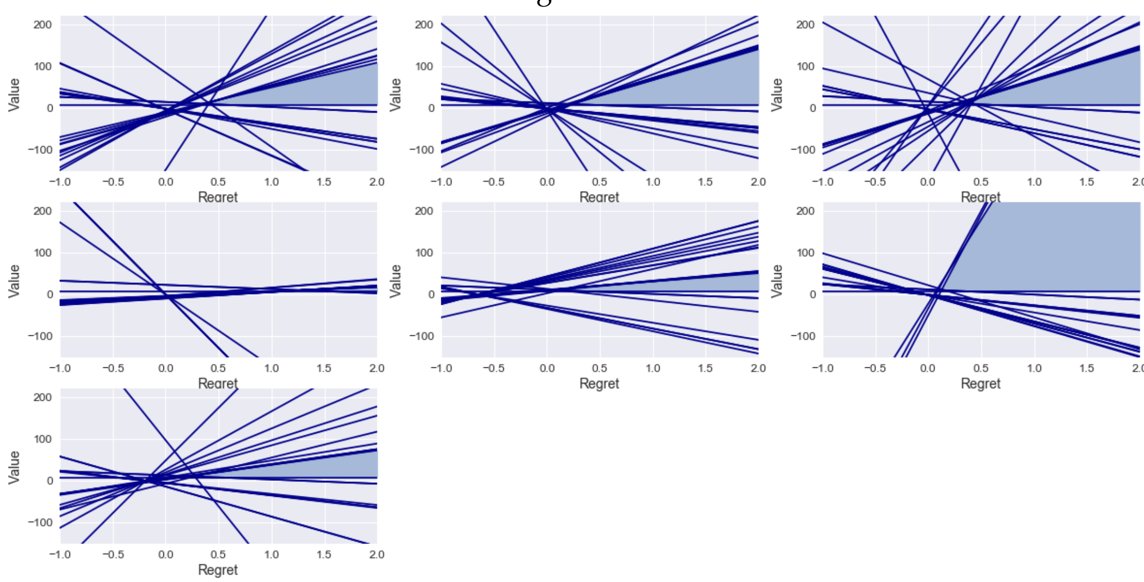

Region 4

Figure B.6: Rationalizable set for 9 agents most frequently changing bids after removing the deviations with at most 2% payment per impression or impression gain. The changed rationalizable sets are marked by green.

