ABSTRACT

Title of dissertation: REVENUE EFFICIENT MECHANISMS FOR ONLINE ADVERTISING

Mohammad Reza Khani, Doctor of Philosophy, 2015

Dissertation directed by: Professor MohammadTaghi Hajiaghayi
Department of Computer Science

Online advertising is an essential part of Internet and the main source of revenue for lots of web-centric companies such as search engines, news websites, Internet social networks, and other types of publishers. Online advertising happens in different settings and includes many challenges and constraints. A key component in each setting is the mechanism which selects and prices the set of winning ads. In this thesis, we consider the mechanism related issues arises in online advertising and propose candidate solutions with a special focus on the revenue aspect.

Generalized Second Price (GSP) auction (the current mechanism of choice in online advertising) has appealing properties when ads are simple (text based and identical in size). But GSP does not generalize to richer ad settings, whereas truthful mechanisms, such as VCG do. Hence there are incentives for search platforms to migrate to truthful

mechanisms, but a straight switch from GSP to VCG either requires all bidders instantly bid truthfully or incurs significant revenue loss. We propose a transitional mechanism which encourages advertisers to update their bids to their valuations, while mitigating revenue loss[1]. The mechanism is equivalent to GSP when nobody has updated her bid, is equivalent to VCG when everybody has updated, and it has the same allocation and payments of the original GSP if bids were in the minimum symmetric Nash equilibrium. In settings where both GSP ads and truthful ads exist, it is easier to propose a payment function than an allocation function. We give a general framework for these settings to characterize payment functions which guarantee incentive compatibility of truthful ads, by requiring that the payment functions satisfy two properties.

Next, we discuss about revenue monotonicity (revenue should go up as the number of bidders increases) of truthful mechanisms in online advertising. This natural property comes at the expense of social welfare - one can show that it is not possible to get truthfulness , revenue monotonicity, and optimal social welfare simultaneously. In light of this, we introduce the notion of Price of Revenue Monotonicity (PoRM) to capture the loss in social welfare of a revenue monotone mechanism. We design truthful and revenue monotone mechanisms for important online advertising auctions with small PoRM and prove a matching lower bound.

Finally, we study how to measure revenue of mechanisms in the prior free settings. One of the major drawbacks of the celebrated VCG auction is its low (or zero) revenue

[1] It is the candidate mechanism to make the transition from GSP to VCG in Microsoft's Bing.

even when the agents have high values for the goods and a *competitive* outcome would have generated a significant revenue. A competitive outcome is one for which it is impossible for the seller and a subset of buyers to 'block' the auction by defecting and negotiating an outcome with higher payoffs for themselves. This corresponds to the well-known concept of *core* in cooperative game theory where designing *core-selecting auctions* is well studied [AM02, AM06, DM08, DC12]. While these auctions are known for having good revenue properties, they lack incentive-compatibility property desired for online advertising. Towards this, we define a notion of *core-competitive* auctions. We say that an incentive-compatible auction is α-core-competitive if its revenue is at least $1/\alpha$ fraction of the minimum revenue of a core-outcome. We study designing core-competitive mechanisms for a famous online advertising scenario.

REVENUE EFFICIENT MECHANISMS FOR ONLINE ADVERTISING

by

Mohammad Reza Khani

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:
Professor MohammadTaghi Hajiaghayi, Chair/Advisor
Professor V.S. Subrahmanian
Professor Andrew Childs
Professor Lawrence M. Ausubel, Department of Economics
Professor Peter Cramton, Department of Economics
Doctor Gagan Goel, Google Inc.

Dedication

To the memory of my father and my lovely mother.

Acknowledgments

First and foremost I would like to thank my advisor, Prof. MohammadTaghi Hajiaghayi for his sincere guidance over the last 4 years. He has been always available to help and a great source of motivation. I have used his helps and advices not only on academic subjects but also on other aspects of my life. Thanks Mohammad for being such a good adviser, and I am hopeful that we maintain a fruitful collaboration in the future as well.

I owe a significant part of my academic achievements to all the great teachers and mentors I sought advice from. I would like to thank my masters advisor professor Mohammad Salavatipour. I also would like to thank my internship mentors Gagan Goel, Ian Kash, Sumit Das, and Zutao Zhu. I have learned a lot by collaborating with Mohammad Hajiaghayi, Mohammad Salavatipour, Gagan Goel, Ian A. Kash, Renato Paes Leme, Yoram Bachrach, Sofia Ceppi, Peter Key, Theodore Johnson, Barna Saha, Hossein Esfandiari, Vahid Liaghat, Hamid Mahini, Harald Racke, Rohit Khandekar, and Guy Kortsarz. Many thanks to Peter Cramton, Lawrence Ausubel, V.S. Subrahmanian, Andrew Childs, and Gagan Geol for serving on my dissertation committee and providing me with informative comments.

I also gratefully acknowledge the support of the following grants during my PhD studies: NSF CAREER award 1053605, NSF grant CCF-1161626, NSF CAREER award 0844872, and Google US/Canada Faculty Research Award.

