DESIGN AND ANALYSIS OF SPONSORED SEARCH MECHANISMS

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

Renato Paes Leme

December 2012

DESIGN AND ANALYSIS OF SPONSORED SEARCH MECHANISMS

Renato Paes Leme, Ph.D.

Cornell University 2012

Auctions have become the standard way of allocating resources in electronic markets. Two main reasons why designing auctions is hard are the need to cope with strategic behavior of the agents, who will constantly adjust their bids seeking more items at lower prices, and the fact that the environment is highly dynamic and uncertain. Many market designs which became de-facto industrial standards allow strategic manipulation by the agents, but nevertheless display good behavior in practice. In this thesis, we analyze why such designs turned out to be so successful despite strategic behavior and environment uncertainty. Our goal is to learn from this analysis and to use the lessons learned to design new auction mechanisms; as well as fine-tune the existing ones.

We illustrate this research line through the analysis and design of Ad Auctions mechanisms. We do so by studying the equilibrium behavior of a game induced by Ad Auctions, and show that all equilibria have good welfare and revenue properties. Next, we present new Ad Auction designs that take into account richer features such as budgets, multiple keywords, heterogeneous slots and online supply.

BIOGRAPHICAL SKETCH

Renato Paes Leme was born in December 1st, 1984 in Rio de Janeiro, where he spent most of his life. There in the tropics, he got a BEng in Computer Engineering from Instituto Militar de Engenharia in 2007 and a MSc in Mathematics from Instituto Nacional de Matematica Pura e Aplicada in the same year. He then moved to Ithaca (not quite in the tropics, but beautiful nevertheless) where he expects to get a PhD in Computer Science with a minor in Operations Research from Cornell University in December 2012.

I dedicate this thesis to

my parents Irio and Sylvia

and to my brother Pedro.

ACKNOWLEDGEMENTS

I want thank Éva Tardos for guiding me through this amazing journey that was my PhD education. Éva was the main force shaping my view of science and my taste for research problems. The researcher I am today, I owe to her. What I'll miss more about Cornell will undoubtely be the conversations I had with her. I remember countless visits to her office when Éva was always patient and took time to walk with me over many angles of the question - technical and non-technical. Above all, I thank her for the care and attention she had for me in the past four years.

I am lucky to have been at such a wonderful place as Cornell, where I found many great mentors, who were sources of inspiration and example for me. I'd like to thank David Shmoys for suggesting interesting problems and for teaching me a lot about linear programming and scheduling. Also, thanks David for being my co-advisor and for being always a great source of advice. Many thanks to Bobby Kleinberg for being a great mentor, a great inspiration and for always being able to hint me in the right direction. I am grateful to Jon Kleinberg for his "Structure of Information Networks" class, which was the first class I took at Cornell and filled my mind with all sorts of interesting questions to think about. To David Williamson for his "Approximation Algorithms" class from which I learned a great deal of useful techniques that I later applied in my research. Thanks to Emin Gün Sirer for serving on my committee and for general advice. To many others faculty members that I had the pleasure to be in contact with during my PhD years, many thanks.

I was very fortunate to be able to spend some part of my PhD years travelling and visiting various institutions and research labs. A special thanks for my industry mentors Moshe Babaioff (who hosted me for a semester in Microsoft

Research Sillicon Valley), Jennifer Chayes and Christian Borgs (for hosting me in Microsoft Research New England for one summer and then for many following visits) and Vahab Mirrokni (for hosting me in Google NYC for two summers and countless many visits during the semester). I was also fortunate to spend two months in the Institute of Advanced Studies in the Hebrew University of Jerusalem (thanks to Noam Nisan and Michal Feldman for organizing the AGT semester, bringing together all the great AGT researchers to the same place and creating such a stimulating research environment). I also had the pleasure to visit MIT-CSAIL for one semester (thanks to Silvio Micali and Costis Daskalakis for hosting me).

I made several friends among my research collaborators, and I'd like to thank them all for all the fun moments doing research together. In special, I'll like to thank Vasilis Syrgkanis for being a great friend and such a fun research collaborator. I specially remember solving the main question in our SODA paper while on the beach in Netanya and the many times we discussed game theory while eating shakshouka for breakfast in Jerusalem.

Thanks also to Brendan Lucier for a very fruitful collaboration on GSP auctions and for teaching me many interesting techniques. Thanks to Gagan Goel with whom I learned a lot about truthful auctions and submodular functions and teaching me how to appreciate NYC (in particular showing me where to get good southern Indian food).

One other gift I got from Cornell is many good friends who made my life here fun and interesting. They formed an important part of my PhD: Bruno Abrahao, Daniel Hauagge, Guilherme Pinto, Hussam Abu-Libdeh, Ashwin Baranidyuru, Daniel Fleischman, Shahar Dobzinski, Sigal Oren, Georgios Piliouras and Marcos Salles.

A special thanks to my friends Thiago Pereira and Thiago Costa for being my great friends for a very long time, since we were all undergrads in Brazil. They both shared with me moments during our undergrad and masters and the anxiety applying for a PhD. We went together through moments of the happiness and dispair in our doctoral studies and sharing them along the way was very important to me. I thank Thiago Costa and Andrei Roman for making my year in Boston so fun and for providing me warmth with your friendship. I also thank Thiago Pereira for introducing me to the ICPC contest during our undergrad years, which was my initial contact with algorithms. It was very fun to spend many nights programming with him and Rafael Paulino. Thiago Pereira should also get the credit for introducing me to travelling - thanks for many interesting trips we had through Europe.

I want to thank my friends from Brazil: Rodrigo, Alexandre, Nathalia, Cechin and Nicodemos.

There are many great things that happened while I was in Boston, but nothing compared to meeting my dear Carol. Thanks for all the beautiful moments we had since then and for always inspiring me in my research and beyond. Thanks for filling my life with joy.

Above all, I'd like to thank my parents Irio and Sylvia, who always loved me and supported me in all my decisions. They taught me to be persistent and never give up on problems, which was definitely very useful in research and in life overall. The love of my parents, my brother Pedro and all my family have been my constant support all this years. To you I dedicate this thesis.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Mechanism Design deals with the problem of designing algorithms that will be deployed in settings where agents are *strategic*. In such settings, the algorithm initially collects reports from a set of agents and uses those reports to compute an outcome. Standard examples are internet marketplaces: in the sale of internet advertisement, bids are elicited from potential advertisers and the choice of which ad to display is based on the bids. Agents may have incentives to overbid, or shade (decrease) their bids depending on the system used. The main difficulty in such situation is how to set proper incentives for the agents.

Simplicity is an important goal in designing mechanisms and the lack thereof is usually the main reason why mechanisms do not get adopted in practice. There are two ways in which we desire mechanisms to be simple: (i) the description of the mechanism itself is simple and (ii) it is simple for the agents participating in the mechanism.

This thesis focuses on designing simple and efficient mechanisms. We focus on sponsored search auctions, that is, auctions to sell advertisement space next to search results, which have become the standard way to monetize web-services and are responsible for moving tens of billions of dollars every year. Besides being practically relevance, sponsored search is also a technically challenging setting and techniques developed for it generalize to many other scenarios.

1.1 Comparing Algorithms in Strategic Settings (or Mechanism Design for Algorithmists)

Traditional algorithm design consists of designing a procedure that maps an input to an output (in a possibly randomized way) satisfying some constraints, say polynomial running time. Given different algorithms, we compare them by how they perform with respect to some objective function. This can be objectively measured by the *approximation ratio*, which is the ratio between the objective the algorithm produces and the best possible objective. The closer to 1 this ratio is, the better the algorithm performs.

One underlying assumption in this model is that the algorithm has direct access to the input. In other words, the choice of the algorithm to be used does not affect the input. While reasonable in most applications, this assumption breaks in market settings. In such settings the input is distributed across many agents who derive some benefit from the output. Each agent will respond to the algorithm by manipulating the input to their benefit. Therefore, once an algorithm is deployed, it will execute not on the real input but on a "manipulated" input, which corresponds to the way agents will respond to the mechanism.

In order to study such settings, it is necessary to describe the way agents will respond to a generic algorithm once it is deployed. This is done by describing a *solution concept*, which maps a pair of input and algorithm to a set of "manipulated inputs". The most common solution concept in the literature is the concept of "Nash equilibrium", which assumes that players will reach an equilibrium where no agent can change his report and become better off. Other popular solution concepts that will be discussed in the thesis are “Bayes-Nash equilibrium”, which incorporates uncertainty and “outcomes of no-regret dynamics” which instead of assuming players converge to a particular equilibrium, studies the outcomes during the learning process.

After the solution concept is fixed, how does one characterize a good algorithm ? A natural generalization of the approximation ratio for strategic settings is called the *price of anarchy*, which measures the worse ratio between the optimal objective for a certain input and the objective of the output of the algorithm on a corresponding “manipulated” input under the solution concept. The price of anarchy measures the robustness of the algorithm with respect to strategic behavior: in the worst case, how much can strategic manipulation harm the objective considered. The closer it is to 1, the more robust the algorithm is.

Below, we will discuss twomain approaches to mechanism design:

The first approach is called *truthful mechanism design*, in which the designer restricts its attention to a subclass of mechanisms (algorithms) that have the truthtelling property under the considered solution concept, i.e., the manipulated input is expected to be equal to the true input. If we find an algorithm with this property, then the concepts of price of anarchy and approximation ratio collapse, since the algorithm is guarantee to run in the correct input.

Truthful mechanisms are often complicated and it is often desirable to sacrifice truthfulness for simplicity, which is crucial for practical applications. This leads us to the second approach, which is to design a mechanism that does not have the truthtelling property and accept the algorithm will not run on the true input, but on a manipulated input. This is interesting if one can provide guarantees that even though we are running the algorithm in the "wrong" input, the output is a good solution with respect to the original input. We call it a *non-truthful mechanism*.

1.2 The Main Inquiry of this Thesis

Our main inquiry in this thesis is to design mechanisms that are simple, yet have good efficiency and revenue properties. We will look at this from two perspectives: first we will analyze the Generalized Second Price (GSP) mechanism, a design that has become incredibly successful in practice and widely adopted in industry. We want to build a theory that explains why such mechanism turned out to be so successful. Understanding why GSP is successful is important for two reasons: first, it will allow us to apply the lessons learned to other settings. Second, we will be able to fine-tune those mechanisms.

Secondly, we look at it from the design perspective. We will seek to design new mechanisms for sponsored search incorporating features that are not natively handled by the original GSP mechanism, such as the relation between auctions for different search queries, budgets (financial constraints for the advertisers) or online supply (the fact that we need to allocate and price ads before knowing the entire supply). We want to design new mechanisms keeping the simplicity goal as much as we can. This perspective also aims to bring new insights to the existing designs.

We believe those perspectives are two sides of the same coin: from one side, we observe a successful story in practice and build a theory around it, and from the other, we try to incorporate real-life features in our theory to bring it closer to practice.

Analysis perspective: understanding the GSP auction The GSP mechanism starts by collecting bids from the agents, which represent the amount they are willing to pay for ad slots. Then it uses those bids to decide on how to allocate ad slots and how much agents will pay for slots they get. This mechanism has the interesting feature that agents might benefit from mis-reporting their value. In other words, it is not a truthful mechanism. Yet, GSP is widely adopted and seems to display good revenue and welfare behavior in practice.

It is intriguing that GSP is adopted even despite the existence of another mechanism for the same setting — the VCG mechanism — which is truthful and fully efficient. One can only speculate the reasons behind such choice. Historically, GSP was devised as an improvement over the Generalized First Price Auction (GFP), also not truthful, used by Overture/Yahoo [34] to prevent the bid volatility that was characteristic of GFP. This engineering decision produced a very successful system that generates billions of dollars in revenue. Changing a system of that scope implies a lot of risk, specially if it is in the core of the company's business model. This, combined with the practical success of GSP, provides a compelling reason why GSP instead of the more theoretically sound VCG mechanism.

The GSP mechanism might be preferable because is has a very simple description of the payment rules. Clarity in the description of the mechanism is a desirable goal in large markets, since most participants are not auction experts and should feel comfortable in participating nevertheless.

In fact, the adoption of a non-truthful mechanism instead of a truthful counterpart is not an isolated phenomenon. There is a beautiful and elegant theory developed by Vickrey, Clarke and Groves [77, 23, 44] called the VCG-theory on how to design truthful mechanism. Although elegant, it rarely finds its way to practical applications, as it was remarked by Ausubel and Milgrom [10]:

> *"it is useful to think of the VCG theory as a lovely and elegant reference point – but not as a likely real-world auction design. Better, more practical procedures are needed."* (Ausubel and Milgrom, 2007, in **The Lovely but Lonely Vickrey Auction**)

In many settings besides sponsored search, non-truthful mechanisms are the standard. Most auctions in financial markets can be modeled by variants of double auctions and uniform-price auction, which are non-truthful. In such auctions, agents submit demand and supply at each given price and a price that clears the market is computed. Truthtelling in such settings is sacrificed in favor of more important properties such as budget balance [77, 65]. In settings like combinatorial auctions, the truthful VCG design suffers from a series of problems discussed in Ausubel and Milgrom [10] and, in practice, designs like Core Selecting Combinatorial Auctions are favored. See [28] for an overview and [27] for an example in the UK Spectrum Auctions. The sale of electricity, which is a multi-billion dollar market, is another good example. Nontruthful mechanisms, such as Locational Marginal Pricing Mechanism (LMP), are preferred to their truthful counterparts. See [78] for technical details and [25] for a detailed exposition on the economics of such markets.

From this viewpoint, the study of the GSP mechanism provides a glimpse on what seems to be a more general and widespread phenomenon.

Design perspective: adding realistic constraints Traditional designs of sponsored search, such as GSP and VCG, do not explicitly handle budget constraints. They also assume that different keywords are independent, i.e., the way an advertiser behaved while bidding for the keyword "car" is independent of the way he bids for the keyword "tires", even if the bids come from the same advertiser. This independence breaks if advertisers have budget constraints, i.e., the total amount of money they have to spend across all keywords is limited. Dollars that are spent on the "car" keyword cannot later be spent on "tires". This motivates us to analyze a richer model where advertisers have budget constraints and bid across many different keywords.

We will also consider the fact that advertisers participate not in one but in many auctions for the same keyword. In fact, each time a search is issued for a particular keyword, the search engine runs an auction for that keyword. So, the advertisers expect to participate in many such auctions but it is not clear how many. On top of that, if advertisers have budget constraints, dollars spent on one auction cannot be spent in future auctions.

We believe that it is important to design auctions that take those issues explicitly into account. On the one hand, techniques developed along the way apply to a variety of settings, as routing, video on demand, matching markets, ... On the other hand, this research line produces insights that can feed into existing designs and make them more robust to budgets and uncertainty in the supply.

1.3 Analysis Approach to Sponsored Search

1.3.1 Sponsored Search Auctions and GSP

The sale of advertising space on the Internet, or Ad Auctions, is the primary source of revenue for many providers of online services and corresponds. According to a recent report [36], $\$25.8$ billion dollars were spent in online advertisement in the US in 2010. The main part of this revenue comes from search advertisement, in which search engines display ads alongside organic search results. The success of this approach is due, in part, to the fact that providers can tailor advertisements to the intentions of individual users, which can be inferred from their search behavior. A search engine, for example, can choose to display ads that synergize well with a query being searched. However, such dynamic provision of content complicates the process of selling ad space to potential advertisers. Each search query generates a new set of advertising space to be sold, each with its own properties determining the applicability of different advertisements, and these ads must be placed near-instantaneously.

The now-standard mechanism for resolving online search advertisement requires that each advertiser places a *bid* that represents the maximum she would be willing to pay if a user clicked her ad. These bids are then resolved in an automated auction whenever ads are to be displayed. By far the most popular bid-resolution method currently in use is the Generalized Second Price (GSP) auction. Edelman et al. [34] and Varian [75] observe that truthtelling is not a dominant strategy under GSP, and GSP auctions do not generally guarantee the most efficient outcome (i.e., the outcome that maximizes social welfare). Nevertheless, the use of GSP auctions has been extremely successful in practice.

This leads to the main inquiry stated in the previous section: *are there theoretical properties of the Generalized Second Price auction that would explain its prevalence?* Edelman et al. [34] and Varian [75] provide a partial answer to this question by showing that, in the full information setting, a GSP auction always has a Nash equilibrium that has same allocation and payments as the VCG mechanism, and therefore is efficient.

However, there are no guarantees that the particular equilibrium studied by Edelman et al and Varian will be selected, since GSP has potentially many other equilibria, many of which are not fully efficient. In such settings, the priceof anarchy analysis is a powerful tool for quantifying the potential loss of efficiency at equilibrium. In a setting without uncertainty, the price of anarchy is surprisingly small, indicating a loss of at most 22% of the welfare. One should note that, remarkably, the welfare loss of these auctions is bounded by a value that does not depend on the number of players or the number of advertisements for sale.

While the results on the full information model provide important insight into the structure of the GSP auction, we will argue that the Generalized Second Price auction is best modeled as a Bayesian game of partial information. Modeling GSP as a full information game assumes that each auction is played repeatedly with the same group of advertisers, and during such repeated play the bids stabilize. The resulting stable set of bids is well modeled by a full information Nash equilibrium. However, the set and types of players can vary significantly between rounds of a GSP auction. Each query is unique, in the sense that it is defined not only by the set of keywords invoked but also by the time the query was performed, the location and history of the user, and many other factors. Search engines take this into account by computing for each advertiser and each query a *quality score* (or factor), which is a number that measures how relevant each ad is to that particular query. The bids are then multiplied by the quality scores and the advertisers are ranked according to this product. Search engines use complex machine learning algorithms to select the ads. We will represent it as an exogenous random process. This results in uncertainty both about the competing advertisers, and about quality factors. We model this uncertainty by viewing the GSP auction as a Bayesian game, and ask: what are the theoretical properties of the Generalized Second Price auction *taking into account the uncertainty that the advertisers face?* Surprisingly, we show that in settings with uncertainty, the price of anarchy is still bounded by a small constant.

GSP and its sources of uncertainty. There are two main sources of uncertainty: the first and main one is about the quality factors that the search engine attributes to each advertiser and the second is about the valuations (*types*) of the players. These sources are different in nature: each advertiser has knowledge of (and can condition her behavior on) her own type, whereas quality factors are fully exogenous and are only revealed ex post. In the rest of the section, we discuss in detail these sources of uncertainty.

Each query is unique, in the sense that it is defined not only by the set of keywords invoked but also the time the query was performed, the location and history of the user, and many other factors. This *context* is taken into account by an underlying *ad allocation algorithm*, which is controlled by the search engine. The ad allocation algorithm not only selects which advertisers will participate in an auction instance, but also assigns a *quality factor* to each advertiser. As a first approximation we can think of the quality factor as a score that measures how likely that participant's ad will be clicked for that query. These quality factors are then used to scale the bids of the advertisers. These scaled bids are known as *effective bids*, which can be viewed as bids derived from a similarly-modified *effective type*. Under our assumption that quality factors measure clickability, the effective type of an advertiser is the expected valuation of displaying the ad (valuation of the ad times its likelihood of getting a click). The effective bid and effective type of a player are therefore random variables, which can be thought of as the original valuations multiplied by quality scores computed exogenously by the search engine. Athey and Nekipelov [8] point out that the uncertainty in quality factors produces qualitative changes in the structure of the game. Thus, even if players converge to a stationary bidding pattern, the resulting equilibrium cannot be described as the outcome of a full information game.

We model the uncertainty about the effective types of advertisers as a Bayesian, partial information game. That is, the inherent uncertainty due to context and the ad allocation algorithm can be captured via prior distributions over effective types, even when the true types of all potential competitors are fully known. The appropriate equilibrium notion is then the Bayes-Nash equilibrium with respect to these distributions. Our model allows arbitrary correlations between the types and quality factors. The uncertainty of ad quality and allocation mostly comes from the query context, and hence is best modeled by correlated distributions of types and ad quality. Search engines use complex machine learning algorithms to compute quality factors based on all available information about the context, whose outcome is hard to predict for the advertisers. We assume that the advertisers are aware of the distribution of quality factors, and that the quality factors computed by the search engine correspond to the clickability of the ad. Our results deteriorate gracefully if the outcome of the machine learning algorithm is not exact, but rather gives only an approximation to the clickability of an ad.

1.3.2 Social Welfare

Our main result is a bound on the Bayesian price of anarchy for the GSP auction. Specifically, we show that the price of anarchy is at most $2(1-\frac{1}{e})^{-1} \approx 3.16$, meaning that the social welfare in any Bayes-Nash equilibrium is at least $1/3.16$ of the optimal social welfare. Notice that this is an unconditional bound, as we make no assumptions on the distribution on valuation profiles and quality factors (it can, for example, be correlated) or on the number of players or slots.

Perhaps just as important as the bound, however, is the straightforward and robust nature of the GSP auction. In particular, our results extend to provide the same welfare guarantees for outcomes of no-regret learning. Also, this bound continues to hold even if players have asymmetric access to distributional information about types and quality scores, in the form of exogenously provided signals. It also degrades gracefully in the presence of approximately rational players or a small fraction of irrational players.

We achieve these bounds by identifying a property that encapsulates many of the insights that drive our bounds. Roughgarden [73] identified a class of games that he termed *smooth* games, defined via a similar property that is used to bound the price of anarchy. The smoothness criterion is quite strong, and does not apply in general to the GSP mechanism. We identify a weaker property, *semi-smoothness*, that is satisfied by the GSP auction, and is strong enough to also imply price of anarchy bounds.

We provide improved results for the case where there is no uncertainty, which is the traditional setting studied in [34, 75]. If valuations and quality factors are fixed, we prove that the social welfare in any *pure Nash equilibrium* is within a factor of $\frac{1}{2}(\sqrt{5}+1) \approx 1.618$ (the golden ratio) of the optimal one and we show a lower bound of 1.259.

1.3.3 Robustness

One feature of our results is that they hold for a variety of models regarding the rationality and the beliefs of the players. This robustness is particularly important in large-scale auctions conducted over the Internet, where assumptions of full information and/or perfect rationality of the participants are unreasonably strong.

Previously, we discussed the Bayes-Nash equilibrium as a main solution concept to analyze auction games when we want to take into account all the sources of uncertainty present in the environment. Now, we discuss relaxations of this concept that allow us to model features present in real-world ad auctions as assymmetric information, learning outcomes and approximate rationality.

1. Asymmetric information. There are different types of players in advertising markets, which may have differing levels of information about their competitors. We assume all players know their own valuations correctly, but some smaller players (such as individual advertisers) might be clueless about the valuations of the other players and expected behavior of quality scores, while others (say bidding agencies or large companies with web advertising departments) may have a much better understanding of how individual rounds of the auction will proceed. Even among this latter group, different advertisers may have access to different information. We can model such information asymmetries by giving each player access to an arbitrary player-specific signal that can carry information about the effective types of the auction participants. Our bounds on social efficiency in the Bayesian model hold in settings with such asymmetry in information.

2. Learning players. So far we have considered equilibria of the auction game. Instead of assuming that players who have played long enough will be in equilibrium, one can model the entire learning process more explicitly. One natural model is that players employ strategies that give them vanishingly small *regret* over time. Roughly speaking, such a model assumes that players observe the bidding patterns of others and modify their own bids in such a way that their long-term performance is at least as good as a single optimal strategy chosen in hindsight. Many simple bidding strategies yield low regret, such as Hart and Mas-Collel's regret matching strategy [47] or the multiplicative weight updating strategy of [55] (see also [6]). These strategies are not necessarily in equilibrium, but capture the intuition that players attempt to learn beneficial bidding strategies over time.

The players can use these standard learning algorithms to learn how to best bid given their valuation and signal. In other words, for each possible valuation and signal, repeated auctions allow players to learn how to best bid taking into account the varying bids of other players, and the uncertainty about quality factors, other players' valuations, and bidding strategies. We will consider the quality degradation of the average social outcome when all players employ strategies with small regret. Blum et al. [15] introduced the term *Price of Total Anarchy* for this analog of the price of anarchy.

3. Approximate rationality. One of the fundamental assumptions in auction analysis is that all players are perfectly rational utility optimizers. However, in reality (and especially in large online settings), it is natural to assume that some fraction of the players participating in an advertising auction might have unsophisticated bidding strategies. In fact, some players may not even play at equilibrium in the single-shot approximation of the GSP auction, or may only be able to find strategies that are approximately utility-maximizing. We discuss the robustness of our bounds to the presence of players bidding with limited (or no) rationality. As we shall see, the GSP auction has the property that its social welfare guarantees degrade gracefully when our assumptions about the rationality of the players are relaxed.

1.3.4 Revenue

After establishing how far the welfare of the GSP auction is from the optimal achievable welfare, we turn our attention to revenue. For the Bayesian setting, we compare the revenue extracted in GSP with the revenue that is extracted by the optimal mechanism in this setting. We consider the setting where valuations are drawn from identical and independent distributions that satisfy the regularity condition. We show that if we allow the auctioneer to include reserve prices the GSP auction always obtains a constant fraction (1/6th) of the optimal VCG revenue, in expectation. This means in particular, that for this setting, GSP with the appropriate reserve is competitive against *any* Bayes-Nash equilibrium of *any* mechanism.

One might also wish to bound the revenue of GSP with respect to the revenue of VCG without reserve prices, but we show that this is not possible: there are cases in which the VCG revenue is unboundedly greater than the GSP revenue. However, if the slot CTRs satisfy a certain well-separatedness condition — namely, that the click-through-rates of adjacent slots differ by at least a certain constant factor — then we prove that GSP always obtains a constant fraction of the VCG revenue even in settings of partial information, extending a result of Lahaie [53], who considered welfare under this assumption on the CTRs.

We consider the full information game. We prove that at any Nash equilibrium, the revenue generated by GSP is at least half of the VCG revenue, *excluding the single largest payment of a bidder*. Thus, as long as the VCG revenue is not concentrated on the payment of a single participant, the worst-case GSP revenue approximates the VCG revenue to within a constant factor. This result also holds with an arbitrary reserve price. We also provide an example illustrating that the factor of 2 in our analysis is tight, and the revenue of GSP at equilibrium may be arbitrarily less than the full revenue of VCG (without excluding a bidder).

Finally, we analyze the tradeoffs of the maximum revenue attainable by the full information GSP mechanism under different equilibrium notions. We demonstrate that there can exist inefficient, non-envy-free equilibria that obtain greater revenue than any envy-free equilibrium. However, we prove that if CTRs are *convex*, meaning that the marginal increase in CTR is monotone in slot position, then the optimal revenue always occurs at an envy-free equilibrium. This implies that when click-through rates are convex, the GSP auction optimizes revenue at an equilibrium that simultaneously maximizes the social welfare. The convexity assumption we introduce is quite natural and may be of independent interest. Note that this assumption is satisfied in the case when CTRs degrade by a constant factor from one slot to the next.

1.3.5 Related work

Here we provide a high-level review of the most important related work. We defer an in-depth survey for the subsequent chapters. There has been considerable amount of work on the economic and algorithmic issues behind sponsored search auctions – see the survey of Lahaie et al. [52] for an overview of the early work and the survey of Maille et al. [59] for recent developments.

The two seminal papers who proposed the main model of GSP adopted in our work are Edelman, Ostrovsky and Schwarz [34] and Varian [75]. The authors notice that even though truthtelling is not a dominant strategy under GSP, the full information game always has a Nash equilibrium that has the same allocation and payments as the VCG mechanism – and therefore is efficient. They focus on a subclass of Nash equilibria which is called *envy-free equilibria* in [34] and *symmetric equilibria* in [75]. They show that such equilibria always exist and are always efficient. In this class, an advertiser would not be better off after switching bids with the advertiser just above her. Note that this is a stronger requirement than Nash, as an advertiser cannot unilaterally switch to a position with higher click-through-rate by simply increasing her bid. In [34, 35, 75], informal arguments are presented to justify the selection of envy-free equilibria, but no formal game-theoretical analysis is done. We believe it is an important question to go beyond this and prove efficiency guarantees for all Nash equilibria. Lahaie [53] also considers the problem of bounding the social welfare obtained at equilibrium, but restricts attention to the special case that click-through-rates decay exponentially along the slots with a factor of $\frac{1}{\delta}$. Under this assumption, Lahaie proves a price of anarchy of $\min\{\frac{1}{\delta}, 1 - \frac{1}{\delta}\}$.

Gomes and Sweeney [42] study the GSP auction in the Bayesian setting, where player types are drawn from independent and identical distributions (without considering the uncertainty due to quality factors). They show that, unlike the full information case, there may not exist symmetric or socially optimal equilibria in this model, and obtain sufficient conditions on click-through-rates that guarantee the existence of a symmetric and efficient equilibrium. Athey and Nekipelov [8] study the effect of uncertainty of quality factors both from a theoretical and an empirical perspective.

1.4 Designing Sponsored Search Mechanisms

In previous sections we focused on analyzing the Generalized Second Price Auction, which is the de-facto mechanism for sponsored search. In the second part of this thesis we take a different approach and try to design a new mechanism for sponsored search incorporating new features which are absent from the standard GSP model: budgets and online supply.

Budgets refer to the fact that advertisers are financially constrained, i.e., they have a cap on how much they can spend across all items they buy. It turns out to be one of the most important practical features of practical advertisement systems. In fact, in the Google AdWords [1] interface, the first question the advertisers are asked when creating a new campaign is: "what is your budget?". This is asked even before they select keywords or bids. In modern advertisement systems, there is an option where the budget is the only thing chosen by the advertiser and the bids are optimized automatically. In the face of this practical consideration, it makes sense to study mechanisms that explicitely take budgets into consideration.

A second feature that is absent in most traditional models of sponsored search is the fact that the supply (pageviews) arrive in an online fashion. If budget constraints are present, this becomes an issue, since the mechanism needs to allocate them to advertisers and charge for it in real-time, without knowing of the entire supply.

1.4.1 Budget Constraints

Satisfying budget constraints while keeping incentive compatibility and efficiency is a challenging problem, and it becomes even harder in the presence of complex combinatorial constraints over the set of feasible allocations. In the presence of budgets, individual rationality and truthfulness cannot be satisfied at the same time as maximizing social welfare [31], and thus the goal of maximizing efficiency can be achieved mainly through *Pareto-optimal* auctions[2]. Therefore, a desirable goal under budget constraints is to design incentive-compatible (IC) and individually-rational (IR) auctions while producing Pareto-

[1] `http://adwords.google.com/`

[2] An auction is Pareto-optimal if it outputs an allocation and payments such that no alternative set of allocation and payments improves the utility of at least one agent and keeps the other agents at least as happy as before. Here, agents include bidders and the auctioneer, where the auctioneer's utility is its revenue.

optimal outcomes. The first successful example of such mechanisms was developed in the seminal paper of Dobzinski, Lavi and Nisan [31], where the authors adapt the clinching auction framework of Ausubel [9] to give a truthful mechanism that achieves Pareto-optimality. Their setting, however, captures only a simple allocation constraint: there is a limited supply of k items and each player has a value of v_i for each item (and hence value of $v_i \cdot t$ for getting t items) and budget B_i.

As for more general allocation constraints, there have been a couple of subsequent work capturing special families of allocation constraints, e.g., unit demands [4], or multi-unit demands with matching constraints [38]. In this thesis we extend those results to a large class of constraints, namely polymatroidal constraints, which includes sponsored search as a special case. We also ask the question of for which polyhedral environments one can design such auctions, and also identify simple environments for which designing such auctions is not possible.

Our results. Our results are mainly inspired by Sponsored Search Ad Auctions, but extend to a large class of environmenta which can be modelled by *polymatroids*. This class of environments also includes spanning tree auctions, bandwidth markets and video on demand. For polymatroidal environments give an auction that achieves all the desired properties, i.e., the auction satisfies incentive-compatible, individually-rationality, and produces Pareto-optimal outcomes while satisfying the budget constraints. We assume that the budgets are public, which was shown in [31] to be a necessary assumption[3].

[3]Dobzinski et. al. [31] showed that with private budgets, truthfulness and Pareto-optimality cannot be achieved using deterministic mechanisms - not even for multi-unit auctions.

While following Ausubel's framework to design this auction, we need to invent the main component of the mechanism, i.e., the *clinching step* that copes with the polyhedral allocation constraints. Our clinching step uses submodular minimization as a subroutine and only needs a value oracle access to the submodular function corresponding to the polymatroid. As a result, our mechanism has a clean geometric description that abstracts away the combinatorial complications of previous designs. This leaves the auctioneer free to focus on modeling the environment, and then use our mechanism as a black-box. This general technique not only generalizes (and simplifies) the previously known results like multi-unit auctions with matching constraints [31, 38], but also extend clinching auctions to many other applications like the AdWords Auction and settings like spanning tree auctions and video on demand [14]. Our main application is in sponsored search auctions where we model the AdWords Auction with multiple keywords and multiple position slots per keyword as a polymatroid called the *AdWords polytope* (See Section 2.7.1 for details).

In order to extend this result to more general polyhedral constraints, we turn our attention to 2-player auctions with budget constraints and prove several structural properties of Pareto-optimal truthful auctions for polyhedral environments In particular, we present a characterization of such auctions that results in various impossibility results and one positive result. On the positive side, we present a truthful individually rational Pareto-optimal auction for any environment if only one player is budget-constrained. On the other hand, if more than one player is budget-constrained, we illustrate simple polytope constraints for which it is impossible to achieve a truthful Pareto-optimal auction even for two players. Moreover, as a byproduct of this characterization, we get an impossibility result for multi-unit auctions with decreasing marginal utilities. This impossibility result disproves an implied conjecture by Ausubel [9] which has been reinforced by follow-up papers [31, 54].

1.4.2 Online Supply

The problem of selling advertisement on the web is essentially an online problem — the supply (pageviews) arrives dynamically and decisions on how to allocate ads to pageviews and price these ads need to be taken instantaneously, without full knowledge of the future supply. What makes these decisions complex is the fact that buyers have budget constraints, which ties the allocation and pricing decisions across different time steps. Another complicating feature of the online advertisement markets is that buyers are strategic and can misreport their values to their own advantage.

These observations have sparkled a fruitful line of research in two different directions. First is that of designing online algorithms where one assumes that the supply arriving dynamically, but makes a simplifying assumption that buyers are non-strategic. This line of research has led to novel tools and techniques in the design of online algorithms (see for example [62, 17, 30, 3]). The second line of research considers the design of incentive-compatible mechanisms that assumes that buyers are strategic, but that the supply is known beforehand. Handling budget constraints using truthful mechanisms is non-trivial since standard VCG-like techniques fail when the player utilities are not quasi-linear. In a seminal work, Dobzinski, Lavi and Nisan [32] showed that one can adapt Ausubel's clinching auction [9] to achieve Pareto-optimal outcomes for the case of multi-unit supply. In settings with budget constraints, the goal of maximizing social welfare is unattainable and efficiency is achieved through Pareto-optimal outcomes. In fact, if budgets are sufficiently large, Pareto-optimal outcomes are exactly the ones that maximize social welfare [40].

From a practical standpoint, it is important to understand what can be done when both the above scenarios are present at the same time. Motivated by this, we study the following question in this paper: *Can one design efficient incentive-compatible mechanisms for the case when agents have budget constraints and the supply arrives online?*

A closely related question was studied by Babaioff, Blumrosen and Roth [11], who asked weather it was possible to obtain efficient incentive-compatible mechanisms with online supply, but instead of budget constraints, they considered capacity constraints, i.e., each agent wants at most k items (capacity) rather than having at most B dollars to spend (budget). They showed that no such mechanism can be efficient and proved lower bounds on the efficiency that could be achieved.

Such lower bounds seem to offer a grim perspective on what can be done with budget constraints, since typically, budget constraints are less well-behaved than capacity constraints. On the contrary, and somewhat surprisingly, we show that, for budget constraints, it is possible to obtain incentive compatible and Pareto-optimal auctions that allocate and charge for items as they arrive, by showing that the Adaptive Clinching Auction in [32] for multi-unit supply can be implemented in an online manner. More formally, we show that the clinching auction for the multi-unit supply case satisfies the following *supply-monotonicity* property: Given the allocation and payments obtained by running the auction for initial supply s, one can obtain the allocation and payments for any other supply $s' \geq s$ by *augmenting* to the auction outcome for supply s. In other words, it is possible to find an allocation for the extra $s' - s$ items and extra (non-negative) payments such that when added to clinching auction outcome for the supply s, we obtain the clinching auction outcome for supply s'. Moreover, we show that each agent's utility is also monotone with respect to the supply, i.e., agents do not have incentive to leave the auction prematurely.

Our results. In the online supply setting, we study a restricted sponsored search model where there is only a single slot per page, but this page might receive multiple pageviews. Our main result is an online variant of the Adaptive Clinching Auction of [32] that allocated and charges for items as they arrive in an incentive-compatible way, without knowing the entire supply in advance. We do so by showing that the Adaptive Clinching Auction satisfies a *supply-monotonicity* property.

From a technical perspective, proving the above result requires a deeper understanding of the structure of the clinching auction, which in general is difficult to analyze because it is described using a differential ascending price procedure rather than a one-shot outcome like VCG. In order to do so, we study the description of the clinching auction given by Bhattacharya, Conitzer, Munagala and Xia [13] by means of a differential equation. At its heart, the proof of the supply monotonicity is a coupling argument. We analyze two parallel differential procedures whose limits correspond to the outcome of the clinching auction with the same values and budgets but different initial supplies. We prove that either one stays ahead of the other or they meet and from this point on they evolve identically (for carefully chosen concepts of 'stay ahead' and 'meet'). We identify many different invariants in the differential description of the auction, and use tools from real analysis to show that these invariants hold.

Towards better heuristics for ad-allocation. One of the main goals of this research program is to provide insights for the design of better heuristics to deal with budget-constrained agents in real ad auctions. Most heuristics in practice are based on bid-throttling or bid-lowering. Bid-throttling probabilistically removes a player from the auction based on her spent budget (throttling). Bid-lowering runs a standard second price auction with modified bids. While sound from an algorithmic perspective, bid-throttling and bid-lowering are not integrated with the underlying auction from the perspective of incentives. We believe clinching auctions provide better insights into designing heuristics that are more robust to strategic behavior.

1.4.3 Related Work

Alternative Ad Auctions design. We would like to survey some related work on the design of alternative Ad Auction mechanisms that take into account budget constraints. Feldman et al. [37] design an auction for the environment with one keyword and multiple slots. Their model is, however, different from the standard utilitarian utility model. Instead of being profit maximizers, the players are click maximizers, i.e., the players want to get as many clicks as possible without exhausting their budget and without paying more per click than their value, which is a simpler setting than ours. In order to design their auction, they describe the structure of the set of possible randomized allocations of players to slots. We note that the structure they identify is in fact a polymatroid and use this fact to apply our auction to this setting. We further extend this characterization to the setting with multiple keywords.

Also for one keyword and multiple slots, Ashlagi et al [7] design an auction for the usual utility model but relax the truthfulness requirement and get an auction that is Pareto-optimal for all ex-post Nash equilibria. The main weakness in the setting of [7] is that the agents are allowed to be allocated only to one slot position for all the different queries of the given keyword. However, in reality, agents can be allocated to different slot positions for different queries of a given keyword. In the restricted setting of [7], the Pareto-optimality requirement becomes easier to satisfy.

Independently of our work, Colini-Baldeschi et al [24] also study the problem of designing incentive compatible, individually rational, budget feasible and Pareto-optimal auctions for sponsored search. The authors present two auctions satisfying those properties: one for the case with a single keyword but multiple slots with different click-through-rates and one for the case of multiple keywords and multiple slots with homogeneous click-through-rates (i.e. all slots are identical).

Auctions with budgets. On the generic question of designing truthful auctions with budgets with the goal of achieving Pareto-optimal outcomes was by Dobzinski, Lavi and Nisan [31], followed by Fiat et al [38], that generalized the previous work to matching settings. Bhattacharya et al [13] show a budget-monotonicity property for the clinching auction of [31], therefore arguing that no player can improve his utility by under-reporting his budget. For the case of unit-demand players, Aggrawal et al [4] design auctions for unit-demand players with budget constraints.

On the question of maximizing revenue, Borgs et al [16] gave a truthful auction whose revenue is asymptotically within a constant factor of the optimal revenue. These results were improved by Abrams [1]. Subsequently, Hafalir, Ravi and Sayedi [45] relax the truthfulness requirement, moving to ex-post Nash equilibrium as a solution concept, and give an auction that, in equilibrium, has good efficiency and revenue properties. More recently, Pai and Vohra [71] gave a revenue-optimal auction for the Bayesian version of the problem. We would like to highlight that the above work focused on the multi-unit setting only.

Our polyhedral clinching auction also generalizes the ascending auction of Bikhchandani et al [14]. The authors consider environments where the set of allocations is defined by a polymatroid, but don't consider budget constraints.

On impossibility results for this setting, Fiat et al [38] gave an impossibility result for achieving Pareto-optimality for heterogeneous goods in the budgeted setting. It remained an open problem whether an auction was possible if goods where identical, i.e., utilities depended only on the number items acquired and not on which items they were. A very recent result by Lavi and May [54] shows an impossibility result for the case where the valuation can be an arbitrary function of the number of items - i.e. players are allowed to express complementarities. Since their setting is more expressive, an impossibility result is easier. Our impossibility result for multi-unit auctions can be seen as a stronger version of their result, since we allow players only to express valuations with diminishing marginals. This came as a surprise to us, since it was generally believed that such a positive result could be achieved using a variation of [31].