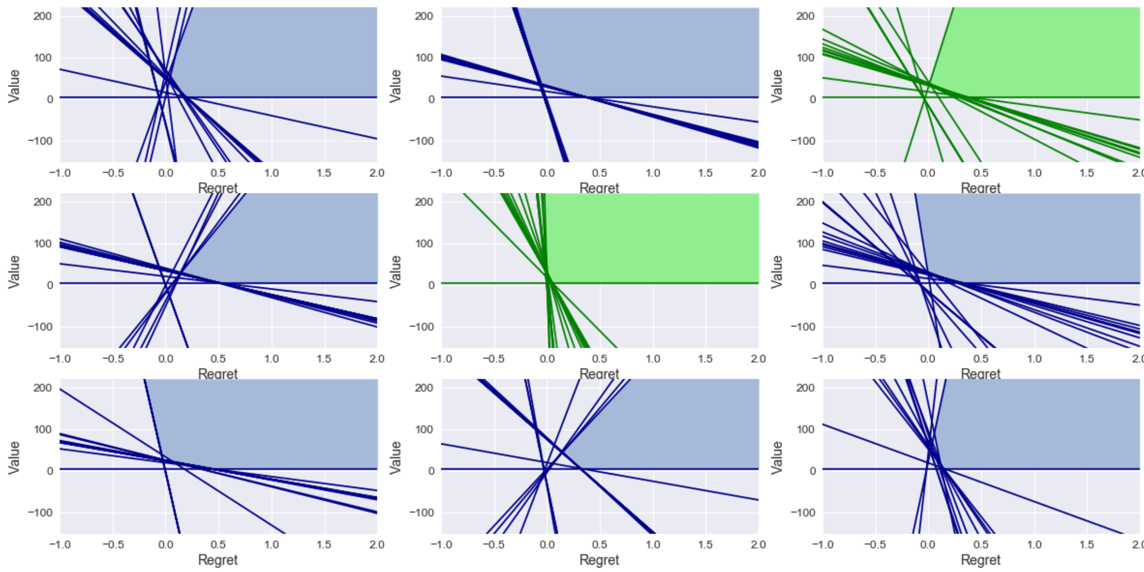

Region 5

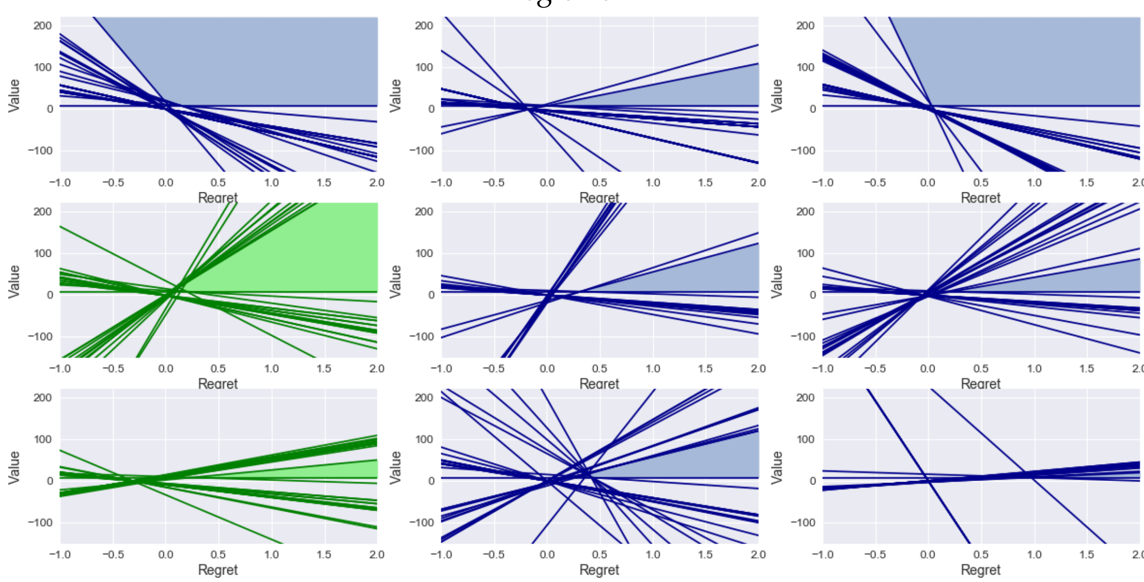

Region 6

Figure B.7: Rationalizable set for 9 agents most frequently changing bids after removing the deviations with at most 2% payment per impression or impression gain. The changed rationalizable sets are marked by green.

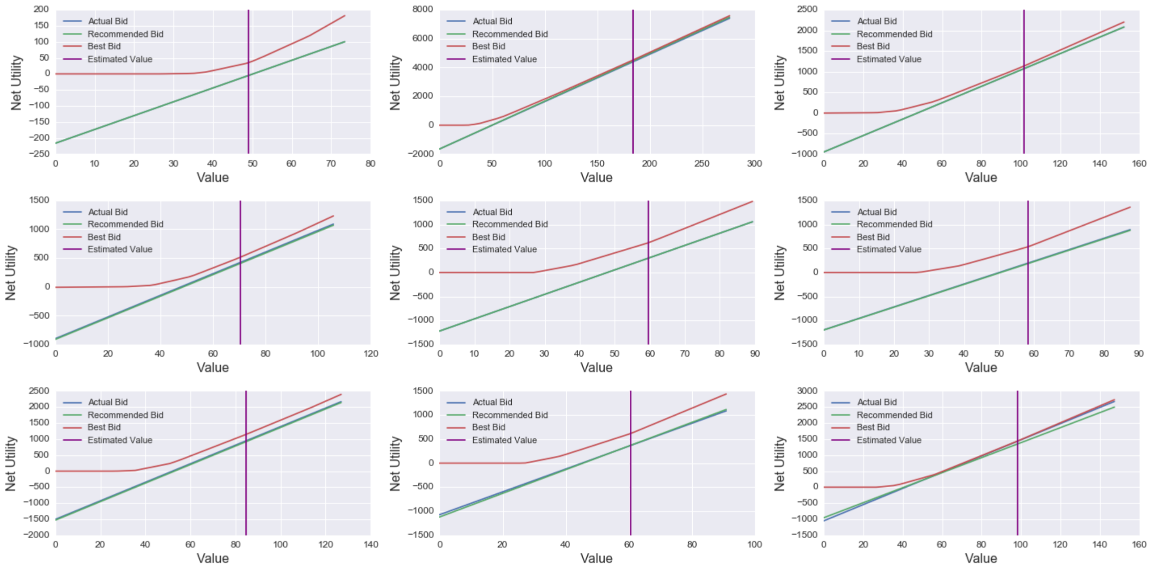

Region 1

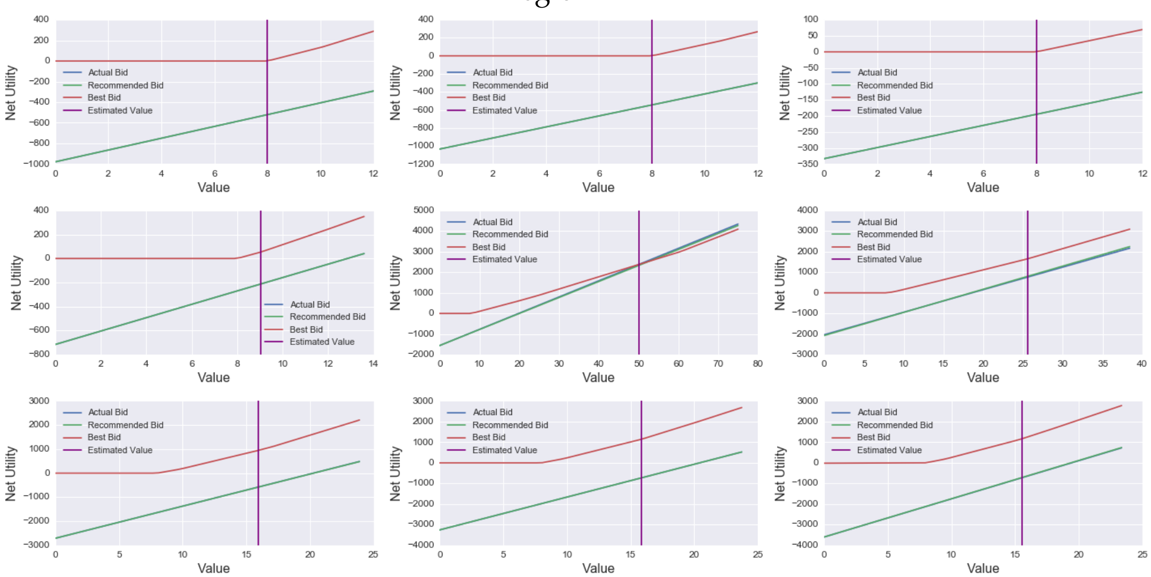

Region 2

Figure B.8: Net value over value for agent's bid, recommended bid and highest fix bid of 9 agents with the most bid changes