Table of Contents

List of Figures

CHAPTER 1

Introduction

1.1 Online Advertising

Online advertising is a growing marketing method which uses Internet to deliver promotional messages. The promotional messages are shown in different forms such as text, image (banners), or video ads along with the publishers' contents. Online advertising is the main source of revenue for many Internet firms and the main method of attracting customers for lots of businesses. With the advent of newer online technologies such as game consoles and mobile computing, this market is expanding and evolving even more quickly. The total revenue of online advertising in U.S. was near 50 billion dollars at 2014 [IP15].

A central component of online advertising is the underlying mechanism that selects and prices the winning ads for a given ad slot. When a user queries to access a publisher's content, the publisher uses a mechanism to select and price a few ads to show along with its content. There are several challenges involved in this process. The mechanisms have to consider incentives and objectives of at least the three main parties involved: the publisher, the advertisers, and the users. For example, the publisher may want to gain more revenue and show high quality ads in its website, the advertisers want to increase their profits by paying less while getting higher click through rates, and the users want to see related and high quality ads. For each user's query and its associated content there are a specific template for showing ads, a special set of interested advertisers, and other constraints specific to that query. This means that the mechanisms have to work for various auctions each with different settings. There are also several other challenges for mechanisms which will be discussed later in this thesis.

Online advertising firms prefer to have a mechanism which has the following properties. It should select high quality ads whose probabilities of getting clicked are high. It has to charge the advertisers upon click high enough to increase the publisher's revenue but not so high that the advertisers decrease their bids significantly. The process of selecting winning ads should be clear so that advertisers can compute their best bids. This process also have to be very fast as users want to see contents without delay. The main purpose of this thesis is to design mechanisms for different online advertising scenarios having the mentioned properties with special focus on revenue. Each of the following chapters of this thesis is kept standalone and readable independent of other chapters.

1.2 Organization of the Thesis

1.2.1 Chapter 2

Chapter 2 is based on Bachrach *et al.* [BCK+15]. The motivation of this chapter is that the Generalized Second Price (GSP) auction (the current mechanism of choice) has appealing properties when ads are simple (text based and identical in size). But GSP does not generalize to richer ad settings, whereas truthful mechanisms, such as VCG do. Hence there are incentives for search platforms to migrate to truthful mechanisms, but a straight switch from GSP to VCG either requires all bidders instantly bid truthfully or incurs significant revenue loss. Therefore, we propose a hybrid auction mechanism for sponsored search, where bidders can be truthful or not, and are accordingly treated differently. Our class of hybrid mechanisms give incentives for non-truthful bidders to bid truthfully, while behaving as a non-truthful auction if no bidders are truthful.

We introduce a transitional mechanism which encourages advertisers to update their bids to their valuations, while mitigating revenue loss. The mechanism is equivalent to GSP when nobody has updated her bid, is equivalent to VCG when everybody has updated, and it has the same allocation and payments of the original GSP if bids were in the minimum symmetric Nash equilibrium. In settings where both GSP ads and truthful (TF) ads exist, it is easier to propose a payment function than an allocation function. We give a general framework for these settings to characterize payment functions which guarantee incentive compatibility of truthful ads, by requiring that the payment functions satisfy two properties.

Finally, we compare the revenue of our transitional mechanism with revenues of GSP and VCG mechanisms when run on a sample of Bing data.

1.2.2 Chapter 3

This chapter (based on [GK14]) is inspired by one of the biggest practical drawbacks of the widely popular Vickrey-Clarke-Groves (VCG) mechanism, which is the unique incentive-compatible mechanism that maximizes social welfare. It is known that VCG lacks a desired property of revenue monotonicity - a natural notion which states that the revenue of a mechanism shouldn't go down as the number of bidders increase or if the bidders increase their bids. Most firms which depend on online advertising revenue have a large sales team to attract more bidders on their inventory as the general belief is that more bidders will increase competition, and hence revenue. However, the lack of revenue monotonicity of VCG conflicts with this general belief and can be strategically confusing for the firm's business.

In this chapter, we seek incentive-compatible mechanisms that are revenue-monotone. This natural property comes at the expense of social welfare - one can show that it is not possible to get incentive-compatibility, revenue-monotonicity, and optimal social welfare simultaneously. In light of this, we introduce the notion of Price of Revenue Monotonicity (PoRM) to capture the loss in social welfare of a revenue-monotone mechanism.

We further study revenue-monotonicity for two important online advertising scenarios. First one is the text vs image ad auction where in an ad slot, one can either show a single image ad or a few text ads. Second one is the video-pod auction where we have a video advertising slot of k seconds which can be filled with multiple video ads. For the image-text auction, we give a mechanism that satisfy both RM and IC and achieve PoRM of $\sum_{i=1}^{k} \frac{1}{i} \approx \ln k$. We also show that the PoRM of our mechanism is the best possible by proving a matching lower bound of $\sum_{i=1}^{k} \frac{1}{i}$ on the PoRM of any deterministic mechanism under some mild assumptions. For the video-pod auction, we give a mechanism that achieves a PoRM of $(\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)$.