Online supply. The study of auctions with online supply was initiated in Mahdian and Saberi [58] who study multi-unit auctions with the objective of maximizing revenue. They provide a constant competitive auction with the optimal offline single-price revenue. Devanur and Hartline [29] study this problem in both the Bayesian and prior-free model. In the Bayesian model, they argue that there is no separation between the online and offline problem. This discussion is then extended to the prior-free setting. The results in [29] assume that the payments can be deferred until all supply is realized, while allocation needs to be done online.

Our work is more closely related to the work by Babaioff, Blumrosen and Roth [11], which study the online supply model with the goal of maximizing social welfare. Unlike previous work, they insist (as we also do) that payments are charged in an online manner. This is a desirable property from a practical standpoint, since it allows players to monitor their spend in real-time. Their results are mainly negative: they prove lower bounds on the approximability of social welfare in setting where the supply is online. Efficiency is only recovered when stochastic information on the supply is available.

We should also note that there is a long line of research at the intersection of online algorithms and mechanism design, mostly dealing with agents arriving and departing in an online manner. We refer to Parkes [72] for a survey.

1.5 Roadmap

I tried to make this thesis as self-contained as possible so that no pre-requisite is required other then basic familiarity with calculus, probability and linear algebra. Results in game theory, combinatorics and algorithms, when needed, are stated and explained, and either have a proof in this thesis or a reference.

Most of the background needed in Game Theory and Mechanism Design is presented in Chapter 2, although the reader can always benefit from a broader exposition, as the by-now-classic AGT book [69] or Hartline's Lecture Notes on Mechanism Design [48].

The first part of this thesis is composed of Chapter 3 which covers the welfare in GSP auctions and Chapter 4 which covers revenue bounds. The second part is composed of Chapter 5, which covers the design of auctions with budgets and Chapter 6 which covers online supply.

1.6 Bibliographic Notes

Equilibrium Analysis of GSP The resuls in this thesis appeared in the following papers: Paes Leme and Tardos [70] (FOCS'10), Lucier and Paes Leme [57] (EC'11), Lucier, Paes Leme and Tardos [56] (WWW'12). Some of the price of anarchy bounds presented in this thesis were improved by Caragiannis, Kaklamanis, Kanellopoulos and Kyropoulou in [18] and [19]. The material in the previous papers on social welfare bounds was combined (and further improved) in [20].

Now we briefly describe the best bounds currently known for GSP. For social welfare, the best known bound for the Bayes-Nash Price of Anarchy is 2.927, for Mixed Nash Equilibria, Coarse-Correlated Equilibria, and Outcomes of No-Regret Learning the best known bound is 2.310 and for pure Nash equilibrium is 1.282, nearly matching a 1.259 lower bound. We refer the reader to [20] for the details. For revenue, the best current bounds for the revenue of GSP in a Bayes-Nash equilibrium are 4.72 for regular distributions and 3.46 for MHR distributions, due to Caragiannis et al [19].

Auction Design for Sponsored Search The second part of this thesis on Auction Design for Sponsored Search with budgets and online supply appeared published in Goel, Mirrokni and Paes Leme [40] (STOC'12) and Goel, Mirrokni and Paes Leme [41] (SODA'13).

CHAPTER 2

TECHNICAL PRELIMINARIES

2.1 Basic Notation

Throught this thesis we will denote by $\mathbb{R}$ the set of real numbers, $\mathbb{Z}$ the set of integers, $\mathbb{Z}_+ = \{z \in \mathbb{Z}; z \geq 0\}$ the set of non-negative integers and $[n] = \{1, 2, \ldots, n\}$. Given two vectors $x, y \in \mathbb{R}^n$ we represent their dot-product as $x^t y$.

Also, given any base set Θ, we define the Θ^n as the set of n-dimensional vectors over Θ. For a vector $v \in \Theta^n$ and $S \subseteq [n]$, we denote by v_S the vector restricted to the S components. In particular we abbreviate $v_{[n]\setminus\{i\}}$ by v_{-i}. This will allow us to represent the vector v as $v = (v_i, v_{-i})$. Whenever we refer to (u, v_{-i}) for some $u \in \mathbb{R}$, we mean a vector in Θ^n that has u in the i-th component and v_j in any j-th component for $j \neq i$.

Also, given any set Ω, we represent by $\Delta(\Omega)$ or simply $\Delta\Omega$ the set of measures (distributions) over Ω. We sometimes abuse notation and denote by $\Delta\Omega$ the set of random variables assuming values in Ω. Which of them we mean will be clear from the context.

Given a random variable $X \in \Delta\mathbb{R}^n$ we define its expectation by $\mathbb{E}[X] \in \mathbb{R}^n$. Also, given $X \in \Delta\Omega$ and a measurable subset $S \subset \Omega$, we define $\mathbb{P}(X \in S)$ as the probability that $X \in S$. Similarly given a measure μ over Ω we define $\mathbb{E}_{x\sim\mu}[x]$ and $\mathbb{P}_\mu(S)$ in the natural way. We drop subscripts when obvious from the context.

2.2 Mechanism Design Basics

We begin by reviewing the basic concepts in mechanism design. For a more extensive treatment, we refer to Chapter 9 of [69] for a more CS-oriented text and to classic economics references [63, 51, 39]. We restrict our presentation to settings which will be useful in reading this thesis.

The essential goal in mechanism design is to implement a desirable outcome in a setting where agents are strategic and will not report their private information unless they are incentized to do so. In order to describe the problem formally, consider a set of n agents and a set of outcomes X. The value of each agent for each outcome $x \in X$ is private information of the agent and is encapsulated in his *type* $\theta_i \in \Theta_i$. Now, we can represent a value of an agent i of type θ_i for an outcome x as $v_i : \Theta_i \times X \to \mathbb{R}$.

A *direct-revelation mechanism* ask agents to report their types using the reports, chooses an outcome and charges payments. In other words, a mechanism is a pair of mappings $x : \times_i \Theta_i \to X$ and $\varphi : \times_i \Theta_i \to \mathbb{R}^n_+$. Notice that we *do not* assume that agents report their true type, for this reason, we will use θ_i to denote true types and $\tilde{\theta}_i$ to denote reported types. Also we denote by $\Theta = \times_i \Theta_i$ the set of type profiles.

The preferences of agents over the outcomes are expressed by means of utility functions. We will consider in this thesis two types of utility functions: the quasi-linear utility function, where:

$$u_i(\theta_i, \tilde{\theta}) = v_i(\theta_i, x(\tilde{\theta})) - \varphi_i(\tilde{\theta})$$

and the budgeted quasi-linear utility function, which is equal to the quasi-linear if $\varphi_i(\tilde{\theta}) \leq B_i$ for some publicly known value B_i and $-\infty$ otherwise.

2.3 Solution Concepts

A *solution concept* is a model of the players rationality that is used to describe how agents will respond to the mechanism once it is deployed. For each mechanism, they describe a set of reported inputs $\tilde{\theta}$ for each true type profile θ. More formally, a solution concept maps for each mechanism (x, φ) a true type profile to a set of possible reported type profiles. More generally, it maps a type profile to a set of distributions over type profiles. Formally, given $\theta \in \Theta$ a solution concept $\mathbf{Sol}$ associates a set $\mathbf{Sol}(\theta) \subseteq \Delta\Theta$. Now, we discuss the most common solution concepts:

Dominant Strategies. We say that $\tilde{\theta}_i$ is a dominant strategy for a player i of type θ_i if:

$$u_i(\theta_i, (\tilde{\theta}_i, \hat{\theta}_{-i})) \geq u_i(\theta_i, (\hat{\theta}_i, \hat{\theta}_{-i})), \forall \hat{\theta}$$

This allows us to define the Dominant Strategies solution concept, which associates each $\theta \in \Theta$ with $\mathbf{DomStr}(\theta) \subseteq \Theta$ such that $\tilde{\theta} \in \mathbf{DomStr}(\theta)$ iff for each i, $\tilde{\theta}_i$ is a dominant strategy for player i.

The dominant strategy solution concept is very powerful and it seems almost too good to be true. In fact it rarely exists, which motivated us to look at weaker solution concepts. On the other hand, it plays an important role in designing games. In fact, one very popular game/mechanism design strategy is to design a mechanism such that it has a dominant strategy for each type. We say that the mechanism is *dominant strategy truthful* if $\theta \in \mathbf{DomStr}(\theta)$ for all $\theta \in \Theta$.

Pure Nash Equilibrium. The Pure Nash Solution concept associates each $\theta \in \Theta$ with $\mathbf{Nash}(\theta) \subseteq \Theta$ such that $\tilde{\theta} \in \mathbf{Nash}(\theta)$ iff:

$$u_i(\theta_i, (\tilde{\theta}_i, \tilde{\theta}_{-i})) \geq u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i})), \forall i, \tilde{\theta}'_i \in \Theta_i$$

The concept of Pure Nash equilibrium is perhaps the most popular and natural among the solution concepts: the strategies chosen by each player are such that there is no player that can deviate and be better-off. Although conceptually simple, Nash equilibria do not always exist and even when they exist, they are not always easy to find.

Mixed Nash equilibrium. The Mixed Nash Solution concept associates each $\theta \in \Theta$ with $\mathbf{mNash}(\theta) \subseteq \Delta\Theta$ such that $\tilde{\theta} \in \mathbf{mNash}(\theta)$ iff the components $\tilde{\theta}_i$ are independent random variables and :

$$\mathbb{E}_{\tilde{\theta}} u_i(\theta_i, (\tilde{\theta}_i, \tilde{\theta}_{-i})) \geq \mathbb{E}_{\tilde{\theta}_{-i}} u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i})), \forall i, \tilde{\theta}'_i \in \Theta_i$$

The concept of Mixed Nash equilibrium is an extension of the Pure Nash equilibrium allowing for randomization. The celebrated Nash's Theorem states that for mild conditions on the strategy space (as Θ_i being finite for each i) a Mixed Nash equilibrium always exists. However, there is strong evidence that suggests that Mixed Nash Equilibria are hard to find algorithmically [26].

The solution concepts described until now are "equilibria" concepts, which assumes that the players somehow reach a stable point in which no profitable deviation exists. Another approach is to study the outcomes of a dynamic in which players repeated play the game using some sort of no-regret algorithm. This motivates the folllowing solution concept.

Outcomes of no-regret learning. Given a true type profile $\theta \in \Theta$, an infinite sequence $\tilde{\theta}^1, \tilde{\theta}^2, \ldots, \tilde{\theta}^t, \ldots$ is said to have the no-regret learning property if there is $R(t)$ with $\lim R(t)/t = 0$ as $t \to \infty$ such that for all $\tilde{\theta}'_i$:

$$\sum_{t=1}^{\tau} u_i(\theta_i, \tilde{\theta}^t) \geq \sum_{t=1}^{\tau} u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}^t_i)) - R(\tau)$$

This means that player i doesn't (much) regret playing his strategy rather then choosing the best single option in hindsight $\tilde{\theta}'_i$ and playing it throughout the game.

For each τ we can associate the random variable $\hat{\theta}^\tau$ that takes value $\tilde{\theta}^t$ with probability $1/\tau$ for $t = 1..\tau$. Now, we say that $\tilde{\theta} \in \mathbf{noRegret}(\theta) \subseteq \Delta\Theta$ iff there is a subsequence of $\hat{\theta}^\tau$ that converges in distribution to $\tilde{\theta}$.

The study of no-regret learning outcome motivates the following definition:

Coarse Correlated equilibrium. The Coarse Correlated Solution concept associates each $\theta \in \Theta$ with $\mathbf{ccNash}(\theta) \subseteq \Delta\Theta$ such that $\tilde{\theta} \in \mathbf{ccNash}(\theta)$ iff:

$$\mathbb{E}_{\tilde{\theta}} u_i(\theta_i, (\tilde{\theta}_i, \tilde{\theta}_{-i})) \geq \mathbb{E}_{\tilde{\theta}} u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i})), \forall i, \tilde{\theta}'_i \in \Theta_i$$

This solution concept is motivated by the study of outcomes of no-regret learning. It easily follows from the definitions that: $\mathbf{noRegret}(\theta) = \mathbf{ccNash}(\theta)$. Besides, Coarse Correlated equilibria can be easily found by Linear Programming approximated using very straightforward algorithms as Multiplicative Weight Updated [6] or Regret Matching [47]. The concept of Coarse Correlated equilibrium is also a generalization of another traditional solution concept:

Correlated Nash equilibrium. The Correlated Nash Solution concept associates each $\theta \in \Theta$ with $\mathbf{cNash}(\theta) \subseteq \Delta\Theta$ such that $\tilde{\theta} \in \mathbf{cNash}(\theta)$ iff:

$$\mathbb{E}_{\tilde{\theta}}[u_i(\theta_i, (\tilde{\theta}_i, \tilde{\theta}_{-i}))|\tilde{\theta}_i] \geq \mathbb{E}_{\tilde{\theta}}[u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i}))|\tilde{\theta}_i], \forall i, \tilde{\theta}'_i \in \Theta_i$$

This solution is motivated by a stronger form of regret, called *swap regret*. It means that for all players i and all strategies $\tilde{\theta}_i$, player i does not regret playing i rather then swapping it for some other $\tilde{\theta}'_i$ instead. More sophisticated Learning Dynamics converge to this type of equilibrium and, moreover, correlated Nash equilibria can be found using Linear Programming.

It follows from the definitions that for every $\theta \in \Theta$:

$$\mathbf{DomStr}(\theta) \subseteq \mathbf{Nash}(\theta) \subseteq \mathbf{mNash}(\theta) \subseteq \mathbf{cNash}(\theta) \subseteq \mathbf{ccNash}(\theta)$$

so properties proved for all Coarse Correlated equilibrium naturally carry out to all other solution concepts.

2.4 Bayesian Solution Concepts

Bayesian games or games of partial information model the scenario where agents are uncertain about the types of other agents. Each agent i knows his type θ_i but not the type θ_j for $j \neq i$. Instead, they know the distribution $\mathbf{F}_{-i}$ where θ_{-i} is drawn from. The strategy of each player will be a mapping that associates for each of this type θ_i a reported type $\tilde{\theta}_i(\theta_i)$.

Formally, a Bayesian-setting is one where the space of type profiles is equipped with a distribution $\mathbf{F} \in \Delta\Theta$. In this model it is assumed that types of the agents are drawn by nature from the distribution $\mathbf{F}$. A solution concept in the Bayesian setting maps a true distribution over the type space into a distribution over reported types.

Bayes Nash equilibrium. The Bayes-Nash solution concept associates each $\mathbf{F} \in \Delta\Theta$ with a reported distribution $\tilde{\mathbf{F}} \in \Delta\Theta$ which is described by means of a "bidding function" $\tilde{\theta}_i : \Theta_i \to \Delta\Theta_i$ such that $(\tilde{\theta}_1(\theta_1), \ldots, \tilde{\theta}_n(\theta_n)) \sim \tilde{\mathbf{F}}$ whenever $\theta \sim \mathbf{F}$. We say that $\tilde{\theta}_i(\cdot) \in \mathbf{BayesNash}(\mathbf{F})$ iff:

$$\mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, \tilde{\theta}(\theta))|\theta_i] \geq \mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i}(\theta_{-i})))|\theta_i], \forall i, \tilde{\theta}'_i \in \Theta_i$$

Correlated Bayes Nash equilibrium. The Correlated Bayes-Nash solution concept associates each $\mathbf{F} \in \Delta\Theta$ with a reported distribution $\tilde{\mathbf{F}} \in \Delta\Theta$ which is described by means of a joint "bidding function" $\tilde{\theta} : \Theta \to \Delta\Theta$ such that $\tilde{\theta}(\theta) \sim \tilde{\mathbf{F}}$ whenever $\theta \sim \mathbf{F}$. We say that $\tilde{\theta}_i(\cdot) \in \mathbf{cBayesNash}(\mathbf{F})$ iff:

$$\mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, \tilde{\theta}(\theta))|\theta_i, \tilde{\theta}_i(\theta_i)] \geq \mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i}(\theta_{-i})))|\theta_i, \tilde{\theta}_i(\theta_i)], \forall i, \tilde{\theta}'_i \in \Theta_i$$

We refer to Bergemann and Morris [12] for an extensive discussion of this solution concept. In speacial, for a discussion on how this solution concept is affected by the presence of additional information signals received by the players.

Coarse Correlated Bayes Nash equilibrium. The Coarse Correlated Bayes-Nash solution concept associates each $\mathbf{F} \in \Delta\Theta$ with a reported distribution $\tilde{\mathbf{F}} \in \Delta\Theta$ which is described by means of a joint "bidding function" $\tilde{\theta} : \Theta \to \Delta\Theta$ such that $\tilde{\theta}(\theta) \sim \tilde{\mathbf{F}}$ whenever $\theta \sim \mathbf{F}$. We say that $\tilde{\theta}_i(\cdot) \in \mathbf{ccBayesNash}(\mathbf{F})$ iff:

$$\mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, \tilde{\theta}(\theta))|\theta_i] \geq \mathbb{E}_{\theta\sim\mathbf{F}}[u_i(\theta_i, (\tilde{\theta}'_i, \tilde{\theta}_{-i}(\theta_{-i})))|\theta_i], \forall i, \tilde{\theta}'_i \in \Theta_i$$

2.5 Price of Anarchy and Revenue

Given a mechanism (x, φ) for a certain environment and a solution concept, we measure how good this mechanism is. We do so, by defining an objective function and measuring the outcomes of the mechanism with respect to such objective. The most natural such objective is called *Social Welfare*:

$$\mathbf{SW}(\theta, x) = \sum_i v_i(\theta_i, x)$$

one would like to compare the welfare of $\mathbf{SW}(\theta, x(\tilde{\theta}))$ of $\tilde{\theta} \in \mathbf{Sol}(\theta)$, or $\mathbb{E}_{\tilde{\theta}}[\mathbf{SW}(\theta, x(\tilde{\theta}))]$ if the solution concept if randomized, with the maximum achievable welfare $\max_{x^*} \mathbf{SW}(\theta, x^*)$. The worse case of this ratio is called the Price of Anarchy. Formally:

$$\mathbf{PoA}[\mathbf{Sol}, (x, \varphi)] = \max_{\theta \in \Theta, \tilde{\theta} \in \mathbf{Sol}(\theta), x^* \in X} \left[\frac{\mathbf{SW}(\theta, x^*)}{\mathbf{SW}(\theta, x(\tilde{\theta}))} \right]$$

If the elements of $\mathbf{Sol}(\theta)$ are distributions in $\Delta\Theta$, say when $\mathbf{Sol} = \mathbf{mNash}$ or $\mathbf{ccNash}$ then, simply add expectations:

$$\mathbf{PoA}[\mathbf{Sol}, (x, p)] = \max_{\theta \in \Theta, \tilde{\theta} \in \mathbf{Sol}(\theta), x^* \in X} \left[\frac{\mathbf{SW}(\theta, x^*)}{\mathbb{E}_{\tilde{\theta}}\, \mathbf{SW}(\theta, x(\tilde{\theta}))} \right]$$

For a Bayesian Solution concept, one does the same analysis, but taking expectations over the distributions:

$$\mathbf{PoA}[\mathbf{Sol}, (x, \varphi)] = \max_{\mathbf{F} \in \Delta\Theta, \tilde{\mathbf{F}} \in \mathbf{Sol}(\mathbf{F})} \left[\frac{\mathbb{E}_{\theta \sim \mathbf{F}} \max_{x^*} \mathbf{SW}(\theta, x^*)}{\mathbb{E}_{\theta \sim \mathbf{F}, \tilde{\theta} \sim \tilde{\mathbf{F}}}\, \mathbf{SW}(\theta, x(\tilde{\theta}))} \right]$$

We also want to draw the attention to the reader to the fact that the payments

φ do not appear explicitly in the definition of the price of anarchy. If agents have quasi-linear utilities, the social welfare sums the utility of all the agents and the utility of the auctioneer (which is the revenue), i.e.:

$$\mathbf{SW}(\theta, x) = \sum_i u_i(\theta_i, x) + \sum_i \varphi_i$$

and the payments cancel out. We would like to point out that even though different payment functions do not change the social welfare, they change the solution concept and therefore impact the equilibria being selected.

We are also interested in measuring the *revenue* of each mechanism, which is simply:

$$\mathbf{Rev} = \sum_i \varphi_i$$

We would like to measure $\min_{\tilde{\theta} \in \mathbf{Sol}(\theta)} \mathbf{Rev}(\varphi(\tilde{\theta}))$, adding expectations if the solution concept so requires. Unlike social welfare, there is not a clear benchmark against which to compare to evaluate if a mechanism is good or not. We defer this discussion to Chapter 4.

2.6 VCG Mechanism: a generic welfare-maximizing mechanism

Surprisingly, given any set of outcomes and valuations, if agents have quasi-linear utility functions, it is possible to design a mechanism that is dominant strategy truthful and always produces social welfare optimal outcomes. This is the Vickrey-Clarke-Groves mechanism [77, 23, 44]. Its description is quite simple: given reported types $\theta = (\theta_1, \ldots, \theta_n)$, the mechanism picks an outcome

$x^*(\theta)$ that maximizes the reported welfare. Formally:

$$x^*(\theta) \in \text{argmax}_{x \in X} \sum_{i=1}^{n} v_i(\theta_i, x)$$

and charges the following payments:

$$\varphi_i(\theta) = \left[\max_{x \in X} \sum_{j \neq i} v_j(\theta_j, x) \right] - \left[\sum_{j \neq i} v_j(\theta_j, x^*(\theta)) \right]$$

The VCG mechanism has two important properties: (i) individual rationality, which means that each player derives non-negative utility from the mechanism, i.e.:

$$v_i(\theta_i, x^*(\theta)) - \varphi_i(\theta) \geq 0$$

(ii) dominant strategy truthful, discussed in Section 2.2, which means that reporting the truth maximizes the player utility. We recall that mathematically, it means:

$$v_i(\theta_i, x^*(\theta)) - \varphi_i(\theta) \geq v_i(\theta_i, x^*(\theta_i', \theta_{-i})) - \varphi_i(\theta_i', \theta_{-i})$$

Theorem 2.6.1 *The VCG mechanism is invidually rational and dominant-strategy truthful. Moreover, the payments are always non-negative.*

Proof : Individual rationality comes easily from the definition of payments:

$$\begin{aligned} v_i(\theta_i, x^*(\theta)) - \varphi_i(\theta) &= \left[\sum_{j} v_j(\theta_j, x^*(\theta)) \right] - \left[\max_{x \in X} \sum_{j \neq i} v_j(\theta_j, x) \right] = \\ &= \left[\max_{x \in X} \sum_{j \neq i} v_j(\theta_j, x) \right] - \left[\max_{x \in X} \sum_{j \neq i} v_j(\theta_j, x) \right] \geq 0 \end{aligned}$$

To show dominant strategy truthfulness, note that:

$$[v_i(\theta_i, x^*(\theta)) - \varphi_i(\theta)] - [v_i(\theta_i, x^*(\theta_i', \theta_{-i})) - \varphi_i(\theta_i', \theta_{-i})] =$$

$$\left[\sum_{j} v_j(\theta_j, x^*(\theta)) \right] - \left[\sum_{j} v_j(\theta_j, x^*(\theta_i', \theta_{-i}) \right] \geq 0$$

by the definition of $x_i^*(\cdot)$. Non-negativity of the payments is obvious from the definition. ■

For more details and generalizations of VCG mechanism, we refer the reader to the survey by Nisan [69].

2.7 Sponsored Search Environment

2.7.1 Single-keyword environment

We consider an auction with n advertisers and n slots[1]. Each advertiser i is an agent whose private type is a real value $v_i \in \Theta_i = \mathbb{R}_+$ representing his value per click. The sequence $v = (v_1, \ldots, v_n)$ is referred to as the *type profile* (or *valuation profile*).

An *outcome* is an assignment of advertisers to slots. An outcome can be viewed as a permutation π with $\pi(j)$ being the advertiser assigned to slot j. We define for notational convenience $\sigma = \pi^{-1}$ and use $\sigma(i)$ to denote the slot where player i is allocated. The probability of a click depends on the slot as well as the advertiser shown in the slot. We use the model of separable click probabilities. We assume slots have associated *click-through-rates* $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$, and each advertiser i has a *quality factor* γ_i that reflects the clickability of the ad. When advertiser i is assigned to the k-th slot, she gets $\alpha_k \gamma_i$ clicks.

If advertiser i is assigned to slot j at a price of p_i per click (i.e. $\varphi_i = \alpha_j \gamma_i p_i$),

[1]We note that we can handle unequal numbers of slots and advertisers by adding virtual slots with click-through-rate zero or virtual advertisers with zero valuation per click.

then her *utility* is $\alpha_j \gamma_i (v_i - p_i)$, which is the number of clicks received times profit per click. Therefore the social welfare of an outcome π is $\mathbf{SW}(\pi, v, \gamma) = \sum_j \alpha_j \gamma_{\pi(j)} v_{\pi(j)}$, the total value of the solution for the participants. The social welfare also depends on the click-through-rates α_j, but throughout this thesis we will assume they are fixed and common knowledge, and as a result we suppress them in the notation. The optimal social welfare is $\mathbf{Opt}(v, \gamma) = \max_\pi SW(\pi, v, \gamma)$, the welfare generated by the socially efficient outcome. Note that the efficient outcome sorts advertisers by their *effective values* $\gamma_i v_i$, and assigns them to slots in this order. The effective value can be thought of as the expected value of showing the ad in a slot with click-through-rate equal to 1.

2.7.2 Geometric Representation

We can represent an outcome π by a vector $x \in \mathbb{R}^n_+$ where x_i is the number of clicks that player i gets, in other words, x is such that: $x_{\pi(j)} = \alpha_j \gamma_{\pi(j)}$. Therefore, we can represent the set of possible outcomes by the set $X \subseteq \mathbb{R}^n_+$ defined above.

This notation allows us, for example, to write the Social Welfare as a dot-product: $\mathbf{SW}(v, x) = v^t x$. In the case where $\gamma_i = 1$ for all i, then X is the set of all the permutations of the coordinates of the vector $(\alpha_1, \ldots, \alpha_n)$.

Notice that X is a finite set of n-dimensional vectors. It is easy to see that the set of vectors of expected number of clicks that are feasible from a randomized allocation, i.e. a distribution over assignments from players to slots, can be represented as the convex-hull of X. In the next section, we show that this convex-hull has a nice and useful description.

2.7.3 Multi-keyword environment

Consider now n advertisers and m keywords. Each advertiser i is interested in a subset of the keywords $\Gamma(i) \subseteq [m]$. We assume this is public information. For a keyword k, we denote by $\Gamma(k)$, the set of advertisers interested in this keyword. With each keyword k, we associate $|\Gamma(k)|$ positions. Position j for keyword k has click-through-rate α_j^k (possibly zero) such that $\alpha_1^k \geq \alpha_2^k \geq \ldots \geq \alpha_{|\Gamma(k)|}^k$ for each k.

Assuming that each keyword gets a large amount of queries, we see α_j^k as the sum of number of clicks that the j-th position of keyword k gets across all queries that it matches. We assume that the number of clicks a player gets in slot j of keyword k is a product $\gamma_i^k \cdot \alpha_j^k$.

Let $\mathcal{A}_k = \{\sigma_k : \Gamma(k) \hookrightarrow [|\Gamma(k)|]\}$ be the set of all allocations (one-to-one maps) from players to slots for keyword k. A deterministic outcome, is a set of allocations $\sigma_k \in \mathcal{A}_k$ for each keyword k. Again, we can represent this geometrically: given $\{\sigma_k\}_k$, consider x such that:

$$x_i = \sum_{k \in \Gamma(i)} \gamma_i^k \cdot \alpha_{\sigma_k(i)}^k$$

For a randomized allocation, let $\Delta(\mathcal{A}_k)$ be the distributions of such allocations. Given that, we can define the AdWords polytope in the following way: an allocation of clicks x is feasible if there is a distribution over allocations of players to slots for each keyword such that player i gets x_i clicks in expectation. More formally:

Definition 2.7.1 *The **AdWords polytope** is the set of feasible allocations* [2] *of clicks*

[2]Since the number of clicks is typically very large we treat them as divisible goods and con-

$x = (x_1, \ldots, x_n)$ *such that there are distributions* $\mathcal{D}_k \in \Delta(\mathcal{A}_k)$ *for each keyword, and*

$$x_i \leq \sum_{k \in \Gamma(i)} \mathbb{E}_{\sigma_k \sim \mathcal{D}_k}[\gamma_i^k \cdot \alpha_{\pi_k(i)}^k]$$

2.8 Polyhedral and Polymatroidal Environments

We call *polyhedral environments*, settings where the set of possible outcomes can be represented by a packing polytope $P = \{x \in \mathbb{R}_+^n; Ax \leq b\}$ for some $m \times n$ matrix with $A_{ij} \geq 0$ and $b \in \mathbb{R}_+^m$. Examples of polyhedral environments are ubiquitous in game theory (see [67, 68] for many examples).

A rich subclass of packing polytopes is the class of polymatroids, which are polytopes that can be written as $P = \{x \in \mathbb{R}_+^n; \sum_{i \in S} x_i \leq f(S)\}$ where $f : 2^{[n]} \to \mathbb{R}_+$ is a monotone submodular function, i.e., a function such that:

$$f(S) + f(T) \geq f(S \cup T) + f(S \cap T), \forall S, T \subseteq [n]$$

and

$$f(S) \leq f(T), \forall S \subseteq T \subseteq [n]$$

Such polymatroidal environments generalize many previously studied environments: the multi-unit auctions environment of Dobzinski et al [31] corresponds to the *uniform matroid* and the matching markets studied in Fiat et al [38] correspond to the *transversal matroid*. Bikhchandani et al [14] give many examples of polymatroid environments including scheduling with due dates, network planning, pairwise kidney exchange, spatial markets, bandwidth markets and multi-class queueing systems [14]. We exemplify some of those applications below:

sider also fractional allocations.

- **Multi-unit auctions:** $P = \{x; \sum_i x_i \leq Q\}$ where Q is the total supply.
- **Combinatorial auctions with matching constraints:** The auction considered in [38], there is a bipartite graph $([m], [n], E)$ between items $[m]$ and bidders $[n]$ and each buyer i has additive value 1 for each item j such that $(i, j) \in E$ and value 0 for each item not connected to him. We can represent this setting by a polymatroid where $f(S)$ is the number of items connected to some player in S. This is called the *transversal matroid*.
- **Video on demand [14]:** Consider company that provides video on demand that is located on a node s of a direct network with capacities on the edges $G = (V, E, c)$. Each buyer corresponds to a node in the network. An allocation x is feasible if it is possible to transmit at rate x_i for each player i simultaneously. This is possible if for each subset $S \subseteq [n]$ of players, $\sum_{i \in S} x_i$ is smaller then the min-cut from s to S. Using the submodularity of the cut-function, it is easy to see that the environment is a polymatroid.
- **Spanning tree auctions:** Consider the abstract setting where the agents are edges of a graph G and the auctioneer is allowed to allocate goods to a set only if it has no cycles. This corresponds to the graphical matroid of graph G. A more practical setting is when a telecommunication company owns a network that contains cycles and decides to auction their redundant edges. This setting corresponds to the dual-graphical matroid of G.

2.8.1 Sponsored Search as a Polymatroidal Environment

In this section we will show that the AdWords Polytope, for the case where $\gamma_i^k = 1$ for all i, k, is a polymatroidal environment. This will enable us to explore the combinatorial properties of polymatroids once designing mechanisms for this setting.

First we start by recalling the definition of the AdWords polytope for a single keyword: an allocation σ maps each player i to a position $\sigma(i)$. Therefore, an allocation vector x is feasible iff there is a probability distribution over allocations such that:

$$0 \leq x_i \leq \mathbb{E}_\sigma[\alpha_{\sigma(i)}]$$

Feldman et al [37] relate the problem of deciding if a vector x is feasible to a classical problem in scheduling theory - scheduling in related machines with preemptions ($Q|pmtn|C_{\max}$ [43]). What follows a re-statement of their characterization in a format that makes it clear it is a polymatroidal environment.

Lemma 2.8.1 (Feldman et al [37]) *An allocation vector x is feasible iff for each S, $x(S) \leq \sum_{j=1}^{|S|} \alpha_j$, where $x(S) = \sum_{i \in S} x_i$ for each set $S \subseteq [n]$.*

Notice that $f_k(S) = \sum_{j=1}^{|S|} \alpha_j^k$ is a submodular function, so the set of feasible allocations for the single-keyword setting is a polymatroid.

For the multiple-keyword setting, we say that an allocation vector x is feasible if we can write $x_i = \sum_{k \in \Gamma(i)} x_i^k$ in such a way that the vector $(x_i^k)_{i \in \Gamma(k)}$ is feasible for keyword k, i.e., $x^k(S) \leq f_k(S)$ for every $S \subseteq \Gamma(k)$.

For the multiple-keyword setting, we say that an allocation vector x is feasible if we can write $x_i = \sum_{k \in \Gamma(i)} x_i^k$ in such a way that the vector $(x_i^k)_{i \in \Gamma(k)}$ is feasible for keyword k, i.e., $x^k(S) \leq f_k(S)$ for every $S \subseteq \Gamma(k)$.

The fact that this allocation set is a polymatroid is a direct consequence of the following theorem, which is a polymatroidal version of Rado's Theorem due to

McDiarmid [61].

Theorem 2.8.2 (McDiarmid [61]) *Given a bipartite graph* $([n] \cup [m], E)$*, its neighborhood map* $\Gamma(\cdot)$*,* m *submodular functions* $f_1, \dots, f_m$ *and their respective polymatroids* $P_1, \dots, P_m$*, then the set:*

$$P^* = \{x \in \mathbb{R}^n_+; x_i = \sum_{k \in \Gamma(i)} x_i^k \text{ and } x^k \in P_k\}$$

is a polymatroid defined by the function

$$f^*(S) = \sum_k f_k(S \cap \Gamma(k))$$

.

2.8.2 Quality factors

So far, we assumed that the click-through-rate of player i allocated to slot j of keyword k depends solely on k and j. More generally, we would like to consider the click-through-rate of a slot depending also on the player allocated in that slot. Let $\alpha^k_{j,i}$ be the click-through-rate of position j of keyword j when player i is placed there. Traditionally, we consider the click-through-rates in a product form, i.e., $\alpha^k_{j,i} = \alpha^k_j \cdot \gamma^k_i$ where γ^k_i is called *quality factor*. Assuming quality factors are public information, one can, in a similar way, define a polytope of feasible allocations. In general it will not be a polymatroid.

If the quality factors are uniform among all queries, i.e., $\gamma^k_i = \gamma_i$, the the set of feasible allocations is given by $P^*_\gamma = \{x; (\frac{x_i}{\gamma_i})_i \in P^*\}$ where P^* is the AdWords polytope defined as a function of α^k_j. We call such polytopes *scaled polymatroids*.

2.9 Generalized Second Price Auction

For the first part of this thesis, we focus on a particular mechanism, the Generalized Second Price auction, which works as follows: the mechanism begins by eliciting players types. They are asked to submit a *bid*, which is his reported valuation. Given a bid profile b, we define the *effective bid* of advertiser i to be $\gamma_i b_i$, which is her bid modified by her quality factor, analogous to the effective value defined above. The auction sets $\pi(k)$ to be the advertiser with the kth highest effective bid (breaking ties arbitrarily). That is, the GSP mechanism assigns slots with higher click-through-rate to advertisers with higher effective bids. Prices per click are then set according to critical value: the smallest bid that guarantees the advertiser the same slot. When advertiser i is assigned to slot k (that is, when $\pi(k) = i$), this critical value is defined as

$$p_i = \frac{\gamma_{\pi(k+1)}}{\gamma_i} b_{\pi(k+1)}$$

where we take $b_{n+1} = 0$. We will write $u_i(b, \gamma)$ for the utility derived by advertiser i from the GSP mechanism when advertisers bid according to b:

$$u_i(b, \gamma) = \alpha_{\pi^{-1}(i)} \gamma_i (v_i - p_i) = \alpha_{\pi^{-1}(i)} [\gamma_i v_i - \gamma_{\pi(\pi^{-1}(i)+1)} b_{\pi(\pi^{-1}(i)+1)}].$$

Notice that π is a function of b, γ as well. In places where we need to be more explicit, we will write $\pi(b, \gamma, j)$ to be the advertiser assigned to slot j by GSP when quality factors are γ and the advertisers bid according to b. We will also write $\sigma(b, \gamma, i)$ for the slot assigned to advertiser i, again when advertisers bid according to b and quality factors are γ. In other words, $\sigma(b, \gamma, \cdot) = \pi^{-1}(b, \gamma, \cdot)$. We write $\pi^i(b_{-i}, \gamma, j)$ to be the advertiser that would be assigned to slot j if advertiser i did not participate in the auction. When b and γ are clear from the context, we write $\pi(i)$ and $\sigma(i)$ instead of $\pi(b, \gamma, i)$ and $\sigma(b, \gamma, i)$. We will also write $\nu(v, \gamma)$ for the optimal assignment of slots to advertisers for valuation profile v, so that $\nu(v, \gamma, i)$ is the slot that would be allocated to advertiser i in the optimal assignment[3].

2.9.1 No overbidding

It is important to note that, in both the full information and Bayesian settings, any bid $b_i > v_i$ is dominated by the bid $b_i = v_i$ in the GSP auction. If by bidding $b_i > v_i$, the next highest effective bid is greater than $\gamma_i v_i$, then the player gets negative utility. If on the other hand, the next highest effective bid is smaller or equal than $\gamma_i v_i$, then bidding $b_i = v_i$ would get the same slot and payment. Based on this, we make the following assumption for the rest of the paper:

Assumption: Players are *conservative* and do not employ *dominated* strategies in GSP auctions. This means that for pure strategies $b_i \leq v_i$, for mixed strategies $\mathbb{P}(b_i > v_i) = 0$, and for Bayesian strategies $\mathbb{P}(b_i(v_i) > v_i) = 0$ for all v_i.

We use this assumption to rule out unnatural equilibria in which advertisers apply dominated strategies such as bidding $b_i > v_i$. We remark that, in these equilibria, the social welfare may be arbitrarily worse than the optimal. It is therefore necessary to exclude dominated strategies in order to obtain meaningful bounds on the price of anarchy. We note, however, that this phenomenon is not specific to the GSP auction: such degenerate equilibria exist even in the

[3]We note that, since GSP makes the optimal assignment for a given bid declaration, we actually have that $\nu(v, \gamma, i)$ and $\sigma(v, \gamma, i)$ are identically equal. We define ν mainly for use when emphasizing the distinction between an efficient assignment for a valuation profile and the assignment that results from a given bid profile.

Vickrey auction for a single good, where truthful bidding is a weakly dominant strategy. Since the Vickrey auction is a special case of GSP auctions (where one slot has $\alpha_1 = 1$, all other slots have $\alpha_i = 0$ and all quality factors have $\gamma_i = 1$), this issue carries over to our setting. Consider the example of a single-item Vickrey auction, where truthful bidding of $b_i = v_i$ is a weakly dominant strategy. Yet with overbidding, there are equilibria where an arbitrary player with low valuation bids excessively high (and hence wins), while everyone else bids 0. Note, however, that this Nash equilibrium seems very artificial as it depends crucially on the low valuation player using the dominated strategy of overbidding. Indeed, such an advertiser is exposed to the risk of negative utility (if some other advertiser submits a new bid between her valuation and bid) without any benefit. We therefore take the position that advertisers will avoid such dominated strategies when participating in the GSP auction.

CHAPTER 3

EFFICIENCY OF EQUILIBRIA IN GSP

3.1 Price of Stability

Before we start discussing the main results in this thesis we begin by reviewing what is known from the structure of Nash and Bayes-Nash equilibria of the GSP auction. The results discussed in this section are due to Edelman, Ostrovsky and Schwarz [34], Varian [75] and Gomes and Sweeney [42].

A classical result in [34, 75] establishes the existence of an efficient Nash equilirium for GSP. More specifically, they define a concept called *envy-free equilibria* in [34] and *symmetric equilibria* in [75], which is a subset of the set of pure Nash equilibria. They both prove that this class is non-empty and that every equilibrium in this class maximizes social welfare.

Definition 3.1.1 *Given a full information GSP game with n players defined by click-through-rates $\alpha_1 \geq \ldots \geq \alpha_n$, quality scores $\gamma_1, \ldots, \gamma_n$ and valuations $v_1, \ldots, v_n$, we say that a set of bids $b_1, \ldots, b_n$ is an envy-free equilibrium iff for any pair j, k of players, player j wouldn't prefer player k's allocation and payments rather then his own. Mathematically:*

$$\alpha_j \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_k \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(k+1)} b_{\pi(k+1)})$$

where $\pi(j)$ is the players witht the j-th highest effective bid (i.e. highest $\gamma_i b_i$ value) and $b_{\pi(n+1)} = 0$.

It is simple to see that every envy-free equilibrium is a Nash equilibrium. In fact, a set of bids $b_1, \ldots, b_n$ is a Nash equilibrium by the definition iff:

$$\alpha_j \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_k \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(k)} b_{\pi(k)}), \qquad \forall k < j$$
$$\alpha_j \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_k \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(k+1)} b_{\pi(k+1)}), \quad \forall k \geq j$$

since a player can easily decrease his bid and acquire a lower slot paying the price the player that is occupying that slot is paying, but in order to acquire a higher slot it needs to pay *not* the price the player occupying that slot is *paying*, but the price the player occupying that slot is *bidding*.

Notice that for $k \geq j$ the inequalities describing Nash equilibria and envy-free equilibrium are the same, however, for $k < j$, the inequalities describing envy-free equilibrium are stricter, since $\gamma_{\pi(k+1)} b_{\pi(k+1)} \leq \gamma_{\pi(k)} b_{\pi(k)}$.