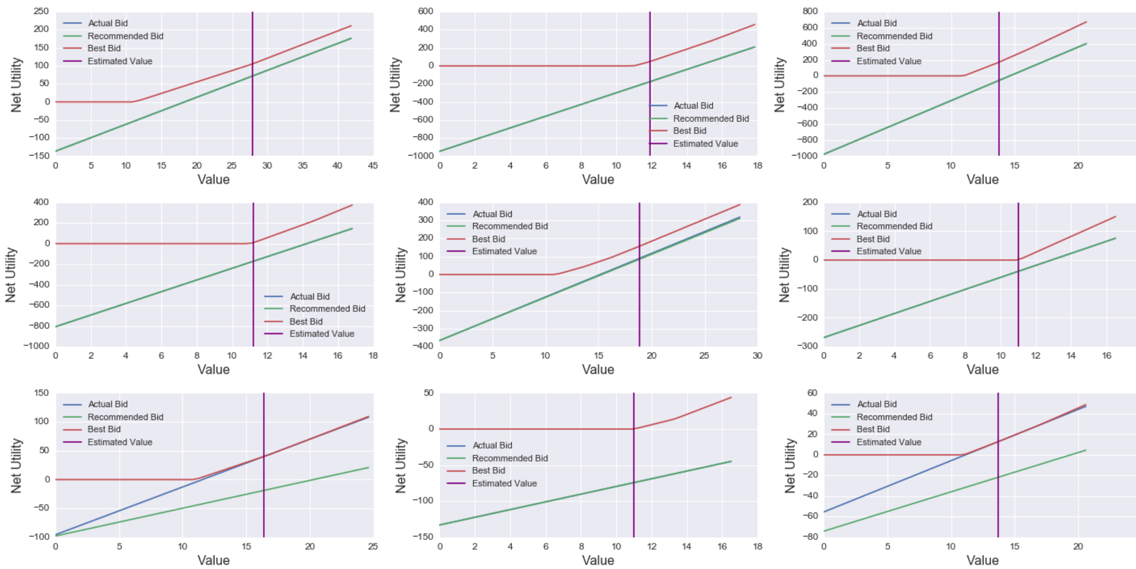

Region 3

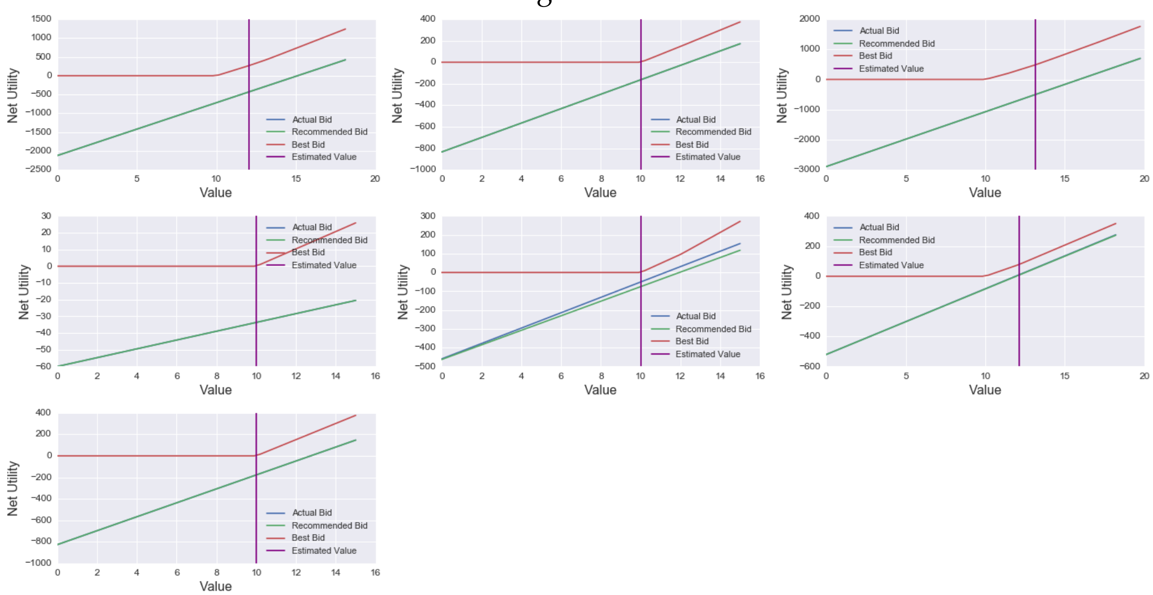

Region 4

Figure B.9: Net value over value for agent's bid, recommended bid and highest fix bid of 9 agents with the most bid changes