1.2.3 Chapter 4

In Chapter 3, it is argued that RM is a desired property which popular mechanisms like VCG do not have. We show that no deterministic RM mechanism can attain PoRM of less than $\ln(k)$ for video-pod auction, *i.e.*, no deterministic mechanism can attain more than $\frac{1}{\ln(k)}$ fraction of the maximum social welfare. We also design a mechanism with PoRM of $O(\ln^2(k))$ for video-pod.

In this chapter (based on [GHK14]) we study designing a mechanism for the Combinatorial Auction with Identical Items (CAII) in which we are interested in selling k identical items to a group of bidders each demanding a certain number of items between 1 and k. CAII generalizes online advertising scenarios such as image-text and video-pod auctions (also are discussed in Chapter 3). We seek to overcome the impossibility result of Chapter 3 for deterministic mechanisms by using the power of randomization. We show that by using randomization, one can attain a constant PoRM. In particular, we design a randomized RM mechanism with PoRM of 3 for CAII.

1.2.4 Chapter 5

In Chapters 3 and 4 we argue that revenue monotonicity is a desired property for mechanisms in online advertising scenario. RM is a tool for measuring revenue by comparing it to the revenue of the same mechanism for a smaller type profile. In this chapter we focus on evaluating the revenue by looking solely at the bids of participants.

One of the major drawbacks of VCG auction is its low (or zero) revenue even when the agents have high value for the goods and a *competitive* outcome would have generated a significant revenue. A competitive outcome is one for which it is impossible for the seller and a subset of buyers to 'block' the auction by defecting and negotiating an outcome with higher payoffs for themselves. This corresponds to the well-known concept of *core* in cooperative game theory.

In particular, VCG revenue is known to be not competitive when the goods being sold have *complementarities*. One important research direction is that of the design of *core-selecting auctions* for goods with complementarities (See Ausubel and Milgrom, Day and Milgrom, Day and Cramton, Ausubel and Baranov). Core-selecting auctions are combinatorial auctions whose outcome implements competitive prices even when the goods are complements. While these auction designs have been implemented in practice in various scenarios and are known for having good revenue properties, they lack the desired incentive-compatibility property of the VCG auction. A bottleneck here is an impossibility result showing that there is no auction that simultaneously achieves competitive prices (a core outcome) and incentive-compatibility.

In this chapter (based on [GKL15]) we try to overcome the above impossibility result by asking the following natural question: is it possible to design an incentive-compatible auction whose revenue is comparable (even if less) to a competitive outcome? Towards this, we define a notion of *core-competitive* auctions. We say that an incentive-compatible auction is α-core-competitive if its revenue is at least $1/\alpha$ fraction of the minimum revenue of a core-outcome. We study Text-and-Image setting and design an $O(\ln \ln k)$ core-competitive randomized auction and an $O(\sqrt{\ln k})$ competitive deterministic auction. We also show that both factors are tight.

CHAPTER 2

Mechanism Design for Mixed Ads

2.1 Introduction

Sponsored search is the main source of revenue for most search engines, such as Google, Yahoo! or Bing. In the classic online ad auction or ‘position auction’, all ad slots and sizes are the same. Search engines typically use a variant of the Generalized Second Price (GSP) mechanism to select and price ads. In GSP, advertisers are rank-ordered by decreasing bids (more generally, by expected revenue or rank score) and slots are assigned in this order. The price of a slot is the minimum bid an advertiser has to make in order to maintain that position, which equates to the next highest bid in the simplest form of GSP. Payment is made when an ad is clicked. The GSP auction’s equilibria, bidding strategies, and other properties are well studied (see, e.g., [EOS05, Var07]).

However, online advertising is becoming more complex. There may be multiple page-templates for search results, different ad formats (*e.g.*, text-ads or image-ads) with different sizes, and several other constraints on showing ads. For these settings, GSP is not well defined and if generalized can be ill-behaved [BCK+14]. Therefore, there is an incentive for migrating from GSP to another mechanism. Truthful mechanisms (such as VCG) are attractive (see [VH14]) because they allow externalities [1] to be calculated which makes handling more complex ad scenarios easier. Truthful mechanisms also remove the computational burden of calculating the optimum bid from advertisers, make the whole system more transparent, and the same bid (valuation) of an advertiser can be used across multiple auctions. Moreover, analyzing the market is easier with true valuations as opposed to bids.

One big obstacle to migration from GSP to VCG is the requirement that advertisers update their bids by increasing them up to their true value. Indeed, as Varian and Harris note, Google "thought very seriously about changing the GSP auction to a VCG auction during the summer of 2002." However, there were several problems, including that "the VCG auction required advertisers to raise their bids above those they had become accustomed to in the GSP auction." As GSP has only gotten more entrenched over the past decade, this issue has only grown larger. If such a switch were made today, advertisers may not update their bids quickly and even if they do update, it might not be their true valuation. Thus, for a long time there will be a mixed set of ads: (i) advertisers who have updated their bids to true valuations and (ii) advertisers with GSP bids. Therefore, it would be desirable to have a transitional mechanism which selects and prices the set of

[1]The externality of each advertiser is the decrease in social welfare of others with and without her (see [NRTV07, Chapter 9]).

winners from such a mixed set of ads.