Lemma 3.1.2 ([34, 75]) *Every envy-free equlibrium is efficient.*

Proof : Let $\pi' : [n] \to [n]$ be any allocation. We will show that the welfare under π is at least as large as the welfare under π', i.e.:

$$\sum_j \alpha_j \gamma_{\pi(j)} v_{\pi(j)} \geq \sum_j \alpha_j \gamma_{\pi'(j)} v_{\pi'(j)}$$

Let $\sigma(j) = \pi^{-1}(j)$ and $\sigma'(j) = (\pi')i^{-1}(j)$. We can re-write the envy-free inequality in terms of σ by taking $j = \sigma(i)$:

$$\alpha_{\sigma(i)} \cdot (\gamma_i v_i - \gamma_{\pi(\sigma(i)+1)} b_{\pi(\sigma(i)+1)}) \geq \alpha_k \cdot (\gamma_i v_i - \gamma_{\pi(k+1)} b_{\pi(k+1)})$$

Now, consider the case where $k = \sigma'(i)$:

$$\alpha_{\sigma(i)} \cdot (\gamma_i v_i - \gamma_{\pi(\sigma(i)+1)} b_{\pi(\sigma(i)+1)}) \geq \alpha_{\sigma'(i)} \cdot (\gamma_i v_i - \gamma_{\pi(\sigma'(i)+1)} b_{\pi(\sigma'(i)+1)})$$

Summing the expression above for $i \in [n]$ gets us:

$$\sum_i \alpha_{\sigma(i)} \cdot \gamma_i v_i - \sum_i \alpha_{\sigma'(i)} \cdot \gamma_i v_i \geq \sum_i \alpha_{\sigma(i)} \gamma_{\pi(\sigma(i)+1)} b_{\pi(\sigma(i)+1)} - \sum_i \alpha_{\sigma'(i)} \gamma_{\pi(\sigma'(i)+1)} b_{\pi(\sigma'(i)+1)} = 0$$

which can be re-written as:

$$\sum_j \alpha_j \cdot \gamma_{\pi(j)} v_{\pi(j)} \geq \sum_j \alpha_j \cdot \gamma_{\pi'(j)} v_{\pi'(j)}$$

■

A less direct but perhaps more insightful proof is that the concept of envy-free equilibrium provides the dual variables for the natural LP that solves the following optimization problem: $\max_\pi \sum_j \alpha_j \gamma_{\pi(j)} v_{\pi(j)}$. Consider the following primal-dual pair of linear programs:

$$\begin{array}{ll|ll}
\max \sum_{ij} \alpha_j \gamma_i v_i x_{ij} & & \min \sum_i u_i + \sum_j \varphi_j & \\
\quad \sum_i x_{ij} \leq 1, & \forall j \quad (\varphi_j) & \quad u_i + \varphi_j \geq \alpha_j \gamma_i v_i, & \forall i, j \quad (x_{ij}) \\
\quad \sum_j x_{ij} \leq 1, & \forall i \quad (u_i) & \quad u_i \geq 0, & \forall i \\
\quad x_{ij} \geq 0, & \forall i, j & \quad \varphi_j \geq 0, & \forall j
\end{array}$$

Notice that the envy-free equilibrium condition provides the dual variables that certify the optimality of the solution $x_{ij} = 1$ when $\pi(j) = i$ and 0 otherwise. Taking:

$$u_{\pi(j)} = \alpha_j \cdot \left(\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} v_{\pi(j+1)}\right) \quad \text{and} \quad \varphi_k = \alpha_k \gamma_{\pi(k+1)} v_{\pi(k+1)}$$

we have that the dual equations correspond exactly to the envy-freeness conditions.

In light of the previous observation it is rather unsurprising that the set of envy-free equilibria is non-empty. Next we prove this fact in a more direct way and we present the connection established in [34] between the envy-free equilibria of GSP and the outcome of the VCG mechanism.

Lemma 3.1.3 *The set of envy-free equilibria of GSP is non-empty.*

Proof : We can sort the agents such that $\gamma_1 v_1 \geq \ldots \geq \gamma_n v_n$. Consider the following set of bids: $b_1 = v_1$ and for $i \neq 1$ we have:

$$b_i = \frac{1}{\alpha_{i-1}\gamma_i}\left[\sum_{j=i}^{n}(\alpha_{j-1} - \alpha_j)\gamma_j v_j\right]$$

It is straightforwards to check that $b_1 \geq b_2 \geq \ldots \geq b_n$ and that the envy-free conditions hold. ■

Indeed, if one wants to inspect the total payment of player i in this mechanism (assuming $\gamma_1 v_1 \geq \ldots \geq \gamma_n v_n$), one will notice that:

$$\varphi_i = \sum_{j=i+1}^{n}(\alpha_{j-1} - \alpha_j) \cdot \gamma_j v_j$$

which corresponds to the externality that player i imposes on the other players, in other words, that is the gain in welfare for all the other players except i that would be incurred in case i decided not to participate in the mechanism. This shows that this equilibrium of the GSP mechanism exactly mimics the outcome of the VCG mechanism discussed in Section 2.6, in the sense that it produces the same allocation and payments.

A consequence of the previous lemmas is:

Theorem 3.1.4 *The (pure) Price of Stability of the GSP mechanism is* 1*, i.e., there is always a fully efficient pure Nash equilibrium.*

However, it is not true that all equilibria of the GSP mechanism are fully efficient. Consider for example one instance with 2 players and 2 slots such that: $\alpha_1 = 1, \alpha_2 = \frac{1}{2}$, $v_1 = 1$, $v_2 = \frac{1}{2}$ and $\gamma_1 = \gamma_2 = 1$. In this case, the bids $b_1 = 0$, $b_2 = \frac{1}{2}$ constitute a pure Nash equilibrium and its efficiency is $1 \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = 1$ while the optimal has efficiency $1 \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{5}{4}$. In section 3.3.2 we discuss this example in more detail.

Price of Stability with Uncertainty In the Bayesian setting, where there is uncertainty about the valuations of the players, Gomes and Sweeney [42] show that GSP might fail to have an efficient Bayes-Nash equilibrium. In what follows we reproduce the argument in [42]:

Lemma 3.1.5 ([42]) *Consider a Bayesian GSP instance with 3 iid players with $v_i \sim$ Uniform$[0,1]$ and quality factors $\gamma_1 = \gamma_2 = \gamma_3 = 1$. Also, let the click-through-rates be $(1, \alpha, 0)$. For sufficiently high-values of α there is no efficient Bayes-Nash equilbrium, in other words, for any equilibrium:*

$$\mathbb{E}[1 \cdot v_{\pi(1)} + \alpha \cdot v_{\pi(2)}] < \mathbb{E}[1 \cdot v^{(1)} + \alpha \cdot v^{(2)}]$$

where $v^{(1)}$ and $v^{(2)}$ indicate the maximum and second maximum values in v_1, v_2, v_3 respectively.

The proof above is based on the revenue equivalence principle [64]. Such result has appeared many times in the literature with increasingly general forms - we refer to Milgrom [63] or Hartline [48] for an extensive discussion. Here we prove a version of it tailored to the proof of Lemma 3.1.5.

Lemma 3.1.6 (Revenue Equivalence) *Given n apriori-identical agents, i.e, $\gamma_1 = \ldots = \gamma_n$ and v_i are drawn iid from the same distribution $\mathbf{F}$. Let $b(v)$ be a symmetric equilibrium of GSP and a monotone increasing bidding function. Also, let $\mathbb{E}[\varphi_i^{GSP}(v)]$ be* *the expected payment of player i if his value is v and his bid is $b(v)$, with the expectation conditioned on the values of all the other agents drawn from $\mathbf{F}$. Also, let $\mathbb{E}[\varphi_i^{VCG}(v)]$ be his expected payment when his value is v on the VCG mechanism (and he bids v since the mechanism is truthful). Then:*

$$\mathbb{E}[\varphi_i^{GSP}(v)] = \mathbb{E}[\varphi_i^{VCG}(v)]$$

Proof : Let $\mathbb{E}[x_i^{\text{GSP}}(v)]$ and $\mathbb{E}[x_i^{\text{VCG}}(v)]$ be the expected number of clicks player i gets on the GSP and VCG mechanism respectively. In GSP by bidding $b(v)$ and in the VCG by bidding v. Notice that for any valuation vector $(v_1, \ldots, v_n)$, the allocations from players to slots is the same, therefore: $\mathbb{E}[x_i^{\text{GSP}}(v)] = \mathbb{E}[x_i^{\text{VCG}}(v)]$.

Notice that since $b(v)$ is an equilibrium, a player with value v doesn't want to bid as a player with value u, which means that:

$$v \cdot \mathbb{E}[x_i^{\text{GSP}}(v)] - \mathbb{E}[\varphi_i^{\text{GSP}}(v)] \geq v \cdot \mathbb{E}[x_i^{\text{GSP}}(u)] - \mathbb{E}[\varphi_i^{\text{GSP}}(u)]$$

Making $u = v + \epsilon$, re-arranging terms, dividing by ϵ and taking ϵ to zero, we get: $v \cdot \partial_v \mathbb{E}[x_i^{\text{GSP}}(v)] - \partial_v \mathbb{E}[\varphi_i^{\text{GSP}}(v)] \leq 0$. Taking $u = v - \epsilon$ and doing the same, we get: $v \cdot \partial_v \mathbb{E}[x_i^{\text{GSP}}(v)] - \partial_v \mathbb{E}[\varphi_i^{\text{GSP}}(v)] \geq 0$. Therefore:

$$v \cdot \partial_v \mathbb{E}[x_i^{\text{GSP}}(v)] - \partial_v \mathbb{E}[\varphi_i^{\text{GSP}}(v)] = 0$$

therefore:

$$\mathbb{E}[\varphi_i^{\text{GSP}}(v)] = \int_0^v u \cdot \partial_v \mathbb{E}[x_i^{\text{GSP}}(u)] du = \int_0^v u \cdot \partial_v \mathbb{E}[x_i^{\text{VCG}}(u)] du = \mathbb{E}[\varphi_i^{\text{VCG}}(v)]$$

■

Proof of Lemma 3.1.5 : If there is an efficient equilibrium, then it corresponds to a symmetric monotone bidding function $b(v)$, since if the bidding function is not symmetric or if it is not monotone, then with some probability the players will produce innefficient outcomes. Now, let's calculate the format of the symmetric function. In order to do that, we use the revenue equivalence theorem. First, we calculate the payment of a player with value v in VCG.

$$\mathbb{E}[\varphi_i^{\text{VCG}}(v)] = 2(1-v)\int_0^v \alpha x dx + 2\int_0^v \int_0^x (1-\alpha)x + \alpha y dy dx$$

where the first term corresponds to the event that among the other players one of them has a value above v (which happens with probability $1-v$) and other has value below v (hence the integral from 0 to v). In this case player i pays the value of the player below him. We multiply by 2 to take into account that their roles might be reversed. Now, in the second case, we consider the case where both players have value below v. In this case player i pays $(1-\alpha)$ times the highest value and α times the second. We again multiply by 2 to take into account that the roles might be reversed. Solving the integral, we get:

$$\mathbb{E}[\varphi_i^{\text{VCG}}(v)] = \alpha v^2 + \frac{2}{3}(1-2\alpha)v^3$$

Now, we get the revenue in GSP by:

$$\begin{aligned}\mathbb{E}[\varphi_i^{\text{GSP}}(v)] &= 2(1-v)\int_0^v \alpha b(x)dx + 2\int_0^v \int_0^x b(x)dydx = \\ &= 2(1-v)\int_0^v \alpha b(x)dx + 2\int_0^v xb(x)dx\end{aligned}$$

Now, we use the revenue equivalence theorem to say that $\mathbb{E}[\varphi_i^{\text{GSP}}(v)] = \mathbb{E}[\varphi_i^{\text{VCG}}(v)]$. By derivating what we get from this expression, we obtain:

$$2\left[-\alpha\int_0^v b(x)dx + \alpha(1-v)b(v) + vb(v)\right] = 2\alpha v + 2(1-2\alpha)v^2$$

$$\int_0^v b(x)dx = b(v)\cdot(1-v+\frac{v}{\alpha}) - v - \frac{1}{\alpha}\cdot(1-2\alpha)v^2$$

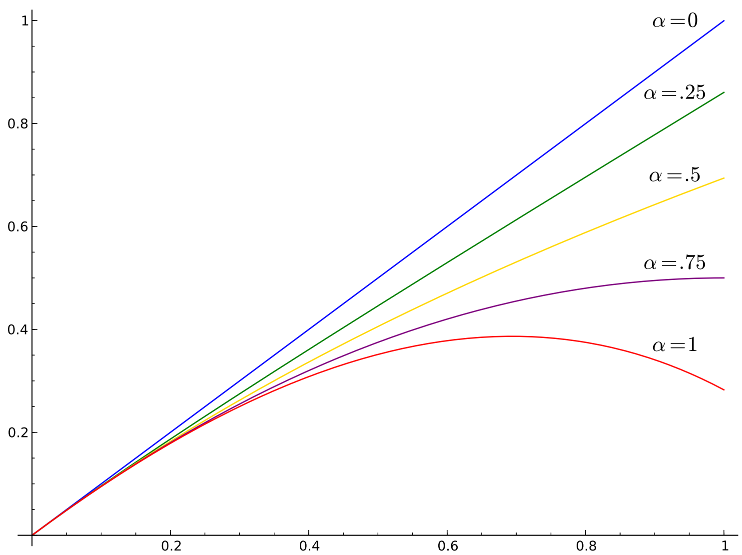

Figure 3.1: Solution $b(v) = B'(v)$ of the ODE for various values of α

Defining $B(v) = \int_0^v b(x)dx$, we get a standard first-order ODE:

$$B'(v) \cdot (1 - v + \frac{v}{\alpha}) - B(v) - v - \frac{1}{\alpha} \cdot (1 - 2\alpha)v^2 = 0 \qquad \text{s.t.} \qquad B(0) = 0$$

This ODE is in the form $f'(x) + p(x) \cdot f(x) + q(x) = 0$ for known p, q functions, so it can be solved by the integrating factor method. Solving it this way and then taking $b(v) = B'(v)$ and plotting for various values of α, we get the result in Figure 3.1.

For high values of α, the unique $b(v)$ that is a solution to the integral equation above is not monotone. Hence, there is no symmetric equilibrium Bayes-Nash equilibrium. Since an efficient equilibrium must be symmetric, there is no efficient Bayes-Nash equilibrium. ∎

3.2 Price of Anarchy with Uncertainty

Our main result is a bound on the price of anarchy for the Generalized Second Price auction with uncertainty. Recall that our model captures two types of uncertainty: uncertainty for player types and uncertainty about quality factors. Further, our result holds even in the presence of information asymmetry in the form of personalized signals available to the players. For simplicity of presentation, we focus on the setting where there are no signals and player valuations and quality factors are drawn from a known joint distribution $(\mathbf{F}, \mathbf{G})$.

Theorem 3.2.1 *The price of anarchy of the Generalized Second Price auction with uncertainty is at most* 3.164*. That is, for any fixed click-through-rates* $\alpha_1, \ldots, \alpha_n$*, any joint distribution* $(\mathbf{F}, \mathbf{G})$ *over valuation profiles and quality factors, and any Bayes-Nash equilibrium* b*,*

$$\mathbb{E}_{v,\gamma,b}[SW(\pi(b,\gamma), v, \gamma)] \geq \frac{1}{3.164}\mathbb{E}_{v,\gamma}[OPT(v,\gamma)].$$

Our proof is based on an extension of a proof technique introduced by Roughgarden [73], which he calls smoothness. We begin by reviewing this notion briefly in the context of a general game. Let $t = (t_1, \ldots, t_n)$ denote the (fixed) player types in a game, and $h = (h_1, \ldots, h_n)$ a pure strategy profile for the players, and let $U_i(t,h)$ denote the utility of player i with player types t, and strategy profile s. Let $sw(t,h)$ denote the social welfare generated by strategy profile s, and $sw^*(t)$ the maximum possible social welfare. Roughgarden defines (λ, μ)-smooth games as games where for all pairs of pure strategy profiles h, h', and any (fixed) vector of types t, we have

$$\sum_i U_i(t, h_i', h_{-i}) \geq \lambda \cdot sw(t, h') - \mu \cdot sw(t, h).$$

Roughly speaking, smoothness captures the property that if strategy profile h' results in a significantly larger social welfare than another strategy profile h, then a large part of this gap in welfare is captured by the marginal increases in the utility of each individual player when unilaterally switching her strategy from h_i to h'_i.

The GSP game is not strictly speaking smooth according to the original definition of smoothness in [73] but, as we show below, it is smooth according to the relaxed definition in Nadav and Roughgarden [66]. Here, we define a related property called *semi-smoothness* which is weaker then the original definition of smoothness in [73] but stronger then the definition in [66]. It has the advantage over the previous definitions that it implies Price of Anarchy bounds for Bayesian games with correlated types. The *semi-smoothness* property goes as follows: there is a *particular* (possibly randomized) strategy profile h' for which a smoothness-like inequality holds for any other pure strategy profile h. In other words, we prove the existence of a single bidding profile h' (depending on the types) that can be used by players unilaterally to improve the efficiency of GSP whenever its allocation is highly inefficient.

Definition 3.2.2 (Semi-Smooth Games) *We say that a game is (λ, μ)-**semi-smooth** if for each player i there exists some (possibly randomized) strategy $h'_i(\cdot)$ (depending only on the type of the player) such that,*

$$\sum_i \mathbb{E}_{h'_i(t_i)}[U_i(t, h'_i(t_i), h_{-i})] \geq \lambda \cdot sw^*(t) - \mu \cdot sw(t, h),$$

for every pure strategy profile h and every (fixed) type vector t. The expectation is taken over the random bits of $h'_i(t_i)$.

Analogous to Roughgarden's [73] proof, semi-smoothness also immediately implies a bound on the price of anarchy with uncertainty.

Lemma 3.2.3 *If a game is (λ, μ)-semi-smooth and its social welfare is at least the sum of the players' utilities, then the price of anarchy with uncertainty (and information asymmetries) is at most $(\mu + 1)/\lambda$.*

Proof : Consider a game in the Bayesian setting where player types are drawn from a joint probability distribution and let h be a Bayes-Nash equilibrium for this game. By the definition of the Bayes-Nash equilibrium, we have that $\mathbb{E}_{t_{-i},h}[U_i(t,h)|t_i] \geq \mathbb{E}_{t_{-i},h}[U_i(t,h'_i(t_i),h_{-i})|t_i]$ for every value the random variable $h'_i(t_i)$ may take. Hence, $\mathbb{E}_{t_{-i},h}[U_i(t,h)|t_i] \geq \mathbb{E}_{t_{-i},h}\mathbb{E}_{h'_i(t_i)}[U_i(t,h'_i(t_i),h_{-i})|t_i]$. Now taking expectation over t_i, we get $\mathbb{E}_{t,h}[U_i(t,h)] \geq \mathbb{E}_{t,h}\mathbb{E}_{h'_i(t_i)}[U_i(t,h'_i(t_i),h_{-i})]$. By summing over all players, and using the fact that the social welfare is at least the sum of the players' utilities, as well as the semi-smoothness property, we have

$$\begin{aligned}
\mathbb{E}_{t,h}[sw(t,h)] &\geq \mathbb{E}_{t,h}[\sum_i U_i(t,h)] \\
&\geq \mathbb{E}_{t,h}[\sum_i \mathbb{E}_{h'_i(t_i)}[U_i(t,h'_i(t_i),h_{-i})]] \\
&\geq \mathbb{E}_{t,h}[\lambda \cdot sw^*(t) - \mu \cdot sw(t,h)] \\
&= \lambda\mathbb{E}_t[sw^*(t)] - \mu\mathbb{E}_{t,h}[sw(t,h)].
\end{aligned}$$

Note that the third inequality follows by applying the semi-smoothness property for every fixed type vector and every pure strategy profile that are simultaneous outcomes of the random vectors t and h. The last inequality implies $\mathbb{E}_t[sw^*(t)] \leq \frac{\mu+1}{\lambda}\mathbb{E}_{t,h}[sw(t,h)]$, as claimed.

We remark that the proof holds without significant changes if we add information asymmetries in the game, i.e., if we assume that each player gets signals that reveal her type and refine her knowledge on the probability distributions of the types of the other players. All we need to change in this case is to replace the expectations over types with expectations over signals. ■

Notice that the usefulness of Lemma 3.2.3 lies in the fact that it can provide bounds on the efficiency loss for Bayesian games (and, as we will see in further sections, under even more general equilibrium concepts) by examining substantially more restricted settings. In the context of GSP auction games, it allows us to focus on identifying a (possibly randomized) deviating bid strategy for each player (i.e., a bid b'_i for each player i) so that the semi-smoothness inequality holds for every fixed valuation vector v and pure bidding profile b. By Lemma 3.2.3, this then immediately implies a bound on the price of anarchy of GSP auction games with uncertainty and information asymmetries.

We note that, technically speaking, the GSP auction does not immediately fit into the framework of semi-smoothness: advertiser payoffs depend on random quality factors which may be correlated with the type profile. However, this notational technicality is easily addressed by expressing advertiser utilities in expectation over quality scores. That is, expressing utilities in the GSP auction in the notation of general games, we have $U_i(v,b) = \mathbb{E}_\gamma[u_i(b,\gamma)|v]$. Since quality factors affect the social welfare as well, we have $sw^*(v) = \mathbb{E}_\gamma[OPT(v,\gamma)|v]$ and $sw(v,b) = \mathbb{E}_\gamma[SW(\pi(b,\gamma),v,\gamma)|v]$.

We are ready to prove that GSP auction games are semi-smooth.

Lemma 3.2.4 *The GSP auction game is $(1-\frac{1}{e},1)$-semi-smooth.*

Proof : We begin by rewriting the definition of semi-smoothness in the notation of GSP auctions. The GSP auction game is $(1 - \frac{1}{e}, 1)$-semi-smooth if and only if, for each valuation profile v, there exists a bid profile b' (with b'_i depending only on the valuation of player i) such that, for every bid profile b,

$$\sum_i \mathbb{E}_\gamma[u_i(b'_i, b_{-i}, \gamma)|v] \geq \left(1 - \frac{1}{e}\right) \mathbb{E}_\gamma[OPT(v, \gamma)|v] - \mathbb{E}_\gamma[SW(\pi(b, \gamma), v, \gamma)|v]. \tag{3.1}$$

We will actually establish the stronger inequality that, for *all* γ,

$$\sum_i u_i(b'_i, b_{-i}, \gamma) \geq \left(1 - \frac{1}{e}\right) OPT(v, \gamma) - SW(\pi(b, \gamma), v, \gamma). \tag{3.2}$$

The desired inequality (3.1) will then follow by taking (3.2) in expectation over the choice of γ (whose distribution may depend on the valuation profile v).

Before establishing inequality (3.2), we will prove the slightly weaker statement that the GSP auction game is $(1/2, 1)$-semi-smooth (which implies a bound of 4 on the price of anarchy with uncertainty). Choose a vector v of fixed valuations, a pure bidding profile b, and quality factors γ. Consider a (deterministic) deviating bid $b'_i = v_i/2$ for each player i. We distinguish between two cases (recalling that $\nu(i)$ is the slot assigned to player i in the efficient allocation given v and γ):

- If by bidding b'_i player i gets slot $\nu(i)$ or better, then $u_i(b'_i, b_{-i}, \gamma) \geq \alpha_{\nu(i)}\gamma_i v_i/2$, as the payment p_i cannot exceed her effective bid.
- If by bidding b'_i player i gets a slot lower than $\nu(i)$, then the effective value of the player $\pi(\nu(i))$ in slot $\nu(i)$ is at least $\gamma_i v_i/2$, as we assume no overbidding.

We conclude that, in either case,

$$u_i(b'_i, b_{-i}, \gamma) \geq \alpha_{\nu(i)}\gamma_i v_i/2 - \alpha_{\nu(i)}\gamma_{\pi(\nu(i))}v_{\pi(\nu(i))}.$$

Summing over all players, and noticing that $\sum_i \alpha_i \gamma_{\pi(i)} v_{\pi(i)} = SW(\pi(b,\gamma), v, \gamma)$, while $\sum_i \alpha_{\nu(i)} \gamma_i v_i = OPT(v,\gamma)$, we arrive at the claimed bound that the GSP auction game is $(1/2, 1)$-semi-smooth:

$$\sum_i u_i(b'_i, b_{-i}, \gamma) \geq \frac{1}{2} OPT(v,\gamma) - SW(\pi(b,\gamma), v, \gamma).$$

To improve the bound to $(1 - \frac{1}{e}, 1)$ we consider a randomized bid b' rather than the deterministic bid of $v_i/2$ considered above. Bid b'_i is a random variable on $[0, v_i]$ with density $f(y) = \frac{1}{v_i - y}$ for $y \in [0, v_i(1 - \frac{1}{e})]$ and $f(y) = 0$ otherwise. We will show that

$$\mathbb{E}_{b'_i}[u_i(b'_i, b_{-i}, \gamma)] \geq \left(1 - \frac{1}{e}\right) \alpha_{\nu(i)} \gamma_i v_i - \alpha_{\nu(i)} \gamma_{\pi(\nu(i))} b_{\pi(\nu(i))}. \tag{3.3}$$

By summing expression (3.3) for all i and using the fact that $b_{\pi(i)} \leq v_{\pi(i)}$ by the non-overbidding assumption, we obtain that the game is $(1 - \frac{1}{e}, 1)$-semi-smooth.

It remains to derive equation (3.3). We have that

$$\begin{aligned}
\mathbb{E}_{b'_i}[u_i(b'_i, b_{-i}, \gamma)] &\geq \mathbb{E}_{b'_i}[\alpha_{\nu(i)} \gamma_i (v_i - b'_i) \mathbb{1}\{\gamma_i b'_i \geq \gamma_{\pi(\nu(i))} b_{\pi(\nu(i))}\}] \\
&= \int_0^{v_i(1-\frac{1}{e})} \alpha_{\nu(i)} \gamma_i (v_i - y) \mathbb{1}\{\gamma_i y \geq \gamma_{\pi(\nu(i))} b_{\pi(\nu(i))}\} \frac{1}{v_i - y} dy \\
&= \alpha_{\nu(i)} \gamma_i \left[v_i \left(1 - \frac{1}{e}\right) - \frac{\gamma_{\pi(\nu(i))}}{\gamma_i} b_{\pi(\nu(i))} \right]^+ \\
&\geq \left(1 - \frac{1}{e}\right) \alpha_{\nu(i)} \gamma_i v_i - \alpha_{\nu(i)} \gamma_{\pi(\nu(i))} b_{\pi(\nu(i))}
\end{aligned}$$

which implies (3.2), completing the proof of Lemma 3.2.4. ∎

Combining Lemmas 3.2.3 and 3.2.4, we get the claimed bound on the price of anarchy.

Theorem 3.2.5 *The price of anarchy of the Generalized Second Price auction with uncertainty (and with information asymmetries) is at most* $2(1 - 1/e)^{-1} \approx 3.164$.

3.3 Pure Nash Equilibria in the Full Information Setting

In this section we turn our attention to the full information setting, where the quality factors γ are fixed and common knowledge. Without loss of generality we can assume that $\gamma_1 v_1 \geq \gamma_2 v_2 \geq \ldots \geq \gamma_n v_n$. In this setting the strategy of a player is a single bid $b_i \in [0, v_i]$, again assuming that players do not overbid. Our main result in this setting is the following:

Theorem 3.3.1 *The (pure) price of anarchy of the Generalized Second Price auction in the full information setting is at most the golden ratio $\frac{1}{2}(1+\sqrt{5}) \approx 1.618$. In other words, for any fixed click-through-rates α, valuation profile v, and quality factors γ, if b is a bid profile in pure Nash equilibrium, then $SW(\pi(b), v) \geq \frac{1}{1.618} \cdot OPT(v) \approx 0.618 \cdot OPT(v)$.*

In the special case of 3 players, we show a matching upper and lower bound of 1.259 for the Price of Anarchy. This provides a lower bound of 1.259 for the general Price of Anarchy. Progress on closing the gap between the general lower bound of 1.259 and the upper bound of 1.618 has been recently made in [21, 20], which improve the upper bound from 1.618 to 1.282.

3.3.1 Weakly feasible allocations

A central concept in the proof of the results stated above is the idea of *weakly feasible allocations*. Recall that each bid profile b defines an allocation π that is a mapping from slots to players $\pi : [n] \rightarrow [n]$.

Definition 3.3.2 (Weakly feasible allocations) *We say that an allocation π is* weakly feasible *if the following holds for each pair i, j of slots:*

$$\frac{\alpha_j}{\alpha_i} + \frac{\gamma_{\pi(i)} v_{\pi(i)}}{\gamma_{\pi(j)} v_{\pi(j)}} \geq 1. \tag{3.4}$$

We also use the term *weak feasibility condition* to refer to inequality (3.4).

Lemma 3.3.3 *If b is a Nash equilibrium of the GSP auction game, then the induced allocation π satisfies the weak feasibility condition.*

Proof : If $j \leq i$ the inequality is obviously true. Otherwise consider the player $\pi(j)$ in slot j. Since b is a Nash equilibrium, the player in slot j is happy with her outcome and does not want to increase her bid to take slot i, so: $\alpha_j(\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_i(\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(i)} b_{\pi(i)})$ since $b_{\pi(j+1)} \geq 0$ and $b_{\pi(i)} \leq v_{\pi(i)}$ then: $\alpha_j \gamma_{\pi(j)} v_{\pi(j)} \geq \alpha_i(\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(i)} v_{\pi(i)})$. ■

The concept of *weakly feasible allocations* encapsulates the fact that an allocation in equilibrium cannot be too far from the optimal. The optimal allocation is such that $\pi(i) = i$, since both $\{\alpha_i\}$ and $\{\gamma_i v_i\}$ are sorted. If an allocation is not optimal, then two slots $i < j$ have advertisers assigned to them such that $\pi(i) > \pi(j)$, i.e., they are assigned in the wrong order. Equation (3.4) implies that at least one of the two ratios is at least $1/2$, and hence whenever advertisers are assigned in the non-optimal order, then either (i) the two advertisers have similar effective values for a click, or (ii) the click-through-rates of the two slots are not very different; in either case their relative order does not affect the social welfare very much.

From Lemma 3.3.3, we can almost directly obtain a bound of 2 for the Price of Anarchy:

Theorem 3.3.4 *The (pure) price of anarchy of the Generalized Second Price auction in the full information setting is at most* 2*.*

Proof of Theorem 3.3.4 : Taking $j = \sigma(i)$ in the definition of weakly feasible allocations, we get that: $\alpha_{\sigma(i)}\gamma_i v_i + \alpha_i \gamma_{\pi(i)} v_{\pi(i)} \geq \alpha_i \gamma_i v_i$. Now, summing this for each player i, we get

$$2 \cdot SW(\pi(b), v) = \sum_i \alpha_{\sigma(i)}\gamma_i v_i + \sum_i \alpha_i \gamma_{\pi(i)} v_{\pi(i)} \geq \sum_i \alpha_i \gamma_i v_i = OPT(v).$$

■

3.3.2 Two and Three slots case

Theorem 3.3.5 *For* 2 *players and* 2 *slots, the price of anarchy is exactly* 1.25*. For* 3 *players and* 3 *slots, the price of anarchy is exactly* 1.259*. By* exactly *we mean that there is a particular GSP auction game with an equilibrium matching this bound.*

Proof : For two slots: consider an example with two players with valuations 1 and $1/2$ respectively, quality factors $\gamma_1 = \gamma_2 = 1$, and two slots with $\alpha_1 = 1$ and $\alpha_2 = 1/2$. The bids $b_1 = 0$ and $b_2 = 1/2$ are at equilibrium, resulting in a social welfare of 1, while the optimal social welfare is 1.25.

We can easily get a matching lower bound, by bounding the inefficiency of weakly feasible allocations. We will assume $\gamma_1 = \gamma_2 = 1$, but we do so only not to over-pollute the notation. The exact same proof works with generic quality factos. We consider a GSP auction game with two slots with click-through-rates $\alpha_1 \geq \alpha_2 = \beta\alpha_1$, for $\beta \in [0, 1]$ and two advertisers with valuations $v_1 \geq v_2 =$

λv_1, for $\lambda \in [0,1]$. The only non-optimal weakly feasible allocation π assigns advertiser 1 to slot 2 and advertiser 2 to slot 1. Its social welfare is $SW(\pi, v) = \alpha_1 v_2 + \alpha_2 v_1 = \alpha_1 v_1(\beta + \lambda)$, while the optimal social welfare is $OPT(v) = \alpha_1 v_1 + \alpha_2 v_2 = \alpha_1 v_1(1 + \beta\lambda)$. Furthermore, the weak feasibility condition for advertiser 1 implies that $\alpha_2 v_1 \geq \alpha_1(v_1 - v_2)$, i.e., $\beta \geq 1 - \lambda$. We have that

$$\frac{OPT(v)}{SW(\pi, v)} = \frac{1 + \beta\lambda}{\beta + \lambda} \leq \frac{1 + (\beta + \lambda)^2/4}{\beta + \lambda} \leq 5/4$$

where the first inequality holds since the product $\beta\lambda$ is maximized when $\beta = \lambda = (\beta+\lambda)/2$ and the second inequality holds since $\beta+\lambda \in [1,2]$ and the function $\frac{1+x^2/4}{x}$ is non-increasing in $x \in [1,2]$.

For three slots: Fix one permutation π. If there is an i s.t. $\pi(i) = i$ then it is easy to show the Price of Anarchy is bounded by 1.25. This excludes all but two allocations which we analyze below. They are: (i) $\pi = [2,3,1]$ and (ii) $\pi = [3,1,2]$.

Case (i): $\pi = [2,3,1]$. We can write the price of anarchy as:

$$PoA = \frac{\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3}{\alpha_3 v_1 + \alpha_1 v_2 + \alpha_2 v_3}$$

Now, notice that the coefficient of v_2 is smaller in the numerator than in the denominator. The same is true for v_3. Now, we use the following simple observation about ratios: if $a \leq b$ and $v \geq v'$ then: $\frac{x+av}{y+bv} \leq \frac{x+av'}{y+bv'}$, which is natural, because decreasing v we decrease the denominator more than the numerator. Now, we use that technique to bound v_2 and v_3 in terms of v_1:

- $v_2 \geq \frac{\alpha_1 - \alpha_3}{\alpha_1} v_1$
- $v_3 \geq \frac{\alpha_2 - \alpha_3}{\alpha_2} v_1$

The first inequality comes from the Nash inequalities $\alpha_3(v_1 - 0) \geq \alpha_1(v_1 - b_2) \geq \alpha_1(v_1 - v_2)$ and the second comes from the fact that $\alpha_3(v_1 - 0) \geq \alpha_2(v_1 - b_3) \geq \alpha_2(v_1 - v_3)$. Now, we get:

$$PoA \leq \frac{\alpha_1 v_1 + \alpha_2 \left[\frac{\alpha_1 - \alpha_3}{\alpha_1} v_1\right] + \alpha_3 \left[\frac{\alpha_2 - \alpha_3}{\alpha_2} v_1\right]}{\alpha_3 v_1 + \alpha_1 \left[\frac{\alpha_1 - \alpha_3}{\alpha_1} v_1\right] + \alpha_2 \left[\frac{\alpha_2 - \alpha_3}{\alpha_2} v_1\right]} \tag{3.5}$$

Which allows us to eliminate v_1 and optimize for α. By standard techniques one can prove that the optimum is 1.25913 which is the root of a fourth degree equation. The values for which it is maximized are $\alpha_1 = 1, \alpha_2 = 0.55079, \alpha_3 = 0.4704$.

Case (ii): $\pi = [3, 1, 2]$. We can write the price of anarchy as:

$$PoA = \frac{\alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3}{\alpha_2 v_1 + \alpha_3 v_2 + \alpha_1 v_3}$$

and again we use the same trick of realizing that $v_1 \leq \frac{\alpha_1}{\alpha_1 - \alpha_2} v_3$ by the fact that player 1 doesn't want to get the first slot, and $v_2 \leq \frac{\alpha_1}{\alpha_1 - \alpha_3} v_3$ by the fact that player 2 doesn't want to take the first slot. That gives us:

$$PoA \leq \frac{\alpha_1 \left[\frac{\alpha_1}{\alpha_1 - \alpha_2} v_3\right] + \alpha_2 \left[\frac{\alpha_1}{\alpha_1 - \alpha_3} v_3\right] + \alpha_3 v_3}{\alpha_2 \left[\frac{\alpha_1}{\alpha_1 - \alpha_2} v_3\right] + \alpha_3 \left[\frac{\alpha_1}{\alpha_1 - \alpha_3} v_3\right] + \alpha_1 v_3}$$

which has the same solution 1.25913 when maximized. Now, it is maximized for $\alpha_1 = 1, \alpha_2 = 0.5295, \alpha_3 = 0.1458$. In fact, it is not hard to see that those two PoA expressions have the same maximum: given a point $(1, \alpha_2, \alpha_3)$ (wlog we can consider $\alpha_1 = 1$ because the expression is homogeneous), the second expressions evaluates to the same value in the point $(1, 1 - \alpha_3, \frac{\alpha_2 - \alpha_3}{\alpha_2})$. ■

We proved that 1.259 is the tight Price of Anarchy for 3 slots (we can use the optimization results in Case(i) to generate a tight example). We also conjecture that this is the correct Price of Anarchy for any $n \geq 3$. Moreover, we conjecture that the allocation maximizing the Price of Anarchy for n slots is $\pi = [2, 3, 4, \ldots, n, 1]$, i.e., the player with higher value takes the bottom slot and all players $i > 1$ take slot $i - 1$. Then, if this is the case, we can prove our desired theorem by showing the following result:

Lemma 3.3.6 *If an equilibrium with n players and n slots is such that $\sigma(1) = n$ and $\sigma(i) = i - 1$ for the other players, then the Price of Anarchy is 1.25913.*

Proof : Following a proof scheme similar to used in the previous Theorem we can write:

$$PoA = \frac{\alpha_1 v_1 + \sum_{i>1} \alpha_i v_i}{\alpha_n v_1 + \sum_{i>1} \alpha_{i-1} v_i} \leq \frac{\alpha_1 + \sum_{i>1} \alpha_i \left[\frac{\alpha_{i-1} - \alpha_n}{\alpha_{i-1}}\right]}{\alpha_n + \sum_{i>1} \alpha_{i-1} \left[\frac{\alpha_{i-1} - \alpha_n}{\alpha_{i-1}}\right]}$$

This boils down to optimizing a function on multiple variables. It can be shown using standard techniques from optimization that the optimum is the same of equation 3.5. In fact, if $(\alpha_1, \alpha_2, 1)$ is a solution to 3 slots, then $(\alpha_1, \alpha_2, 1, \ldots, 1)$ is a solution for n slots. ∎

3.3.3 Golden Ratio Upper Bound for the Price of Anarchy

Now, we are ready to prove the main result in this section, Theorem 3.3.1.

Proof of Theorem 3.3.1 : As before, we prove the desired bound for all weakly feasible permutations. We also assume here $\gamma_i = 1$ for all i not to overcomplicate the notation. The exact same proof works for generic quality factors.

We define a sequence of values r_k so that we can prove that for k slots social welfare is at least an r_k fraction of the optimum, and prove that r_k converges to the desired bound. Let $r_2 = 1.25$ and suppose we have $r_2, r_3, ..., r_{n-1}$ and that this property holds for them. Let's calculate some "small" value of r_n so that the property still holds.

Again, consider parameter α, v, a weakly feasible permutation π and let's assume $i = \pi^{-1}(1)$ and $j = \pi(1)$. If $i = j = 1$, this is an easy case and it is straightforward to see that in this case the price of anarchy can be bounded by r_{n-1}. If not, assume without loss of generality that $i \leq j$ (since equation 3.4 is symmetric in α and v we can just interchange the roles of them in the proof if $i > j$). Let $\beta = \frac{\alpha_1}{\alpha_i}$ and $\gamma = \frac{v_1}{v_j}$. We know that $\frac{1}{\beta} + \frac{1}{\gamma} \geq 1$. Following the lines of the proof of the last theorem we have:

$$\begin{aligned}\sum_k \alpha_k v_{\pi(k)} &= \alpha_i v_1 + \sum_{k \neq i} \alpha_k v_{\pi(k)} \geq \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} \left(\sum_{k=2}^{i} \alpha_{k-1} v_k + \sum_{k=i+1}^{n} \alpha_k v_k \right) \geq \\ &= \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} \left[\sum_{k=2}^{i} (\alpha_{k-1} - \alpha_k) v_k + \sum_{k>1} \alpha_k v_k \right] \geq \\ &\geq \frac{1}{\beta}\alpha_1 v_1 + \frac{1}{r_{n-1}} (\alpha_1 - \alpha_i) v_i + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k\end{aligned}$$

Now, we can use $i \leq j$ to say: $v_i \geq v_j = \frac{1}{\gamma} v_1 \geq \left(1 - \frac{1}{\beta}\right) v_1$.

$$\sum_k \alpha_k v_{\pi(k)} \geq \left[\frac{1}{\beta} + \frac{1}{r_{n-1}} \left(1 - \frac{1}{\beta}\right)^2 \right] \alpha_1 v_1 + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k$$

So, we would like to find some r_n such that we can say that $\sum_k \alpha_k v_{\pi(k)} \geq \frac{1}{r_n} \sum_k \alpha_k v_k$ for all $\beta \geq 1$, so we would like to have: $\frac{1}{r_n} \leq \min \left\{ \frac{1}{r_{n-1}}, \frac{1}{\beta} + \frac{1}{r_{n-1}} \left(1 - \frac{1}{\beta}\right)^2 \right\}$ for any $\beta \geq 1$. But notice some other bound we can get is:

$$\sum_k \alpha_k v_{\pi(k)} \geq \frac{1}{\gamma}\alpha_1 v_1 + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k \geq \left(1 - \frac{1}{\beta}\right) \alpha_1 v_1 + \frac{1}{r_{n-1}} \sum_{k>1} \alpha_k v_k$$

by following the lines of the proof of last theorem, but removing slot 1 and advertiser j in the inductive step. So another alternative is to get: $\frac{1}{r_n} \leq \min\left\{\frac{1}{r_{n-1}}, 1 - \frac{1}{\beta}\right\}$ for every $\beta \geq 1$. So if we can get $1/r_n$ bounded by the maximum of those two quantities, we are done. Summarizing that, we need:

$$r_n \geq \max\left\{ r_{n-1}, \left[\max\left\{1 - \frac{1}{\beta}, \frac{1}{\beta} + \frac{1}{r_{n-1}}\left(1 - \frac{1}{\beta}\right)^2\right\}\right]^{-1}\right\}$$

for all $\beta \geq 1$.