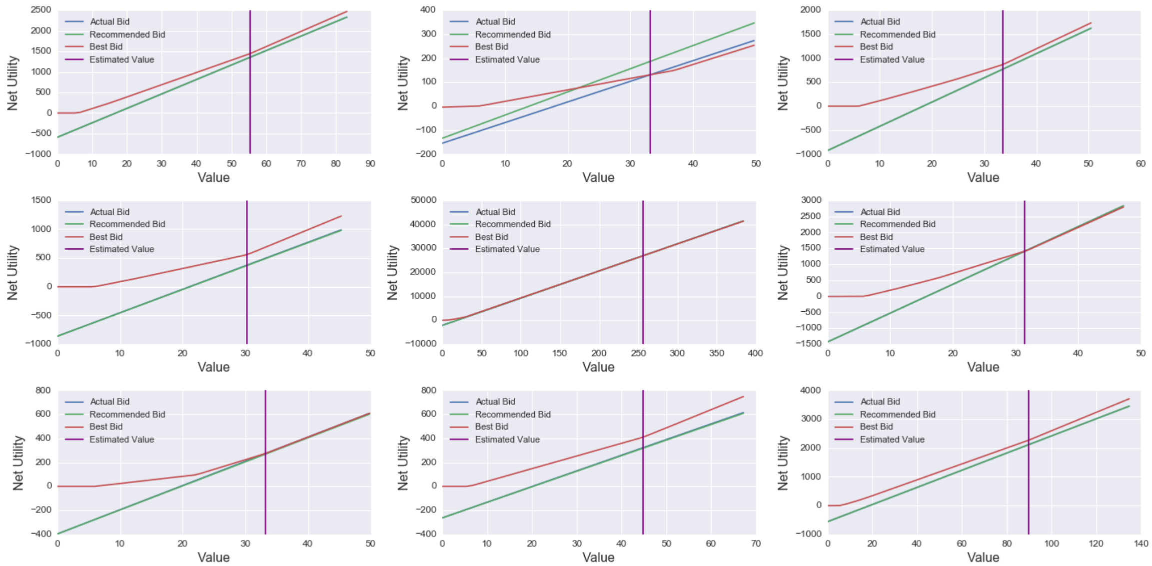

Region 5

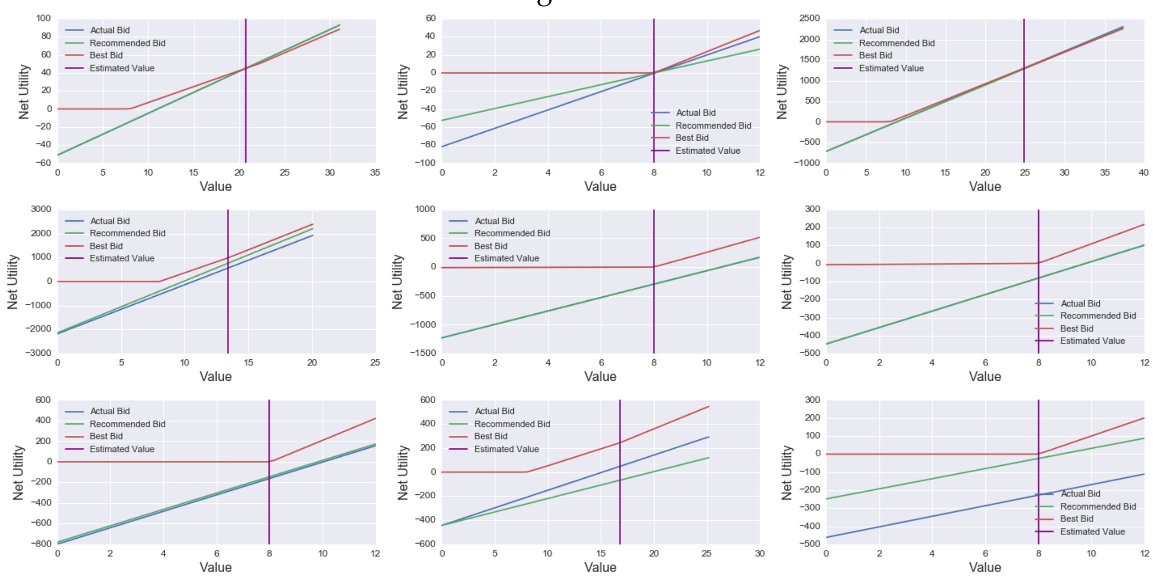

Region 6

Figure B.10: Net value over value for agent's bid, recommended bid and highest fix bid of 9 agents with the most bid changes

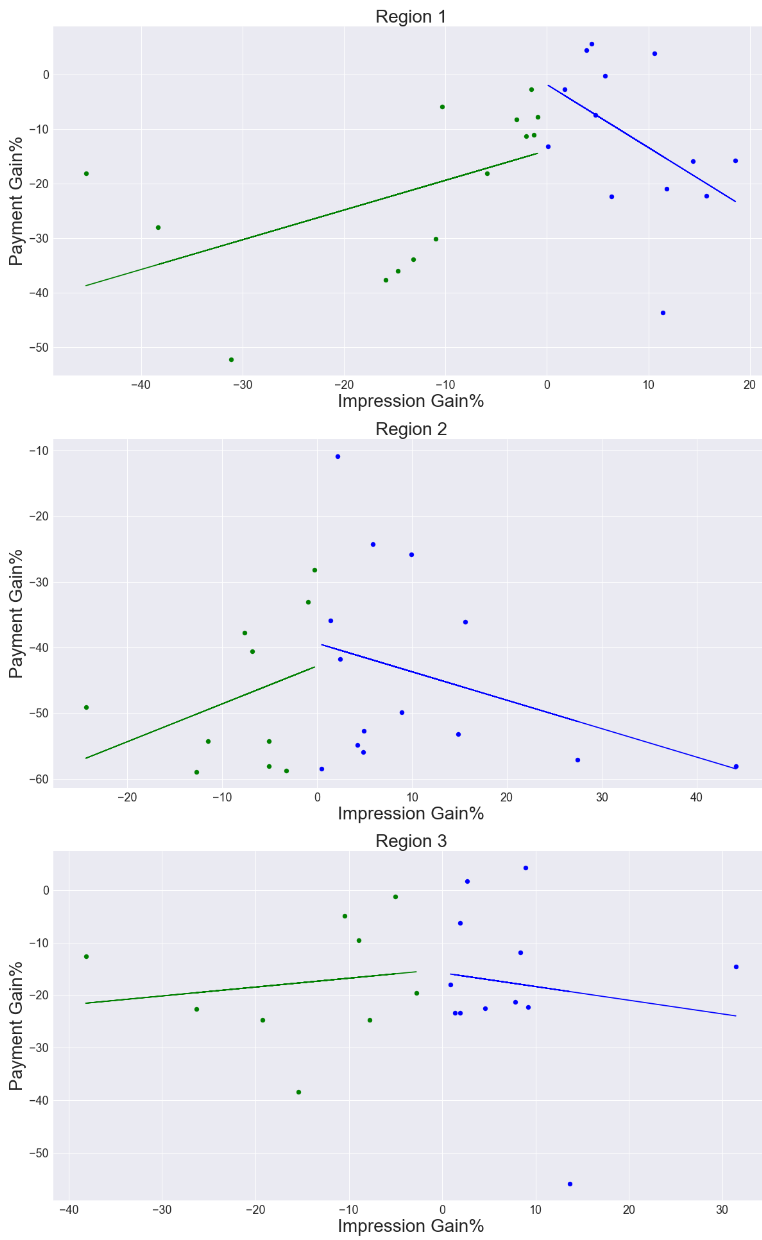

Figure B.11: Each point corresponds to a deviation that determines the value of an agent. The X and Y axis are payment per impression and impression gains respectively. The green and blue points correspond to the deviations that imply a lower bound and upper bound on the value respectively. The Green and blue line correspond to the linear least-squares regression for the green and blue points.