In this chapter we introduce a transitional mechanism in order to migrate GSP bids to true valuations. At the start of the transition the mechanism behaves the same as GSP. Then (in one implementation of our mechanism) when an advertiser tries to update her bid, we notify her that the mechanism will optimize her allocation and payment assuming that her new bid is her true valuation. This means that the best response for the advertiser is to bid her true valuation. We show that since this optimization will not happen for old GSP bids, the mechanism actually encourages advertisers to update their bids to their valuations as soon as possible (this is exactly what we want them to do). In the middle of the transition, the mechanism maintains a classification of bids as either truthful (updated bids) or GSP so that they can be treated differently, respecting incentives for truthful ads, but not GSP ads. Finally when all advertisers update their bids, the mechanism behaves as VCG and the transition finishes.

The key difficulty for a transition mechanism is to decide how bids of one group affect allocation and payment of the other group. For example, if the bid of a truthful ad is higher than the bid of a GSP ad, it does not necessarily mean that the truthful ad has to have a better allocation, since we know the GSP ad is shading its bid downward. Moreover, if a truthful ad is given an ad slot with click probability f and the ad below him is a GSP ad with bid g, we may want the expected payment of the truthful ad to be at least $g \cdot f$. The usual approach for designing truthful mechanisms is to give a monotone allocation function and then derive the unique payments using Myerson's lemma [Mye81].

However, in our setting it is not easy to first give an allocation function, as the bid of a GSP ad implies that *payments* of ads above him should be altered. Hence, designing an allocation function which takes into account the payment constraints is harder.

In this chapter we give a general framework, not limited to only GSP or VCG ads[2], for designing mechanisms in settings similar to the above where it is easier to specify a payment function than an allocation function. In practice, this flexibility allows our framework to apply in the presence of some of the additional complexity that exists in actual auction systems. To describe this framework more formally, assume ad slots are indexed from 1, the top slot, to n, the bottom one, and there are n advertisers. The designer just has to specify a payment function $p^i : \mathbb{R}^{n-i} \mapsto \mathbb{R}$ for each $i \in [n]$ which specifies the payment of the advertiser assigned to position i given types of advertisers assigned to positions below him. Our framework requires that any payment function satisfies two simple properties: (i) Minimum Marginal Increase (MMI): the payment has to be high enough so that truthful bidders assigned to lower slots do not envy the winner of a higher slot and (ii) Exact Marginal Increase (EMI): the marginal payment increase of the slot directly above a truthful ad has to be equal to the truthful ad's bid. Given a set of payment functions satisfying MMI and EMI, our framework shows how to construct an allocation rule with payments that are exactly those given by the payment functions applied to the realized allocation.

While we primarily focus on deriving a practical rule for transitioning from GSP

[2]The set of ads can have arbitrary number of groups. For example in addition to GSP and truthful ads, it may contain advertisers who have pre-established contracts and pay a fixed amount if their ad gets clicked.

to VCG, the set of ad auction mechanisms that fit in our framework is quite general. In fact, we prove that by using our framework one can design any truthful mechanism in which the payment of an ad is derived solely from ads below that ad (subject to a few additional requirements). Equivalently, our framework encapsulates mechanisms where raising the bid of an ad does not affect the allocation of ads that were previously allocated below it. More broadly, while the question of what properties of allocation rules lead to truthful mechanisms has been intensively studied (see, e.g., [AK08, ABHM10, FK14]), the question of what properties of payment rules lead to truthful mechanisms has not. Indeed, the only prior characterization we know of is the taxation principle for single-agent mechanisms. Since EMI and MMI are more inspired by Myerson than by the details of our auction setting, this approach may be of independent interest.

Having designed a large class of candidate mechanisms, and selected a representative member of the class, we analyze it both theoretically and in simulations based on Bing data. On the theory side, we show that transitional behavior is particularly nice if the system starts in the lowest symmetric Nash equilibrium. In particular, allocations and prices do not change regardless of the order of updates, in principle leading to a painless transition.

This conclusion relies on a number of strong (and false) assumptions. In our simulations, we analyze the consequences of relaxing them. We find that essentially all the costs of a transition are in terms of revenue—the welfare effects for both advertisers and users are small. We also see that the hybrid mechanism can have significant revenue benefits if bidders directly update to true valuations, but these benefits decrease if the bidders fail to do that for various reasons such as not knowing the true valuations, not being utility maximizers, or not trusting the system. So, rather than simply telling advertisers that they will be treated differently once they change their bid (which may not even be feasible in practice), it may make sense to attempt to use a learning approach to identify which bidders to treat as truthful. We simulate a few behaviors, and consequently identification models, and analyze the performance of the hybrid, GSP, and VCG mechanisms.

To summarize, our three main contributions in this chapter are:

1. a new framework for deriving truthful mechanisms from payment rules that satisfy MMI and EMI (Section 2.4),
2. a specific hybrid mechanism to enable transitioning from GSP to VCG (Section 2.5), and
3. an evaluation of the mechanism based on Bing data (Section 2.6).