Now we need to evaluate for which value of $\frac{1}{\beta} \in (0,1]$ we have the minimum for $\max\left\{1 - \frac{1}{\beta}, \frac{1}{\beta} + \frac{1}{r_{n-1}}\left(1 - \frac{1}{\beta}\right)^2\right\}$. The minimum can be in two points: the minimum of the quadratic function or the intersection between those two functions. They intersect at $\frac{1}{\beta} = -r + 1 + \sqrt{r^2 - r}$ (where r stands for r_{n-1}) and the quadratic minimum is at $1 - \frac{1}{2}r$. So, for $r \geq \frac{4}{3}$, the minimum occurs in the intersection and for $r < \frac{4}{3}$, it occurs in the quadratic minimum. So:

$$r_n = \begin{cases} \left(1 - \frac{r_{n-1}}{4}\right)^{-1} & , r_{n-1} < \frac{4}{3} \\ \left(r_{n-1} - \sqrt{r_{n-1}^2 - r_{n-1}}\right)^{-1} & , r_{n-1} \geq \frac{4}{3} \end{cases}$$

since we want the smallest possible ratio. This allows to define r_k recursively from $r_2 = 1.25$ and it is easy to see that the sequence monotonically converges to the fixed point of that function which is the golden ration $\varphi = \frac{1+\sqrt{5}}{2} \approx 1.618$. This happens because the function that maps r_{n-1} to r_n is non-decreasing and has a fixed point in φ, so if $r_{n-1} \leq \varphi$ then $r_n \leq \varphi$.

■

3.4 Computational search for the GSP Price of Anarchy

In this section we describe how we can computationally search for lower bounds on the pure Price of Anarchy for GSP. In fact, this approach is not only useful to produce lower-bounds, but was very useful in providing us intuition to prove the results in the previous sections.

It is easy to see that we can focus our attention to the case where $\gamma_1 = \ldots = \gamma_n = 1$, since given any instance with α, γ, v, b we can substitute by an instance α, γ', v', b' where $\gamma_i' = 1$, $v_i' = \gamma_i v_i$ and $b_i' = \gamma_i b_i$. If b is a pure Nash equilibrium of the game induced by α, γ, v then b' is a Nash equilibrium of the game induced by α, γ', v'. Moreover, the ratio between the social welfare and the optimum is the same.

Consider the following problem: what is the worse Price of Anarchy one can obtain for an instance with n agents and click-through-rates $\alpha_1 \geq \ldots \geq \alpha_n \geq 0$? The Price of Anarchy should look like:

$$\beta = \frac{\sum_i \alpha_i v_i}{\sum_i \alpha_i v_{\pi(i)}} \tag{3.6}$$

for some permutation $\pi : [n] \to [n]$. This permutation should be such that there is a bid profile that supports it, i.e, a set of bids $b_1, \ldots, b_n$ such that:

$$\begin{aligned}
& b_i \leq v_i, \forall i \\
& b_{\pi(1)} \geq b_{\pi(2)} \geq \ldots \geq b_{\pi(n)} \geq 0 \\
& \alpha_j \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_k \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(k)} b_{\pi(k)}), \forall k < j \\
& \alpha_j \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(j+1)} b_{\pi(j+1)}) \geq \alpha_k \cdot (\gamma_{\pi(j)} v_{\pi(j)} - \gamma_{\pi(k+1)} b_{\pi(k+1)}), \forall k \geq j
\end{aligned} \tag{3.7}$$

where the first set of constraints is the no-overbidding condition, the second set of constraints is the fact that the bids generate that permutation and the third set indicates that b form a Nash equilibrium. We should also impose the following condition on the valuation vector such that $\sum_i \alpha_i v_i$ is the optimum:

$$v_1 \geq v_2 \geq \ldots \geq v_n \geq 0 \tag{3.8}$$

Essentially, we want to solve the problem:

$$\max \beta \quad \text{s.t.} \quad \text{constraints (3.6), (3.7), (3.8)}$$

This is not quite a linear program yet, since equation 3.6 is non-linear. However, this is easy to solve. Notice that the problem is homogeneous in v, b, i.e., given any scalar $s > 0$, if (β, v, b) is a solution, then $(\beta, s \cdot v, s \cdot b)$ is also a solution. So, we can substitute equation 3.6 by:

$$\begin{aligned} \beta &= \sum_i \alpha_i v_i \\ 1 &= \sum_i \alpha_i v_{\pi(i)} \end{aligned} \tag{3.9}$$

Now, we can write it as the LP:

$$\max \beta \quad \text{s.t.} \quad \text{constraints (3.9), (3.7), (3.8)}$$

Call $\mathbf{LP}(\alpha, \pi)$ the solution of the LP above for some $\alpha = (\alpha_1, \ldots, \alpha_n)$ and a permutation π of $[n]$. We note that this program is always feasible taking $v = (1, 1, \ldots, 1)$, $b_{\pi(j)} = 1 - \frac{\alpha_n}{\alpha_j}$ and $\beta = 1$, this gives us a feasible solution.

Also, it is not hard ot see that the feasible region is bounded. Therefore, there is always a bounded optimal solution.

Together with this thesis we provide a code (in Octave) to solve this problem. This is in the file `gsp_poa.m` which implements the function

```
[beta, v, b] = gsp_poa(pi, alpha)
```

which takes a permutation and a vector of click-through-rates and returns the worse Price of Anarchy one can get with those click-through-rates and with this permutation.

The pure Price of Anarchy of GSP for n slots is given by:

$$\mathbf{PoA}(n) = \max_{\alpha \in \mathbb{I}_n} \max_{\pi \in \mathbb{S}_n} \mathbf{LP}(\alpha, \pi)$$

where $\mathbb{I}_n$ is the set of vectors $\alpha = (\alpha_1, \ldots, \alpha_n)$ such that $1 \geq \alpha_1 \geq \ldots \geq \alpha_n \geq 0$ and $\mathbb{S}_n$ is the set of permutations on n elements.

This approach is less useful for proving upper-bounds, but more for computationally searching for lower bounds on the actual Price of Anarchy. In fact, we we take $\mathbb{I}_n(\epsilon)$ to be the set of $\alpha \in \mathbb{I}_n$ where all coordinates are integral multiples of ϵ.

Since there are around $\epsilon^{-n}/n!$ vectors in $\mathbb{I}_n(\epsilon)$, we can compute the lower bound: $\max_{\alpha \in \mathbb{I}_n(\epsilon)} \max_{\pi \in \mathbb{S}_n} \mathbf{LP}(\alpha, \pi)$ by solving the problem $\mathbf{LP}(\alpha, \pi)$ roughly ϵ^{-n} times. By observing that the problem is homogeneous in α and we can take $\alpha_1 = 1$ w.l.o.g., we can reduce it to ϵ^{-n+1}. In the file `lower_bound_poa.m` we provide a code in Octave to compute this lower bound. In the following table, we provide the results for running this code for $n = 2, 3, 4, 5$:

n	ϵ used	**PoA**	α	v	b
2	0.01	1.2500	$[1.0, 0.5]$	$[1.0, 0.5]$	$[0.0, 0.5]$
3	0.01	1.2591	$[1.0, 0.55, 0.47]$	$[0.925, 0.490, 0.134]$	$[0.0, 0.491, 0.135]$

3.5 Quality of Learning Outcomes in GSP

In this section, we bound the average quality of outcomes in a repeated play of a GSP auction game where players employ strategies that guarantee no external regret. In both the full information setting and the setting with uncertainty, we can reduce the problem over declaration sequences to a problem over distributions. This will allow us to adapt our earlier bounds on the price of anarchy from Sections 3.2 and 3.3 to bound the price of total anarchy.

3.5.1 Learning in the full information setting

We will first focus upon the full information setting of the GSP auction. Recall that, in this model, the valuation profile v and quality factors γ are fixed and common knowledge. As in the previous section, we will assume that $\gamma_1 v_1 \geq \gamma_2 v_2 \geq \ldots \geq \gamma_n v_n$.

We will begin by proving a relationship between the price of total anarchy and the set of *coarse correlated equilibria* for the GSP auction in the full information model. Given a valuation profile v, a distribution $\mathbf{D}$ over bid profiles is called a coarse correlated equilibrium if

$$\mathbb{E}_{b\sim\mathbf{D}}[u_i(b)] \geq \mathbb{E}_{b\sim\mathbf{D}}[u_i(b'_i, b_{-i})], \forall i, b'_i.$$

As we shall show, the price of total anarchy can be bounded by considering the social welfare generated at any coarse correlated equilibrium.

Lemma 3.5.1 *The price of total anarchy in the full information setting is at most*

$$\sup_{v,\mathbf{D}\in\mathbf{ccNash}} \frac{OPT(v)}{\mathbb{E}_{b\sim\mathbf{D}}[SW(\pi(b),v)]}$$

where **ccNash** *is the set of coarse correlated equilibria.*

Proof : Consider a declaration sequence $D=(b^1,\ldots,b^t,\ldots)$ in the full information case. For each T let $\mathbf{D}^T$ be the distribution over bid profiles where each b^t for $t\leq T$ is drawn with probability $\frac{1}{T}$. Proving that the price of total anarchy is bounded by η is equivalent to showing that:

$$\liminf_T \mathbb{E}_{b\sim\mathbf{D}^T}[SW(\pi(b),v)] \geq \frac{1}{\eta}OPT(v).$$

Since the set of all possible bid profiles is compact, one needs to prove that for all distributions $\mathbf{D}$ such that there is a subsequence of $\{\mathbf{D}^T\}_T$ converging in distribution to $\mathbf{D}$ we have:

$$\mathbb{E}_{b\sim\mathbf{D}}[SW(\pi(b),v)] \geq \frac{1}{\eta}OPT(v).$$

It is therefore sufficient to show that such a $\mathbf{D}$ is a coarse correlated equilibrium. We note that the fact that the declaration sequence D minimizes external regret implies that, for each distribution $\mathbf{D}$ which can be written as the limit of a subsequence of $\{\mathbf{D}^T\}_T$, it holds that

$$\mathbb{E}_{b\sim\mathbf{D}}[u_i(b)] \geq \mathbb{E}_{b\sim\mathbf{D}}[u_i(b_i',b_{-i})], \forall i, b_i'$$

as required. ■

Using this connection to coarse correlated equilibria, we are able to obtain a bound of 3.16 on the price of total anarchy of the GSP auction using semi-smoothness techniques in Section 3.2.

Theorem 3.5.2 *The price of total anarchy of the Generalized Second Price auction in the full information setting is at most* 3.164*.*

3.5.2 Learning with uncertainty

Let us now turn to the model of learning outcomes with uncertainty. As in the full information model, we can define a Bayesian version of the coarse correlated equilibrium. A *Bayesian coarse correlated equilibrium* is a mapping from valuation profiles v to distributions over bid profiles $\mathbf{D}(v)$ such that

$$\mathbb{E}_{v,\gamma,b\sim\mathbf{D}(v)}[u_i(b,\gamma)|v_i] \geq \mathbb{E}_{v,\gamma,b\sim\mathbf{D}(v)}[u_i(b'_i,b_{-i},\gamma)|v_i], \forall i, v_i, b'_i.$$

Similarly to Lemma 3.5.1, we can show that the price of total anarchy with uncertainty can be bounded by considering the social welfare generated at any Bayesian coarse correlated equilibrium.

Lemma 3.5.3 *The price of total anarchy with uncertainty is at most*

$$\sup_{\mathbf{F},\mathbf{G},\mathbf{D}(\cdot)\in\mathbf{ccBayesNash}} \frac{\mathbb{E}_{v,\gamma}OPT(v,\gamma)}{\mathbb{E}_{v,\gamma,b\sim\mathbf{D}(v)}[SW(\pi(b),v,\gamma)]}$$

where $\mathbf{ccBayesNash}$ *is the set of Bayesian coarse correlated equilibria.*

The arguments in the proof of Lemma 3.2.3 can be used with essentially no change to show that (λ,μ)-semi-smoothness implies a bound of $(\mu+1)/\lambda$ to the price of total anarchy with uncertainty. From this, we know that:

Theorem 3.5.4 *The price of total anarchy of the Generalized Second Price auction with uncertainty is bounded by* 3.164.

CHAPTER 4

REVENUE OF EQUILIBRIA IN GSP

In this chapter we are mainly concerned with bounding the revenue that is extracted by the GSP mechanisms. We will essentially be concerned with comparing the revenue extraction power of GSP in comparison to the VCG mechanism described in section 2.6. For the (single-keyword) sponsored search environment defined in section 2.7.1, it is easy to check that the VCG mechanism assumes the following form.

1. the auction elicits bids b_i from each agent, which correspond to their *reported* value per click
2. the agents are sorted by bid, i.e., the highest bid is assigned to the top slot and so on...
3. player i's payment per click is: $p_i = \frac{1}{\alpha_{\sigma(i)}} \left[\sum_{j=\sigma(i)+1}^{n} (\alpha_{j-1} - \alpha_j) b_{\pi(j)} \right]$

Notice that above we ignore quality scores, i.e., we assume they are uniform $\gamma_i = 1$. We will do this assumption throughout this section.

VCG is a truthful mechanism: regardless of what the other players are doing, it is a weakly dominant strategy for player i to report his true valuation. The resulting outcome of VCG is therefore social-welfare optimal and the revenue is

$$\mathbf{Rev}^{VCG}(v) = \sum_i \sum_{j>i} (\alpha_{j-1} - \alpha_j) v_j = \sum_{i=2}^{n} (i-1)(\alpha_{i-1} - \alpha_i) v_i.$$

We will also consider the comparison between VCG and GSP in the presence of a reserve price. Let VCG$_r$ be the VCG mechanism with reserve price r, where we discard all players with bids smaller then r and run the VCG mechanism on the remaining players, who then pay price per click $\max\{p_i, r\}$. In the analogous variant of the GSP mechanism, which we call GSP with reserve price r (GSP$_r$), we also discard all players with bids smaller then r, the remaining players are allocated using GSP, and the last player to be allocated pays price r per click.

Below, we represent the special classes of equilibria that have been studied in the literature, which we call **equilibrium hierarchy** for GSP. We define and discuss them in Section 4.3 :

$$\left\{\begin{array}{c}\text{VCG}\\ \text{outcome}\end{array}\right\} \subseteq \left\{\begin{array}{c}\text{envy-free}\\ \text{equilibria}\end{array}\right\} \subseteq \left\{\begin{array}{c}\text{efficient}\\ \text{Nash eq}\end{array}\right\} \subseteq \left\{\begin{array}{c}\text{all}\\ \text{Nash}\end{array}\right\}$$

4.1 Revenue with Uncertainty

In this section we consider the revenue properties of GSP at Bayes-Nash equilibrium. We prove that if agent values are drawn iid from a sufficiently nice distribution (i.e. a distribution satisfying regularity, which we will soon define) and GSP is paired with an appropriate reserve price, the revenue generated at equilibrium will be within a constant factor of the VCG revenue with optimal reserve, the revenue-optimal mechanism over all Bayes-Nash implementations. So our result implies that GSP revenue is within a constant factor of the optimal. We will first consider a special case of regular distributions, the so called MHR distributions, then prove our result in the more general setting where values are drawn from regular distributions.

In what follows, we start by providing a general set of tools for studying revenue in auctions. In particular, we precisely define regular and MHR distributions and show the format of the revenue-optimal auctions for the sponsored search setting. Then we we show the use of reserve prices is crucial: there are instances in which the GSP auction without reserve generates no revenue, whereas the VCG auction generates positive revenue.

After we present the necessary set of tools (section 4.1.1) and motivate the need for reserve prices (section 4.1.2), we proceed to our main results, which are bounds on the revenue extraction power of GSP (sections 4.1.3, 4.1.4 and 4.1.5).

4.1.1 Useful set of tools

A useful tool for studying revenue in the Bayesian setting is Myerson's Lemma, which can be rephrased in the AdAuctions setting as follows. Given a distribution F over agent values, the **virtual valuation** function is defined by $\phi(x) = x - \frac{1-F(x)}{f(x)}$.

Lemma 4.1.1 (Myerson's Lemma [64]) *At any Bayes-Nash equilibrium of an AdAuction mechanism, we have that, for all i, $\mathbb{E}[\alpha_{\sigma(i)} p_i] = \mathbb{E}[\alpha_{\sigma(i)} \phi(v_i)]$ where p_i is the payment per click of player i and $\alpha_{\sigma(i)}$ is the number of clicks received by agent i, and expectation is with respect to $v \sim \mathbf{F}$.*

We say that a distribution is **regular** if $\phi(x)$ is a monotone non-decreasing function. For regular distributions, it follows directly from Myerson's Lemma that the revenue-optimal mechanism for AdAuctions corresponds to running VCG with Myerson's reserve price r, which is the largest value such that $\phi(r) = 0$. We will refer to this as *Myerson's mechanism*, VCG_r.

Running GSP (or VCG) with **reserve price** r means not allocating any user with value $v_i < r$ and running GSP (or VCG) with the remaining agents. For the allocated agents, the mechanism charges per click the maximum between the GSP (VCG) price and r.

A special class of regular distributions is the **monotone hazard rate** distributions (**MHR**), which are the distributions for which $f(x)/(1-F(x))$ is non-decreasing.

4.1.2 Revenue without Reserves: Bad Examples

We start by providing an example in the Bayesian setting where VCG generates positive revenue and GSP has a Bayes-Nash equilibrium that generates zero revenue. Consider three players with iid valuations drawn uniformly from $[1,2]$ and three slots with $\alpha = [1, 0.5, 0.5]$. Let $v^{(i)}$ be the i^{th} largest valuation (which is naturally a random variable defined by v). We have

$$\mathbb{E}[\mathbf{Rev}^{VCG}(v)] = \mathbb{E}[0.5 v^{(2)}] = \frac{3}{4}.$$

Now, consider the following equilibrium of GSP: $b_i(v_i) = 0$ for $i = 2,3$ and $b_1(v_1) = v_1$. Clearly player 1 is in equilibrium. To see that players $i = 2,3$ are in equilibrium, suppose player i has valuation $v_i > 0$. Then his expected utility when bidding any value in $[0,1]$ is $0.5 v_i$, whereas if he changed his bid to some $b > 1$ his utility would be

$$\begin{aligned}\mathbb{E}[u_i(b', b_{-i})|v_i] &= 0.5 v_i + 0.5 v_i \mathbb{P}(v_1 \leq b') - \int_0^{b'} v_1 d\mathbb{P}(v_1) = \\ &= 0.5 v_i + 0.5 v_i (b' - 1) - \frac{(b')^2 - 1}{2} \leq \\ &\leq 0.5 v_i.\end{aligned}$$

Thus agent i cannot increase his expected utility by placing a non-zero bid.

4.1.3 Warmup: MHR Valuations

We now show that if valuations are drawn from a MHR distribution and GSP is paired with the Myerson reserve price, the resulting mechanism extracts a constant fraction of the optimal revenue.

In what follows we will write x^+ to denote $\max\{x, 0\}$.

Theorem 4.1.2 *If valuations are drawn iid from a MHR distribution F and r is the Myerson reserve price for F, then the expected revenue of GSP$_r$ at any Bayes-Nash equilibrium is at least $\frac{1}{6}$ of the optimal revenue.*

Our proof will make use of the fact that, for MHR distributions, $\phi(x) \geq x - r$ for any $x \geq r$. To see this, note that $x - \phi(x) = \frac{1-F(x)}{f(x)} \leq \frac{1-F(r)}{f(r)} = r$ by monotonicity and the definition of Myerson's reserve price.

Proof : Let b be a Bayes-Nash equilibrium of GSP$_r$, and let $\mathbf{Rev}_r(v)$ be the expected revenue of GSP$_r$ at this equilibrium. Let $\mathbf{Rev}_r^{VCG}(v)$ be the VCG$_r$ revenue. Let random variable $\mu(i)$ denote the slot occupied by player i in the optimal (i.e. efficient) allocation. By Myerson's Lemma, $\mathbb{E}[\mathbf{Rev}_r^{VCG}(v)] = \mathbb{E}[\sum_i \alpha_{\mu(i)}\phi(v_i)^+]$. For each player i, let E_1^i denote the event that $b_{\pi(\mu(i))} < v_i/2$, and let E_2^i denote the event that $b_{\pi(\mu(i))} \geq v_i/2$. We will consider each of these events in turn. For the first event, we'll show that player i contributes to the revenue at least 1/2 his contribution in the optimum. Consider a player i with value v_i. We have

$$\mathbb{E}_{v_{-i}}\left[\alpha_{\mu(i)}\frac{v_i}{2}\mathbb{1}\{E_1^i\}\right] \le \mathbb{E}_{v_{-i}}\left[u_i\left(\frac{v_i}{2}, b_{-i}\right)\right] \le \mathbb{E}_{v_{-i}}[u_i(b)] \le \mathbb{E}_{v_{-i}}[\alpha_{\sigma(i)}v_i]$$

where the first inequality is due to the definition of E_1^i implying that a bid of $v_i/2$ would win slot $\mu(i)$ (or better) at price no more than $v_i/2$; the second follows since b is a Bayes-Nash equilibrium, and the third comes from the definition of utility. Notice that all the expectations are taken over v_{-i} and v_i is a constant, so we can divide by v_i, multiply by $\phi(v_i)^+$, take expectations over v_i and sum over all players i to obtain

$$\sum_i \mathbb{E}_v[\alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_1^i\}] \le 2\sum_i \mathbb{E}_v[\alpha_{\sigma(i)}\phi(v_i)^+] = 2\mathbb{E}_v[\mathbf{Rev}_r(v)].$$

For the second event, consider again a player i with value v_i. We will show that the player who gets slot $\mu(i)$ contributes to the revenue. We have

$$\begin{aligned}
\mathbb{E}_{v_{-i}}\left[\alpha_{\mu(i)}\frac{\phi(v_i)^+}{2}\mathbb{1}\{E_2^i\}\right] &\le \mathbb{E}_{v_{-i}}\left[\alpha_{\mu(i)}\frac{v_i}{2}\mathbb{1}\{E_2^i\}\right] \\
&\le \mathbb{E}_{v_{-i}}[\alpha_{\mu(i)}v_{\pi(\mu(i))}] \\
&\le \mathbb{E}_{v_{-i}}[\alpha_{\mu(i)}(r + \phi(v_{\pi(\mu(i))})^+)]
\end{aligned}$$

where we used the fact that $x \ge \phi(x)^+ \ge x - r$ for all x. Taking expectations over v_i, summing over all players, and noting that event E_2^i implies that $v_{\pi(\mu(i))} \ge r$, we obtain

$$\sum_i \mathbb{E}_v[\alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_2^i\}] \le 2\sum_i \mathbb{E}_v[\alpha_{\sigma(i)}\phi(v_i)^+] + 2\sum_i \alpha_{\sigma(i)}r\mathbb{1}\{v_i \ge r\}.$$

Since GSP$_r$ extracts a revenue of at least r per click from every bidder with $v_i > r$, we have $\mathbb{E}_v[\mathbf{Rev}_r] \ge \mathbb{E}_v[\sum_i \alpha_{\sigma(i)}r\mathbb{1}\{v_i \ge r\}]$. We conclude that $\sum_i \mathbb{E}_{v_{-i}}[\alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_2^i\}] \le 4\mathbb{E}[\mathbf{Rev}_r]$. Combining our analysis for the two events, we have

$$\mathbb{E}[\mathbf{Rev}_r^{VCG}(v)] = \mathbb{E}\left[\sum_i \alpha_{\mu(i)}\phi(v_i)^+(\mathbb{1}\{E_1^i\} + \mathbb{1}\{E_2^i\})\right]$$
$$\leq 2\mathbb{E}[\mathbf{Rev}_r(v)] + 4\mathbb{E}[\mathbf{Rev}_r(v)] = 6\mathbb{E}[\mathbf{Rev}_r(v)].$$

■

4.1.4 Regular valuations

We now show that if player valuations are drawn from a regular distribution, then there exists an r' such that running GSP with reserve r' extracts a constant fraction of the optimal revenue. The bound for the MHR bounding the contribution of the player at slot $\mu(i)$ took advantage of the fact that in a MHR distribution $\phi(x) \geq x - r$, which may not be true in a regular distribution. Instead, we will use that the player in slot $\mu(i) - 1$ pays at least the bid in slot $\mu(i)$. This leaves us with the added difficulty in bounding the revenue generated by the first slot. To address this issue, we make use of the well-studied Prophet Inequalities [49, 50, 46].

A simplified version of the Prophet Inequality is as follows. Suppose z_i are independent non-negative random variables. Given any $t \geq 0$, write y_t for the value of the first z_i (by index) satisfying $z_i > t$ (or 0 if there is no such z_i). Then the Prophet Inequality states that there exists some $t \geq 0$ such that $\mathbb{E}[y_t] \geq \frac{1}{2}\mathbb{E}[\max_i z_i]$. Since the proof is of this fact is very short, we include it here for completness.

Theorem 4.1.3 (Prophet Inequality [46]) *Given independent random variables $z_1, \ldots, z_n$ if one defines t as the solution of the equation $t = \sum_{i=1}^n \mathbb{E}(z_i - t)^+$, then by* *defining* $y_t = z_i$*, where* i *is the smallest index such that* $z_i \geq t$*, and zero if* $\max z_i < t$*, then:*

$$\mathbb{E}[y_t] \geq \frac{1}{2}\mathbb{E}[\max_i z_i]$$

Proof : We can upper bound $\mathbb{E}[\max_i z_i]$ as :

$$\mathbb{E}[\max_i z_i] \leq t + \mathbb{E}[\max_i (z_i - t)^+] \leq t + \mathbb{E}[\sum_i (z_i - t)^+] = 2t$$

and lower bound $\mathbb{E}[y_t]$ as:

$$\mathbb{E}[y_t] = t\mathbb{P}(\max_i z_i \geq t) + \sum_i \mathbb{E}[(z_i - t)^+ | \max_{j=1..i-1} z_j < t]\mathbb{P}(\max_{j=1..i-1} z_j < t) \geq$$

$$\geq t\mathbb{P}(\max_i z_i \geq t) + \sum_i \mathbb{E}[(z_i - t)^+]\mathbb{P}(\max_i z_i < t) = t$$

Notice that since t is increasing and $\sum_{i=1}^n \mathbb{E}(z_i - t)^+$ is decreasing, a solution always exists if each z_i has a distribution that has positive density everywhere. If this is not the case, the prophet inequality still holds by taking t to be either the supremum of $\{t \,:\, t \leq \sum_{i=1}^n \mathbb{E}(z_i - t)^+\}$ or the infimum of $\{t \,:\, t \geq \sum_{i=1}^n \mathbb{E}(z_i - t)^+\}$ (whichever results in larger $\mathbb{E}[y_t]$). ■

As has been noted elsewhere [22], the Prophet Inequality has immediate consequences for the revenue of auctions with anonymous reserve prices. The following lemma encapsulates the observation we require.

Lemma 4.1.4 *If* v_i *are drawn iid from a regular distribution then there exists* $r_2 \geq 0$ *such that, writing* Z *for the event that* $\max_i v_i \geq r_2$, $\mathbb{E}[\max_i \phi(v_i)^+ | Z]\mathbb{P}(Z) \geq \frac{1}{2}\mathbb{E}[\max_i \phi(v_i)^+]$.

Proof : (sketch) This follows by applying the Prophet Inequality to virtual values $z_i = \phi(v_i)$ and noting that regularity implies that $v_i \geq r_2$ iff $\phi(v_i) \geq \phi(r_2)$. ■

It is important to remark that the proof of the Prophet Inequality is constructive. If one is able to efficiently compute $\mathbb{E}[(v_i - t)^+]$ for every t, then we can compute r_2 exactly using binary search.

Our approach will now be to analyze the revenue of GSP under two different reserve prices. An argument similar to Theorem 4.1.2 shows that GSP with Myerson reserve obtains a constant fraction of the optimal revenue for all slots other than the first slot. On the other hand, GSP with reserve r_2 from Lemma 4.1.4 will obtain at least half of the optimal revenue generated by the first slot. One of these two reserve prices must therefore generate a constant fraction of the optimal revenue.

Theorem 4.1.5 *If valuations v_i are drawn iid from a regular distribution F, then there is a reserve price r such that the expected revenue of GSP$_r$ at any Bayes-Nash equilibrium is at least $\frac{1}{6}$ of the optimal revenue.*

Proof : Define $\mathbf{Rev}_r^{VCG}(v)$, $\mathbf{Rev}_r(v)$, and $\mu(i)$ as in Theorem 4.1.2. Let r_1 denote the Myerson reserve price for F. By Myerson's Lemma, $\mathbb{E}[\mathbf{Rev}_r^{VCG}(v)] = \mathbb{E}[\sum_i \alpha_{\mu(i)}\phi(v_i)^+]$. For each player i, we define the following three events:

- $E_1^i = \{\ b_{\pi(\mu(i))} < v_i/2 \text{ and } \mu(i) \neq 1\ \}$
- $E_2^i = \{\ b_{\pi(\mu(i))} \geq v_i/2 \text{ and } \mu(i) \neq 1\ \}$
- $E_3^i = \{\ \mu(i) = 1\ \}$

We wish to bound the virtual value of the optimal allocation, conditioning on each of these events in turn. For the first event, we proceed precisely as in Theorem 4.1.2 to obtain

$$\sum_i \mathbb{E}_{v_{-i}}[\alpha_{\mu(i)}\phi(v_i)^+ \mathbb{1}\{E_1^i\}] \leq 2\mathbb{E}[\mathbf{Rev}_{r_1}].$$

For the second event, we use the revenue from slot $\mu(1)-1$. Let random variable p_i denote the payment per click of the player in slot i. Then for all v,

$$\alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_2^i\} \leq \alpha_{\mu(i)}v_i\mathbb{1}\{E_2^i\} \leq 2\alpha_{\mu(i)-1}p_{\mu(i)-1}\mathbb{1}\{E_2^i\}$$

where the second inequality follows since E_2^i implies $p_{\mu(i)-1} = b_{\mu(i)} \geq v_i/2$. Therefore, summing over all agents i and taking expectations, we get

$$\mathbb{E}_v\left[\sum_i \alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_2^i\}\right] \leq 2\mathbb{E}_v\left[\sum_i \alpha_i p_i\right] = 2\mathbb{E}[\mathbf{Rev}_{r_1}].$$

Finally, for event E_3^i, consider setting the reserve price to be r_2 from the statement of Lemma 4.1.4 (with distribution F). Note that

$$\mathbb{E}\left[\sum_i \alpha_{\mu(i)}\phi(v_i)^+\mathbb{1}\{E_3^i\}\right] = \alpha_1\mathbb{E}[\max_i \phi(v_i)^+].$$

On the other hand, setting reserve price r_2 for GSP we get

$$\mathbb{E}[\mathbf{Rev}_{r_2}] \geq \alpha_1\mathbb{E}[\max_i \phi(v_i)^+ \mid \max_i v_i \geq r_2]\mathbb{P}(\max_i v_i \geq r_2) \geq \frac{1}{2}\alpha_1\mathbb{E}[\max_i \phi(v_i)^+]$$

where the first inequality follows by considering only the expected virtual value due to the first slot and the last inequality follows from Lemma 4.1.4. Combining our analysis for each of the three cases, we have

$$\begin{aligned}\mathbb{E}[\mathbf{Rev}_r^{VCG}] &= \mathbb{E}\left[\sum_i \alpha_{\mu(i)}\phi(v_i)^+(\mathbb{1}\{E_1^i\} + \mathbb{1}\{E_2^i\} + \mathbb{1}\{E_3^i\})\right] \\ &\leq 4\mathbb{E}[\mathbf{Rev}_{r_1}] + 2\mathbb{E}[\mathbf{Rev}_{r_2}]\end{aligned}$$

and hence $\max\{\mathbb{E}[\mathbf{Rev}_{r_1}], \mathbb{E}[\mathbf{Rev}_{r_2}]\} \geq \frac{1}{6}\mathbb{E}[\mathbf{Rev}_r^{VCG}]$. ■

4.1.5 Bayesian Revenue with Well-separated CTRs

Another way to bound the revenue of GSP in settings of incomplete information, without imposing reserve prices, is to assume that the slot click-through-rates are well separated, in the sense of [53]. We say that click-through-rates are δ-well separated if $\alpha_{i+1} \leq \delta \alpha_i$ for all i.

Lemma 4.1.6 *If click-through-rates are δ-well separated, then bidding $b_i(v_i) < (1 - \delta)v_i$ is dominated by bidding $(1-\delta)v_i$.*

Proof : Suppose player i bids $b_i < (1-\delta)v_i$. If he increases his bid to $b_i' = (1-\delta)v_i$ then with some probability he still gets the same slot (event S) and with some probability he gets a better slot (event B). Then clearly $\mathbb{E}[u_i(b_i, b_{-i})|v_i] \leq \mathbb{E}[u_i(b_i', b_{-i})|v_i]$ since the expectation conditioned to S is the same and conditioned to B it can only increase by changing the bid to b_i'. To see that, let $\alpha_{\pi(i)}$ be the slot player i gets under b_i and $\alpha_{\pi'(i)}$ the slot he gets under b_i'. Conditioned on B we know that $\alpha_{\pi'(i)} \geq \delta^{-1}\alpha_{\pi(i)}$, and this generates value for bidder i of at least $\alpha_{\pi'(i)}(v_i - b_i')$, while the value with bid b_i was at most $\alpha_{\pi(i)}v_i$, which implies the claim:

$$
\begin{aligned}
E[u_i(b_i, b_{-i})|v_i, B] \leq \mathbb{E}[\alpha_{\pi(i)}v_i|v_i, B] \leq \mathbb{E}[\delta\alpha_{\pi'(i)}v_i|v_i, B] = \\
= \mathbb{E}[\alpha_{\pi'(i)}(v_i - (1-\delta)v_i)|v_i, B] \leq \\
\leq \mathbb{E}[u_i(b_i', b_{-i})|v_i, B].
\end{aligned}
$$

■

Recall that under truthful bidding, the revenue of GSP is at least the revenue of VCG. If one eliminates the strategies $b_i(v_i) < (1-\delta)v_i$ from the players strategy set, then it is easy to see that any Bayesian-Nash equilibrium b has high revenue.

Corollary 4.1.7 *If click-through-rates are δ-well separated, and all players play undominated strategies, then*

$$\mathbb{E}_v[\mathbf{Rev}(b(v))] \geq (1-\delta)\mathbb{E}_v[\mathbf{Rev}^{VCG}(v)].$$

Further, for any reserve price r, we also get

$$\mathbb{E}_v[\mathbf{Rev}_r(b)] \geq (1-\delta)\mathbb{E}_v[\mathbf{Rev}_r^{VCG}(v)].$$

Next we consider whether this bound on GSP revenue, with respect to the expected GSP revenue when all players report truthfully, continues to hold if agents do not eliminate dominated strategies. That is, we consider settings of limited rationality in which players may not be able to find dominated strategies. If we allow players to use dominated strategies, then we might have equilibria with very bad revenue compared to the expected revenue when agents bid truthfully, as one can see in the following example:

Example 4.1.8 *Consider two players with iid valuations $v_i \sim$ Uniform$([0,1])$ and two slots with $\alpha = [1, 1-\epsilon]$. Then VCG generates revenue $\mathbb{E}[\mathbf{Rev}^{VCG}(v)] = \mathbb{E}[\epsilon \min\{v_1, v_2\}] = O(\epsilon)$, and if agents report truthfully the GSP auction generates revenue $\mathbb{E}[\min\{v_1, v_2\}] = O(1)$. However, consider the following equilibrium:*

$$b_1(v_1) = \begin{cases} \epsilon(1-\delta), & v_1 \geq \epsilon(1-\delta) \\ \epsilon v_1, & v_1 < \epsilon(1-\delta) \end{cases}$$

$$b_2(v_2) = \begin{cases} \epsilon, & v_2 \geq 1-\delta \\ \epsilon^2(1-\delta), & \epsilon(1-\delta) \leq v_2 < 1-\delta \\ \epsilon v_2, & v_2 < \epsilon(1-\delta) \end{cases}$$

It is not hard to check that this is an equilibrium. In fact, for two player GSP in the Bayesian setting, playing $(\alpha_1 - \alpha_2)v_i/\alpha_1$ is a best reply - and any bid that gives the *player the same outcome is also a best reply. So, in the above example, one can simply check that the bids generate the same utility as bidding $b_i(v_i) = \epsilon v_i$. This example generates revenue $\mathbb{E}\mathbf{Rev}(b) = O(\epsilon(\epsilon + \delta))$, so taking $\delta = O(\epsilon)$ in the above example give us $O(\epsilon^2)$ revenue.*

However, this is a feature of having only 2 players, as shown in the following theorem, which is a version of Corollary 4.1.7 that doesn't depend on eliminating dominated strategies.

Theorem 4.1.9 *With n players with iid valuations v_i and δ-well separated click-through-rates, then for all Bayes-Nash equilibria b in which agents do not overbid,*

$$\mathbb{E}[\mathbf{Rev}(b)] \geq \frac{n-2}{n}(1-\delta)\mathbb{E}[\mathbf{Rev}^{VCG}(v)].$$

Proof : We will prove the stronger result that the expected GSP revenue at equilibrium is within a factor of $\frac{n-2}{n}(1-\delta)$ of the expected GSP revenue when agents report truthfully. We first claim that, for a profile b in Bayesian-Nash equilibrium and any two players i and j, we have that

$$\mathbb{P}_{v \sim F}[b_i(v) < (1-\delta)v - \epsilon, b_j(v) < (1-\delta)v - \epsilon] = 0.$$

To see this, suppose the contrary. Then there is $\epsilon' \ll \epsilon$ such that if we take $F' = F|_{[v^0 - \epsilon', v^0 + \epsilon']}$ then

$$\mathbb{P}_{v \sim F'}[b_i(v) < (1-\delta)v - \epsilon, b_j(v) < (1-\delta)v - \epsilon] > 0.$$

For ϵ' small enough $\underline{v_0} = v^0 - \epsilon$ and some $\epsilon'' < \epsilon$, we have

$$\mathbb{P}_{v \sim F'}[b_i(v) < (1-\delta)\underline{v_0} - \epsilon'', b_j(v) < (1-\delta)\underline{v_0} - \epsilon''] > 0.$$

Now pick v^i, v^j in this interval such that $\mathbb{P}_{v \sim F'}[b_i(v^i) \leq b_i(v) < (1-\delta)\underline{v_0}] > 0$ and the same for j. By lemma 4.1.6, playing $(1-\delta)v^i$ is a best response, then for player j for example, it can't be the case that any of the other players play between $b_j(v^j)$ and $(1-\delta)v^j$ with positive probability. Therefore

$$\mathbb{P}_{v \sim F'}[b_j(v) \in [b_i(v^i), (1-\alpha)v^i)] = 0$$

$$\mathbb{P}_{v \sim F'}[b_i(v) \in [b_j(v^j), (1-\alpha)v^j)] = 0$$

but notice this is a contradiction. This completes the proof of the claim.

Now, we can think of the procedure of sampling v iid from F in the following way: sample $v_i'' \sim F$ iid, let v_i' be the sorted valuations, and then apply a random permutation $\tau \in S_n$ to the values so that $v_i = v'_{\tau(i)}$. Notice that v is iid and now, notice that with $\geq 1 - \frac{2}{n}$ probability, v_i' and v'_{i+1} will generate $(1-\delta)v_i'$ and $(1-\delta)v'_{i+1}$ bids producing $(1-\delta)\alpha_i v'_{i+1}$ revenue, therefore

$$\mathbb{E}[\mathbf{Rev}(v)] \geq \mathbb{E}\left[\sum_i \left(1 - \frac{2}{n}\right)(1-\delta)\alpha_i v'_{i+1}\right] \geq \frac{n-2}{n}(1-\delta)\mathbb{E}\left[\mathbf{Rev}^V(v)\right].$$

■

4.2 Revenue in Full Information GSP

We now wish to compare the revenue properties of GSP and VCG in the full information setting. We start by giving examples showing that there are no universal constants that bound these two quantities. Then we introduce a new benchmark related to VCG, and show that the GSP revenue is not too low relative to this benchmark.

4.2.1 Full Information Revenue: Examples

Unfortunately, there are no universal constants $c_1, c_2 > 0$ such that for every full information AdAuctions instance α, v and for all equilibria b of GSP it holds that

$$c_1 \cdot \mathbf{Rev}^{VCG}(v) \leq \mathbf{Rev}(b) \leq c_2 \cdot \mathbf{Rev}^{VCG}(v).$$

In fact, GSP can generate arbitrarily more revenue than VCG and vice-versa. For example, consider two players with $\alpha = \{1, 0\}$, $v = \{2, 1\}$. Then VCG generates revenue 1, but GSP has the Nash equilibrium $b = [2, 0]$ that generates no revenue.