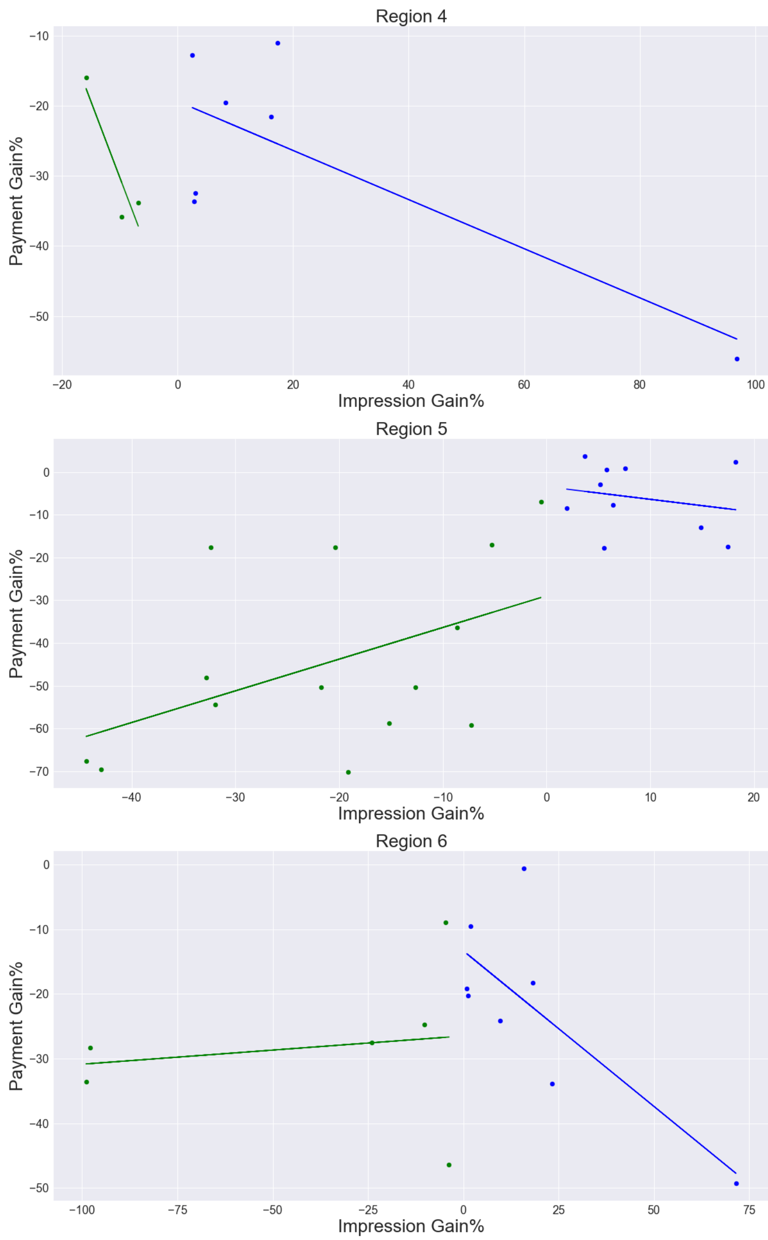

Figure B.12: Each point corresponds to a deviation that determines the value of an agent. The X and Y axis are payment per impression and impression gains respectively. The green and blue points correspond to the deviations that imply a lower bound and upper bound on the value respectively. The Green and blue line correspond to the linear least-squares regression for the green and blue points.

APPENDIX C

ADDITIONAL EXAMPLES AND DISCUSSION FOR CHAPTER 4

C.1 Examples

Interim Pareto-efficiency of RSD It is well known that RSD is ex post Pareto-efficient (efficient after the random choices are made), but not interim Pareto-efficient. From the perspective of a fixed agent, ordinal preferences are not always sufficient for ranking the interim outcomes of a mechanism (i.e., ranking of distributions over outcomes), even though they are sufficient for ranking the ex-post outcomes. However, in the example below, ordinal preferences are enough to show that the interim outcomes are not Pareto-efficient. Consider the following example with agents $A_1, \ldots, A_4$ and items $I_1, \ldots, I_4$ in which the preferences of the agents are indicated on the edges (only the first two top choices of each agent).

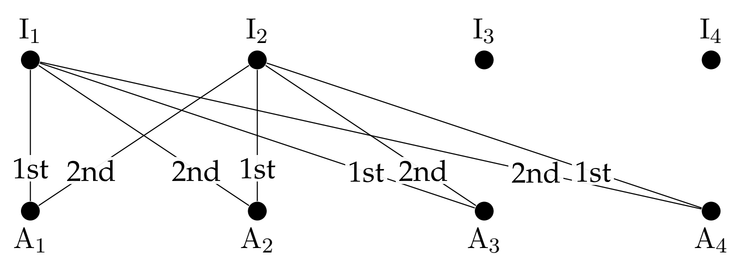

Let x_{ij} denote the probability that agent i gets item j assuming the items are allocation using RSD. Notice that A_1 gets I_2 when the agents are served in the order A_3, A_1, A_2, A_4. Similarly A_2 gets I_1 when the agents are served in the order

A_4, A_2, A_1, A_3. Therefore both x_{12} and x_{21} are strictly positive. However that implies the allocation is not interim Pareto-efficient because A_1 and A_2 would both strictly benefit by exchanging some fraction of their allocation of I_1 and I_2.

C.2 Pareto-efficiency with cardinal versus ordinal preferences

Cardinal and ordinal preferences are different in evaluating outcomes. In the next example the RSD interim allocations is interim Pareto inefficient with cardinal preferences, even though with the corresponding ordinal preferences would not have been interim Pareto ifefficient. Consider the following example in which the valuation of each agent for each item is specified on the corresponding edge and $\epsilon < 0.5$.