2.2 Related Work

Sponsored search auctions are arguably the most successful recent application of auction theory to a business environment. As a result, much research has been conducted regarding the influence of the mechanism used for the auction on social welfare and the generated revenue. In the case where the VCG mechanism is used, truthfulness is the dominant bidding behavior. However, the same does not hold for the GSP auction and predicting bids in this case is trickier.

A complete information analysis of GSP auctions is discussed by Edelman et al. [EOS05], Varian [Var07], and Aggarwal et al. [AGM06]. A common theme in this line of work is the equivalence between the auctioneer's revenue and bidders' utility under a VCG auction and under the lowest symmetric Nash equilibrium of a GSP auction (which is sometimes referred to as the "bidder-optimal locally envy free equilibrium"). Ashlagi et al. [AMT07] generalize this, showing that in many auction types in which the payments are a function of the lowest ranked bids, there exists an equilibrium in which bidders' utility is equivalent to their utility under the VCG auction. Roberts et al. [RGKK13] generalize this along a different axis, showing that this result also holds for a variety of rank score functions other than simply ranking by highest bid.

Much of the research on equilibria in GSP auctions has focused on symmetric equilibria. Edelman and Schwarz [ES10] examined the revenue of different symmetric Nash equilibria, noting that under a certain comparison to optimal revenue possible under the Bayesian setting, the "lowest" equilibrium is the reasonable one. A generalized auction proposed by Aggarwal et al. [AFM07] allows advertisers to specify not only a bid but also the positions they are interested in, ruling out the bottom positions. They show that this auction has a symmetric Nash equilibrium implementing the same outcome (i.e., allocation and pricing) as the VCG auction.

Complementary to studies on symmetric equilibria, several researchers have studied the inefficiency that can result from asymmetric equilibria [LT09, CKKK11, LPLT12].

Some studies of auction tuning have also explored the full set of equilibria [TLB13]

Taking a Bayesian perspective, Gomes and Sweeney [GS14] examined the existence and uniqueness of Bayes-Nash equilibria in a GSP auction. Several models have also been proposed for inferring the valuations of advertisers based on the observed bid data [PK11, AN10]. The model by Pin and Key [PK11] considers advertisers best responding in an uncertain environment in a repeated auction setting, relating the bidding behavior to scenarios when the Bayes-Nash Equilibria of Gomes and Sweeney [GS14] are known to exist. The model of Athey and Nekipelov [AN10] starts directly from the Bayes-Nash Equilibria, but has a different model of the information available to the bidders. Instead, Vorobeychik [Vor09] proposed a framework based on agent simulation to approximate the Bayes-Nash equilibria in GSP auctions, which relies on restricting the space of allowed bidding strategies.

The dynamics leading to the equilibrium outcomes in GSP auctions are less studied. Cary et al. [CDE+07] consider dynamics under a greedy bidding strategy, where each bidder chooses the optimal bid for the next round assuming the other bidders do not change their bids. They show this bidding strategy has a unique fixed point, with payments identical to those of the VCG mechanism.

Closest to the work of this chapter, Aggarwal et al. [AMPP09] propose a framework that frames both GSP and VCG in terms of the assignment game with appropriate models of bidder utility. Their framework can be applied to derive a hybrid auction that incorporates both GSP and VCG bidders that is a special case of our more general framework.

However, they do not explore this and their framework lacks the flexibility ours provides.

2.3 Preliminaries

We study the standard model of a sponsored search auction. There is a set of $\{1, \ldots, n\}$ of ads, denoted by $[n]$. We assume that there are k ad positions ($[k]$), also called slots, where position i has CTR (Click-Through-Rate) f_i. Without loss of generality we assume that the first position has the highest CTR and the k-th position has the lowest ($0 < f_k < \ldots < f_1 \leq 1$). For notational convenience, and without loss of generality, we take $k = n$ and assume that there are no ad quality scores—the probability of any ad being clicked in slot i is exactly f_i.

In our setting ads can have different characteristics. For example they could be,

- Truthful (TF) ads which have updated their bids to their valuations;
- Generalized Second Price (GSP) ads which assume that they are participating in a GSP auction.
- First price ads which have a contract to pay a fixed amount upon being clicked.

We are going to provide a framework for mechanism designers which only requires a payment rule to be specified, and which places special constraints on the payments of TF ads. Therefore, we specially differentiate between TF and non-TF ads and assume each ad has a type taken from set $\mathcal{T} = \{TF, \text{non-}TF\} \times \mathbb{R}^+$ that specifies the truthfulness attribute and bid. Here, non-TF ads can be of any nature and our framework does not limit the designer's ability for deciding the allocation and payments of them. This can be thought of as modeling a situation where we are designing a system for the TF bidders, and therefore care about their incentives. However, there will be some "legacy" bidders whose behavior does not reflect the new system (and perhaps cannot, because inherently changing their bid will result in them being reclassified as TF). Since the right way to treat these legacy bidders will depend on their exact nature, we do not put any constraints on what they are charged (and indeed neglect modeling this entirely in the proving the theoretical results of Section 2.4).