As a counter-example for the second inequality, consider the following instance: $\alpha = \{1, 1-\epsilon\}$, $v = \{\epsilon^{-1}, 1\}$. Notice that the revenue produced by VCG is ϵ, while GSP has the equilibrium $b = [1, 1]$ generating revenue 1.

4.2.2 Revenue Bound in Full Information

Next, we will prove that the GSP revenue cannot be much less than a revenue benchmark based on the VCG auction. Intuitively, the difficulty behind our bad examples is in extracting revenue from the player with the largest private value. Motivated by this, we consider the following benchmark:

$$\mathcal{B}(v) = \sum_{i=2}^{n} p_i^{VCG} \alpha_{\sigma(i)} = \sum_{i=2}^{n} \sum_{j>i} (\alpha_{j-1} - \alpha_j) v_j = \sum_{i=2}^{n} (i-2)(\alpha_{i-1} - \alpha_i) v_i$$

which is the VCG revenue from players $2, 3, \ldots, n$. Recall that in the full information setting we assumed that players are numbered such that $v_1 \geq v_2 \geq \ldots$. We show that the GSP revenue is always at least half of this benchmark at any equilibrium. Thus, unless VCG gets most of its revenue from a single player, GSP revenue will be within a constant factor of the VCG revenue.

Theorem 4.2.1 *Given an AdAuctions instance α, v, and a Nash equilibrium b of GSP, we have $\mathbf{Rev}(b) \geq \frac{1}{2}\mathcal{B}(v)$, and this bound is tight.*

We prove Theorem 4.2.1 in two steps. First we define the concept of up-Nash[1] equilibrium for GSP, then we show that any inefficient Nash equilibria can be written as an efficient up-Nash equilibrium. In the second step, we prove the desired revenue bound for all efficient up-Nash equilibria.

Definition 4.2.2 *Given a bid profile b, we say it is **up-Nash** for player i if he can't increase his utility by taking some slot above, i.e.*

$$\alpha_{\sigma(i)}(v_i - b_{\pi(\sigma(i)+1)}) \geq \alpha_j(v_i - b_{\pi(j)}), \forall j < \sigma(i).$$

*Analogously, we say that b is **down-Nash** for player i if he can't increase his utility by taking some slot below, i.e.*

$$\alpha_{\sigma(i)}(v_i - b_{\pi(\sigma(i)+1)}) \geq \alpha_j(v_i - b_{\pi(j+1)}), \forall j > \sigma(i).$$

A bid profile is up-Nash (down-Nash) if it is up-Nash (down-Nash) for all players i. Clearly a bid profile b is a Nash equilibrium iff it is both up-Nash and down-Nash.

Lemma 4.2.3 *If a bid profile b is a Nash equilibrium, then the bid profile b' where $b'_i = b_{\pi(i)}$ is up-Nash.*

Proof : We will prove the lemma by modifying bid profile b in a sequence of steps. Fix some $k \leq n$, and suppose that b is a bid profile (with corresponding allocation π) such that

[1]Our concepts of up-Nash and down-Nash equilibria are very similar to the concepts of upwards stable and downwards stable equilibria in Markakis and Telelis [60]

- players $j = 1, \ldots, k$ satisfy the Nash conditions (i.e. both up-Nash and down-Nash) in b,
- players $j = k+1, \ldots, n$ are such that $\sigma(j) = j$ and they satisfy the up-Nash conditions in b,
- $\sigma(k) < k$.

We then define b' by swapping the bids of players k and $\pi(k)$, that is setting $b'_i = b_i$ for $i \neq k, \pi(k)$, $b'_k = b_{\pi(k)}$, and $b'_{\pi(k)} = b_k$. We claim that b' is up-Nash for players $k, \ldots, n$ and Nash for the remaining players. This then implies the desired result, since we can apply this operation for $k = n$, followed by $k = n-1, \ldots, 2$, resulting in the required bid profile.

Since our transformation does not alter the bids associated with given slots, we just need to check three things: the up and down-Nash inequalities for player $\pi(k)$, and the up-Nash inequality for player k.

Under bid profile b', player $\pi(k)$ gets slot $\sigma(k)$. This player doesn't want to change his bid to win any slot $j > \sigma(k)$ since in the bid profile b player k with lower value didn't want to get these slots. We therefore have $\alpha_{\sigma(k)}(v_k - b_{\pi(\sigma(k)+1)}) \geq \alpha_j(v_k - b_{\pi(j+1)})$ and since $v_{\pi(k)} \geq v_k$, we conclude

$$\alpha_{\sigma(k)}(v_{\pi(k)} - b_{\pi(\sigma(k)+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j+1)}). \tag{4.1}$$

To see that player $\pi(k)$ would not prefer to take any slot $j < \sigma(k)$, notice that $\pi(k)$ didn't want to move to a higher slot in b, so $\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$. This, combined with equation (4.1) for $j = k$ (stating that $\pi(k)$ prefers slot $\sigma(k)$ to k) gives the up-Nash inequality for player $\pi(k)$.

Next consider player k in bid profile b', where we gets slot k. We wish to prove the up-Nash inequality for k. Notice that, in b, $\pi(k)$ had slot k and didn't want to switch to a higher slot, so we know $\alpha_k(v_{\pi(k)} - b_{\pi(k+1)}) \geq \alpha_j(v_{\pi(k)} - b_{\pi(j)})$.

Now, since $v_{\pi(k)} \geq v_k$, we have $\alpha_k(v_k - b_{\pi(k+1)}) \geq \alpha_j(v_k - b_{\pi(j)})$ which is the desired inequality. ■

Now to prove Theorem 4.2.1 we use the up-Nash profile b'.

Proof of Theorem 4.2.1 : Given any Nash equilibrium b, consider the bid profile b' of Lemma 4.2.3, which is an up-Nash equilibrium in which each player k occupies slot k. By the up-Nash inequalities, for each k we have

$$\alpha_k(v_k - b'_{k+1}) \geq \alpha_{k-1}(v_k - b'_{k-1}).$$

We can rewrite this as

$$\alpha_{k-1}b'_{k-1} \geq (\alpha_{k-1} - \alpha_k)v_k + \alpha_k b'_{k+1}.$$

Then, since $\alpha_k \geq \alpha_{k+1}$,

$$\alpha_{k-1}b'_{k-1} \geq \sum_{j \in k+2\mathbb{N}} (\alpha_{j-1} - \alpha_j)v_j$$

where $k + 2\mathbb{N} = \{k, k+2, k+4, \ldots\}$. Now we can bound the revenue of b:

$$\begin{aligned}\mathbf{Rev}(b) = \mathbf{Rev}(b') = \sum_k \alpha_k b'_{k+1} \geq \sum_k \alpha_{k+1} b'_{k+1} \geq \\ \geq \sum_k \sum_{j \in k+2+2\mathbb{N}} (\alpha_{j-1} - \alpha_j)v_j \geq \\ \geq \sum_{k=2}^{n} \frac{k-2}{2}(\alpha_{k-1} - \alpha_k)v_k = \frac{1}{2}\mathcal{B}(v).\end{aligned}$$

■

To show that the bound in Theorem 4.2.1 is tight, consider the following example with n slots and n players, parametrized by $\delta > 0$:

$$\alpha = [1, 1, \ldots, 1, 1 - \delta, 0],$$

$$v = [1, 1, \ldots, 1, 1, \delta],$$

$$b = [\delta, \delta, \ldots, \delta, \delta, 0].$$

In this case $\mathbf{Rev}(b) = (n-2)\delta + \delta(1-\delta)$ and $\mathbf{Rev}^{VCG}(v) = (2\delta - \delta^2)(n-3) + \delta(1-\delta)$. Therefore $\lim_{n\to\infty} \frac{\mathbf{Rev}(b)}{\mathcal{B}(v)} = 2 - \delta$ and it tends to 2 as $\delta \to 0$.

Notice that those bounds also carry for the case where there is a reserve price r. We compare the revenue $\mathbf{Rev}_r(b)$ with reserve price r, against a slightly modified benchmark: $\mathcal{B}_r(v)$ which is the revenue VCG$_r$ extracts from players $2, \ldots, n$.

Corollary 4.2.4 *Let b be a Nash equilibrium of the GSP$_r$ game, then*

$$\mathbf{Rev}_r(b) \geq \frac{1}{2}\mathcal{B}_r(v)$$

Proof : We can assume wlog that $v_i, b_i \geq r$ (otherwise those players don't participate in any of the auctions). We can define an upper-Nash bid profile b' as in Lemma 4.2.3. Now, notice that all players in b' are paying at least r per click. We can divide the players in two groups: players $1 \ldots k$ are paying more than r in VCG$_r$ and players $k+1 \ldots n$ are paying exactly r. It is trivial that for the players $k+1 \ldots n$ we extract at least the same revenue under VCG$_r$ then under GSP$_r$. For the rest of the players we need to do the exact same analysis as in the proof of Theorem 4.2.1. ∎

4.3 Tradeoff between Revenue and Efficiency

In this section we consider the tradeoff between efficiency and revenue, and ask if optimal efficiency and optimal revenue can always be achieved in the same equilibrium. We give a negative answer to this question, showing that for some AdAuction instances, one can increase revenue by selecting inefficient equilibria. First we recall the equilibrium hierarchy briefly discussed in the introduction.

Then we characterize the maximum revenue possible for envy free equilibrium (that is always efficient). Does this equilibrium class generate more or less revenue than other classes, such as efficient equilibria or all pure equilibria? This question of comparing the revenue of VCG and envy-free equilibria of GSP was addressed by [34], who show that the revenue in any envy-free equilibrium is at least that of the VCG outcome (i.e. the VCG outcome is the envy-free equilibria generating smallest possible revenue). Moreover, as we've shown, an envy-free equilibrium can generate arbitrarily more revenue than the VCG outcome. Varian [76] shows how to compute the revenue optimal envy free Nash equilibrium, if we assume that agents will overbid. Allowing overbidding can result in very high revenue (eg.., the maximum valuation in a single item auction). Here we determine the maximum revenue that can be obtained if we do not assume that agents bid at envy-free equilibria, and without requiring that agents apply the dominated strategy of overbidding.

Finally, we use this characterization to we give a natural sufficient condition under which there is a revenue-optimal equilibrium that is efficient.

4.3.1 Equilibrium hierarchy for GSP

Edelman, Ostrovsky and Schwarz [34] and Varian [75] showed that the full information game always has a Pure Nash equilibrium, and moreover, there is a pure Nash equilibrium which has same outcome and payments as VCG. The authors also define a class of equilibria called *envy-free* or *symmetric equilibria*. All envy-free equilibria are Nash equilibria, though not all Nash equilibria are envy-free. We refer the reader to Section 3.1 for such results. Relevant for the discussion following is the following result on the revenue of envy-free equilibria:

Lemma 4.3.1 ([34, 75]) *That is, if b is an envy-free equilibrium of the game induced by α and v, $\mathbf{Rev}(b) \geq \mathbf{Rev}^{VCG}(v)$.*

Proof : From Lemma 3.1.2, all envy free are efficient, so if the players are sorted such that $v_1 \geq v_2 \geq \ldots \geq v_n$ we get: $\mathbf{Rev}(b) = \sum_{i=2}^{n} \alpha_{i-1} b_i$. By envy-freeness, $\alpha_i(v_i - b_{i+1}) \geq \alpha_{i-1}(v_i - b_i)$ so: $\alpha_{i-1} b_i \geq (\alpha_{i-1} - \alpha_i) v_i + \alpha_i b_{i+1}$. Telescoping this inequality, we get: $\alpha_{i-1} b_i \geq \sum_{j=i}^{n} (\alpha_{j-1} - \alpha_j) v_j$. So:

$$\mathbf{Rev}(b) = \sum_{i=2}^{n} \alpha_{i-1} b_i \geq \sum_{i=2}^{n} \sum_{j=i}^{n} (\alpha_{j-1} - \alpha_j) v_j = \mathbf{Rev}^{VCG}(v)$$

■

Also, although all envy-free equilibria are efficient, there are efficient equilibria that are not envy-free, as one can see for example in Figure 4.1, as well as inefficient equilibria. We therefore have the following hierarchy:

$$\left\{ \begin{array}{c} \text{VCG} \\ \text{outcome} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{envy-free} \\ \text{equilibria} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{efficient} \\ \text{Nash eq} \end{array} \right\} \subseteq \left\{ \begin{array}{c} \text{all} \\ \text{Nash} \end{array} \right\}$$

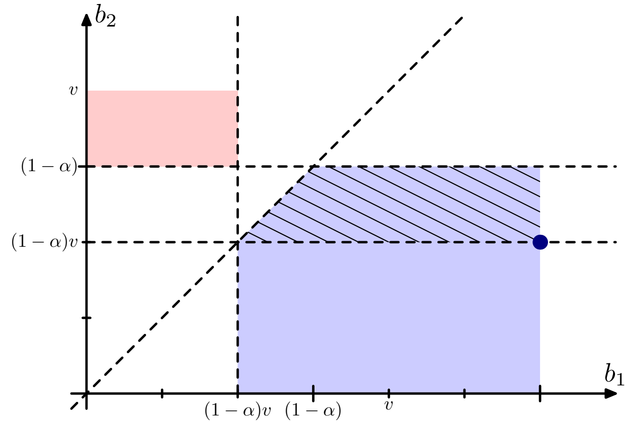

Figure 4.1: Equilibria hierarchy for GSP for $\alpha = [1, 1/2]$, $v = [1, 2/3]$: the strong blue dot represents the VCG outcome, the pattern region the envy-free equilibria, the blue region all the efficient equilibria and the red region the inefficient equilibria

4.3.2 Envy-free and efficient equilibrium

As shown in the example of Figure 4.1, there are efficient equilibria that generate arbitrarily less revenue then any envy-free equilibrium. For the other direction, we show that all revenue-optimal efficient equilibria are envy-free.

Theorem 4.3.2 *For any AdAuctions instance such that $\alpha_i > \alpha_{i+1} \forall i$, all revenue-optimal efficient equilibria are envy-free. Moreover, we can write the revenue optimal efficient equilibrium explicitly as function of α, v recursively as follows:*

$$b_n = \min\left\{v_n, \frac{\alpha_{n-1} - \alpha_n}{\alpha_{n-1}} v_{n-1}\right\},$$
$$b_i = \min\left\{v_i, \frac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \frac{\alpha_i}{\alpha_{i-1}} b_{i+1}\right\} \ \forall i < n.$$

Proof : Given an efficient equilibrium b, if it is not envy-free, we show that we can improve revenue by slightly increasing one of the bids. If the equilibrium is not envy-free, there is at least one player that envies the player above, i.e.

$$\alpha_i(v_i - b_{i+1}) < \alpha_{i-1}(v_i - b_i).$$

As pointed out in [34], if in an efficient equilibrium no player envies the above slot (i.e. no player i wants to take the above slot $i-1$ by the price per click player i is paying) then the equilibrium is envy-free.

Let i be the player with the smallest index that envies slot $i-1$. Consider the bid profile b' such that $b'_j = b_j$ for $j \neq i$ and $b'_i = b_i + \epsilon$. We will verify that the Nash inequalities for player $i-1$ still hold when $\epsilon > 0$ is sufficiently small. In other words, we will show that no Nash inequality for player $i-1$ holds with equality in b.

For slots $j > i-1$, notice that

$$\alpha_j(v_i - b_{j+1}) \leq \alpha_i(v_i - b_{i+1}) < \alpha_{i-1}(v_i - b_i)$$

where the first is a standard Nash inequality and the second is the hypothesis that player i envies the above slot. Now, since $v_{i-1} > v_i$ in an efficient equilibrium, we have

$$\alpha_j(v_{i-1} - b_{j+1}) < \alpha_{i-1}(v_{i-1} - b_i).$$

For slots $j < i-1$, we use the fact that player i is the first envious player. Also, without loss of generality, we can assume player 1 bids v_1. Therefore we need to verify the Nash inequalities only for $j = 2, 3, \ldots, k-1$. We have

$$\alpha_{i-1}(v_i - b_i) \geq \alpha_j(v_i - b_{j+1}) > \alpha_j(v_i - b_j)$$

where the first inequality comes from the fact that player $i-1$ doesn't envy any player j above him and the second inequality comes from the fact that $b_j > b_{j+1}$, since otherwise the player in slot j would envy the player in slot $j-1$. This shows that the revenue optimal equilibrium is envy free.

To see that the bid profile defined in the theorem is optimal, we need to show the following things about this bid profile b: (i) it is in Nash equilibrium,

(ii) it is envy free, and (iii) no other efficient Nash equilibrium generates higher revenue. Begin by noticing that if b is Nash, then player $i-1$ doesn't want to take slot i, for all i, and therefore $\alpha_{i-1}(v_{i-1} - b_i) \geq \alpha_i(v_{i-1} - b_{i+1})$ and this is satisfied by definition by the bid vector presented. Notice also that this series of inequalities implies an upper bound on the maximum revenue in an efficient equilibrium and this bound is achieved exactly by the bid profile defined above.

Furthermore, for all $j \leq i-1$ we have $\alpha_{i-1}(v_j - b_i) \geq \alpha_i(v_j - b_{i+1})$ therefore by composing this expression with different values of i and j, it is straightforward to show that no player can profit by decreasing his bid. We prove that no player can profit by overbidding as a simple corollary of envy-freeness. For that, we need to prove that

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_{i-1}(v_i - b_i).$$

If $b_i = v_i$ than this is trivial. If not, then substitute the expression for b_i and notice it reduces to $v_{i-1} \geq v_i$. Now, this proved local envy-freeness, what implies that no player wants the slot above him by the price he player above him is paying. This in particular implies that no player wants to increase his bid to take a slot above. ■

4.3.3 Cost of efficiency: definition and example

Next we will analyze the relation between revenue and efficiency in GSP auctions.

We define the *cost of efficiency* for a given profile of click-through-rates as

$$\mathbf{CoE}(\alpha) = \max_v \frac{\max_{b \in \mathbf{Nash}(\alpha,v)} \mathbf{Rev}(b)}{\max_{b \in \mathbf{EffNash}(\alpha,v)} \mathbf{Rev}(b)}$$

where **Nash** is the set of all bid profiles in Nash equilibrium and **EffNash** is the set of all efficient Nash equilibrium.

First we give examples in which $\mathbf{CoE}(\alpha) > 1$, in which case all revenue-optimal equilibria occur at inefficient equilibria. Our example will have $n = 3$ slots and advertisers. The click-through rates are given by $\alpha = [1, \frac{2}{3}, \frac{1}{6}]$ and the agent types are $v = [1, \frac{7}{8}, \frac{6}{8}]$. In this case, the best possible revenue generated by an efficient outcome is given by $\frac{1}{3} + \frac{7}{8} \approx 1.20833$ (this can be calculated using the formula in Theorem 4.3.2). However, for the (inefficient) allocation $\pi = [2, 1, 3]$, there is an equilibrium that generates revenue 1.21528.

In Figure 4.2 we calculate this value empirically for each $\alpha = [1, \alpha_2, \alpha_3]$, where each α_i is an integer multiple of 0.01 in $[0, 1]$. In all cases we found that $1 \leq \mathbf{CoE}(\alpha) < 1.1$. The color of (α_1, α_2) in the graph corresponds to $\mathbf{CoE}(1, \alpha_2, \alpha_3)$, where blue represents 1 and red represents 1.1. By solving a constrained non-linear optimization problem, one can show that the worst **CoE** for 3 slots is 1.09383.

Computing the Cost of Efficiency In the manner of Section 3.4, we will show how to compute $\mathbf{CoE}(\alpha)$ for fixed click-through-rates $\alpha = (\alpha_1, \ldots, \alpha_n)$ using Linear Programming. The main ingredient will be Theorem 4.3.2 in the previous section. Given α, v the cost of efficiency is the ratio between the highest revenue Nash-equilibrium, which we call it b^1, and the best revenue efficient Nash equilibrium, which we call b^2.

Like we did in Section 3.4, we can fix a permutation $\pi : [n] \to [n]$ representing the allocation of players to slots under b^1, and re-define b^1 as the equilibrium with highest revenue among the equilibria that allocate players to slots according to π. Therefore, what we want to maximize is:

$$\max \frac{\sum_{i=2}^{n} \alpha_{i-1} b^1_{\pi(i)}}{\sum_{i=2}^{n} \alpha_{i-1} b^2_i}$$

where (v, b^1) follow the constraints (3.7) and (3.8) with b substituted by b^1 and b^2 can be obtained from v using Theorem 4.3.2. Even after fixing v, we don't have an LP yet, since the conditions in Theorem 4.3.2 involve a minimum operator. From that theorem, we have:

$$b^2_i = \min \left\{ v_i, \frac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \frac{\alpha_i}{\alpha_{i-1}} b^2_{i+1} \right\}, \quad \forall i$$

where $b^2_{n+1} = 0$. In order to go around this, we break this problem in 2^n linear programs. For each $X = (X_1, \ldots, X_n) \in \{0, 1\}^n$ add the following constraints:

$$\begin{aligned} \text{if } X_i = 0 : \quad & \begin{cases} v_i \geq \dfrac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \dfrac{\alpha_i}{\alpha_{i-1}} b^2_{i+1} \\ b^2_i = \dfrac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \dfrac{\alpha_i}{\alpha_{i-1}} b^2_{i+1} \end{cases} \\ \text{if } X_i = 1 : \quad & \begin{cases} v_i \leq \dfrac{\alpha_{i-1} - \alpha_i}{\alpha_{i-1}} v_{i-1} + \dfrac{\alpha_i}{\alpha_{i-1}} b^2_{i+1} \\ b^2_i = v_i \end{cases} \end{aligned} \tag{4.2}$$

Following the same argument as in Section 3.4, we reduce it to the problem of computing for each (π, X) the solution of $\max \beta$ where $\beta = \sum_{i=2}^{n} \alpha_{i-1} b^1_{\pi(i)}$, $1 = \sum_{i=2}^{n} \alpha_{i-1} b^2_i$ where (v, b^1) are subject to constraints (3.7) and (3.8) and (v, b^2) is subject to constraints $(4.2)_X$. Therefore, we can compute $\mathbf{CoE}(\alpha)$ by solving $2^n \cdot n!$ linear programs with $O(n)$ variables and $O(n^2)$ constraints each. In the file `cost_of_efficiency.m` we provide a code that computes the cost of efficiency for each α.

The utility of such LP formulation is two-fold: one is to produce lower bounds on the value of $\max_{\alpha\in\mathbb{I}_n} \mathbf{CoE}(\alpha)$. But a perhaps more interesting use is to be able to plot and get intuition on how the **CoE** function behaves. This allowed us to plot the $\mathbf{CoE}(1, \alpha_2, \alpha_3)$ in Figure 4.2 and observe the phenomenon proved in the next section.

4.3.4 Efficiency Versus Revenue when Click-Through-Rates are Convex

We now present a condition on α that implies $\mathbf{CoE}(\alpha) = 1$. Our condition is that the click-through-rates are *convex*, meaning that $\alpha_i - \alpha_{i+1} \geq \alpha_{i+1} - \alpha_{i+2}$ for all i. We note that most models for CTRs studied in the literature satisfy convexity, such as exponential CTRs [53] and Markovian user models [2].

Theorem 4.3.3 *If click-through-rates α are convex then there is a revenue-maximizing Nash equilibrium that is also efficient.*

Our proof follows from a local improvement argument: given an instance with convex click-through-rates and an equilibrium that is not efficient, we show how to either improve it revenue or its welfare. A key step of the proof is bounding the maximum revenue possible in equilibrium for a given allocation, extending Theorem 4.3.2 to inefficient allocations.

Proof : Let b be the revenue maximizing efficient Nash equilibrium. Fix an allocation π and let b' be an equilibrium under allocation π. We say that b is **saturated** for slot i if $b_i = v_i$. We start by presenting the proof of the theorem under

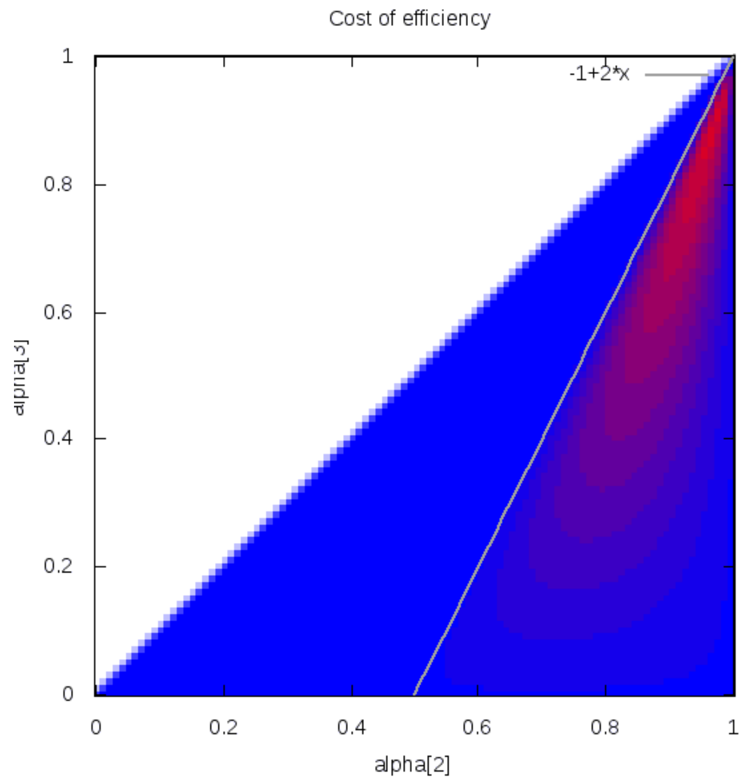

Figure 4.2: Cost of efficiency for $\alpha = [1, \alpha_2, \alpha_3]$: in the plot, blue means 1.0 and red means 1.1.

the simplifying assumption that no slot is saturated in the maximum revenue equilibrium..

Under the no-saturation assumption, Theorem 4.3.2 implies

$$\mathbf{Rev}(b) = \sum_i \alpha_i b_{i+1} = \sum_i \sum_{j \geq i} (\alpha_j - \alpha_{j+1}) v_j. \tag{4.3}$$

Notice that we can view this expression as a dot product of two vectors where one has elements of the form v_i and other has elements in the form $\alpha_j - \alpha_{j+1}$. Notice also that due to the convexity assumption, this is a dot product of two sorted vectors. Now, for b', we will bound revenue as follows. Define $m(\pi, i, j) = \max\{\pi(i), \pi(i+1), \pi(i+2), \ldots, \pi(j)\}$. Let p be such that the $k = i, i+1, \ldots, i+p$ are all the indices such that $m(\pi, i, k) = \pi(i)$. Now, notice that the player in slot i doesn't want to take slot $i+p+1$, so

$$\alpha_i(v_{\pi(i)} - b'_{\pi(i+1)}) \geq \alpha_{i+p+1}(v_{\pi(i)} - b'_{\pi(i+p+2)}).$$

This implies

$$\begin{aligned}\alpha_i b'_{\pi(i+1)} &\leq \alpha_{i+p+1} b'_{\pi(i+p+2)} + (\alpha_i - \alpha_{i+p+1}) v_{\pi(i)} \\ &= \alpha_{i+p+1} b'_{\pi(i+p+2)} + \sum_{j=i}^{i+p} (\alpha_j - \alpha_{j+1}) v_{m(\pi,i,j)}.\end{aligned}$$

We can now apply recursion to conclude that $\alpha_i b'_{\pi(i+1)} \leq \sum_{j\geq i}(\alpha_j - \alpha_{j+1}) v_{m(\pi,i,j)}$, and hence

$$\mathbf{Rev}(b') = \sum_i \alpha_i b'_{\pi(i+1)} \leq \sum_i \sum_{j\geq i} (\alpha_j - \alpha_{j+1}) v_{m(\pi,i,j)}. \tag{4.4}$$

Notice that equation (4.4) can also be written as a dot product between two vectors of type v_i and $\alpha_j - \alpha_{j+1}$. If we sort the vectors, we see that the $(\alpha_j - \alpha_{j+1})$-vector is the same in both (4.3) and (4.4). Moreover, the sorted vector of v_j for equation (4.4) is dominated by that of equation (4.3), in the sense that it is pointwise smaller. To see this, simply count how many times we have one of $v_1, \ldots, v_i$ appear in both vectors for each index i: for equation (4.3) they appear $\sum_{j=1}^{i} j$ times, whereas for equation (4.4) they appear at most $\sum_{j=1}^{i} 1 + \max\{p \mid m(\pi(j, j+p)) \leq i\} \leq \sum_{j=1}^{i} j$ times. Since the $(\alpha_j - \alpha_{j+1})$-vectors are the same in both equations, the v_i vector in the first equation dominates the order and in the first equation both vectors are sorted in the same order, so it must be the case that $\mathbf{Rev}(b) \geq \mathbf{Rev}(b')$.

It remains to remove our simplifying assumption about saturation and prove the general result. Let b be the optimal efficient equilibrium and let $S \subseteq [n+1]$ be the set of saturated bids, including $n+1$ (where we consider a "fake" player $n+1$ with $b_{n+1} = v_{n+1} = 0$), i.e., $i \in S$ iff $b_i = v_i$. Let $S(i) = \min\{j \in S; j > i\}$.

Given an allocation π, we wish to define an upper bound, $\overline{\mathbf{Rev}}_\pi$, on the revenue of a bid profile that induces allocation π at equilibrium. To this end, we define

$$B_\pi(j) = \begin{cases} \alpha_{S(j)-1} v_{S(j)} + \sum_{i=\sigma(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi,\sigma(j),i)} \\ \qquad \text{if } \sigma(j) \le S(j)-1 \\ \alpha_{S(j)-1} v_{S(j)} - v_j(\alpha_{S(j)-1} - \alpha_{\sigma(j)}) \\ \qquad \text{if } \sigma(j) \ge S(j)-1 \end{cases}$$

We then define

$$\overline{\mathbf{Rev}}_\pi = \sum_j B_\pi(j).$$

We claim that this is, indeed, an upper bound on revenue. Moreover, this bound is tight for revenue at efficient equilibria (i.e. when π is the identity id).

Claim 4.3.4 *If bid profile b induces allocation π at equilibrium, then $\mathbf{Rev}(b) \le \overline{\mathbf{Rev}}_\pi$.*

Claim 4.3.5 *There exists an efficient equilibrium with revenue $\overline{\mathbf{Rev}}_{id}$.*

Using these two claims we want to argue that id is the permutation that maximizes $\overline{\mathbf{Rev}}_\pi$ and therefore we can show that for all inefficient bid profile b' we have

$$\mathbf{Rev}(b') \le \overline{\mathbf{Rev}}_\pi \le \overline{\mathbf{Rev}}_{id} = \mathbf{Rev}(b).$$

To show this, consider some permutation π. Let $j = \max\{k : \pi(k) \neq k\}$ and define a permutation π' such that $\pi'(k) = k$ for $k \ge j$ and $\pi'(k) = \pi(k)$ for $k < \sigma(j)$ and $\pi'(k) = \pi(k+1)$ for $\sigma(j) \le k < j$. Essentially this is picking the last player that is not allocated to his correct slot and bring him there. Now, if we prove that $\overline{\mathbf{Rev}}_{\pi'} \ge \overline{\mathbf{Rev}}_\pi$, then we are done, since we can repeat this procedure many times and get to id.

Claim 4.3.6 $\overline{\mathbf{Rev}}_{\pi'} \geq \overline{\mathbf{Rev}}_{\pi}$.

This completes the proof, subject to the proof of our claims, which we do next.

■

Proof of Claim 4.3.4 : We will show that for all b' inducing allocation π, we have $\alpha_{\sigma(j)} b'_{\sigma(j)+1} \leq B_\pi(j)$. For $\sigma(j) = S(j) - 1$, we use the fact that $b'_{\sigma(j)+1} = b'_{S(j)} \leq v_{S(j)}$. For $\sigma(j) < S(j) - 1$ the result follows in the same way as in the unsaturated case. For $\sigma(j) > S(j) - 1$, we use the fact that player j doesn't want to take slot j and therefore

$$\alpha_{\sigma(j)}(v_j - b'_{\sigma(j)+1}) \geq \alpha_{S(j)-1}(v_j - b'_{S(j)-1}) \geq \alpha_{S(j)-1}(v_j - v_{S(j)})$$

since

$$b'_{S(j)} \leq \min\{v_{\pi(1)}, \ldots, v_{\pi(S(j)-1)}\} \leq v_{S(j)}$$

and $\sigma(j) > S(j) - 1$ so one of the players with value $\leq v_{S(j)}$ must be among the first $S(j) - 1$ slots. Reordering the Nash inequalities above gives us the desired result. ■

Proof of Claim 4.3.5 : This claim follows from the formula defining the optimal-revenue efficient equilibrium in the previous section. ■

Proof of Claim 4.3.6 : Note first that $B_\pi(k) = B_{\pi'}(k)$ for all $k > j$. Moreover, for any k with $\sigma(k) < \sigma(j)$, we will have $\sigma'(k) = \sigma(k)$. In this case, either

$S(k) < \sigma(k)$ in which case $B_{\pi'}(k) = B_\pi(k)$, or else

$$
\begin{aligned}
B_\pi(k) &= \alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi,\sigma(k),i)} \geq \\
&\geq \alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma'(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi',\sigma'(k),i)} \\
&= B_{\pi'}(k).
\end{aligned}
$$

It remains to consider k is such that $\sigma(j) \leq \sigma(k) \leq j$; that is, those players k such that $\sigma(k) \neq \sigma'(k)$. For each such player, we will consider the difference between $B_\pi(k)$ and $B_{\pi'}(k)$. First note that, for player j, we have

$$
\begin{aligned}
&B_\pi(j) - B_{\pi'}(j) \\
&= \left(\alpha_{S(j)-1} v_{S(j)} + \sum_{i=\sigma(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi,\sigma(j),i)} \right) \\
&\quad - \left(\alpha_{S(j)-1} v_{S(j)} + \sum_{i=\sigma'(j)}^{S(j)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi',\sigma'(j),i)} \right) \\
&= \sum_{i=\sigma(j)}^{j-1} (\alpha_i - \alpha_{i+1}) v_j
\end{aligned}
$$

For $k \neq j$, we claim that $B_{\pi'}(k) - B_\pi(k) \geq v_j(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})$. We proceed by two cases. First, if $S(k) \leq \sigma(k)$, we have

$$
\begin{aligned}
B_{\pi'}(k) - B_\pi(k) &= \big(\alpha_{S(k)-1} v_{S(k)} - v_k(\alpha_{S(k)-1} - \alpha_{\sigma'(k)})\big) \\
&\quad - \big(\alpha_{S(k)-1} v_{S(k)} - v_k(\alpha_{S(k)-1} - \alpha_{\sigma(k)})\big) \\
&= v_k(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)}) \geq v_j(\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})
\end{aligned}
$$

Second, if $S(k) - 1 > \sigma(k)$, then we have

$$\begin{aligned}
&B_{\pi'}(k) - B_\pi(k) \\
&= \left(\alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma'(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi',\sigma'(k),i)} \right) \\
&\quad - \left(\alpha_{S(k)-1} v_{S(k)} + \sum_{i=\sigma(k)}^{S(k)-2} (\alpha_i - \alpha_{i+1}) v_{m(\pi,\sigma(k),i)} \right) \\
&= (\alpha_{S(k)-2} - \alpha_{S(k)-1}) v_{m(\pi',\sigma'(k),S(k)-2)} \\
&\quad + \sum_{i=\sigma'(k)}^{S(k)-3} v_{m(\pi',\sigma'(k),i)} [(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2})] \\
&\geq v_j (\alpha_{S(k)-2} - \alpha_{S(k)-1}) \\
&\quad + \sum_{i=\sigma'(k)}^{S(k)-3} v_j [(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2})] \\
&= v_j (\alpha_{\sigma(k)-1} - \alpha_{\sigma(k)})
\end{aligned}$$

Notice that we strongly use the fact that click-through-rates are convex in the last inequality to ensure that $(\alpha_i - \alpha_{i+1}) - (\alpha_{i+1} - \alpha_{i+2}) \geq 0$.

Therefore, taking the sum over all k with $\sigma(j) \leq \sigma(k) \leq j$, we have

$$\begin{aligned}
\sum_{k:\sigma(j)<\sigma(k)\leq j} (B_{\pi'}(k) - B_\pi(k)) &\geq \sum_{i=\sigma(j)}^{j-1} v_j (\alpha_i - \alpha_{i+1}) \\
&= B_\pi(j) - B_{\pi'}(j)
\end{aligned}$$

so that

$$\sum_{k:\sigma(j)\leq\sigma(k)\leq j} (B_{\pi'}(k) - B_\pi(k)) \geq 0.$$

Combining this with the fact that $B_{\pi'}(k) \geq B_\pi(k)$ for all k with $\sigma(k) < \sigma(j)$ or $\sigma(k) > j$, we conclude

$$\overline{\mathbf{Rev}}_{\pi'} = \sum_k B_{\pi'}(k) \geq \sum_k B_\pi(k) = \overline{\mathbf{Rev}}_\pi$$

as desired. ■

CHAPTER 5

SPONSORED SEARCH WITH BUDGETS: A DESIGN APPROACH

In the previous chapter, we studied a mechanism that works well in practice and provided a theoretical analysis of its revenue and welfare in equilibrium, seeking to explain its prevalence on real life applications.

In this and the next chapter we look at sponsored search auctions from a design perspective. Our main goal is to identify important real-life features that are usually ignored in traditional models and propose mechanisms that take those features explicitly into account. As we discussed in Section 1.4, the fact that agents are financially constrained in central to real world ad systems. In fact, the first information Google Ad Words collects from advertisers when creating a campaign is their budget. We also study the problem of designing sponsored search mechanisms that take into account that users are bidding on multiple keywords. In setting with budgets, this becomes very relevant, since budgets tie different keywords together, after all, money spent on a certain keyword cannot be spent on another.

Unlike in previous sections where we study Nash and Bayes-Nash equilibria of auction games, here we adopt the truthful mechanism design methodology, i.e., we design mechanisms that incentivize agents to report their true value. We also tackle the most general problem of designing auctions with budgets for polymatroidal environment. In section 2.8.1 we demonstrated that a very general model of sponsored search can be captured by polymatroids. This will allow our design to be applicable to a wide range of problems in sponsored search, dealing with multiple keywords and multiple slots simultaneously.

5.1 Desiderata

Consider n players, where player i has a positive value v_i per unit of some good g and a budget of B_i. We will consider polyhedral environments (section 2.8), where the set of possible allocations is given by a packing polytope $P \subseteq \mathbb{R}^n$ and we will consider budgeted quasi-linear utilities (section 2.2), which mean that if player i receives x_i amount of good g and pays φ_i, her utility u_i is equal to $v_i x_i - \varphi_i$ if $\varphi_i \leq B_i$ and $-\infty$ otherwise. However, since we will require the mechanism to never charge more than the budgets, we won't have to deal with the latter case. Our goal is to design an auction mechanism that elicits valuations v from the players and outputs a feasible allocation $x(v) \in P$ and a feasible payment vector $\varphi(v) \leq B$ that satisfies the following three properties:

- *Individual Rationality* (a.k.a. voluntary participation): Each player has net non-negative utility from participating in the auction, i.e., $u_i \geq 0$.
- *Incentive compatibility* (a.k.a. truthfulness) : It is a dominant strategy for each player to participate in the auction and report their true value, i.e., $v_i x_i(v_i, v_{-i}) - \varphi_i(v_i, v_{-i}) \geq v_i x_i(v_i', v_{-i}) - \varphi_i(v_i', v_{-i})$. The characterization of single-parameter truthful mechanisms in [64, 5] states that this is equivalent to x_i being a non-decreasing function of v_i (for a fixed v_{-i}) and payments being calculated by $\varphi_i(v_i, v_{-i}) = v_i x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(u, v_{-i}) du$.
- *Pareto-optimality*: An allocation $x(v) \in P$ and payments $\varphi(v) \leq B$ is Pareto-optimal if and only if there is no alternative allocation and payments where all players' utilities and the revenue of the auctioneer do not

decrease, and at least one of them increases. In other words, there is no alternative (x', φ') such that $v_i x'_i - \varphi'_i \geq v_i x_i(v) - \varphi_i(v)$, $\sum_i \varphi'_i \geq \sum_i \varphi_i(v)$ and at least one of those inequalities is strict.

Next we prove a useful lemma about the structure of Pareto-optimal outcomes.

Lemma 5.1.1 *A feasible outcome (x, φ), i.e. $x \in P$ and $\varphi \leq B$, is Pareto-optimal iff there is no $d \in \mathbb{R}^n$ in a dominated direction at x (i.e. $x + d \in P^0 := \{x' \in P; \exists \hat{x} \in P \setminus x', \hat{x} \geq x'\}$) such that $d^t v \geq 0$ and $d_i \leq 0$ for all i that have $\varphi_i = B_i$.*

Proof : In order to show the $\Rightarrow$ direction, assume there is a dominated direction d such that $d^t v \geq 0$ and $d_i \leq 0$ for all i that have $\varphi_i = B_i$. Then define $x'_i = x_i + d_i$ and $\varphi'_i = \varphi_i + v_i d_i$ and we obtain same utilities and the total payment didn't decrease, since $\sum_i \varphi'_i - \sum_i \varphi_i = d^t v \geq 0$. Now, since $x + d$ is not in the boundary of the polytope, we can give some more of good g to some players without charging extra payments and increase their utility. Therefore (x, φ) is not Pareto-optimal.