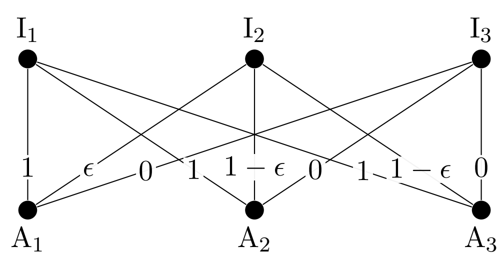

All of the agents in the above example have the same ordinal preferences, but different cardinal preferences. Running the RSD mechanism would yield an allocation of $x_i = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ for each agent $i \in [3]$ which would have been Pareto-efficient for the corresponding ordinal preferences. However this allocation is not Pareto efficient for the above cardinal preferences because it is strictly Pareto dominated by the following allocation for any $\lambda \in [0, 1]$:

$$\mathbf{x}_1 = (\frac{1}{3} + \frac{\lambda(1-\epsilon) + (1-\lambda)\epsilon}{3}, 0, \frac{1}{3} + \frac{\lambda\epsilon + (1-\lambda)(1-\epsilon)}{3}),$$
$$\mathbf{x}_2 = \mathbf{x}_3 = (\frac{1}{3} - \frac{\lambda(1-\epsilon) + (1-\lambda)\epsilon}{6}, \frac{1}{2}, \frac{1}{3} - \frac{\lambda\epsilon + (1-\lambda)(1-\epsilon)}{3}).$$

Basically the interim RSD allocation can be Pareto improved by agent 1 trading its whole share of item 2 with agent 2 and 3 in return for a smaller share of item 1 where item 3 is just traded as a *dummy item* to allow the total allocation of each agent to remain equal to 1.

C.3 Existence of Equilibria

In this section, we prove the existence of a market equilibrium by constructing a corresponding concave game with convex externality constraints as formally defined in C.3.1 such that any equilibrium of the concave game corresponds to a market equilibrium. We then show that this game satisfies the requirements of 46 which implies it has an equilibrium, hence a market equilibrium exists.

Given a market with n unit demand agents and m items with C_j copies for each item $j \in [m]$, we define an $n+1$ player concave game as follows. For each $i \in [m]$, player i corresponds to agent i in the market and chooses $\mathbf{x}_i \in \mathbb{R}^m_{\geq 0}$ to optimize the following linear program:

$$
\begin{aligned}
\text{maximize} \quad & \sum_j v_{ij} x_{ij} && \text{(C.3.1)} \\
\text{subject to} \quad & \sum_j x_{ij} p_j \leq 1 \\
& \sum_j x_{ij} \leq 1 \\
& x_{ij} \geq 0 && \forall j \in [m] \,.
\end{aligned}
$$

Player 0 corresponds the market maker and chooses the price vector $\mathbf{p} \in \mathbb{R}^m_{\geq 0}$ to optimize the following linear program:

$$
\begin{aligned}
\text{maximize} \quad & \sum_j (\sum_i x_{ij} - C_j) p_j && \text{(C.3.2)} \\
\text{subject to} \quad & \sum_j C_j p_j \leq n \\
& p_j \geq 0 && \forall j \in [m] \,.
\end{aligned}
$$

Lemma 44. *Any equilibrium of the above game corresponds to an equilibrium of the original market.*

Proof. It is easy to see that in any equilibrium of the game the action of player i, $\mathbf{x}_i$, is an optimal allocation for agent i of the original market with respect to price vector $\mathbf{p}$. So we only need to verify that the market clearing conditions are also satisfied. It is easy to see that if there are over-allocated items, then player 0 (market maker) should have chosen a price vector $\mathbf{p}$ that would make the object value of C.3.2 strictly positive and would also make the first constraint tight. But a strictly positive objective means the total amount of money spend by all agents, $\sum_j (\sum_i x_{ij} p_j)$, is more than $\sum_j (\sum_i C_j p_j)$ which is itself equal to n by tightness of the first constraint, hence a contradiction. On the other hand, if there is any under-allocated item, the market maker should have assigned a price of 0 to that item, hence all agents must be fully allocated. ■

Theorem 45. *A market equilibrium for the original market always exists.*

Proof. By 44 an equilibrium for the original market exists if the corresponding concave game has an equilibrium. By 46 a concave game has an equilibrium if every player has a default action that lies strictly inside the feasible set of actions for that player for all possible actions of other players. Let $\epsilon = 1/(n \sum_j C_j)$. A default action for player 0 is given by $p_j^0 = \epsilon$ for every $j \in [m]$. A default action for each player $i \in [n]$ is given by $\mathbf{x}_{ij}^0 = \epsilon$ for all $j \in [m]$. ■

C.3.1 Multi Agent Concave Games with Externality Constraints

In this section we prove more generally the existence of an equilibrium for a general class of games in which both the utility function and the set of feasible actions for a player may depend on the actions of the other players. Formally, suppose there are n players and the optimal utility of player $i \in [n]$ is captured by the following convex program

$$\begin{aligned} &\text{maximize} && v_i(\mathbf{x}_i, \mathbf{x}_{-i}) && && \text{(C.3.3)} \\ &\text{subject to} && h_{ik}(\mathbf{x}_i, \mathbf{x}_{-i}) \leq 0 && \forall k \in [d_i] \\ & && \mathbf{x}_i \in \mathrm{S}_i \subseteq \mathbb{R}^{m_i} \end{aligned}$$

in which $\mathbf{x}_i$ is the action of player i; $\mathbf{x}_{-i}$ is the vector of actions of players other than i; v_i is the utility function of player i which is concave in $\mathbf{x}_i$ and continuous in all arguments; $\mathrm{S}_i \subset \mathbb{R}^{m_i}$ is a compact and convex set representing the feasible actions of player i; and h_{ik} is convex in $\mathbf{x}_i$ and continuous in all arguments. Define $\mathbf{S} = \mathrm{S}_1 \times \cdots \times \mathrm{S}_n$. A vector of actions $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n) \in \mathbf{S}$ is an *equilibrium* iff $\mathbf{x}_i$ is an optimal assignment for C.3.3 for each $i \in [n]$. The next theorem establishes that an equilibrium always exists under a mild assumption.