In section 2.5, we examine what happens when we look to maintain parity with existing GSP prices for non-TF bidders, and where the non-TF bidders behave as GSP bidders in equilibrium.

We denote an assignment of ads to slots by the permutation $\Pi = (\pi_1, \ldots, \pi_n)$ where ad π_i is assigned to position i. The (expected) payment of ad π_i is the cost per click for being in position i multiplied by the CTR of position i. Throughout the chapter we work with expected payments $p_i, \forall i \in [n]$ as opposed to cost per click. Further, we assume that all ads have quasilinear utilities, *i.e.*, if ad π_i is assigned to position i and pays p_i then its utility will be $u_i(\Pi, p) = f_i * v(\pi_i) - p_i$, where $v(\pi_i)$ is the valuation of ad π_i. Throughout the chapter, we use $b(i)$ for the bid of ad i and if a TF ad, we also use $v(i) = b(i)$ to emphasize the fact that the bid and valuation are the same.

In this work we assume that the auctioneer can distinguish between TF ads and non-TF ads, i.e., the nature of each ad is known. Since we are aware that this is a strong assumption in practice, in Section 2.6 we elaborate on this point and show some empirical results where the auctioneer is uncertain about the ads nature. Additionally, our model assumes away the growing richness that is part of the ad ecosystem that is part of the motivation for a switch. However, at this point basic text ads still represent the bulk of ad impressions, so a solution that works well for them would be useful to aid a near-term transition in anticipation of this future richness.

2.4 Top Interference Free Payment Framework

In this section we introduce a framework for designing mechanisms in ad auction-like settings where it is easier to provide a payment rule than to give an allocation rule. In other words, we do not know exactly what we want the final allocation of ads to be, but we do know what we want the payment for ad π_i assigned to position i in the overall assignment Π to be. More formally, we explore the space of payment rules which are a set of n functions $\mathcal{P} = \{p^{(i)}\}_{i\in[n]}$, where function

$$p^{(i)} : \overbrace{\mathcal{T} \times \ldots \times \mathcal{T}}^{n-i} \longmapsto \mathbb{R}^{+}.$$

That is, if the ad in slot i is TF its payment will be $p^{(i)}(\pi_{i+1}, \ldots, \pi_n)$. (Recall that we allow the payments of non-TF bidders to be arbitrary.) Note that $p^{(i)}$ is the expected payment of the ad assigned to position i and $p^{(i)}/f_i$ is its cost per click.

This formulation implicitly restricts the set of payment rules we consider. Without loss of generality, the payment of a truthful ad does not depend on its own bid. However, with loss of generality, we also assume that payment does not depend on the bids of ads assigned to slots above it. This is a natural restriction in a setting without externalities, and indeed one that is satisfied by both the GSP and VCG payment rules.

Our framework specifies two further intuitive properties which we require the payment rule satisfy. We show that these two properties are necessary in the sense that any anonymous mechanism whose payment rule for an ad depends only on ads below it can be implemented using our framework. By anonymous, we mean that permuting the input to the mechanism simply permutes the output (up to tie breaking among ads with identical bids). In order to specify these properties we need to restrict the domain of the payment rules to exclude nonsensical inputs where the TF bidders are mis-ordered, e.g., where TF ads are not assigned to slots monotonically with respect to their bid.

Let $\Pi = (\pi_1, \ldots, \pi_n)$ denote the assignment of ads to positions. We use the notation $\Pi^{(k)} = (\pi_k, \ldots, \pi_n)$ to show the partial assignment of the last $n - k + 1$ ads to positions k to n. In the following we define what partial assignments are valid and thus form the domain of payment rule $\{p^{(i)}\}_{i \in [n]}$.

Definition 2.1 (valid ordering). *A partial assignment of $n - k + 1$ ads $\Pi^{(k)}$ is valid if and only if for any $i, j \in \{k, \ldots, n\}$ such that π_i and π_j are TF ads and $i < j$, we have $v(\pi_i) \geq v(\pi_j)$.*

Recall that Myerson's characterization of truthful mechanisms asks a designer to give a monotone allocation rule and the payments are then uniquely derived from the area above the curve. A monotone allocation rule, when seen from the payment perspective, implies monotone marginal increases of the payments (see Figure 2.1).