For the $\Leftarrow$ direction, suppose (x', φ') is a Pareto improvement. Define $d = x' - x$. First we claim that $d^t v > 0$. By the definition of Pareto optimality, $v_i x_i - \varphi_i \leq v_i x'_i - \varphi'_i$, $\sum_i \varphi_i \leq \sum_i \varphi'_i$ and at least one inequality is strict. Summing them all, we get that $\sum_i v_i x_i < \sum_i v_i x'_i$, which implies that $d^t v > 0$. Now, consider two cases:

If $x'_i \leq x_i$ for all i with $\varphi_i = B_i$, then $d_i \leq 0$ for all such i. Simply pick some i for which $d_i > 0$ and decrease d_i slightly. The result will be a dominated direction d' (since $x + d' \leq x + d$ and $x + d' \neq x + d$) with $d'^t v > 0$ and $d'_i \leq 0$ for all i with $\varphi_i = B_i$.

If $x'_i > x_i$ for some i with $\varphi_i = B_i$. Then define d' such that $d'_i = d_i$ if $\varphi_i < B_i$ and $d'_i = \min\{0, d_i\}$ if $\varphi_i = B_i$. Now, consider $x'' = x + d'$ and $\varphi'' = \varphi'$. Clearly d' is a dominated direction (since $x + d' \leq x + d$ and $x + d' \neq x + d$) and $d'_i \leq 0$ for $\varphi_i = B_i$ by definition. Now, we will show that $d'^t v \geq 0$. Notice that $\sum_i p''_i = \sum_i \varphi'_i \geq \sum_i \varphi_i$. Also, except for i with $x'_i > x_i$ and $p_i = B_i$ we have: $v_i x''_i - \varphi''_i = v_i x'_i - \varphi'_i \geq v_i x_i - \varphi_i$. For theremaining i, one has: $v_i x''_i - \varphi''_i = v_i x_i - \varphi'_i \geq v_i x_i - \varphi_i$. Summing all those inequalities, we get $\sum_i v_i x''_i \geq \sum_i v_i x_i$, implying that $d'^t v \geq 0$. ∎

Another simple observation is that if (x, φ) is a Pareto-optimal outcome in which no budget is fully exhausted, then $x = \text{argmax}_{x \in P} v^t x$. For small valuations, any Pareto-optimal mechanism that satisfies individually rationality cannot exhaust budgets, so it must behave like VCG.

5.2 Clinching Auction for Polymatroids

In this section, we describe our main positive result, i.e., an auction with all the desirable properties for polymatroidal environments. This auction is based on the clinching auctions framework of Ausubel [9]. Before we study more complicated constraints, let's recall the clinching auction [9, 31] for the multi-unit setting, i.e., $P = \{x; \sum_i x_i \leq s_0\}$. We begin by setting the supply $s = s_0$ and B_i the budget available to each agent. We maintain a price clock p that begins at zero and gradually ascends. For each price p, the agents are asked how much of the good they demand at the current price. Their demand will be $d_i = \frac{B_i}{p}$ (how much they can afford with their remaining budget) if $p \leq v_i$, and zero otherwise (the case where the price exceeds their marginal value). Then agent i is able to *clinch* an amount $\delta_i = [s - \sum_{j \neq i} d_j]^+$, which is the minimum amount we can give to player i while we are still able to meet the aggregate demands of the other players. *Clinching* means that player i gets δ_i amounts of the good, and $\delta_i p$ is subtracted from his budget. The price increases and we repeat the process until the supply is completely sold.

The heart of the mechanism is the clinching step and generalizing it for more complicated environments involves various challenges: how does one define the notion of supply and aggregate demand (it is not a single number anymore, since there are constraints restraining the possible allocation)? Finally, we need to make sure the clinching step doesn't violate feasibility.

Clinching Framework. First, in Algorithm 1, we consider a slightly modified version of the clinching framework: we maintain a price vector $p \in \mathbb{R}^n_+$ and increase the prices one player at a time. The vector $\rho \in \mathbb{R}^n_+$ contains the promised allocations in each step and its final value is the final allocation of the mechanism. The payment of each agent is the total amount that was deducted from their budget during the execution [1].

For each price, we calculate the demand d_i of each player, which is the amount of the good they would like to get for price p. Then we invoke a procedure called **clinch** which decides the amount to grant to each player at that price. We update the promises, remaining budget and adjust demands[2]. Then

[1]Note that the details of the main procedure **clinch** is not described in this algorithm.

[2]Clearly updating demands is not necessary at this point, but we do in order to make the analysis cleaner.

Algorithm 1: Polyhedral Clinching Auction

Input: P, v_i, B_i

$p_i = 0$, $\rho_i = 0$, $\hat{i} = 1$

do

$d_i = B_i/p_i$ if $p_i < v_i$ and $d_i = 0$ otherwise,

$\delta = \textbf{clinch}(P, \rho, d)$,

$\rho_i = \rho_i + \delta_i, \quad B_i = B_i - p_i\delta_i$,

$d_i = B_i/p_i$ if $p_i < v_i$ and $d_i = 0$ otherwise,

$p_{\hat{i}} = p_{\hat{i}} + \epsilon\,, \quad \hat{i} = \hat{i} + 1 \mod n$

while $d \neq 0$

we increase the price.

In order to define clinching, we need to define analogues of the remnant supply and to demands for the case where the the environment is a generic polytope. Instead of being a single number as in the multi-unit auctions case, the remnant supply and aggregate demands will be polytopes:

Definition 5.2.1 (aggregate demands) *Given P, a vector of promised allocation $\rho \in P$, the remnant supply is described by the polytope $P_\rho = \{x \geq 0; \rho + x \in P\}$. If $d \in \mathbb{R}^n_+$ is the demand vector, the aggregate demand is defined by $P_{\rho,d} = \{x \geq 0; \rho + x \in P, x \leq d\}$.*

In the multi-unit auctions case, the amount player i clinched was the maximum amount we could give him while still being able to meet the demands of the other players. We generalize this notion to polyhedral environments (the concepts are depicted in Figure 5.1):

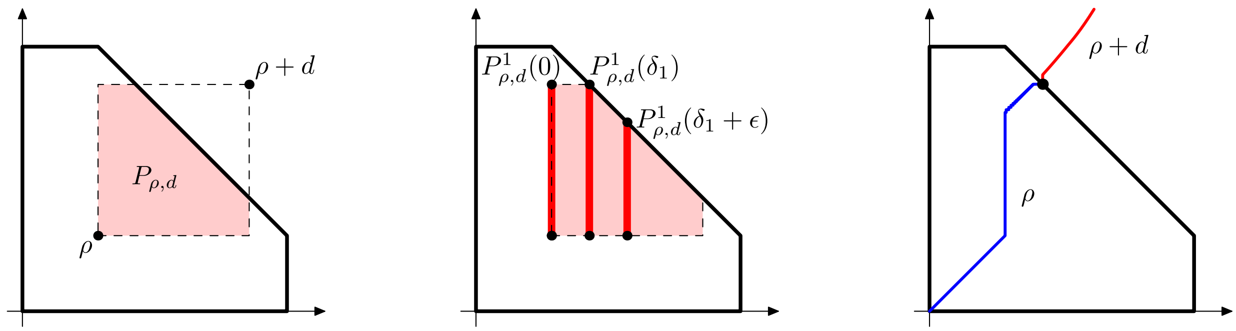

Figure 5.1: Illustration of polyhedral clinching: the first two figures depict the polytopes defined in Definitions 5.2.1 and 5.2.2. The third depicts the mechanism running on polytope P: during the execution, the vector ρ walks inside the polytope (blue line) and the vector $\rho + d$ walks outside it (red line), until they meet at the boundary. The point they meet corresponds to the final allocation.

Definition 5.2.2 (polyhedral clinching) *The* demand set *of players $[n] \setminus i$ if one allocates x_i to player i is represented by the polytope $P^i_{\rho,d}(x_i) = \{x_{-i} \in \mathbb{R}^{[n]\setminus i}_+; (x_i, x_{-i}) \in P_{\rho,d}\}$. Since P is a packing polytope, clearly $P^i_{\rho,d}(x_i) \supseteq P^i_{\rho,d}(x'_i)$ if $x_i \leq x'_i$. The amount player i is able to **clinch** is the maximum amount we can give him without making any allocation for the other players infeasible. More formally, $\delta_i = \sup\{x_i \geq 0; P^i_{\rho,d}(x_i) = P^i_{\rho,d}(0)\}$.*

We need to ensure that the clinching step is well-defined, i.e., that after clinching is performed, the vector of promised allocations is still feasible. This is done by the following lemma:

Lemma 5.2.3 *For each step of the auction above, if $\rho \in P$, then $\rho + \delta \in P$.*

Proof : Let χ^i be the i-th coordinate vector. Note that $\delta_1 \chi^1 \in P_{\rho,d}$ by definition of δ. Now, notice that $\delta_1 \chi^1 \in P^2_{\rho,d}(0) = P^2_{\rho,d}(\delta_2)$, so: $\delta^1 \chi_1 + \delta_2 \chi^2 \in P_{\rho,d}$. By induction, we can show that $\sum_{i=1}^{j} \delta_i \chi^i \in P_{\rho,d}$. The induction is easy: $\sum_{i=1}^{j} \delta_i \chi^i \in P_{\rho,d}^{j+1}(0) = P_{\rho,d}^{j+1}(\delta_{j+1})$, so $\sum_{i=1}^{j+1} \delta_i \chi^i \in P_{\rho,d}^{j+1}$. ■

This auction is clearly truthful, since each player i reports only v_i, and she can stop her participation earlier (which she doesn't want, since she will potentially miss items she are interested in) or later (which will potentially give her items for a price higher than her valuation). It is also individually rational, since players only get items for prices below their valuation and respect budgets by the definition. Notice that those facts are true regardless of the trajectory of the price vector: any process that increases prices (in a potentially non-uniform way) has this property.

Lemma 5.2.4 *The auction in Algorithm 1 along with the clinching step described in Definition 5.2.2 is truthful, individually-rational and budget-feasible.*

Clinching for polymatroids. Notice that we haven't used anything from polymatroids yet, so Lemma 5.2.4 holds for any polytope P. However, two things are left to be shown: (i) that amount clinched can be computed efficiently and (ii) that the outcome is Pareto optimal. To show both of these properties, we use the fact that P is a polymatroid.

Lemma 5.2.5 *If the environment is a polymatroid P defined by a submodular function f, then the amount player i clinches in Algorithm 1 is given by:*

$$\delta_i = (\max_{x \in P_{\rho,d}} \mathbb{1}^t x) - (\max_{x \in P_{\rho,d}} \mathbb{1}_{-i}^t x_{-i}).$$

Moreover, this can be calculated efficiently using submodular minimization.

The main ingredients of the proof are the following two facts about polymatroids:

Fact 5.2.6 (Schrijver [74], sections 44.1 and 44.4) *If P is a polymatroid defined by the submodular function f, then $P_{\rho,d}$ is also a polymatroid defined by the following submodular function:*

$$\hat{f}(S) = \min_{T \subseteq S} \{f(T) - \rho(T) + d(S \setminus T)\}.$$

Notice that $\hat{f}(\cdot)$ might not be monotone. However,

$$\bar{f}(S) = \min_{S' \supseteq S} \hat{f}(S')$$

is a monotone submodular function that defines the same polymatroid.

Fact 5.2.7 *Given two monotone submodular functions $f, \tilde{f}$, then the polymatroids $P, \tilde{P}$ defined by them are equal iff the functions are equal. The ($\Leftarrow$) direction is trivial. For the other direction, notice that if $f(S) < \tilde{f}(S)$ say for $S = \{1, \dots, i\}$, notice that the point x such that $x_j = \tilde{f}(\{1..j\}) - \tilde{f}(\{1..j-1\})$ for $j \leq i$ and zero otherwise is such that $x \in \tilde{P} \setminus P$.*

Proof of Lemma 5.2.5 : Using the fact 5.2.6, we know that $P^i_{\rho,d}(x_i)$ is also a polymatroid defined over $[n] \setminus i$ by the function $\tilde{f}(S) = \min\{\bar{f}(S), \bar{f}(S \cup i) - x_i\}$. Now, we use Fact 5.2.7 to see that $P^i_{\rho,d}(x_i) = P^i_{\rho,d}(0)$ iff $\bar{f}(S) \leq \bar{f}(S \cup i) - x_i, \forall S \subseteq [n] \setminus i$. So, $\delta_i = \min_{S \subseteq [n] \setminus i} \bar{f}(S \cup i) - \bar{f}(S)$. Since $\bar{f}$ is submodular, the smallest marginal can only be

$$\bar{f}([n]) - \bar{f}([n] \setminus i) = \max\{0, \hat{f}([n]) - \hat{f}([n] \setminus i)\}$$

which is exactly the expression in the statement of the lemma. Now, one can easily see that evaluating $\hat{f}$ is a submodular minimization problem. ■

Now we prove that the outcomes are Pareto-optimal in two steps. The first step is to characterize Pareto-optimal allocations for polymatroids. This characterization is stronger than that of Lemma 5.1.1, since it explores the structure of polymatroids. Afterwards, we show that the outcomes of the chinching auction defined in Algorithm 1 satisfy the two conditions in the characterization lemma.

In the following, for a vector $x \in \mathbb{R}^n$ and $S \subseteq [n]$ we denote $x(S) = \sum_{i \in S} x_i$.

Lemma 5.2.8 *For a polymatroidal environment P defined by a submodular function f, an allocation (x, φ) is Pareto optimal iff:*

1. *All items are sold, i.e., $x([n]) = f([n])$, and*
2. *Given a player i with $\varphi_i < B_i$ and player j with $v_j < v_i$, then there exists a set S such that $x(S) = f(S)$, $i \in S$ and $j \notin S$*

The following elementary facts about submodular functions will be useful in the proof of the Lemma 5.2.8:

Fact 5.2.9 *Given a vector $x \in P$, if two sets S, T are tight (i.e. $x(S) = f(S)$ and $x(T) = f(T)$), then $S \cap T$ and $S \cup T$ are also tight. The proof is quite elementary: $x(S \cup T) = x(S) + x(T) - x(S \cap T) \geq f(S) + f(T) - f(S \cap T) \geq f(S \cup T)$. So, all the inequalities must be tight and therefore $x(S \cup T) = f(S \cup T)$ and $x(S \cap T) = f(S \cap T)$.*

Fact 5.2.10 *If $x([n]) < f([n])$ then there is one component x_i that we can increase by $\delta > 0$ such that x is still in P. It follows from the previous fact: if all players i were contained in a tight set, one could take the union of those and $[n]$ would be tight. Then there is some element i which is in no tight set.*

Proof of Lemma 5.2.8 : The ($\Rightarrow$) direction is easy. If $x([n]) < f([n])$ then we can increase some x_i (Fact 5.2.10) and still get point P generating a Pareto improvement. Also, if there is $\varphi_i < B_i$ and $v_j < v_i$ and no tight set separating them, then we can consider another outcome where we increase x_i by some $\delta > 0$, decrease x_j by some $\delta < 0$ and still get a feasible point improving $x^t v$. Now, this would not be Pareto optimal by Lemma 5.1.1.

For the ($\Leftarrow$) direction, let (x, φ) be an outcome satisfying properties 1 and 2 and suppose (x', φ') is a Pareto-improvement. This means that $v_i x'_i - \varphi'_i \geq v_i x_i - \varphi_i$ and $\sum_i \varphi'_i \geq \sum_i \varphi_i$.

Let $\{i_1, \ldots, i_k\} = \{n\} \cup \{i; \varphi_i < B_i\}$, sorted in non-increasing order of v_i. Using property 2 (notice it holds for player n trivially) together with fact 5.2.9, we define the following family of tight sets $S_1 \subseteq S_2 \subseteq \ldots \subseteq S_k = [n]$, tight in the sense that $x(S_i) = f(S_i)$. For all $v_t < v_{i_1}$ there is a tight set S_{1t} that has i_1 but not t. Let S_1 be the intersection of such sets. Now, given $S_1 \subseteq \ldots \subseteq S_{j-1}$, we define S_j in the following way: If $i_j \in S_{j-1}$, take $S_j = S_{j-1}$ (notice can only happen if $v_{i_j} = v_{i_{j-1}}$). If not, for each $v_t < v_{i_j}$ there is a tight set S_{jt} that has i_j but not t. Now, define S_j as the union of S_{j-1} and the intersection of the S_{jt} sets.

By eliminating duplicates and its corresponding elements from $\{i_1 \ldots i_k\}$, we get a family $S_1 \subset \ldots \subset S_k$. Define $T_j = S_j \setminus S_{j-1}$ and it is clear the family obtained has the following properties:

- all $t \in S_j$ have $v_t \geq v_{i_j}$
- $i_j \in T_j$
- for all $i \in T_j$ either $v_i = v_{i_j}$ or $\varphi_i = B_i$.

Let $T'_j = \{i \in T_j ; v_i \neq v_{i_j}\}$ and $T''_j = \{i \in T_j ; v_i = v_{i_j}\}$. Since the players in T'_j have exhausted their budget, $\varphi_i \geq \varphi'_i$. Using that and Pareto-optimality, we get:

$$\begin{aligned}
&\sum_{i \in T'_j} \varphi_i - \varphi'_i \geq \\
&\quad \geq \sum_{i \in T'_j, x_i \geq x'_i} \varphi_i - \varphi'_i \geq \sum_{i \in T'_j, x_i \geq x'_i} v_i(x_i - x'_i) \overset{*}{\geq} \\
&\quad \geq \sum_{i \in T'_j, x_i \geq x'_i} v_{i_j}(x_i - x'_i) \overset{**}{\geq} \sum_{i \in T'_j} v_{i_j}(x_i - x'_i)
\end{aligned} \tag{5.1}$$

Now, we can add the inequality $\varphi_i - \varphi'_i \geq v_{i_j}(x_i - x'_i)$ for $i \in T''_j$ and obtain:

$$\sum_{i \in T_j} \varphi_i - \varphi'_i \geq \sum_{i \in T_j} v_{i_j}(x_i - x'_i) \tag{5.2}$$

Summing those for all j and get:

$$\sum_i \varphi_i - \varphi'_i \geq \sum_j \sum_{i \in T_j} v_{i_j}(x_i - x'_i) = \sum_j (v_{i_j} - v_{i_{j+1}}) \sum_{i \in S_j} (x_i - x'_i) \geq 0$$

since $x(S_j) = f(S_j) \geq x'(S_j)$. Therefore $\sum_i \varphi_i \geq \sum_i \varphi'_i$ and therefore equal. This means in particular all of the inequalities in (5.1) and (5.2) must be tight. Therefore for all $i \in T'_j$ we need to have $x_i = x'_i$, since if $x_i > x'_i$ then inequality $*$ in (5.1) would be strict. If $x_i < x'_i$, then inequality $**$ would be strict. We use this fact to show that $\sum_{i \in S_j} v_i(x_i - x'_i) \geq 0$ by induction on j. If we show that, we can take $j = k$ and then we are done, since this will imply that $\sum_i v_i x_i \geq \sum_i v_i x'_i$ and therefore (x', φ') cannot be a Pareto-improvement.

For $j = 1$, this is trivial, since we can write:

$$\sum_{i \in S_1, v_i \neq v_{i_1}} v_i(x_i - x'_i) \geq \sum_{i \in S_1, v_i \neq v_{i_1}} v_{i_1}(x_i - x'_i)$$

since both terms are zero, and then sum $v_i(x_i - x'_i)$ for the rest of the elements in

S_1 and use the fact that S_1 is tight. For other j, we use that:

$$\sum_{i \in S_j} v_i(x_i - x'_i) \geq$$
$$\geq v_{i_{j-1}} \sum_{i \in S_{j-1}} (x_i - x'_i) + \sum_{i \in T'_j} v_i(x_i - x'_i) + \sum_{i \in T''_j} v_{i_j}(x_i - x'_i) \geq$$
$$\geq v_{i_j} \sum_{i \in S_j} (x_i - x'_i) = v_{i_j}(x(S_j) - x'(S_j)) \geq 0,$$

by the fact that S_j is tight. ■

Now, we argue that, for sufficiently small ϵ, the outcome satisfied the two properties in Lemma 5.2.8 and hence is Pareto-optimal. We prove this fact using the following sequence of lemmas:

Lemma 5.2.11 *After the clinching step is executed, and before updating prices,* $\hat{f}([n]) \leq \hat{f}([n] \setminus j), \forall j \in [n]$.

Proof : In the clinching step, given an initial $\hat{f}_0$, we define $\delta_i = \max\{0, \hat{f}_0([n]) - \hat{f}_0([n] \setminus i)\}$. After we update ρ, B, d, $\hat{f}$ is updated to $\hat{f}_1(S) = \hat{f}_0(S) - \delta(S)$. Now, it is easy to check that:

$$\hat{f}_1([n]) = \hat{f}_0([n]) - \sum_i \delta_i = \hat{f}_0([n]) - \delta_j - \sum_{i \neq j} \delta_i \leq \hat{f}_0([n] \setminus j) - \sum_{i \neq j} \delta_i = \hat{f}_1([n] \setminus j)$$

■

Lemma 5.2.12 *The outcome* (x, φ) *of the clinching auction is such that* $x([n]) = f([n])$.

Proof : We show the following invariant: if we define $\hat{f}$ as in Fact 5.2.6, updating it each round as ρ, d changes, we claim that the value of $\mathbb{1}^t \rho + \hat{f}([n])$ remains constant.

To do so, we consider the events that can cause it to drop:

1. clinching: just after clinching occurs (i.e. ρ_i increases by δ_i, budgets decrease by $p\delta_i$, demands are adjusted, but before the price increases), the amount $\mathbb{1}^t\rho + \hat{f}([n])$ remains the same since $\rho^t\mathbb{1}$ increases by $\delta^t\mathbb{1}$ and for all S, $\hat{f}(S)$ decreases by $\delta(S)$, because to each i, ρ increases by δ_i and d_i decreases by δ_i.

2. price p_i increases and d_i decreases by $\theta, 0 \leq \theta \leq d_i$. If $\mathbb{1}^t\rho + \hat{f}([n])$ decreased then there was some $T, i \notin T$ such that:

$$\hat{f}([n] \setminus i) + d_i - \theta \leq f(T) - \rho(T) + d([n] \setminus T) - \theta < \hat{f}([n])$$

Using Lemma 5.2.11, we know that $\hat{f}([n]) \leq \hat{f}([n] \setminus i)$, so $d_i < \theta$ which is not true.

$\blacksquare$

The proofs of the previous two lemmas intuitively establishes the maximality of the clinching procedure. Lemma 5.2.11 can be interpreted as saying that if we apply the clinching procedure twice, without updating prices, then the second time will have no effect. The proof of Lemma 5.2.12 identifies an invariant that is maintained during the execution of the mechanism.

Lemma 5.2.13 *If $\epsilon < \min_{v_i \neq v_j} |v_i - v_j|$, then property 2 of Lemma 5.2.8 is satisfied.*

Proof : Suppose not and for the final outcome there are $v_j < v_i$, $\varphi_i < B_i$ and all sets S such that $i \in S$, $j \notin S$ are not tight. First, clearly $x_j \neq 0$, otherwise $[n] \setminus j$

would be tight by Lemma 5.2.12. Then consider $\tilde{x}$ where $\tilde{x}_i = x_i + \theta$, $\tilde{x}_j = x_j - \theta$ and $\tilde{x}_k = x_k$ for all $k \neq i, j$. It is feasible for some small θ.

Now, consider the promised allocation ρ and demands d just before the last time player j clinched an amount $\delta_j > 0$. If necessary decrease θ so that it becomes smaller than this last amount clinched, i.e., $\theta < \delta_j$. At this point $\rho \leq x \leq \rho + d$. By the definition of clinching: $P^j_{\rho,d}(\delta_j) = P^j_{\rho,d}(\theta) = P^j_{\rho,d}(0)$.

At this point, $\rho \leq x$ and $\rho_j + \theta < \rho_j + \delta_j = x_j$. Therefore $\tilde{x} \geq \rho$. Also, we have that $x - \rho \leq d$ and $x_i - \rho_i < d_i$, since agent i hasn't dropped his demand to zero yet and his demand never increases and won't be met while $v_i < p_i$. Here we are strongly using that $\epsilon < \min_{v_i \neq v_j} |v_i - v_j|$ to ensure that for the last time player j clinches, player i demand is not zero yet. This implies that $\tilde{x} - \rho \in P_{\rho,d}$ so $(\tilde{x} - \rho)_{-j} \in P^j_{\rho,d}(0)$. Now, the fact that $P^j_{\rho,d}(0) = P^j_{\rho,d}(\delta_i)$ implies that $\hat{x} = (x_j, \tilde{x}_{-j}) \in P$. But $\hat{x}([n]) = x([n]) + \theta = f([n]) + \theta > f([n])$, which is an absurd. ∎

We can summarize the results as:

Theorem 5.2.14 *For a polymatroidal environment, the auction in Algorithm 1 along with the clinching step described in Definition 5.2.2 has all the desirable properties.*

5.2.1 Extensions of the clinching framework

The clinching framework described in Algorithm 1 and Definition 5.2.2 is quite flexible: one can change the way clinching is done or the way prices ascend and obtain an auction that is still truthful, individually rational, and respects budgets. Pareto-optimality, however, is a delicate property to achieve.

Scaled polymatroids: If $P \subseteq \mathbb{R}^n_+$ is a polymatroid, $\gamma \in \mathbb{R}^n_+$, then we call $P_\gamma = \{x; (\frac{x_i}{\gamma_i})_i \in P\}$ a scaled polymatroid. If the environment is P_γ, it is easy to see that a truthful, individually-rational, Pareto-optimal and budget-feasible auction is obtained by running the polymatroid clinching auction on P with inputs $\gamma_i v_i$ instead of v_i. It is simple to see that this is equivalent to a standard clinching auction with input values v_i but price clocks advancing on a different speed for each player. Scaled polymatroids are important since they correspond to the setting of AdWords with Quality Factors discussed in section 2.8.2).

Beyond scaled polymatroids: one could change the way clinching is done and one could change the price trajectories, maybe in a more sophisticated way then the one we did for the scaled polymatroids. If the trajectory is such that it only depends on P and budgets (not on values) and p_i never decreases, the auction retains truthfulness and budget feasibility. Here we argue that none of such changes would generate a Pareto-optimal auction when P is not a scaled polymatroid. Assume $n = 2$ for simplicity and imagine that there is trajectory for the price vector p and a clinching procedure. Also assumes valuations are much smaller than budgets in such a way that the mechanism cannott exhaust budgets and therefore the auction must allocate like VCG. Any such mechanism must decide on the whole allocation the first time $p_i = v_i$ for some component, having only the information that $v_j \geq p_j$ for the other component, since at this point he needs to allocate to the second player. So, the environment must be such that there is a price trajectory $p(t)$, where the optimal allocation is constant for all points (p_1, v_2) for all $v_2 > p_2$. And also, it should be constant for all (v_1, p_2) for all $v_1 > p_1$. Notice that scaled polymatroids are exactly those environments.

5.2.2 Faster clinching subroutines

In Lemma 5.2.5 we showed that for any generic polymatroidal environment P we can perform the clinching step in polynomial time if we have oracle access to the submodular function defining the polymatroid. In order to do so, we solve a submodular minimization problem. For most practical applications, however, one can design much simpler and faster algorithms for clinching. Clinching involves solving the following problem: given an environment P, $\rho \in P$ and $d \in \mathbb{R}^n_+$ we want to compute:

$$\max\{\mathbb{1}^t x; x + \rho \in P, 0 \leq x \leq d\} \tag{5.3}$$

We illustrate how to solve this problem efficiently for the single-keyword AdWords polytope: given $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_n$, consider the environment:

$$P = \{x \in \mathbb{R}^n_+; x(S) \leq \textstyle\sum_{j=1}^{|S|} \alpha_j, \forall S \subseteq [n]\}$$

Lemma 5.2.15 *For the single-keyword AdWords polytope, the optimization problem defined in equation (5.3) can be solved using the following greedy algorithm: we can assume wlog that the components are sorted such that* $\rho_1 + d_1 \geq \rho_2 + d_2 \geq \ldots \geq \rho_n + d_n$. *Now, define inductively*

$$z_i = \min\{\rho_i + d_i, \textstyle\sum_{j=1}^{i} \alpha_j - \sum_{j=1}^{i-1} z_j\}.$$

Then $\sum_i z_i - \rho_i$ *is the solution to the problem.*

Proof : If we drop the restriction that $x \geq 0$ in (5.3), then it is easy to see that $x = z - \rho$ is an optimal solution to this problem using, for example, a local exchange argument. Now, we show that we can fix this problem, by modifying z such that $z \geq \rho$.

In order to fix that, consider the smaller i such that $z_i < \rho_i$. By the definition of z_i, it must be the case that $z_i = \sum_{j=1}^{i} \alpha_j - \sum_{j=1}^{i-1} z_j$, so $\sum_{j=1}^{i} z_j = \sum_{j=1}^{i} \alpha_j$. Then there must be some $k < i$ such that $z_k > \rho_k$, otherwise we would have $\sum_{j=1}^{i} \rho_j > \sum_{j=1}^{i} \alpha_j$ contradicting the fact that $\rho \in P$. Notice we can increase z_i by some small δ and decrease z_k by a small δ. And obtain another vector z which is also such that $z \in P, z \leq \rho + d$ and has the same $\mathbb{1}^t z$ value (the fact that $z \in P$ after this transformation is due to the nature of the constraints). We can repeat this process until we get $z \geq \rho$. ■

5.3 Limitations of auctions for budget-constrained agents

Previously, we argued why simple modifications to the clinching auction would not work for polyhedral environments beyond (scaled) polymatroids. Here, we explore the possibility of designing an auction of a different format achieving those properties and show that this is not possible even for two players. We do so through a general characterization of Pareto-optimal auctions with desirable properties. Before stating the characterization, we study the case with one budget-constrained player and prove some lemmas that are useful in proving the general characterization result later.

5.3.1 One budget-constrained player

For ease of exposition, we first focus on 2 players and assume that the feasible set of allocations P has a smooth and strictly-concave boundary, in the sense that for each $v \in \mathbb{R}^2_+$ there is a single point $x^*(v) \in P$ maximizing $v^t x$ such that $x^*(v)$ is a $\mathcal{C}^\infty$-function. In fact, one can approximate any polytope by such a set using the technique of Dolev et al [33]. Using compactness arguments, it is possible to get an auction for the original environment by taking the limit of the auctions obtained for its $\mathcal{C}^\infty$-approximations.

Assume that player 1 is not budget constrained and player 2 has budget B_2 and let (x^*, φ^*) be the VCG mechanism for this setting. Now, we can define the function:

$$\xi(v_1) = \min\{v_2; \varphi^*_2(v_1, v_2) \geq B_2\}$$

Theorem 5.3.1 *The allocation rule*

$$x(v_1, v_2) = x^*(v_1, \min\{v_2, \xi(v_1)\})$$

is monotone. Moreover, when coupled with the appropriate payment rule, it generates a Pareto-optimal and budget feasible mechanism

Proof : The main part of the proof is to show that the allocation is monotone. If we show that, it is clearly budget feasible for player 2, since we use the VCG-payment rule until the point the budget of player 2 gets exhausted and from that point on, the allocation is constant. When the budgets of the players are not exhausted, the allocation is efficient (since it mimics VCG) and therefore is Pareto-optimal. The allocation when the budget of player 2 is exhausted is equivalent to the VCG allocation of a pair (v_1, v_2') with $v_2' \leq v_2$, so player 1 is getting $x_1^*(v_1, v_2') \geq x_1^*(v_1, v_2)$ by monotonicity of VCG. This implies Pareto-optimality as a consequence of Lemma 5.1.1.

Monotonicity: The allocation rule is clearly monotone for player 2. We need to show it is monotone for player 1, i.e. that the function $t \mapsto x_1(v_1 + t, v_2)$ is monotone non-decreasing. It is clearly so for intervals where $\xi(v_1 + t) \geq v_2$, so let's assume that for $t \in (-\epsilon, +\epsilon)$ we have $\xi(v_1 + t) < v_2$. Our goal is to show that: $\frac{d}{dt} x_1^*(v_1 + t, \xi(v_1 + t)) \geq 0$. Since the VCG-allocation lies in the boundary of P, this is the same as showing that $\frac{d}{dt} x_2^*(v_1 + t, \xi(v_1 + t)) \leq 0$. The crucial observation is that the VCG-payment for player 2 on the curve $(v_1 + t, \xi(v_1 + t))$ is constant, i.e.:

$$B_2 \equiv \varphi_2^*(v_1 + t, \xi(v_1 + t)) = \xi(v_1 + t) x_2^*(v_1 + t, \xi(v_1 + t)) - \int_0^{\xi(v_1+t)} x_2^*(v_1 + t, u) du$$

Now, we can simply derivate it with respect to t. We use the notation $\partial_i f(\cdot)$ for the derivative of f with respect to the i-th variable. We also define $x_2^*(t) = x_2^*(v_1 + t, \xi(v_1 + t))$. Now,

$$0 = \xi'(v_1 + t) x_2^*(t) + \xi(v_1 + t) \frac{d}{dt} x_2^*(t) - \xi'(v_1 + t) x_2^*(t) - \int_0^{\xi(v_1+t)} \partial_1 x_2^*(v_1 + t, u) du$$

$$\xi(v_1 + t) \frac{d}{dt} x_2^*(t) = \int_0^{\xi(v_1+t)} \partial_1 x_2^*(v_1 + t, u) du \leq 0$$

since $x_2^*(v_1, v_2)$ decreases with v_1 by the definition of the VCG allocation. ■

A variant of the proof can be used to show the following result for 2 budget constrained players. This is useful for our general characterization.

Corollary 5.3.2 *If the functions $\xi_1(v_2), \xi_2(v_1)$ are such that the regions $\{v; v_2 \geq \xi_2(v_1)\}$ and $\{v; v_1 \geq \xi_1(v_2)\}$ are disjoint, then one can define*

$$x(v_1, v_2) = x^*(\min\{v_1, \xi_1(v_2)\}, \min\{v_2, \xi_2(v_1)\})$$

$$\varphi_i(v) = v_i x_i(v) - \int_0^{v_i} x_i(u, v_{-i}) du.$$

If $\varphi_2(v_1, \xi_2(v_1)) \equiv B_2$ and $\varphi_1(\xi_1(v_2), v_2) \equiv B_1$, then (x, φ) is a mechanism with the desirable properties.

The above corollary has a strong fixed-point flavour and it is tempting to believe one could get the existence of such a mechanism from this theorem. This is however not true, as shown in the next section. However, this result remains useful as a tool for searching for such mechanisms whenever they exist. For example, one can extend the above theorem to prove the existence of the mechanisms for polyhedral environments when B_1 is much larger then B_2.

5.3.2 Characterization and impossibility

Now we discuss our main negative result: which states an impossibility of extending the auction for polymatroids to general polyhedral environments.

Theorem 5.3.3 (Impossibility) *There is no general auction for every polyhedral environment and every pair of budgets that satisfies the desirable properties.*

We prove it in two steps: first we prove a sequence of lemmas characterizing 2-player auctions for polyhedral environments satisfying all the desirable properties. Then we fix a specific polyhedral environment and argue that no mechanism can possibly satisfy this characterization.

First, we begin by understanding the format of an auction with the desirable properties where the environment is a packing polytope $P \subseteq \mathbb{R}^2_+$. We start by defining a family of VCG auctions.

VCG-family: We say that a mechanism is in the *VCG-family* if its allocation $x(v) \in P$ is such that

$$x(v) \in X^*(v) := \text{argmax}_{x \in P} v^t x.$$

Notice that there might be more than one such mechanism: if v is normal to an edge of the polytope, then the entire edge is in the argmax. Nevertheless, $x(v_1+, v_2) = \lim_{v_1' \downarrow v_1} x(v_1', v_2)$ and $x(v_1, v_2+) = \lim_{v_2' \downarrow v_2} x(v_1, v_2')$ are common for the entire family, as we see in the following lemma. A consequence of this fact is that the payment function might not be unique, but $\varphi(v_1+, v_2)$ and $\varphi(v_1, v_2+)$ are unique.

Lemma 5.3.4 *Given a convex set $P \subseteq \mathbb{R}^2_+$ and two allocation rules $x(v), \tilde{x}(v) \in X^*(v)$:= argmax$_{x \in P} v^t x$, then $x(v_1+, v_2) = \tilde{x}(v_1+, v_2)$, where $x(v_1+, v_2) = \lim_{v_1' \downarrow v_1} x(v_1', v_2)$.*

Proof : Suppose that $x(v_1+, v_2) \neq \tilde{x}(v_1+, v_2)$, then say that $x_1(v_1+, v_2) > \tilde{x}_1(v_1+, v_2)$. Since both are in the boundary of the polytope, it means that for all (v_1', v_2) with $v_1' > v_1$, $\tilde{x}_1(v_1', v_2) \geq x_1(v_1+, v_2)$, since the points to the right of $x(v_1+, v_2)$ are clearly better then the ones to the left of it. So $\tilde{x}_1(v_1', v_2)$ can't converge to $\tilde{x}_1(v_1+, v_2) < x_1(v_1, x_2)$. ■

To illustrate this fact, consider the simple case of $P = \{x \in \mathbb{R}^2_+; x_1 + x_2 \leq 1\}$. Then the VCG mechanism is well-defined for $x_1 \neq x_2$, which is, simply to allocate to the player with the highest value the entire amount. But notice that completing this mechanism with any allocation in the points (v, v) generates a truthful mechanism. The payments of different mechanisms of the VCG family differ on (v, v), for example $\varphi_i(v, v) = vx_i(v)$, but notice that the payments everywhere else are well-defined.

Pareto-optimal mechanisms: Now we turn our attention back to Pareto-optimal mechanisms for two budget-constrained players. Let (x, φ) be such mechanism. As a direct consequence of the characterization of Pareto optimal outcomes (Lemma 5.1.1), we know the following:

- if $\varphi_i(v) < B_i$ then $x_i(v) \geq \min\{x_i; x \in X^*(v)\}$
- if $\varphi_1(v) < B_1, \varphi_2(v) < B_2$ then $x(v) \in X^*(v)$

And a simple consequence of truthfulness:

- if $\varphi_i(v) = B_i$ then for all $v_i' \geq v_i$, $x(v_i', v_{-i}) = x(v_i, v_{-i})$.

Now, we are ready to start proving the characterization theorem. We will characterize the mechanism in terms of the regions in the space of valuations where the budgets get exhausted. For formally, we are interested in understanding the sets:

$$E_i = \{v \in \mathbb{R}^2_+; \varphi_i(v) = B_i\}$$

Lemma 5.3.5 *If mechanism (x, φ) has the desirable properties, then either $E_1 \cap E_2 = \emptyset$ or there exists some $a \in \mathbb{R}^2_+$ such that $\{v; v > a\} \subseteq E_1 \cap E_2 \subseteq \{v; v \geq a\}$.*

Proof : Assume that $E_1 \cap E_2$ is not empty. Then we will prove the lemma in two parts. For the first part we will prove two statements: (i) if that if $v^0 \in E_1 \cap E_2$ and $v \geq v^0$ then $x(v) = x(v^0)$ and (ii) if $v, v' \in E_1 \cap E_2$ and $v_i^* = \min\{v_i, v_i'\}$ then $x(v^*) = x(v) = x(v')$. Then for the second part, we show that this whole region that has constant allocation has budget exhausted for the two players. See figure 5.2 for an illustration of the proof.

For (i), let $v^1 = (v_1, v_2^0)$ and $v^2 = (v_1^0, v_2)$ and notice that $x(v^0) = x(v^1) = x(v^2)$, since budgets are exhausted so the allocation can't increase. By monotonicity, $x_2(v) \geq x_2(v^1) = x_2(v^0)$ and $x_1(v) \geq x_1(v^2) = x_1(v^0)$. Since all allocations lie in the boundary of the polytope, we must have $x(v) = x(v^0)$.

The proof of (ii) is very similar, define $v_i^0 = \max\{v_i, v_i'\}$. Then by (i), $x(v) = x(v^0) = x(v')$ now, by the exact same argument as above we show that $x(v^*) = x(v^0)$.