Theorem 46. *If for each player $i \in [m]$ there exists a* default action $\mathbf{x}_i^0 \in \mathrm{S}_i$ *for which the first set of constraints in C.3.3 hold strictly for all $\mathbf{x}_{-i} \in \mathbf{S}_{-i}$, then an equilibrium always exists.*

Proof. Define the set of optimal actions of player i in response to $\mathbf{x}_{-i} \in \mathbf{S}_{-i}$ as

$$R_i(\mathbf{x}_{-i}) = \{\mathbf{x}_i \in \mathrm{S}_i \mid \mathbf{x}_i \text{ is a solution to C.3.3 with } \mathbf{x}_{-i}\}\,.$$

Define the collective optimal actions in response to $\mathbf{y} \in \mathbf{S}$ as

$$R(\mathbf{y}) = \left\{\mathbf{x} \in \mathbf{S} \,\middle|\, \mathbf{x}_i \in R_i(\mathbf{y}_{-i}) \quad \forall i \in [n]\right\}.$$

It is easy to see that R is a set valued function from $\mathbf{S}$ to $2^{\mathbf{S}}$ whose fixed points correspond to the equilibria of the game. We prove that R satisfies the requirements of Kakutani's fixed point theorem and therefore must have a fixed point which would then imply the existence of an equilibrium. It is easy to see that $\mathbf{S}$ is a non-empty compact and convex set, and that $R(\mathbf{y})$ is a also a non-empty compact convex set for every $\mathbf{y} \in \mathbf{S}$. Therefore we only need to prove that R has a closed graph. Formally, we need to show that for any sequence $(\mathbf{x}^t, \mathbf{y}^t)_{t\in\mathbb{N}}$ where $\mathbf{x}^t \in R(\mathbf{y}^t)$ for all $t \in \mathbb{N}$, if $\lim_{t\to\infty}(\mathbf{x}^t, \mathbf{y}^t) = (\mathbf{x}^*, \mathbf{y}^*)$, then $\mathbf{x}^* \in R(\mathbf{y}^*)$. Consider C.3.3 in which $\mathbf{x}_{-i}$ set to $\mathbf{y}^*_{-i}$. From the continuity of h_{ik} and compactness of S_i, it follows that $\mathbf{x}_i^*$ is a feasible action for player i in response to $\mathbf{y}_{-i}^*$, so we only need to show that it is also optimal. Pick any $\mathbf{x}_i^\dagger \in R_i(\mathbf{y}_{-i}^*)$. We will prove optimality by showing that $v_i(\mathbf{x}_i^*, \mathbf{y}_{-i}^*) \geq v_i(\mathbf{x}_i^\dagger, \mathbf{y}_{-i}^*)$. Recall that $v_i(\mathbf{x}_i^*, \mathbf{y}_{-i}^*) = \lim_{t\to\infty} v_i(\mathbf{x}_i^t, \mathbf{y}_{-i}^t)$, so ideally we would like to show that $v_i(\mathbf{x}_i^t, \mathbf{y}_{-i}^t) \geq v_i(\mathbf{x}_i^\dagger, \mathbf{y}_{-i}^t)$, however $\mathbf{x}_i^\dagger$ may not even be a feasible action in response to $\mathbf{y}_{-i}^t$ for any t and therefore such an inequality may not hold in general. This is where we use the assumption of the theorem to show that there exists an action $\mathbf{x}_i'^t$ between $\mathbf{x}_i^t$ and $\mathbf{x}_i^\dagger$ which is feasible for player i in response to $\mathbf{y}_{-i}^t$ for all t and also converges to $\mathbf{x}_i^\dagger$ as $t \to \infty$ which will imply the desired inequality as follows:

$$\begin{aligned}
v_i(\mathbf{x}_i^*, \mathbf{y}_{-i}^*) &= \lim_{t\to\infty} v_i(\mathbf{x}_i^t, \mathbf{y}_{-i}^t) \\
&\geq \lim_{t\to\infty} v_i(\mathbf{x}_i'^t, \mathbf{y}_{-i}^t) \\
&= v_i(\mathbf{x}_i^\dagger, \mathbf{y}_{-i}^*).
\end{aligned}$$

To define $\mathbf{x}_i'^t$, we first define the following two quantities in which $\mathbf{x}_i^0$ is the default action of player i:

$$\begin{aligned}
\delta^t &= \max\left(0, \max_{i,k} h_{ik}(\mathbf{x}_i^\dagger, \mathbf{y}_{-i}^t)\right) \\
\delta^0 &= -\max_{i,k,\mathbf{x}_{-i}''} h_{ik}(\mathbf{x}_i^0, \mathbf{x}_{-i}'').
\end{aligned}$$

Note that $\delta^0 > 0$ because of the hypothesis of the theorem. Let $\lambda^t = \frac{\delta^0}{\delta^0+\delta^t}$. Define $\mathbf{x}'^t = \lambda^t \mathbf{x}^\dagger + (1-\lambda^t)\mathbf{x}^0$. First notice that $\mathbf{x}_i'^t$ is a feasible response to $\mathbf{y}_{-i}^t$ because

$$\begin{aligned}
h_{ik}(\mathbf{x}_i'^t, \mathbf{y}_{-i}^t) &\leq \lambda^t h_{ik}(\mathbf{x}_i^\dagger, \mathbf{y}_{-i}^t) + (1-\lambda^t) h_{ik}(\mathbf{x}_i^0, \mathbf{y}_{-i}^t) && \text{by convexity of } h_{ik} \text{ in } \mathbf{x}_i'^t \\
&\leq \lambda^t \delta^t - (1-\lambda^t)\delta^0 && \text{by definition of } \delta^t \text{ and } \delta^0 \\
&= 0 && \text{by definition of } \lambda^t.
\end{aligned}$$

Second, notice that continuity of h_{ik} guarantees $\lim_{t\to\infty} \delta^t = 0$, which implies $\lim_{t\to\infty} \lambda^t = 1$ (recall that $\delta^0 > 0$), which implies $\lim_{t\to\infty} \mathbf{x}'^t = \mathbf{x}^\dagger$ as claimed.