Definition 2.2 (marginal operator $\nabla^{(i,j)}$). *For two positions $i, j \in [n]$ where $i < j$, the*

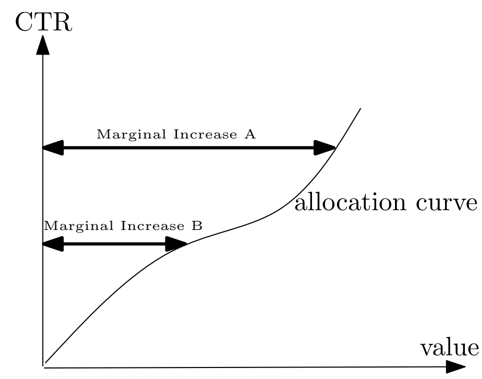

Figure 2.1: The area above the allocation curve of a winner is his payment. The two arrows show the marginal increase of the payment at different points. If the allocation curve is monotone the marginal increases are also monotone.

marginal increase of payment rule $\mathcal{P}$ *for a valid assignment* Π *is*

$$\nabla^{(i,j)}\mathcal{P}(\Pi) = \frac{p^{(i)}(\Pi^{(i+1)}) - p^{(j)}(\Pi^{(j+1)})}{f_i - f_j}.$$

Now we are ready to specify the first property which the payment rule $\mathcal{P}$ should satisfy.

Definition 2.3 (Exact Marginal Increase (EMI))**.** *The payment rule* $\mathcal{P}$ *satisfies EMI, if for any valid assignment* Π *and position* $i \in [n-1]$*, if* π_{i+1} *is a TF ad then*

$$\nabla^{(i,i+1)}\mathcal{P}(\Pi) = v(\pi_{i+1}).$$

The intuition behind the EMI requirement (Definition 2.3) is that, since TF ads are shown in the order of their bid, the minimum bid TF ad π_i needed to get shown above the TF ad π_{i+1} is exactly $v(\pi_{i+1})$. Thus, the marginal payment he should make for being in slot i as opposed to $i+1$ is exactly this minimum bid.

Note that in our setting there are non-TF ads that can be placed between TF ads. Therefore, we need to generalize EMI in order to make sure that the payments for TF ads remain incentive compatible.

Definition 2.4 (Minimum Marginal Increase (MMI))**.** *The payment rule $\mathcal{P}$ satisfies MMI if for any valid assignment Π, position $i \in [n]$ such that π_i is a TF ad, and position $j \in \{1, \ldots, i-1\}$ we have*

$$\nabla^{(j,i)}\mathcal{P}(\Pi) \geq v(\pi_i).$$

Now we give our algorithm to derive the final allocation given payment rule $\mathcal{P}$ which satisfies EMI and MMI. Our algorithm is very simple and intuitive. It starts filling from position n all the way up to position 1. The TF ads get assigned to positions in the increasing order of their valuations. Let us assume that the current TF ad to be assigned to a position is π. Our algorithm tries to fill all the remaining positions by non-TF ads which are not yet assigned, choosing the ads sequentially such that the payment of the next position is minimized. Then, our algorithm puts ad π in the position i for which its profit is maximized. The algorithm then takes the next TF ad and restarts from position $i-1$. The formal description of our Allocation Algorithm (AA) is given in Algorithm 1.

Algorithm 1 Allocation Algorithm (AA)

input : n ads $\{1, \ldots, n\}$ and payment rule $\{p^{(i)}\}_{i \in n}$.

output: assignment $(\pi_1, \ldots, \pi_n)$ of ads to positions and their payments.

$T \leftarrow$ `Extract-TF-ads`$(\{1, \ldots, n\})$ $N \leftarrow$ `Extract-NonTF-ads`$(\{1, \ldots, n\})$

$\ell \leftarrow n$

while $|T| > 0$ **do**

  Let $\pi \in T$ be a TF ad with minimum value

  Remove π from T

  `Fill-With-NonTF-ads` (see Algorithm 2)

  $i \leftarrow \arg\max_{j \in \{\ell - |N|, \ell - |N| + 1, \ldots, \ell\}} f_j \cdot v(\pi) - p^{(j)}(\pi_{j+1}, \ldots, \pi_n)$

  $\pi_i \leftarrow \pi$

  $N \leftarrow N - \{\pi_{i+1}, \ldots, \pi_\ell\}$

  $\ell \leftarrow i - 1$

end

`Fill-With-NonTF-ads` (see Algorithm 2)

Set the payment of π_i to be $p^{(i)}(\pi_{i+1}, \ldots, \pi_n)$

Algorithm 2 Fill with Non-TF Ads

```
/* Sub-procedure Fill-With-NonTF-ads provisionally
   assigns all the remaining non-TF ads to next available
   positions. At each step it selects a non-TF ad which
   makes the next payment as small as possible.          */
14 Fill-With-NonTF-ads:
   begin
15     N' ← N
16     for i ← ℓ downto ℓ − |N| − 1 do
17         π_i ← arg min_{π∈N'} p^{(i−1)}(π, π_{i+1} ..., π_n)
18         N' ← N' − {π_i}
19     end
20 end
```

Description of the algorithms. Set T contains all the TF ads which are not yet assigned. Similarly set N contains all the non-TF ads which are not yet assigned. In Line 2 we initialize the value of ℓ which keeps the index of current position to be filled. In Line 4 we select a TF ad with minimum value in order to assign it to a position. In Line 6 we provisionally fill the next $|N|$ positions with non-TF ads. In Lines 7 and 8 we find and assign a position with the best profit for TF ad π. In Line 9 we remove all the non-TF ads which are assigned permanently (appear after the position i) from N. In Line 12 we fill the remaining positions by the rest of non-TF ads. Finally at Line 13 we set the payments of allocated ads according to $p^{(i)}$.