Now, if $a_i = \inf\{v_i; v \in E_1 \cap E_2\}$, let us show that for $v > a$, $v \in E_1 \cap E_2$. Let us show that $v \in E_2$ and then E_1 is analogous. By definition, there is some $v' \in E_1 \cap E_2$ with $v_1' < v_1$. Then $(v_1', v_2) \in E_2$ since the allocation is constant, for $v > a$, the budget of 2 is exhausted in (v_1', v_2) iff it is exhausted in (v_1', v_2'). Now, note that by monotonicity $x_1(v_1', u) \leq x_1(v_1, u)$ and since allocation is in the boundary of the polytope $x_2(v_1', u) \geq x_2(v_1, u)$. By the payment formula, $p_2(v) = v_2 x_2(v) - \int_0^{v_2} x_2(v_1, u) du$ and by the fact that $x_2(v) = x_2(v_1', v_2)$, we have that $B_2 \geq p_2(v) \geq p_2(v_1', v_2) = B_2$, so $v \in E_2$. ■

The next lemma further describes the regions E_1 and E_2:

Lemma 5.3.6 *If P is a packing polytope, there is a finite set of vectors $\{u^1, u^2, \ldots, u^k\}$ such that for $v \neq tu^i$ for some $t \geq 0$, the VCG family is uniquely defined. Moreover,*

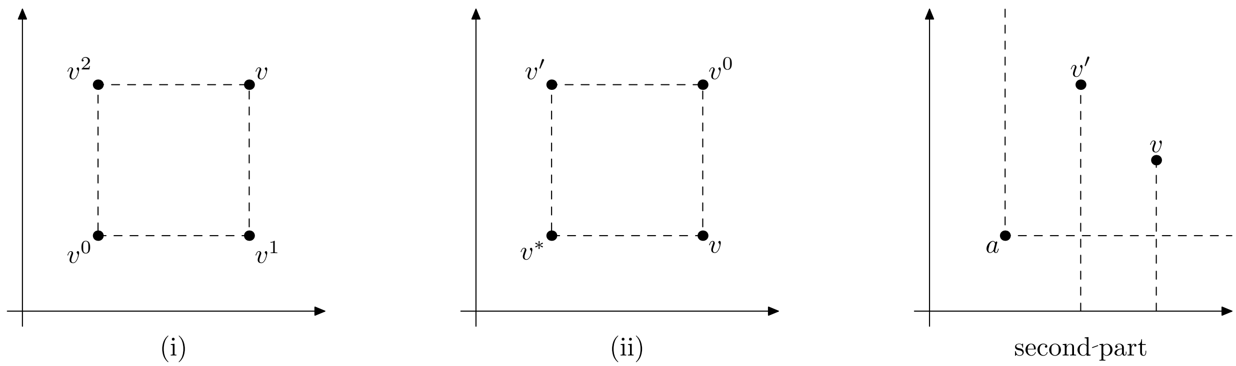

Figure 5.2: Illustration of the proof of Lemma 5.3.5

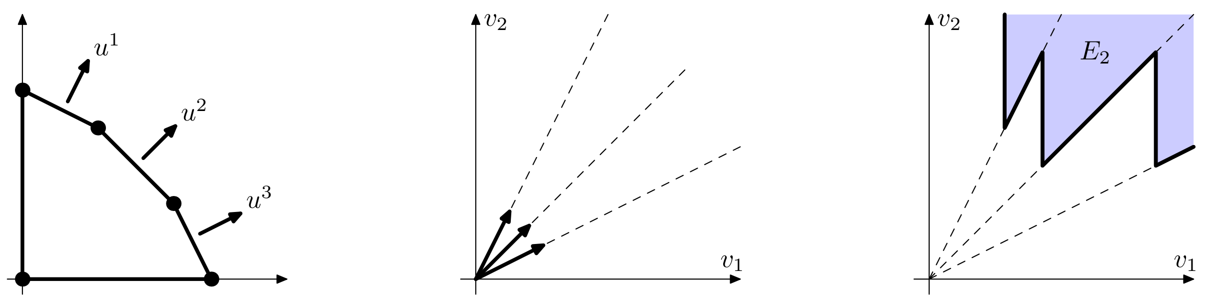

Figure 5.3: Illustration of the proof of Lemma 5.3.6

if (x, φ) *is a mechanism with the desirable properties and* $E_1 \cap E_2 \subseteq \{v; v \geq a\}$ *then if* $\xi_i(v_{-i}) = \inf\{v_i; p_i(v_i, v_{-i}) = B_i\}$ *and* $v_{-i} < a_{-i}$*, then* $(\xi_i(v_{-i}), v_{-i}) = tu^i$ *for some* t, i*.*

Proof : The set of vectors $\{u^1, u^2, \ldots, u^k\}$ is simply the set of normals of the edges of the polytope as depicted in the first part of Figure 5.3. If v is not an edge in the polytope, then the point in P maximizing $v^t x$ is a vertex and therefore uniquely defined. If we draw the lines $u^i \cdot t$ for $t > 0$ we divide the space of all possible valuations in regions (see second part of the figure): the regions correspond to the vertices and the lines to edges of the polytope.

Now, given a certain mechanism (x, φ), suppose that $v_1 < a_1$ the vector $(v_1, \xi_2(v_1))$ is not normal to any edge of P. Then clearly $x^*(v_1, \xi_2(v_1))$ is well-defined and moreover, for some $\delta > 0$ and $v_2' \in [\xi_2(v_1) - \delta, \xi_2(v_1) + \delta]$, $x^*(v_1, v_2')$ is well-defined and constant in the v_2'-range. Also $(v_1, v_2') \notin E_1 \cap E_2$, since $v_1 < a_1$. For such $v_2' > \xi_2(v_1)$, we know that $(v_1, v_2') \in E_2$, thus it is not in E_1 and therefore $x_2(v_1, v_2') = x_2(v_1, \xi_2(v_1)) \leq x_2^*(v_1, \xi_2(v_1))$. For $v_2' < \xi_2(v_1)$, $(v_1, v_2') \notin E_2$ so $x_2(v_1, v_2') \geq x_2^*(v_1, v_2')$. Using that $x_2(v_1, v_2')$ is constant in v_2' in this range and taking $v_2' \uparrow \xi_2(v_1)$, we get: $x_2(v_1, v_2') \geq x_2^*(v_1, \xi_2(v_1))$.

Now, by monotonicity, we have $x_2(v_1, v_2') = x_2^*(v_1, \xi_2(v_1))$ for all v_2' in the interval $(\xi_2(v_1) - \delta, \xi_2(v_1) + \delta)$. Therefore the budget of player 2 could not have been exhausted on $\xi_2(v_1)$. ∎

The third part of Figure 5.3 illustrates how the region E_2 typically looks like. If $u^1, u^2, \ldots, u^k$ are sorted such that u^1 corresponds to the edge that is higher to the left and u^k corresponds to the edge that is lower to the right, then we can divide $[0, a_1)$ in segments $[0, r_1), [r_1, r_2), \ldots, [r_k, a_1)$ such that for $v_1 \in [0, r_1)$, $\xi_2(v_1) = \infty$, and for $v_1 \in [r_i, r_{i+1})$, $(v_1, \xi_2(v_1))$ is in the line $\{t \cdot u^i; t \geq 0\}$. To see why this is true, consider $v_1 < v_1'$ and assume that $(v_1, \xi_2(v_1))$ is in the line $\{t \cdot u^i; t \geq 0\}$. Then $(v_1', \xi_2(v_1'))$ cannot be strictly above this line, by a similar argument used in Lemma 5.3.5: look at the allocation curves $x_2(v_1, v_2) \geq x_2(v_1', v_2)$ and the payment formula – then player 2 must be paying just above the $\{t \cdot u^i; t \geq 0\}$ for v_1' line at least as much as he was paying above this line for v_1 and hence his budgets must be exhausted.

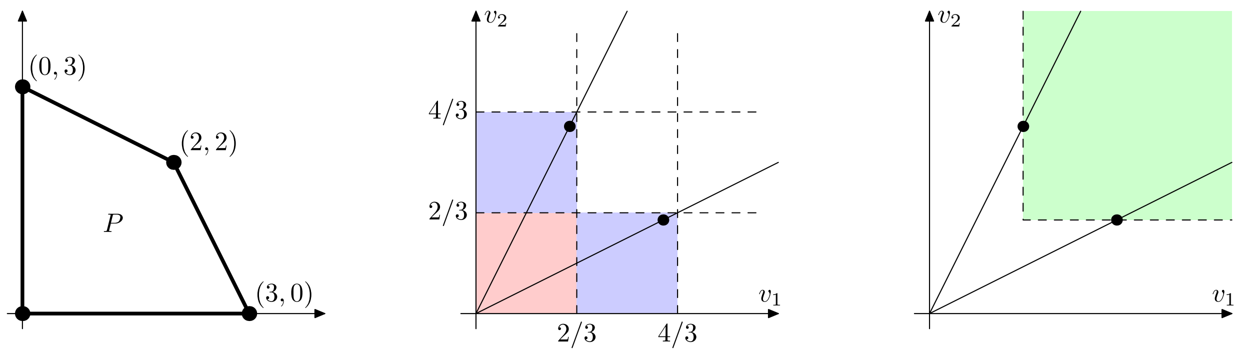

Figure 5.4: Illustration of the proof of the impossibility theorem: in the first part we represent the polytope P, in the second we represent the first main steps: we show that any auction must resemble VCG in the red region (Fact 5.3.7) and extend the definition of the auction to the blue region (Facts 5.3.8 and 5.3.9), showing that the budget of player 1 must get exhausted at the point $(\frac{1.2381}{2}, 1.2381)$ and the budget of player 2 must get exhausted at the point $(1.2381, \frac{1.2381}{2})$. We use this fact to show that the allocation must be constant in the green region (Fact 5.3.10), contradicting Pareto-optimality for allocation of the form (v_1, v_2) where $\frac{1.2381}{2} < v_1 < \frac{2}{3}$ as $v_2 \to \infty$.

5.3.3 Proof of the Impossibility Theorem

Now, we are ready to prove Theorem 5.3.3, which states that there is no general auction with all the desirable properties for all polyhedral environments P. We fix the following setting: a set of feasible allocations

$$P = \{x \in \mathbb{R}^2_+; 2x_1 + x_2 \leq 6, x_1 + 2x_2 \leq 6\}$$

and budgets $B_1 = B_2 = 1$. Assume that (x, φ) is a mechanism with the desirable properties for this setting. We will use the characterization lemmas in Section 5.3.2 to find a contradiction. We illustrate the flow of the proof in Figure 5.4 (a brief summary of the proof is given in the caption of the figure).

Fact 5.3.7 *In the region $[0, \frac{2}{3}) \times [0, \frac{2}{3})$, the mechanism (x, φ) produces an efficient allo-*

cation (i.e., it is equal to some mechanism of the VCG family).

Proof : First notice that no mechanism in the VCG-family for P exhausts the budget in $[0, \frac{2}{3}) \times [0, \frac{2}{3})$. Now, we turn our attention to the mechanism (x, φ), which is a mechanism with the desirable properties for this setting.

In $[0, \frac{1}{3}) \times [0, \frac{2}{3})$, player 1 cannot exhaust his budget, since $x_1 \leq 3$ and $v_1 < \frac{1}{3}$. We claim that in this area the mechanism needs to behave like VCG. If there is a point in this region where $x(v)$ doesn't maximize $v^t x$ for $x \in P$, then the budget of player 2 must be exhausted. So, there is some (v_1, v_2) such that for $v_2' > v_2$, $p_2(v_1, v_2') = B_2$ and for $v_2' < v_2$, $x(v_1, v_2')$ is in the VCG family. Now, notice that the allocation for $v_2' > v_2$ must be $x_2(v_2', v_1) \leq x_2^*(v_2', v_1)$. This contradicts the fact that the budget is exhausted for $v_2' \downarrow v_2$. So, this shows that (x, φ) must allocate efficiently on $[0, \frac{1}{3}) \times [0, \frac{2}{3})$. Now, for $[0, \frac{2}{3}) \times [0, \frac{1}{3})$, we can do the same argument. Now, what remains are the points in $[\frac{1}{3}, \frac{2}{3}) \times [\frac{1}{3}, \frac{2}{3})$. Let v be such a point. Notice that for $(\frac{1}{3} - \epsilon, v_2)$ and $(v_1, \frac{1}{3} - \epsilon)$, the allocation must be $(2, 2)$ because of the previous argument. By monotonicity, $x(v) = (2, 2)$ which is the efficient allocation. ■

Now, we know how any mechanism (x, φ) with the desirable properties should look like in the red region of Figure 5.4. Next, we try to understand how it should look in the blue region. In order to do so, we need some definitions. From Lemma 5.3.5 we know that there is $a \in \mathbb{R}^2_+ \cup \{(\infty, \infty)\})$ such that: $\{v; v > a\} \subseteq E_1 \cap E_2 \subseteq \{v; v \geq a\}$. We also define:

$$\tilde{v}_2 = \min\{v_2; p_1(\frac{v_2}{2}+, v_2) = 1\}$$

$$\tilde{v}_1 = \min\{v_1; p_2(v_1, \frac{v_1}{2}+) = 1\}$$

And we focus on the region R, which we define as the interior of the rectangle between $(0,0)$ and $(\tilde{v}_2/2, \tilde{v}_2)$.

Fact 5.3.8 *The budget of player* 1 *does not get exhausted in* R*. Also,* $\frac{2}{3} < \tilde{v}_2 \leq 1.2381$

Proof : First, assume $\tilde{v}_2 < 4/3$. Then by Fact 5.3.7, the budget of player 1 does not get exhausted in $[0, \frac{2}{3}) \times [0, \frac{2}{3})$ and by the definition of $\tilde{v}_2$ and lemma 5.3.6 it cannot be exhausted for $[0, \frac{\tilde{v}_2}{2}) \times (\frac{1}{3}, \tilde{v}_2)$. Notice that a cannot be in this region, because it would also contradict the definition of $\tilde{v}_2$.

Given this fact, let us analyze how the mechanism should be in this setting. We do so, by fixing v_1 and looking at $x_2(v_1, v_2)$. We can use the same argument as in the previous fact to argue that $x(v_1, v_2)$ must be the efficient allocation for $v_2 < 2v_1$ in R and also for all $v_1 \leq \frac{1}{3}$. For $v_1 > \frac{1}{3}$ and $v_2 > 2v_1$, if we allocate as in VCG, we exceed the budget. So, the allocation for those points must exactly match the budget of player 2. So, the only possible value must be such that:

$$2 \cdot \frac{1}{2} v_1 + (x_2(v) - 2) 2 v_1 = 1$$

and therefore the allocation must be

$$x(v) = \left(3 - \frac{1}{v_1}, 2 + \frac{1 - v_1}{2v_1}\right).$$

Now, this determines the payment of player 1 in the rectangle. We know that for $\delta \downarrow 0$, $\varphi_1(\frac{\tilde{v}_2}{2} - \delta, \tilde{v}_2 - 3\delta) < 1$, which we can write as:

$$\lim_{\delta \downarrow 0} \varphi_1(\frac{\tilde{v}_2}{2} - \delta, \tilde{v}_2 - 3\delta) = 2 \cdot \frac{\tilde{v}_2}{2} - \int_{1/3}^{\tilde{v}_2/2} 3 - \frac{1}{z} dz = 1 - \frac{\tilde{v}_2}{2} + \log\left(\frac{3}{2}\tilde{v}_2\right) \leq 1$$

This implies that $\tilde{v}_2 \leq 1.2381$. Notice that this excludes the fact that $\tilde{v}_2 > 4/3$, otherwise we could have done the same analysis on $[0, \frac{2}{3}] \times [0, \tilde{v}_2]$ and arrived in the same conclusion that $\tilde{v}_2 \leq 1.2381$. ∎

Fact 5.3.9 $\tilde{v}_2 = 1.2381$ *and* $x(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) = (2, 2)$.

Proof : Since $\frac{2}{3} < \tilde{v}_2 \leq 1.2381$, we know that $x_2(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) \geq 2$ by using that $x_2(\frac{\tilde{v}_2}{2}+, \frac{1}{3}) = 2$ (Fact 5.3.7) and monotonicity. Therefore, $x_1(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) \leq 2$. Writing the payment for player 1 at this point we get:

$$1 = \varphi_1(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) = x_1(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) \cdot \frac{\tilde{v}_2}{2} - \int_{1/3}^{\tilde{v}_2/2} 3 - \frac{1}{z} dz \leq 1 - \frac{\tilde{v}_2}{2} + \log\left(\frac{3}{2}\tilde{v}_2\right)$$

which implies that $\tilde{v}_2 = 1.2381$ and $x(\frac{\tilde{v}_2}{2}+, \tilde{v}_2) = (2, 2)$, since all inequalities must be tight. ■

Fact 5.3.10 $\tilde{v}_1 = \tilde{v}_2 = 1.2381$ *and* $x(v) = (2, 2)$ *for all valuation profiles* $v > (\frac{\tilde{v}_1}{2}, \frac{\tilde{v}_2}{2})$.

Proof : We can apply the same argument exchanging 1 and 2 and conclude that $\tilde{v}_1 = \tilde{v}_2$. Now, to see that for $v_1', v_2' > \frac{\tilde{v}_i}{2}$ we have $x(v_1', \tilde{v}_2) = x(\tilde{v}_1, v_2') = (2, 2)$, we analyze four regions. If $v \in (\tilde{v}_i/2, \tilde{v}_i] \times (\tilde{v}_i/2, \tilde{v}_i]$ or $v \in [\tilde{v}_i, \infty) \times [\tilde{v}_i, \infty)$ we can use the standard monotonicity argument to show that $x(v) = (2, 2)$.

For the regions $[\tilde{v}_1, \infty) \times (\frac{\tilde{v}_1}{2}, \tilde{v}_1)$ and $(\frac{\tilde{v}_2}{2}, \tilde{v}_2) \times [\tilde{v}_2, \infty)$ is a little trickier. We do the analysis for the first one. The second is analogous.

Clearly $\varphi_2(\tilde{v}_1, \tilde{v}_2) = 1$. Now, for $v_1' > \tilde{v}_1$, $x(v_1', \tilde{v}_2) = x(\tilde{v}_1, \tilde{v}_2) = (2, 2)$. And for all v_2 we have $x_1(v_1', v_2) \geq x_1(\tilde{v}_1, v_2)$ and therefore $x_2(v_1', v_2) \leq x_2(\tilde{v}_i, v_2)$. Now, since $\varphi_2(v_1, v_2) = v_2 x_2(v) - \int_0^{v_2} x_2(v_1, u) du$ clearly $\varphi_2(v_1', \tilde{v}_2) \geq \varphi_2(\tilde{v}_1, \tilde{v}_2) = 1$. Therefore $\varphi_2(v_1', \tilde{v}_2) = 1$, since the mechanism respects budgets. Notice that the only way it can be true is that if $x_2(v_1', v_2) = x_2(\tilde{v}_i, v_2)$, so we must have $x_2(v_1', v_2) = (2, 2)$ for $v_1' \in [\tilde{v}_i, \infty)$.

We can use the exact same argument for region: $(\frac{\tilde{v}_i}{2}, \tilde{v}_i) \times [\tilde{v}_i, \infty)$. ■

Now we are ready to prove Theorem 5.3.3:

Proof of Theorem 5.3.3 : Now, by putting Fact 5.3.10 and Fact 5.3.7 together, we get a contradiction with the Pareto-optimality: consider $\frac{\tilde{v}_i}{2} < v_1 < \frac{2}{3}$ then by facts 5.3.10 and 5.3.7 combined, we know that $x_2(v_1, v_2) = 0$ for $v_2 < \frac{v_1}{2}$ and $x_2(v_1, v_2) = 2$ for $v_2 > \frac{v_1}{2}$, so the budget of player 2 never gets exhausted even for $v_2 \to \infty$. This contradicts Pareto-optimality for $v_2 > 2v_1$, since if his budget is not exhausted, he should get allocated at least as much as he gets in VCG. ∎

5.3.4 Multi-unit auctions with decreasing marginals

As a by-product of Theorem 5.3.3, we can answer in a negative way the question of the existence of truthful Pareto-optimal auctions for multi-unit auctions with decreasing marginals. Consider the following setting:

Setting: Consider a supply of s of a certain divisible good and two players in such a way that the feasible allocations are (x_1, x_2) such that $x_1 + x_2 \leq s$. Player i has a public budget B_i and a private valuation which is a increasing concave function $V_i : [0, s] \to \mathbb{R}_+$. Upon getting x_i units of the good and paying φ_i, player i has utility $u_i = V_i(x_i) - \varphi_i$.

It is tempting to believe that one could adapt the clinching framework in Algorithm 1 to deal with this setting, by simply redefining the demand function as something like:

$$d_i = \min \left\{ \frac{B_i}{p}, \max\{x_i; \partial V_i(\rho_i + x_i) \leq p\} \right\}$$

where $\partial V_i(x_i)$ is the marginal valuation at x_i. Indeed, if $V_i(x_i) = v_i \cdot x_i$, this recovers the original way of calculating demands. The intuition behind why it doesn't work is that some player can increase his declared value on items he won't get anyway in order to increase the payment of his opponent, exhausting his budget earlier. This way, he is able to get items for cheaper in the end.

In the following theorem, we show that no auction mechanism can satisfy all the desirable properties for this setting:

Theorem 5.3.11 *There is no truthful, Pareto-optimal and budget-feasible auction for this setting.*

Proof : Suppose that $(\hat{x}(V_1, V_2), \hat{\varphi}(V_1, V_2))$ is a mechanism satisfying all the desirable properties for multi-unit auctions with decreasing marginals. Then we can use it as a black-box to construct a mechanism for a general polyhedral environment, contradicting Theorem 5.3.3. Given a certain polyhedral environment, we can describe

$$P = \{x \in [0, \alpha] \times [0, \beta]; x_2 \leq h(x_1)\}$$

where $h : [0, \alpha] \to [0, \beta]$ is a monotone non-increasing concave function, $h(0) = \beta$ and $h(\alpha) = 0$. Now, using $(\hat{x}, \hat{\varphi})$ for $s = 1$, build the following mechanism: if players report valuations v_1, v_2 build the following concave functions: $V_1(x_1) = v_1 \cdot \alpha x_1$ and $V_2(x_2) = v_2 \cdot h(\alpha - \alpha x_2)$.

Now, simply define

$$x(v_1, v_2) = (\alpha \hat{x}_1, h(\alpha \hat{x}_1))$$

$$\varphi(v_1, v_2) = \hat{\varphi}(V_1, V_2).$$

The mechanism is clearly truthful, individually rational and budget-feasible. It is also easy to see that sets of allocation can me mapped 1-1 between those two settings, preserving Pareto-optimality. ■

CHAPTER 6

AUCTIONS WITH ONLINE SUPPLY

We continue the agenda proposed in the last chapter of identifying important features of internet advertisement neglected in traditional models and designing mechanisms taking such features into account. In this chapter we focus on the fact that ad impressions arrive in an online manner and need to be allocated and priced in real time. This gives rise to an interesting question that is simultaneously a mechanism design question and an online algorithms question.

Pricing items without knowing the exact size of the inventory is a tricky problem: if the items turned out to be scarce and the competition for them intense, they should end up having high prices. On the other hand, if the items were abundant and not much competition for them, items should be priced low by any reasonable mechanism. In the online setting, what should we do when we get the first item ? At this point, it is unclear if this is the only item we have to sell or if it is the first of a million items.

Intuitively, the mechanism we will describe in this section will allocate the first item assuming the supply is small, and then update the allocation and payments once new items arrive. Players that acquired items at high prices in the beginning will be compensated by facing cheper prices on future items.

6.1 Online Supply Model

We consider auctions where the feasibility set is not known in advance to the auctioneer. For each time $t \in \{0, \dots, T\}$, we associate an environment $P_t \subseteq \mathbb{R}^n_+$, which keeps track of the allocations done in times $t' = 1..t$ [1]. In each step, the mechanism needs to output an allocation vector $x^t = (x_1^t, \ldots, x_n^t) \in P_t$ and a payment vector $\varphi^t = (\varphi_1^t, \ldots, \varphi_n^t) \geq 0$ by augmenting x^{t-1} and φ^{t-1}. Given a set of desirable properties, we would like to maintain them for all t. To make the problem tractable, have to restrict the set of possible histories $\{P_t\}_{t \geq 0}$. We do so by defining a partial ordering $\preccurlyeq$ on the set of feasibility constraints such that if $t \leq s$ then $P_t \preccurlyeq P_s$.

Our main goal is to design auctions where the auctioneer can allocate and charge payments *'on the fly'*. The auctioneer will face a set of environments $P_1 \preccurlyeq P_2 \preccurlyeq \ldots \preccurlyeq P_t$ and at time t, he needs to allocate $x^t \in P_t$ and charge φ^t, maintaining a set of desirable properties. He doesn't know if P_t will be the final outcome, or if some new environment $P_{t+1} \succcurlyeq P_t$ will arrive, in which case he will need to augment $x^t \in P_t$ to an allocation x^{t+1} in P_{t+1}. It is crucial that his decision at time t doesn't depend the knowledge about P_{t+1}.

Definition 6.1.1 (Online Supply Model) *Consider a family of feasibility allocation constraints indexed by $\mathcal{F}$, i.e, for each $f \in \mathcal{F}$ associate a set of feasible allocation vectors $P^f \subseteq \mathbb{R}_+^n$ (a set P^f is often called environment). Also, consider a partial order $\preccurlyeq$ defined over $\mathcal{F}$ such that if $f \preccurlyeq f'$ then $P^f \subseteq P^{f'}$. An auction for environment P^f consists of functions $x^f : \Theta = \times_i \Theta_i \to P^f$ and $\varphi^f : \Theta \to \mathbb{R}_+^n$.*

*An **auction in the strong online supply model** for $(\mathcal{F}, \preccurlyeq)$ is a family of auctions such that $x^f \leq x^{f'}$ and $\varphi^f \leq \varphi^{f'}$ whenever $f \preccurlyeq f'$. Moreover, we say that the auction satisfies a certain property if it satisfies this property for each f (e.g. the auction is incentive compatible if for each $f \in \mathcal{F}$, (x^f, φ^f) is an incentive compatible auction).*

[1] We want to stress the fact that P_t *doesn't* represent the set of allocations allowed *at* time t, but the set of allocations allowed *until* time t. The set of new possible allocations in time t is the difference between P_t and P_{t-1}

*An **auction in the weak online supply model** for $(\mathcal{F}, \preccurlyeq)$ is essentially the same, except that we drop the requirement of $\varphi^f \leq \varphi^{f'}$. The intuition is that we are required to allocate goods online, but are allowed to charge payments only in the end.*

The main idea behind the definition is that if at some point the auctioneer runs the auction (x^f, φ^f) for some environment P^f and at a later time some more goods arrive perhaps with new constraints such that the environment is augmented to $P^{f'}$ with $f' \succcurlyeq f$, then the auctioneer can run $(x^{f'}, \varphi^{f'})$ and augment the allocation of player i by $x_i^{f'}(v) - x_i^f(v)$ goods and charge him more $\varphi_i^{f'}(v) - \varphi_i^f(v)$.

Example 6.1.2 (Multi-unit auctions) *Let $\Delta_s = \{x \in \mathbb{R}_+^n; \sum_i x_i \leq s\}$ and define $\mathcal{F}^{MU} = \{\Delta_s; s \geq 0\}$ and let $\Delta_s \preccurlyeq^{MU} \Delta_t$ iff $s \leq t$. Let the value of player i for one unit of the good, v_i, lies in $\Theta_i = \mathbb{R}_+$. Now $u_i = v_i x_i - \varphi_i$. Thus, we are in a simple multi-unit auction setting. In this setting, VCG is incentive compatible, individually rational and efficient (in the sense that it has those three properties once run for each Δ_s) auction in the strong online model for $(\mathcal{F}^{MU}, \preccurlyeq^{MU})$.*

Example 6.1.3 (Multi-unit auctions with capacities) *Curiously, if players have capacity constraints, i.e., their utilities are $u_i = v_i \min\{x_i, C_i\} - \varphi_i$, then the VCG allocations for $(\mathcal{F}^{MU}, \preccurlyeq^{MU})$ are still monotone in the supply, but the payments are not. For example, consider two agents with values $v_1 = 1, v_2 = 2$ and capacities $C_1 = C_2 = 1$. With supply 1, one item is allocated to player 2 and he is charged 1. With supply 2, both players get one unit of the item, but the VCG prices are zero. Therefore, there is no incentive compatible, individually rational and efficient in the strong online model. Babaioff, Blumrosen and Roth [11] strengthen this result showing that no $\Omega(\log\log n)$-approximately efficient auction exists in the strong online model.*

Example 6.1.4 (Polymatroidal auctions) *Now, let $\mathcal{F}^{PM}$ be the set of all polymatroidal domains and consider the naive-partial-order $\preccurlyeq^N$ to be such that $f \preccurlyeq f'$ iff $P^f \subseteq P^{f'}$. The VCG is not even online in the weak sense for $(\mathcal{F}^{PM}, \preccurlyeq^N)$. Consider the following example:*

$$P^f = \{x \in \mathbb{R}^2_+; x_1 \leq 2, x_2 \leq 2, x_1 + x_2 \leq 3\} \quad \textit{and} \quad P^{f'} = \{x \in \mathbb{R}^2_+; x_1 + x_2 \leq 4\}$$

then clearly $f \preccurlyeq f'$ but if $v_1 > v_2$. $x^f(v) = (2,1)$ but $x^{f'} = (4,0)$ violating the monotonicity property. But now, let's define a different partial order $\preccurlyeq^{PM}$ such that $f \preccurlyeq^{PM} f'$ if there is a polymatroid P' such that $P^{f'} = P^f + P'$ where the sum is the Minkowski sum. In the following lemma, we show that VCG is an auction in the strong online model for $(\mathcal{F}^{PM}, \preccurlyeq^{PM})$.

Lemma 6.1.5 *VCG is an auction in the strong online model for $(\mathcal{F}^{PM}, \preccurlyeq^{PM})$.*

Proof : Assume that $P^f, P^{f'}, P'$ are defined respectively by the monotone submodular functions f, f', g. If $P^{f'} = P^f + P'$, then by McDiarmid's Theorem [61], $f' = f + g$.

Now, let's remind how VCG allocated for this setting. If the polymatroid is defined by f, VCG begins by sorting the players by their value (and breaking ties lexicographically). So, we can assume $v_1 \geq \ldots \geq v_n$. Then it chooses the outcome:

$$x_i = f([i]) - f([i-1])$$

$$\varphi_i = v_{i+1} \cdot (f([i+1] \setminus i) - f([i-1]) - x_{i+1}) + \sum_{j > i+1} v_j \cdot (f([j] \setminus i) - f([j-1] \setminus i) - x_j)$$

where $[i]$ is an abbreviation for $\{1, \ldots, i\}$. Now, once we do this for f, f' we notice that the allocation and payments for f' are simply the sum of the allocation and payments for f and g, hence they are monotone along $\preccurlyeq^{\text{PM}}$. ■

One interesting property of incentive-compatible auctions in the online supply model is that utilities are monotone with the supply. If bidders have the option of leaving in each timestep collecting their current allocations for their current payment, they still (weakly) prefer to stay until the end of the auction.

Lemma 6.1.6 (Utility monotonicity) *Consider a setting where agents have single-parameter valuations $\Theta_i = \mathbb{R}_+$ and quasilinear utilities $u_i = v_i x_i - \varphi_i$. Given a truthful auction in the weak online supply model and $f \preccurlyeq f'$, then the utility of agent i increased with the supply, i.e.: $u^{f'} = v_i x^{f'} - \varphi_i^{f'} \geq v_i x^f - \varphi_i^f = u^f$*

Proof : The proof follows directly from Myerson's characterization [64] of payments in quasi-linear settings: $u^{f'} = v_i x_i^{f'} - \varphi_i^{f'} = \int_0^{v_i} x_i^{f'}(u)du \geq \int_0^{v_i} x^f(u)du = v_i x^f - \varphi_i^f = u^f$. ■

6.2 Clinching Auctions and Supply Monotonicity

Our main theorem states that the Adaptive Clinching Auction (defined in Dobzinski, Lavi and Nisan [32] and Bhattacharya el al [13]) is an incentive-compatible auction in the strong online supply model for budget constrained agents in the multi-unit setting. Formally:

Theorem 6.2.1 *Given n agents with public budgets B_i and single-dimensional types $v_i \in \mathbb{R}_+$ such that their utility is given by $u_i = v_i x_i - \varphi_i$ if $\varphi_i \leq B_i$ and $u_i = -\infty$ otherwise, the Adaptive Clinching Auction is an auction in the strong online supply model for $(\mathcal{F}^{MU}, \preccurlyeq^{MU})$. In other words, if $x(v, B, s)$ and $\varphi(v, B, s)$ is the outcome* *of the auction for valuation profile v, budgets B and supply s, then if $s \leq s'$, then: $x(v, B, s) \leq x(v, B, s')$ and $\varphi(v, B, s) \leq \varphi(v, B, s')$.*

Notice that this is in sharp contrast with what happens in Example 6.1.3 where getting a Pareto optimal auction in the strong online model is not possible, not even in an approximate way [2]. This is somewhat surprising, since capacity constraints on the allocations are usually more nicely-behaved compared to budget constraints.

Before proving the result, we review the Adaptive Clinching Auction [13, 32], presenting it in a way which will be more convenient for the proof.

6.2.1 Adaptive Clinching Auction

The Adaptive Clinching Auction is essentially the version of the Clinching Auction described in Section 5.2 for the Uniform Matroid (i.e. $f(S) = 1, \forall S$) in the limit as the price increase ϵ goes to zero. In what follows, we review its main properties and present the infinitesimal version.

The clinching auction takes as input the valuation profile v, the budget profile B and the initial supply s, then it runs a procedure based on the ascending price framework to determine final allocation and payments. There is a price clock p, and for each price, the auction mantains[3] $x_i(p)$ denoting the current allocation of player i and $B_i(p)$, which is the current remaining budget of player i. Initially, $x_i(0) = 0$ and $B_i(0) = B_i$, their initial budget. For each p, the auction

[2]for that setting, since there are no budgets, Pareto optimality boils down to efficiency.

[3]note that here we prefer to index the ascending process by the price itself rather then an external variable, like in Bhattacharya el al [13].

defines the values of the right-derivatives $\partial_p x_i(p)$ and $\partial_p B_i(p)$ and described its behavior in the points in which it is discontinuous. Notice we will use $\partial_p f(p)$ to denote the right-derivative of f at p.

For simplicity, we define the auction and prove our results for valuation profiles v such that $v_i \neq v_j$ for each $i \neq j$ (we call it a profile in generic form) This is mainly a technical assumption to avoid over-complicating the statement and the proof. For the definition and its subsequent discussion, we will use the following implicitly defined notation:

- remnant supply: $S(p) = s - \sum_i x_i(p)$
- active players: $A(p) = \{i; v_i > p\}$
- clinching players: $C(p) = \{i \in A(p); S(p) = \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p}\}$
- maximum remaining budget:$B_*(p) = \max_{i \in A(p)} B_i(p)$
- for any function f, let $f(\bar{p}-) = \lim_{p \uparrow \bar{p}} f(p)$ and $f(\bar{p}+) = \lim_{p \downarrow \bar{p}} f(p)$

Definition 6.2.2 (Adaptive Clinching Auction) *Given as input a valuation vector v in generic form, a budget vector B and initial supply s, consider the functions $x_i(p), B_i(p)$ such that:*

(i) $x_i(0) = 0$ *and* $B_i(0) = B_i$.

(ii) $\partial_p x_i(p) = \frac{S(p)}{p}$ *and* $\partial_p B_i(p) = -S(p)$ *if* $i \in C(p)$ *and* $\partial_p x_i(p) = \partial_p B_i(p) = 0$ *otherwise.*

(iii) the functions x_i and B_i are right-continuous at all points p, i.e., $x_i(p) = x_i(p+)$ and $B_i(p) = B_i(p+)$ for all p and it is left-continuous at all points $p \notin \{v_1, \ldots, v_n\}$, i.e., $x_i(p-) = x_i(p)$ and $B_i(p-) = B_i(p)$ for all $p \notin \{v_1, \ldots, v_n\}$

(iv) for $p = v_i$, let $\delta_j = \left[S(v_i-) - \sum_{k \in A(v_i) \setminus j} \frac{B_k(v_i-)}{v_i}\right]^+$. For $j \in A(v_i)$, let $x_j(v_i) = x_j(v_i-) + \delta_j$ and $B_j(v_i) = B_j(v_i-) - v_i \delta_j$ and for $j \notin A(v_i)$, $x_j(v_i) = x_j(v_i-)$ and $B_j(v_i) = B_j(v_i-)$.

The existence and uniqueness of those functions follow from elementary real analysis. The outcome associated with v, B, s is $x_i = \lim_{p\to\infty} x_i(p)$ and $\varphi_i = B_i(0) - \lim_{p\to\infty} B_i(p)$. Notice that this is well defined since x and B are constant for $p > \max_i v_i$.

The verb *clinch* means acquiring goods that are underdemanded at the current price. So clinching a δ_i amount at price p means receiving δ_i amount of the good and paying $\delta_i p$ for it. When we refer to a player clinching some amount, either we refer to the infinitesimal clinching happening in (ii) or the player clinching positive units in (iv).

Theorem 6.2.3 (Bhattacharya et al [13]) *The Adaptive Clinching Auction in Definition 6.2.2 is incentive-compatible, individually-rational, budget-feasible and produces Pareto-optimal outcomes.*

As one can possibly guess, it is possible to solve the differential equation in each interval between two adjacent values of v_i and give an explicit description of the clinching auction. We do so in Section 6.3. Nevertheless, we mostly prove our results using the differential form in Definition 6.2.2 which is more insightful than the explicit version.

Example 6.2.4 *At this point, it is instructive to consider an example of the auction. Consider an auction between $n = 4$ players with valuations $v = [9, 10, 11, 5.7]$ and*

$B = [3, 2, 1, .5]$. *The functions* $x_i(p), B_i(p)$ *are depicted in Figure 6.1. For* $p < p_0^1 = 3.5$, *the clinching set* $C(p)$ *is empty. At this price* $S(p_0^1) = 1 = \frac{2+1+0.5}{3.5} = \sum_{j \neq 1} \frac{B_j}{p}$, *so player* 1 *alone begins clinching.*

Since he is clinching alone for a while, $x_1(p) = s - S(p)$. *Now by derivating this expression we get that* $\frac{S(p)}{p} = \partial_p x_1(p) = -\partial_p S(p)$. *Solving for the supply with the condition that* $S(p_0^1) = 1$, *we get:* $S(p) = \frac{sp_0^1}{p}$ *and* $x_1(p) = s - \frac{sp_0^1}{p}$. *This continues while no other player enters the clinching set. The supply function* $S(p)$ *is illustrated in the first part of Figure 6.2.*

Notice that for this period, the budget of 1 *is being spent while the budgets of the other agents are intact. Eventually, the budget of* 1 *meets the budget of* 2, *and at this point, those two players are indistinguishable from the perspective of the differential procedure as long as both are active. Therefore, both start spending their budget at the same rate and acquiring goods at the same rate* $\partial_p x_1(p) = \partial_p x_2(p) = S(p)$. *Since from this point on* $S(p) = s - x_1(p) - x_2(p)$, *then* $\frac{S(p)}{p} = \partial_p x_1(p) = -\frac{1}{2}\partial_p S(p)$. *Solving again for the supply and the boundary condition* $S(p_0^2) = \frac{sp_0^1}{p_0^2}$ *we get:* $S(p) = \frac{sp_0^1 p_0^2}{p^2}$. *Using this, one can calculate* $x_1(p)$ *and* $x_2(p)$. *Both continue clinching at the same rate until the price reaches* $p = v_4$, *where player* 4 *exits the active set prompting the agents to clinch a positive amount according to (iv).*

Their allocation $x_i(p)$ *and budgets* $B_i(p)$ *are discontinuous at this point, but continue to follow the differential procedure after this point, having their budgets all equal (not coincidentally, as we will see in Lemma 6.2.11), until price* $p = v_1$ *is reached and player* 1 *exits the active set. At this point, all the remaining active players clinch a positive amount according to (iv) that exhausts the supply. Therefore, allocations and budgets are constant from this point on.*

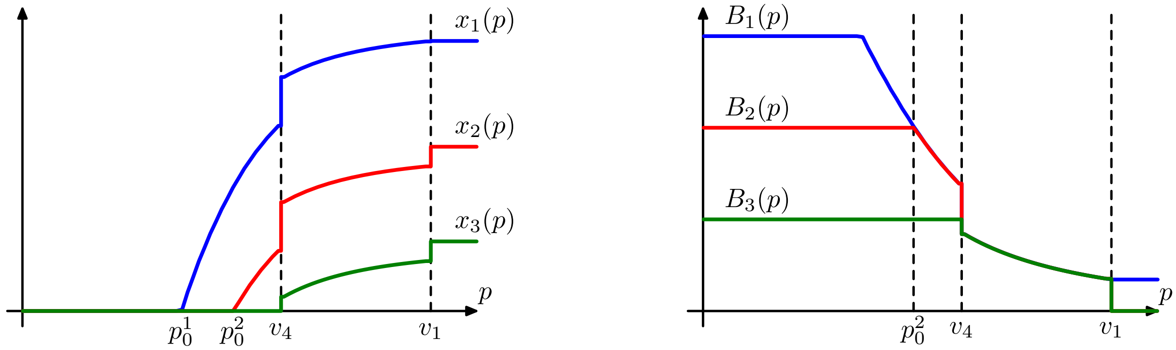

Figure 6.1: The functions $x(p)$ and $B(p)$ for an auction in Example 6.2.4

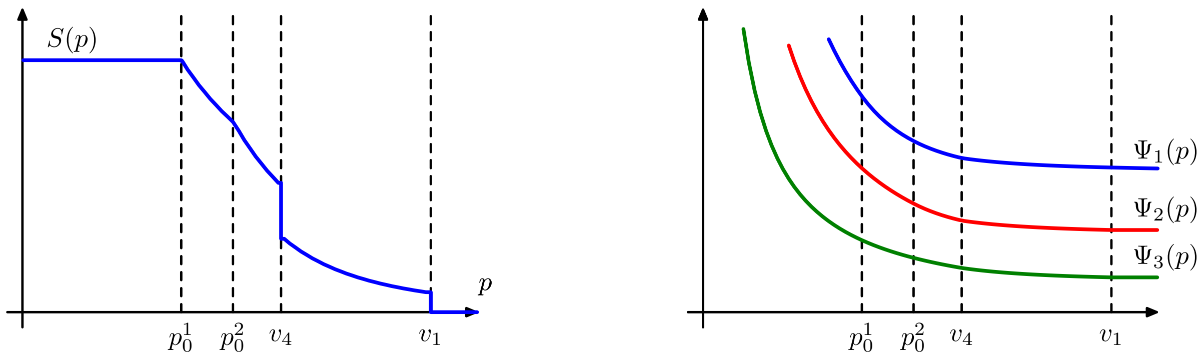

Figure 6.2: Supply $S(p)$ and wishful allocation $\Psi(p)$ for an auction in Example 6.2.4

One important tool in analyzing this auction is the concept of the *wishful allocation*. We define a $\Psi_i(p)$ as a function of $x_i(p)$ and $B_i(p)$ which is continuous even at the points where $x_i(p)$ and $B_i(p)$ are not. It is carefully set up so that the discontinuities from both functions cancel out. Intuitively, it represents a sum of what the player acquired already at the current price $x_i(p)$ with the maximum amount he would like to acquire at this price, which is $\frac{B_i(p)}{p}$.