BIBLIOGRAPHY

[ABM16] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. Coverage, matching, and beyond: New results on budgeted mechanism design. In *International Conference on Web and Internet Economics*, pages 414–428. Springer, 2016.

[ABM17] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On budget-feasible mechanism design for symmetric submodular objectives. *CoRR*, abs/1704.06901, 2017.

[AGN14] Nima Anari, Gagan Goel, and Afshin Nikzad. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 266–275. IEEE, 2014.

[AJKT17] Saeed Alaei, Pooya Jalaly Khalilabadi, and Eva Tardos. Computing equilibrium in matching markets. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 245–261, New York, NY, USA, 2017. ACM.

[AN10] Susan Athey and Denis Nekipelov. A structural model of sponsored search advertising auctions. In *Sixth ad auctions workshop*, volume 15, 2010.

[AS98] Atila Abdulkadiroglu and Tayfun Sonmez. Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems. *Econometrica*, 66(3):689–702, May 1998.

[ASZ14] Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. Efficiency of truthful and symmetric mechanisms in one-sided matching. In *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*, pages 13–24, 2014.

[ATCO16] D. Agarwal, S. Ghos T. Cui, and M. Ostrovsky. Bid suggestions for online advertising auctions: A field experiment. *Working paper*, 2016.

[BCGL12] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of* *the forty-fourth annual ACM symposium on Theory of computing*, pages 449–458. ACM, 2012.

[BCI$^+$05] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 44–51. ACM, 2005.

[BCK11] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Social welfare in one-sided matching markets without money. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings*, pages 87–98, 2011.

[BH16] Eric Balkanski and Jason D Hartline. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 189–203. International World Wide Web Conferences Steering Committee, 2016.

[BM01] Anna Bogomolnaia and Herve Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295 – 328, 2001.

[BPR98] S. Basu, R. Pollack, and M.-F. Roy. A new algorithm to find a point in every cell defined by a family of polynomials. In B. Caviness and J. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, Texts and Mongraphs in Symbolic Computation, pages 341–350. Springer-Verlag, 1998.

[CGL11] Ning Chen, Nick Gravin, and Pinyan Lu. On the approximability of budget feasible mechanisms. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 685–699. SIAM, 2011.

[CV07] Bruno Codenotti and Kasturi Varadarajan. Computation of market equilibria by convex programming. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 6, pages 135–158. Cambridge University Press, 2007.

[DFH$^+$12] Danny Dolev, Dror G. Feitelson, Joseph Y. Halpern, Raz Kupferman, and Nathan Linial. No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 68–75, New York, NY, USA, 2012. ACM.

[DK08] Nikhil R. Devanur and Ravi Kannan. Market equilibria in polynomial time for fixed number of goods or agents. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 45–53, Washington, DC, USA, 2008. IEEE Computer Society.

[DPS11] Shahar Dobzinski, Christos H Papadimitriou, and Yaron Singer. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 273–282. ACM, 2011.

[EOS05] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Working Paper 11765, National Bureau of Economic Research, November 2005.

[EW12] Federico Echenique and Adam Wierman. Finding a walrasian equilibrium is easy for a fixed number of agents. In *EC*, page 495, 2012.

[GN12] Avital Gutman and Noam Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 719–728, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[GNS14] Gagan Goel, Afshin Nikzad, and Adish Singla. Mechanism design for crowdsourcing markets with heterogeneous tasks. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.

[goo] Estimate your results with bid simulators. https://support.google.com/adwords/answer/2470105?hl=en. Google, 2016.

[HZ79] Aanund Hylland and Richard Zeckhauser. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy*, 87(2):293–314, April 1979.

[JNT17] Pooya Jalaly, Denis Nekipelov, and Éva Tardos. Learning and trust in auction markets. *CoRR*, abs/1703.10672, 2017.

[JT17] Pooya Jalaly and Éva Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. *CoRR*, abs/1703.10681, 2017.

[KMN99] Samir Khuller, Anna Moss, and Joseph Seffi Naor. The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45, 1999.

[Mye81] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[NN17] Noam Nisan and Gali Noti. An experimental evaluation of regret-based econometrics. In *Proceedings of the 26th International Conference on World Wide Web*, pages 73–81. International World Wide Web Conferences Steering Committee, 2017.

[NST15] Denis Nekipelov, Vasilis Syrgkanis, and Eva Tardos. Econometrics for learning agents. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 1–18. ACM, 2015.

[RRS11] Francesco Ricci, Lior Rokach, and Bracha Shapira. Introduction to recommender systems handbook. In *Recommender systems handbook*, pages 1–35. Springer, 2011.

[Sin10] Yaron Singer. Budget feasible mechanisms. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 765–774. IEEE, 2010.

[SK13] Adish Singla and Andreas Krause. Incentives for privacy tradeoff in community sensing. In *First AAAI Conference on Human Computation and Crowdsourcing*, 2013.

[Svi04] Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.

[WM15] Xiaodong Wang and José F. Martínez. XChange: A market-based approach to scalable dynamic multi-resource allocation in multicore architectures. In *International Symposium on High-Performance Computer Architecture (HPCA)*, Bay Area, CA, February 2015.

[ZG16] Inc Zillow Group. Zillow group reports fourth quarter and full year 2016 results, 2016.

[Zho90] Lin Zhou. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990.