Definition 6.2.5 (Wishful allocation) *The wishful allocation is defined as* $\Psi_i(p) = x_i(p) + \frac{B_i(p)}{p}$.

Lemma 6.2.6 *The wishful allocation is continuous and right-differentiable for all $p \geq 0$. Moreover, its right-derivative is given by: $\partial_p \Psi_i(p) = -\frac{B_i(p)}{p^2}$.*

Proof : The function $\Psi_i(p)$ is clearly continuous for $p \notin \{v_1, \ldots, v_n\}$ and right-continuous everywhere. Now, we claim that it is also left-continuous at v_j, i.e., $\Psi(v_j-) = \Psi(v_j)$. This fact is almost immediate:

$$\Psi_i(v_j) = x_i(v_j) + \frac{B_i(v_j)}{v_j} = [x_i(v_j-) + \delta_i] + \frac{[B_i(v_j-) - \delta_i v_j]}{v_j} = x_i(v_j-) + \frac{B_i(v_j-)}{v_j} = \Psi_i(v_j-)$$

Calculating its derivative is also easy:

$$\partial_p \Psi_i(p) = \partial_p \left[x_i(p) + \frac{B_i(p)}{p} \right] = \partial_p x_i(p) + \frac{\partial_p B_i(p)}{p} - \frac{B_i(p)}{p^2} = -\frac{B_i(p)}{p^2}$$

since $\partial_p x_i(p) = \frac{S(p)}{p} = -\frac{\partial_p B_i(p)}{p}$. ∎

Since $\Psi_i(p) \geq x_i(p)$ and is a monontone non-increasing function converging to the final allocation as $p \to \infty$, it constantly gives us an upper bound of the final allocation.

Now, we study some other properties of the above auction, which will be useful in the proof of our main theorem. First we prove a Meta Lemma that sets the basic structure for most of our proofs. The lemma is based on elementary facts of real analysis.

Meta-Lemma 6.2.7 *Given a property Λ that depends on p, if we want to prove for all $p \geq p_0$, it is enough to prove the following facts:*

(a) it holds for $p = p_0$.

(b) if Λ holds for p, then there is some $\epsilon_p > 0$ such that Λ holds for $[p, p + \epsilon_p)$

(c) if Λ holds for all p' such that $p_0 \leq p' < p$, then Λ also holds for p.

Proof : Let $F = \{p \geq p_0; \Lambda \text{ doesn't hold for } p\}$. We want to show that if the properties (a),(b),(c) in the statement hold, then $F = \emptyset$. Assume for contradiction that (a),(b),(c) hold but $F \neq \emptyset$. Let $\bar{p} = \inf F$, i.e., the smallest $\bar{p}$ such that for all $\epsilon > 0$, $[\bar{p}, \bar{p} + \delta) \cap F \neq \emptyset$ for all $\delta > 0$.

Now, there are two possibilities:

(1) either $\bar{p} \notin F$, in this case we can invoke (b) to see that there should be an $\epsilon > 0$ such that $[\bar{p}, \bar{p} + \epsilon) \cap F = \emptyset$ which contradicts the fact that $\bar{p} = \inf F$.

(2) or $\bar{p} \in F$. By (a), we know $\bar{p} > p_0$. Then we can use that by the definition of $\inf$, Λ holds for all $p < \bar{p}$, so we can invoke (c) to show that Λ should hold for $\bar{p}$. And again we arrive in a contradiction. ■

For most properties Λ that we want to prove about the Adaptive Clinching Auction, part (a) is easy to show, part (b) requires using the right-continuity of the function and the value of the right-derivatives given in item (ii) of Definition 6.2.2 and part (c) is usually proved using continuity for $p \notin \{v_1, \ldots, v_n\}$ and using part (iv) of Definition 6.2.2.

The first two lemmas (whose proof is based on the Meta-Lemma) state that once a player start acquiring goods (i.e. $\partial_p x_i(p) > 0$), he continues to do so for all the prices until p becomes equal to his value v_i.

Lemma 6.2.8 *Once a player i enters the clinching set, then he is in the clinching set until he becomes inactive, i.e., if $i \in C(p)$ for some p, then $i \in C(p')$ for all $p' \in [p, v_i)$.*

Proof : The proof is based on the Meta Lemma. Part (a) is trivial.

For part (b), there is $\epsilon > 0$ such that in $[p, p+\epsilon)$ the active set is the same as $A(p)$. We will show that if $i \in C(p)$, then $i \in C(p')$ for all $p' \in [p, p+\epsilon)$, or in other words: $S(p') = \sum_{j \in A(p) \setminus i} \frac{B_j}{p'}$. This equality holds for p. Now, we will simply show that the derivative of both sides is the same in the $[p, p+\epsilon)$ interval, i.e.: $\partial_p S(p) = \partial_p \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p}$.

$$
\begin{aligned}
\partial_p \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p} &= \frac{p[\sum_{j \in A(p) \setminus i} \partial_p B_j(p)] - \sum_{j \in A(p) \setminus i} B_j(p)}{p^2} = \\
&= \frac{1}{p} \sum_{j \in C(p) \setminus i} \partial_p B_j(p) - \frac{1}{p} S(p) = - \sum_{j \in C(p) \setminus i} \frac{1}{p} S(p) - \frac{1}{p} S(p) \\
&= - \sum_{j \in C(p)} \partial_p x_j(p) = - \sum_{j \in A(p)} \partial_p x_j(p) = \partial_p S(p)
\end{aligned}
$$

For part (c), it is trivial for $p \notin \{v_1, \ldots, v_n\}$ by left-continuity: if $S(p') = \sum_{j \in A(p') \setminus i} \frac{B_j}{p'}$ for $p' < p$ and the functions involved are left-continuous, then it holds for p. Now, for $p = v_j$, if $S(v_j-) = \sum_{k \in A(v_j-) \setminus i} \frac{B_k(v_j-)}{v_j}$, then for δ_k as defined in (iii) of Definition 6.2.2 we have:

$$
\begin{aligned}
S(v_j) &= S(v_j-) - \sum_{k \in A(v_j)} \delta_k = \\
&= \left[\sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j-) - \delta_k v_j}{v_j} \right] + \frac{B_j(v_j-)}{v_j} - \delta_i = \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j)}{v_j}
\end{aligned}
$$

since:

$$
\begin{aligned}
\delta_i &= \left[S(v_j-) - \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j-)}{v_j} \right]^+ = \\
&= \left[\sum_{k \in A(v_j-) \setminus i} \frac{B_k(v_j-)}{v_j} - \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j-)}{v_j} \right]^+ = \frac{B_j(v_j-)}{v_j}
\end{aligned}
$$

■

Lemma 6.2.9 *For each price p and each active player i, $S(p) \leq \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p}$.*

Proof : Again we prove it using the Meta Lemma. (a) is trivial, for (b) if $S(p) < \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p}$, then by right-continuity the strict inequality continues to hold in some region $[p, p + \epsilon)$. If $S(p) = \sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p}$ we can do the same analysis as in Lemma 6.2.8. For (c) it is again trivial for $p \notin \{v_1, \dots, v_n\}$ by left-continuity and for $p = v_j$ we use the fact that comes directly from the proof of the previous lemma:

$$S(v_j) - \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j)}{v_j} = \left[S(v_j-) - \sum_{k \in A(v_j-) \setminus i} \frac{B_k(v_j-)}{v_j} \right] + \frac{B_j(v_j-)}{v_j} - \delta_i \leq 0 \quad (6.1)$$

by the definition of δ_i, since:

$$\delta_i = \left[S(v_j-) - \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j-)}{v_j} \right]^+ \geq \left[S(v_j-) - \sum_{k \in A(v_j-) \setminus i} \frac{B_k(v_j-)}{v_j} \right] + \frac{B_j(v_j-)}{v_j}$$

■

Corollary 6.2.10 *If at price $p = v_j$, player $i \in A(v_j)$ acquires any positive amount of the good $\delta_i > 0$, then he enters in the clinching set (if he wasn't previously), i.e., $i \in C(v_j)$.*

Proof : If $\delta_i > 0$, then, $\delta_i = S(v_j-) - \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j-)}{v_j}$. Substituting that in equation (6.1) we get that $S(v_j) = \sum_{k \in A(v_j) \setminus i} \frac{B_k(v_j)}{v_j}$ and therefore $i \in C(v_j)$. ■

A crucial observation for our proof is that the evolution of the profile of remaining budgets follows a very structured format. At any given price, the remaining budget of an agent is either his original budget or the maximum budget among all agents. It is instructive to observe that in Figure 6.1.

Lemma 6.2.11 *For each price p, if $C(p) \neq \emptyset$, then $C(p) = \{i \in A(p); B_i(p) = B_*(p)\}$.*

Proof : It is easy to see that all bidders in the clinching set have the same remaining budget, since if $i, i' \in C(p)$, then $\sum_{j \in A(p) \setminus i} \frac{B_j(p)}{p} = S(p) = \sum_{j \in A(p) \setminus i'} \frac{B_j(p)}{p}$ and therefore $B_i(p) = B_{i'}(p)$. Also, clearly, all players with the same budget will be in the clinching set. The fact that the players clinching have the largest budget follows directly from Lemma 6.2.9. ∎

Corollary 6.2.12 *For each $i \in A(p)$, $B_i(p) = \min\{B_i(0), B_*(p)\}$.*

6.2.2 Supply Monotonicity

Now we are ready to prove Theorem 6.2.1 which is our main result. For that we fix a budget profile B and a valuation profile v in generic form (i.e. $v_i \neq v_j$ for $i \neq j$, which is not needed for the proof and is mainly intended to simplify the exposition). Now, we consider two executions of the adaptive clinching auction. One with initial supply s^{b} which we call the *base auction* and one with initial supply $s^{\mathsf{a}} \geq s^{\mathsf{b}}$ which we call the *augmented auction*. Running the base and augmented auction with the same valuations and budgets we get functions $x^{\mathsf{b}}(p), B^{\mathsf{b}}(p)$ and $x^{\mathsf{a}}(p), B^{\mathsf{a}}(p)$. From this point on, we use superscripts b and a to refer to the base and augmented auctions respectively. For the set of active players at a given price, we omit the superscript, since $A^{\mathsf{b}}(p) = A^{\mathsf{a}}(p)$ for all p.

As the first step toward the proof of Theorem 6.2.1, we prove that the payments are monotone with the supply, that the final payment of each agent in the augmented auction is higher than in the base auction:

Proposition 6.2.13 (Payment Monotonicity) *Given the base and augmented auction as defined above, then for all $p \geq 0$ and all agents i, $B_i^{\mathsf{b}}(p) \geq B_i^{\mathsf{a}}(p)$.*

Proof : The first part of the proof consists of showing that clinching starts first in the augmented auction. Then we divide the prices in three intervals: in the first where no clinching happens in both auctions, in the second where clinching happens only in the augmented auction and the third in which clinching happens in both auctions. Then we prove the claim in each of the intervals.

First part of the proof: Clinching starts earlier in the augmented auction

Let $p_0^{\mathsf{b}} = \min\{p; C^{\mathsf{b}}(p) \neq \emptyset\}$ and $p_0^{\mathsf{a}} = \min\{p; C^{\mathsf{a}}(p) \neq \emptyset\}$. We claim that $p_0^{\mathsf{a}} \leq p_0^{\mathsf{b}}$. In order to see that, assume the contrary: $p_0^{\mathsf{b}} < p_0^{\mathsf{a}}$. At p_0^{b}, there is one agent i such that $S^{\mathsf{b}}(p_0^{\mathsf{b}}) = \sum_{k \in A(p_0^{\mathsf{b}}) \setminus i} \frac{B_k^{\mathsf{b}}(p_0^{\mathsf{b}})}{p_0^{\mathsf{b}}}$. If $p_0^{\mathsf{b}} \notin \{v_1, \ldots, v_n\}$, then by Corollary 6.2.10, no budget was spent in neither of the auctions at this price and no goods were acquired, so $S^{\mathsf{b}}(p_0^{\mathsf{b}}) = s^{\mathsf{b}}$, $S^{\mathsf{a}}(p_0^{\mathsf{b}}) = s^{\mathsf{a}}$, $B_k^{\mathsf{b}}(p_0^{\mathsf{b}}) = B_k^{\mathsf{b}}(0)$ and $B_k^{\mathsf{a}}(p_0^{\mathsf{b}}) = B_k^{\mathsf{a}}(0)$. This implies that at this point $S^{\mathsf{a}}(p_0^{\mathsf{b}}) > S^{\mathsf{b}}(p_0^{\mathsf{b}}) = \sum_{k \in A(p_0^{\mathsf{b}}) \setminus i} \frac{B_k^{\mathsf{a}}(p_0^{\mathsf{b}})}{p_0^{\mathsf{b}}}$, which contradicts Lemma 6.2.9 for the augmented auction. Now, the case left to analyze is the one where $p_0^{\mathsf{b}} = v_j$ for some $j \neq i$ and i entered the clinching set after acquiring a positive amount of good $\delta_i^{\mathsf{b}} > 0$ at price v_j. Then: $\delta_i^{\mathsf{b}} = s^{\mathsf{b}} - \sum_{k \in A(v_j) \setminus i} \frac{B_k(0)}{v_j} > 0$. But in this case $\delta_i^{\mathsf{a}} > 0$, contradicting that $p_0^{\mathsf{a}} > p_0^{\mathsf{b}}$.

Second part of the proof: Proof for the first interval $[0, p_0^{\mathsf{a}})$.

For the p in the interval $[0, p_0^{\mathsf{a}})$, no clinching occurs, so $B_i^{\mathsf{a}}(p) = B_i^{\mathsf{a}}(0) = B_i^{\mathsf{b}}(0) = B_i^{\mathsf{b}}(p)$.

Third part of the proof: Proof for the second interval $[p_0^{\mathsf{a}}, p_0^{\mathsf{b}})$.

In the interval $[p_0^{\mathsf{a}}, p_0^{\mathsf{b}})$, some players are acquiring goods in the augmented auction but no player is neither acquiring goods nor paying anything in the base auction, so: $B_i^{\mathsf{a}}(p) \leq B_i^{\mathsf{a}}(0) = B_i^{\mathsf{b}}(0) = B_i^{\mathsf{b}}(p)$.

Fourth part of the proof: Proof for the third interval $[p_0^{\mathsf{b}}, \infty)$.

In this interval, both players are clinching. Now, we use the Meta Lemma to show that for all $p \geq p_0^{\mathsf{b}}$, the property $B_i^{\mathsf{b}}(p) \geq B_i^{\mathsf{a}}(p)$ for all i holds.

For part (a) of the Meta Lemma, we need to show that $B_i^{\mathsf{b}}(p_0^{\mathsf{b}}) \geq B_i^{\mathsf{a}}(p_0^{\mathsf{b}})$. If $p_0^{\mathsf{b}} \notin \{v_1, \ldots, v_n\}$ this follows directly from continuity and the third part of the proof. If $p_0^{\mathsf{b}} = v_j$ for some j, then by the previous cases we know that $B_i^{\mathsf{a}}(v_j-) \leq B_i^{\mathsf{b}}(v_j-)$. We have that $B_i^{\mathsf{a}}(v_j) = B_i^{\mathsf{a}}(v_j-) - \delta_i^{\mathsf{a}} v_j$ and $B_i^{\mathsf{b}}(v_j) = B_i^{\mathsf{b}}(v_j-) - \delta_i^{\mathsf{b}} v_j$. Now we analyze the clinched amounts δ_i^{a} and δ_i^{b}. If $p_0^{\mathsf{a}} = p_0^{\mathsf{b}}$, it is straightforward to see that $\delta_i^{\mathsf{a}} \geq \delta_i^{\mathsf{b}}$ and therefore $B_i^{\mathsf{a}}(v_j) \leq B_i^{\mathsf{b}}(v_j)$. So, let's focus on the case where $p_0^{\mathsf{a}} < p_0^{\mathsf{b}}$. For this case:

$$
\begin{aligned}
\delta_i^{\mathsf{a}} &= \left[S^{\mathsf{a}}(v_j-) - \sum_{k \in A(v_j) \setminus i} \frac{B_k^{\mathsf{a}}(v_j-)}{v_j} \right]^+ = \\
&= \left[\sum_{k \in A(v_j-)} \frac{B_k^{\mathsf{a}}(v_j-)}{v_j} - \frac{B_*^{\mathsf{a}}(v_j-)}{v_j} - \sum_{k \in A(v_j) \setminus i} \frac{B_k^{\mathsf{a}}(v_j-)}{v_j} \right]^+ = \\
&= \left[\frac{B_i^{\mathsf{a}}(v_j-)}{v_j} + \frac{B_j^{\mathsf{a}}(v_j-)}{v_j} - \frac{B_*^{\mathsf{a}}(v_j-)}{v_j} \right]^+ = \\
&= \frac{1}{v_j} \left[\min\{B_i(0), B_*^{\mathsf{a}}(v_j-)\} + \min\{B_j(0), B_*^{\mathsf{a}}(v_j-)\} - B_*^{\mathsf{a}}(v_j-) \right]^+
\end{aligned}
$$

where the last step is an invocation of Corollary 6.2.12. For the base auction we have essentially the same, except that $S^{\mathsf{b}}(v_j-) \leq \sum_{k\in A(v_j-)} \frac{B^{\mathsf{b}}_k(v_j-)}{v_j} - \frac{B^{\mathsf{b}}_*(v_j-)}{v_j}$ holds as an inequality rather than equality, so we get:

$$\delta^{\mathsf{b}}_i \leq \left[\frac{B^{\mathsf{b}}_i(v_j-)}{v_j} + \frac{B^{\mathsf{b}}_j(v_j-)}{v_j} - \frac{B^{\mathsf{b}}_*(v_j-)}{v_j}\right]^+ =$$
$$= \frac{1}{v_j}\left[\min\{B_i(0), B^{\mathsf{b}}_*(v_j-)\} + \min\{B_j(0), B^{\mathsf{b}}_*(v_j-)\} - B^{\mathsf{b}}_*(v_j-)\right]^+$$

In order to prove that $B^{\mathsf{a}}_i(v_j) \leq B^{\mathsf{b}}_i(v_j)$, we study two cases:

- Case A: $B^{\mathsf{a}}_*(v_j-) \leq B^{\mathsf{a}}_j(0)$, i.e. $B^{\mathsf{a}}_j(v_j-) = B^{\mathsf{a}}_*(v_j-)$. In this case, $\delta^{\mathsf{a}}_i = \frac{B^{\mathsf{a}}_i(v_j-)}{v_j}$ and therefore $B^{\mathsf{a}}_i(v_j) = 0$, so, it is trivial that $B^{\mathsf{a}}_i(v_j) = 0 \leq B^{\mathsf{b}}_i(v_j)$.
- Case B: $B^{\mathsf{a}}_*(v_j-) > B^{\mathsf{a}}_j(0)$. Now, consider the function $\Phi(\beta) = [\min\{\beta, \mu\} + \min\{\beta, \gamma\} - \beta]^+$ for $\beta \geq \min\{\mu, \gamma\}$. This function is monotone non-increasing in such range. Now, take $\mu = B_i(0), \gamma = B_j(0)$ and use that $B^{\mathsf{a}}_*(v_j-) \leq B^{\mathsf{b}}_*(v_j-)$ to conclude that $\delta^{\mathsf{a}}_i = \Phi(B^{\mathsf{a}}_*(v_j-)) \geq \Phi(B^{\mathsf{b}}_*(v_j-)) \geq \delta^{\mathsf{b}}_i$. This implies $B^{\mathsf{a}}_i(v_j) \leq B^{\mathsf{b}}_i(v_j)$.

This finishes the proof of part (a) of the Meta Lemma.

Now, for part (b) of the Meta-Lemma, consider two cases:

- $B^{\mathsf{a}}_*(p) < B^{\mathsf{b}}_*(p)$, then by right-continuity of the budget function, there is some $\epsilon > 0$ such $B^{\mathsf{a}}_*(p') < B^{\mathsf{b}}_*(p')$ for any $p' \in [p, p+\epsilon)$.
- $B^{\mathsf{a}}_*(p) = B^{\mathsf{b}}_*(p)$, therefore, $B^{\mathsf{a}}_i(p) = B^{\mathsf{b}}_i(p)$ for all $i \in A(p)$, moreover, $S^{\mathsf{a}}(p) = S^{\mathsf{b}}(p)$, since

$$S^{\mathsf{a}}(p) = \sum_{i\in A(p)} \frac{B^{\mathsf{a}}_i(p)}{p} - \frac{B^{\mathsf{a}}_*(p)}{p} = \sum_{i\in A(p)} \frac{B^{\mathsf{b}}_i(p)}{p} - \frac{B^{\mathsf{b}}_*(p)}{p} = S^{\mathsf{b}}(p)$$

Since the behavior of the function $B(\cdot)$ for $p' \geq p$ just depends on $S(p)$ and $B(p)$, for all $p' \geq p$, then for all $p' \geq p$, $B^{\mathsf{a}}(p) = B^{\mathsf{b}}(p)$. In other words, when $B_*^{\mathsf{b}}(p)$ and $B_*^{\mathsf{a}}(p)$ meet, then the auctions become *fully coupled*.

Part (c) of the Meta-Lemma is essentially the same argument made in item (a). This part is trivial for $p \notin \{v_1, \ldots, v_n\}$ by continuity of $B(p)$. For $p = v_j$ we use that $B_*^{\mathsf{a}}(v_j-) \leq B_*^{\mathsf{b}}(v_j-)$ and study δ_i^{a} and δ_i^{b}. As in (c) we get:

$$\delta_i^{\mathsf{a}} = \frac{1}{v_j}\left[\min\{B_i(0), B_*^{\mathsf{a}}(v_j-)\} + \min\{B_j(0), B_*^{\mathsf{a}}(v_j-)\} - B_*^{\mathsf{a}}(v_j-)\right]^+$$

$$\delta_i^{\mathsf{b}} = \frac{1}{v_j}\left[\min\{B_i(0), B_*^{\mathsf{b}}(v_j-)\} + \min\{B_j(0), B_*^{\mathsf{b}}(v_j-)\} - B_*^{\mathsf{b}}(v_j-)\right]^+$$

Now, by analyzing cases A and B as in part (a) of the Meta-Lemma, we conclude that $B_i^{\mathsf{a}}(v_j) = B_i^{\mathsf{a}}(v_j-) - v_j\delta_i^{\mathsf{a}} \leq B_i^{\mathsf{b}}(v_j-) - v_j\delta_i^{\mathsf{b}} = B_i^{\mathsf{b}}(v_j)$ as desired. ■

Now, we want to establish allocation monotonicity, i.e., that $x_i^{\mathsf{a}}(p) \geq x_i^{\mathsf{b}}(p)$ for all $p \geq 0$. We will prove a stronger claim, that the *wishful allocation* Ψ_i is monotone in the supply, i.e., $\Psi_i^{\mathsf{a}}(p) \geq \Psi_i^{\mathsf{b}}(p)$ for all $p \geq 0$.

Proposition 6.2.14 (Allocation Monotonicity) *For all $p \geq 0$ and all agents i, the following invariant holds: $\Psi_i^{\mathsf{b}}(p) \leq \Psi_i^{\mathsf{a}}(p)$.*

Proof : This proof follows from combining Proposition 6.2.13 and Lemma 6.2.6. For small values of p, $\Psi_i^{\mathsf{b}}(p) \leq \Psi_i^{\mathsf{a}}(p)$ is definitely true, since both are equal to $\frac{B_i}{p}$. Now, if it is true for some small p, then it is true for any $p' \geq p$, since:

$$\Psi_i^{\mathsf{a}}(p') = \Psi_i^{\mathsf{a}}(p) - \int_p^{p'} \frac{B_i^{\mathsf{a}}(\rho)}{\rho^2} d\rho \geq \Psi_i^{\mathsf{b}}(p) - \int_p^{p'} \frac{B_i^{\mathsf{b}}(\rho)}{\rho^2} d\rho = \Psi_i^{\mathsf{b}}(p').$$

■

The proof of our main theorem follows immediately from Propositions 6.2.13 and 6.2.14.

Proof of Theorem 6.2.1 : For the allocation monotonicity, Proposition 6.2.14 implies that $x_i^{\mathsf{a}}(p) + \frac{B_i^{\mathsf{a}}(p)}{p} \geq x_i^{\mathsf{b}}(p) + \frac{B_i^{\mathsf{b}}(p)}{p}$. Since $B_i^{\mathsf{a}}(p) \leq B_i^{\mathsf{b}}(p)$, then clearly: $x_i^{\mathsf{a}}(p) \geq x_i^{\mathsf{b}}(p)$, taking $p \to \infty$ we get that for each player i, the final allocation in the augmented auction and in the base auction are such that $x_i^{\mathsf{a}} \geq x_i^{\mathsf{b}}$.

The monotonicity of the payment function follows directly from Proposition 6.2.13. The remaining budget in the end is larger in the base auction then in the augmented auction for each agent. So, the final payments are such that $\varphi_i^{\mathsf{a}} \geq \varphi_i^{\mathsf{b}}$.

■

6.3 Algorithmic Form of the Adaptive Clinching Auction

We presented the Adaptive Clinching Auction in Definition 6.2.2 as the limit as $p \to \infty$ of a differential procedure following Bhattacharya el al [13]. Here we present the same auction in an algorithmic format, i.e., an $\tilde{O}(n)$ steps procedure to compute (x, φ) from (v, B, s). The idea is quite simple: given a price p and the values of $B(p), x(p)$, we solve the differential equation in item (i) of Definition 6.2.2 and using it, we compute the next point $\bar{p}$ where either a player leaves the active set, or a player enters the clinching set. Given that, we compute $B(\bar{p}-), x(\bar{p}-)$. Then we obtain the values of $B(\bar{p}), x(\bar{p})$ either by the procedure in (iv) if a player leaves the active set on $\bar{p}$ or simply by taking $B(\bar{p}) = B(\bar{p}-)$

and $x(\bar{p}) = x(\bar{p}-)$ otherwise.

Lemma 6.3.1 *Consider the functions $x(p)$ and $B(p)$ obtained in the Adaptive Clinching Auction. If for prices $p' \in [p, \bar{p})$, the clinching and active set are the same, i.e., $C(p') = C(p)$ and $A(p') = A(p)$, then given $k = |C(p)|$, the players i in the clinching set are such that:*

- *if $k = 1$, $S(p') = \frac{pS(p)}{p'}$, $x_i(p') = x_i(p) + [S(p') - S(p)]$ and $B_i(p') = B_i(p) + pS(p)[\log p - \log p']$.*
- *if $k > 1$, $S(p') = \frac{p^k S(p)}{(p')^k}$, $x_i(p') = x_i(p) + \frac{1}{k}[S(p') - S(p)]$ and $B_i(p') = B_i(p) + \frac{p^k S(p)}{k-1}\left[\frac{1}{p'^{k-1}} - \frac{1}{p^{k-1}}\right]$*

Proof : The proof is straightforward. For the case of $k = 1$, we follow the discussion in Example 6.2.2: let i be the only player in $C(p)$, then $S(p') + x_i(p')$ is constant in this range, since all that is subtracted from the supply is added to the allocation of player 1, therefore:

$$\partial S(p') = -\partial_p x_i(p') = -\frac{S(p')}{p'} \Rightarrow S(p') = \frac{\alpha}{p'}.$$

using the boundary condition $S(p) = \frac{\alpha}{p}$, we get the value of $\alpha = pS(p)$. Now, clearly $x(p') = x(p) + [S(p') - S(p)]$, since player i is the only one clinching. For his budget:

$$B_i(p') - B_i(p) = \int_p^{p'} \partial_p B_i(\rho) d\rho = \int_p^{p'} -S(\rho) d\rho = \int_p^{p'} -\frac{pS(p)}{\rho} d\rho = pS(p)[\log p - \log p']$$

For $k > 1$, $S(p') + \sum_{i \in C(p')} x_i(p')$ is constant and therefore:

$$\partial S(p') = -\sum_{i \in C(p')} \partial_p x_i(p') = -k\frac{S(p')}{p'} \Rightarrow S(p') = \frac{\alpha}{(p')^k}.$$

Using the boundary condition $S(p) = \frac{\alpha}{p^k}$, we get the value of $\alpha = p^k S(p)$. We use the observation in Lemma 6.2.11 that players in the clinching set have the same budget, and therefore the auction treats them equally from this point on as long as they remain in the active set, i.e., they will get allocated and charged at the same rate. Therefore: $x(p') = x(p) + \frac{1}{k}[S(p') - S(p)]$. For the budgets:

$$B_i(p') - B_i(p) = \int_p^{p'} -S(\rho)d\rho = \int_p^{p'} -\frac{pS(p)}{\rho^k}d\rho = \frac{p^k S(p)}{k-1}\left[\frac{1}{(p')^{k-1}} - \frac{1}{p^{k-1}}\right]$$

■

Theorem 6.3.2 (Algorithmic Form) *It is possible to compute the allocation and payments of the Adaptive Clinching Auction in $\tilde{O}(n)$ time.*

Proof : Using the lemma above, we just need to compute x and B for the points where one of the following events happen: (a) one leaves the active set and (b) one player enters the clinching set. Clearly there are at most n events of type (a) and by Lemma 6.2.8 also at most n events of type (b).

The algorithm starts at price $p = 0$ and at each time computes the next event. For example, at price $p = 0$, the next event of type (a) occurs in $p = \min_i v_i$. The next event of type (b) occurs at price $p = \frac{1}{s}[\sum_A B_i - \max_A B_i]$ if no event of type (a) happens before. First we compute which one occurs first. Let $\bar{p}$ be such a price. Then, computing $B(\bar{p}-), x(\bar{p}-)$ is trivial, since no clinching happened so far, so at that price: $B(\bar{p}-) = B$ (initial budgets) and $x(\bar{p}-) = 0$. Now, if $\bar{p}$ is an event of type (a), then use step (iv) in Definition 6.2.2 to compute $x(\bar{p}), B(\bar{p})$. If not, simply take $B(\bar{p}) = B(\bar{p}-)$ and $x(\bar{p}) = x(\bar{p}-)$.

From this point on, at each considered price p, the clinching set will be non-empty, so we know the format of $x(p')$ and $B(p')$ for $p' \in [p, p+\epsilon)$. If the next event that happens is of type (a), it happens at $\min\{v_j; v_j > p\}$, if it is of type (b), it happens at $\min\{p'; B_*(p') = \max_{i' \in A(p') \setminus C(p')} B_{i'}\}$, where the expression for $B(p')$ is given in the previous lemma. For example, if $|C(p)| = 1$, then this happens at:

$$p' = \exp\left[\frac{1}{pS(p)}(B_*(p) - \max_{i' \in A(p) \setminus C(p)} B_{i'}(p)) + \log p\right]$$

and if $|C(p)| = k > 1$, it happens at:

$$p' = \left[\frac{1}{p^{k-1}} - \frac{(B_*(p) - \max_{i' \in A(p) \setminus C(p)} B_{i'}(p))}{p^k S(p)} \cdot (k-1)\right]^{-1/(k-1)}$$

Those expressions are easily obtained by taking $B_i(p')$ as calculated in the previous lemma and calculating for which p' it becomes equal to $\max_{i' \in A(p) \setminus C(p)} B_i'(p)$. Now, we simply need to find out which of those events happen first. Let it be $\bar{p}$, then we compute $B(\bar{p}-), x(\bar{p}-)$ using the previous lemma and then update to $B(\bar{p}), x(\bar{p})$ as described above. ∎

Accompaigning this thesis we also provide an implementation of the algorithm described above. The implementation can be found in `clinching.m`, which is an Octave code that takes as an input a vector of valuation v, a vector of budgets B and the supply level s and produces allocation x and payments π.

BIBLIOGRAPHY

[1] Zoë Abrams. Revenue maximization when bidders have budgets. In *SODA*, pages 1074–1082, 2006.

[2] Gagan Aggarwal, Jon Feldman, S. Muthukrishnan, and Martin Pál. Sponsored search auctions with markovian users. In *Proceedings of the 4th International Workshop on Internet and Network Economics*, WINE '08, pages 621–628, Berlin, Heidelberg, 2008. Springer-Verlag.

[3] Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *SODA*, pages 1253–1264, 2011.

[4] Gagan Aggarwal, S. Muthukrishnan, Dávid Pál, and Martin Pál. General auction mechanism for search advertising. In *WWW*, pages 241–250, 2009.

[5] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *FOCS*, pages 482–491, 2001.

[6] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and some applications. *Theory of Computing Journal*, (to appear), 2012.

[7] Itai Ashlagi, Mark Braverman, Avinatan Hassidim, Ron Lavi, and Moshe Tennenholt. Position auctions with budgets: Existence and uniqueness. In *B.E. journal of Theoretical Economics Advances, forthcoming*, 2011.

[8] Susan Athey and Denis Nekipelov. Equilibrium and uncertainty in sponsored search advertising. In *AdAuctions Workshop 2010*, 2010.

[9] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94, 1997.

[10] Lawrence M. Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press, 2006.

[11] Moshe Babaioff, Liad Blumrosen, and Aaron Roth. Auctions with online supply. In *ACM Conference on Electronic Commerce*, pages 13–22, 2010.

[12] Dirk Bergemann and Stephen Morris. Correlated equilibrium in games with incomplete information. Cowles Foundation Discussion Papers 1822, Cowles Foundation for Research in Economics, Yale University, 2011.

[13] Sayan Bhattacharya, Vincent Conitzer, Kamesh Munagala, and Lirong Xia. Incentive compatible budget elicitation in multi-unit auctions. In *SODA*, pages 554–572, 2010.

[14] Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V. Vohra. An ascending vickrey auction for selling bases of a matroid. *Operations Research*, 59(2):400–413, 2011.

[15] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *STOC08*, 2008.

[16] Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce*, pages 44–51, 2005.

[17] Niv Buchbinder, Kamal Jain, and Joseph Naor. Online primal-dual algorithms for maximizing ad-auctions revenue. In *ESA*, pages 253–264, 2007.

[18] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *ACM Conference on Electronic Commerce*, pages 81–90, 2011.

[19] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. Revenue guarantees in sponsored search auctions. In *ESA*, 2012.

[20] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, Maria Kyropoulou, Brendan Lucier, Renato Paes Leme, and Éva Tardos. On the efficiency of equilibria in generalized second price auctions. *CoRR*, abs/1201.6429, 2012.

[21] Ioannis Caragiannis, Panagiotis Kanellopoulos, Christos Kaklamanis, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *EC'11*, 2011.

[22] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing.

In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 311–320, New York, NY, USA, 2010. ACM.

[23] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1), September 1971.

[24] Riccardo Colini-Baldeschi, Monika Henzinger, Stefano Leonardi, and Martin Starnberger. On multiple round sponsored search auctions with budgets. *CoRR*, abs/1112.6361, 2011.

[25] Peter Cramton and Axel Ockenfels. Economics and design of capacity markets for the power sector. *Zeitschrift fr Energiewirtschaft*, 36:113–134, 2012. 10.1007/s12398-012-0084-2.

[26] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.

[27] Robert Day and Peter Cramton. The quadratic core-selecting payment rule for combinatorial auctions. Technical report, 2008.

[28] Robert Day and Paul Milgrom. Core-selecting package auctions. *Int. J. Game Theory*, 36(3-4):393–407, 2008.

[29] Nikhil R. Devanur and Jason D. Hartline. Limited and online supply and the bayesian foundations of prior-free mechanism design. In *ACM EC*, pages 41–50, 2009.

[30] Nikhil R. Devenur and Thomas P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *ACM Conference on Electronic Commerce*, pages 71–78, 2009.

[31] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.

[32] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[33] Danny Dolev, Dror G. Feitelson, Joseph Y. Halpern, Raz Kupferman, and Nathan Linial. No justified complaints: on fair sharing of multiple resources. In *ICTS*, pages 68–75, 2012.

[34] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, March 2007.

[35] Benjamin Edelman and Michael Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, (2), 2010.

[36] eMarketer. Online ad spend continues double-digit growth. `http://www.emarketer.com/Article.aspx?R=1008087`.

[37] Jon Feldman, S. Muthukrishnan, Evdokia Nikolova, and Martin Pál. A truthful mechanism for offline ad slot scheduling. In *SAGT*, pages 182–193, 2008.

[38] Amos Fiat, Stefano Leonardi, Jared Saia, and Piotr Sankowski. Single valued combinatorial auctions with budgets. In *ACM Conference on Electronic Commerce*, pages 223–232, 2011.

[39] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[40] Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Polyhedral clinching auctions and the adwords polytope. In *STOC*, pages 107–122, 2012.

[41] Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Clinching auctions with online supply. In *SODA*, 2013.

[42] Renato D. Gomes and Kane S. Sweeney. Bayes-Nash equilibria of the generalized second price auction. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 107–108, New York, NY, USA, 2009. ACM.

[43] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: a survey. *Ann. Discrete Math.*, 4:287–326, 1979.

[44] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[45] I. Hafalir, R. Ravi, and Amin Sayedi. Sort-cut: A pareto optimal and semi-truthful mechanism for multi-unit auctions with budget-constrained bidders. *Ad Auctions Workshop*, 2009.

[46] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd national conference on Artificial intelligence - Volume 1*, pages 58–65. AAAI Press, 2007.

[47] Sergiu Hart and Andreu M. Colell. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

[48] Jason D. Hartline. Approximation in economic design. Lecture Notes, `http://users.eecs.northwestern.edu/~hartline/amd.pdf`.

[49] U. Krengel and L. Sucheston. Semiamarts and finite values. In *Bull. Amer. Math. Soc. 83*, 1977.

[50] U. Krengel and L. Sucheston. On semiamarts and amarts, and processes with finite value. In *Probability on Banach Spaces*, 1978.

[51] V. Krishna. *Auction Theory*. Academic Press/Elsevier, 2009.

[52] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored search auctions, pages 699–716. Cambridge University Press, 2007.

[53] Sébastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *EC '06: Proceedings of the 7th ACM conference on Electronic commerce*, pages 218–227, New York, NY, USA, 2006. ACM.

[54] Ron Lavi and Marina May. A note on the incompatibility of strategy-proofness and pareto-optimality in quasi-linear settings with public budgets. *WINE*, 2011.

[55] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2), 1994.

[56] Brendan Lucier, Renato Paes Leme, and Éva Tardos. On revenue in the generalized second price auction. In *WWW*, pages 361–370, 2012.

[57] Brendan Lucier and Renato Paes Leme. Improved social welfare bounds for GSP at equilibrium. In *arXiv*, 2011.

[58] Mohammad Mahdian and Amin Saberi. Multi-unit auctions with un-

known supply. In *Proceedings of the 7th ACM conference on Electronic commerce*, EC '06, pages 243–249, New York, NY, USA, 2006. ACM.

[59] P. Maille, E. Markakis, M. Naldi, G. D. Stamoulis, and B. Tuffin. Sponsored search auctions: An overview of research with emphasis on game theoretic aspects. To appear in the Electronic Commerce Research journal (ECR), 2012.

[60] E. Markakis and O. Telelis. Discrete strategies in keyword auctions and their inefficiency for locally aware bidders. In *Workshop on Internet and Network Economics (WINE)*, 2010.

[61] Colin J. H. McDiarmid. Rado's theorem for polymatroids. *Mathematical Proceedings of the Cambridge Philosophical Society*, 78(02):263–281, 1975.

[62] Aranyak Mehta, Amin Saberi, Umesh V. Vazirani, and Vijay V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.

[63] P. Milgrom. *Putting Auction Theory to Work*. Churchill Lectures in Economics. Cambridge University Press, 2004.

[64] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[65] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, April 1983.

[66] Uri Nadav and Tim Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *WINE*, pages 319–326, 2010.

[67] Thành Nguyen and Éva Tardos. Approximately maximizing efficiency and revenue in polyhedral environments. In *ACM Conference on Electronic Commerce*, pages 11–19, 2007.

[68] Thành Nguyen and Milan Vojnovic. The weighted proportional sharing mechanisms. In *SIGMETRICS*, 2011.

[69] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[70] Renato Paes Leme and Éva Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auctions. In *FOCS10: 51st Annual IEEE Symposium on Foundations of Computer Science*, 2010.

[71] Mallesh M. Pai and Rakesh Vohra. Optimal auctions with financially constrained bidders. Discussion papers, Northwestern University, Center for Mathematical Studies in Economics and Management Science, Aug 2008.

[72] David C. Parkes. Online mechanisms. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 16, pages 411–439. Cambridge University Press, 2007.

[73] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 513–522, New York, NY, USA, 2009. ACM.

[74] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

[75] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 2006.

[76] Hal R. Varian. Online ad auctions. *American Economic Review Papers and Proceedings*, 2009.

[77] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[78] Felix Wu, Pravin Varaiya, Pablo Spiller, and Shnmel Oren. Folk theorems on transmission access: Proofs and counterexamples. *Journal of Regulatory Economics*, 10:5–24, 1996.