Note that at Lines 4 and 17, we might have multiple valid choices, in which case we break the ties by the choosing the ad with the smallest index. The only other instance where a tie can happen is at Line 7, when we select the largest feasible j.

The following theorem shows that mechanisms derived from our framework are incentive compatible for TF ads.

Theorem 2.1. *Given a set of payment functions $\mathcal{P} = \{p^{(i)}\}_{i\in[n]}$ that satisfy EMI and MMI, the mechanism derived from applying AA to $\mathcal{P}$ is incentive compatible for TF ads.*

Proof. Let $\mathcal{M}$ be the resulting mechanism after applying AA to set of payment functions. Observe that $\mathcal{M}$ assigns TF ads to positions in the increasing order of their value, *i.e.*, the larger the value of a TF ad is, the higher position he receives. This follows from Line 4 of AA.

In order to prove incentive compatibility of mechanism $\mathcal{M}$, we show that an arbitrary TF ad gets the best utility when he bids his true valuation. Assume that θ is an arbitrary type profile, $\mathcal{M}$ outputs assignment $\Pi = (\pi_1, \ldots, \pi_n)$ for θ, and π_k is a TF ad. We show that utility of π_k does not increase if he bids v' considering three cases: (1) he is considered in the same iteration of Algorithm 1, (2) he is considered in a later iteration, and (3) he is considered in an earlier one.

Case (1): Since he is considered in the same iteration, all that changes is that Line 7 optimizes with respect to v' rather than v, giving him a weakly worse position. Thus, he does not benefit.

Case (2): Since he is considered in a later iteration, some other TF with value $v'' \geq v$ is considered in his original iteration and assigned to slot k''. By MMI, $\nabla^{(k',k'')}\mathcal{P}(\Pi') \geq v'' \geq v$. Thus, his marginal payment for all the clicks he gets beyond what he would get in slot k'' is at least his value, and he is no better off than he would have been originally taking slot k'', which is a contradiction.

Case (3): Without loss of generality, let π_k be the bidder in the lowest slot (according to π) who can benefit from lowering his bid. Let k'' be the highest slot below k such that $\pi_{k''}$ (with value v'') is TF. By the taxation principle, there is a price that $\pi_{k''}$ faced for every slot at or below k'', and at those prices he preferred k''. π_k could have faced those same prices by bidding $v'' - \epsilon$ for sufficiently small ϵ, and as $v \geq v''$ he too prefers slot k'' among all those options. By EMI, $\nabla^{(k''-1,k'')}\mathcal{P}(\Pi) = v'' \leq v$. Thus, he weakly prefers taking slot $k'' - 1$ to taking slot k''. Since slot $k'' - 1$ was one of his options, he weakly prefers slot k to it, a contradiction. □

Having shown that every mechanism derived from our framework is truthful, it is natural to characterize the class of mechanisms that are implementable with our framework. We show that this class is characterized by three natural axioms and one technical one.

First note that payment functions $\{p^{(i)}\}_{i\in[n]}$ only use the bid and nature of the ads and do not use the identity (index) of the ads to determine payments[3]. This means that mechanisms derived from our framework satisfy anonymity, defined formally below.

Definition 2.5. *[Anonymous Mechanism (AM)] A mechanism ($\mathcal{M} = (x, p)$) with allocation function x and payment function p is anonymous if the following holds. Let θ and θ' be two type profiles that are permutations of each other (i.e. the set of natures and bids are the same but the identities of ads are permuted) and have no ties. Say $\theta = \sigma(\theta')$, for some fixed permutation σ. Then we have $x(\theta) = \sigma(x(\theta'))$ and $p(\theta) = \sigma(p(\theta'))$. For type profiles with ties we require permutations to permute the output, except that the payments and allocations of tied bidders can be exchanged arbitrarily.*

Secondly, note that the payment of the ad assigned to position i is specified by looking only at the ads that are assigned to positions below i. Therefore, mechanisms in our framework also satisfy the following property.

Definition 2.6 (Top Interference Free (TIF))**.** *A mechanism $\mathcal{M} = (x, p)$ satisfies TIF, if when an ad changes its type and gets a better position then the allocation of ads assigned to lower positions remains unaltered. More formally, let $x(\theta)$ be the allocation given by x on type profile $\theta = \{\theta_1, \ldots, \theta_n\}$ and $x(\theta')$ the allocation given by x on type profile θ' where $\theta'_h = \theta_h, \forall h \in [1, \ldots, k-1, k+1, \ldots, n]$ and $\theta'_k \neq \theta_k$. Assume that ad k is in position i with allocation $x(\theta)$ and in position j with allocation $x(\theta')$ such that $j < i$.*

[3]This would appear to rule out current systems that incorporate other information such as click probability and other quality measures into a rank score. However, our results still apply in this more general setting as long as the rank scores are treated anonymously.

Mechanism $\mathcal{M}$ satisfies TIF if the ads assigned to positions $i + 1$ to n are the same in both allocations $x(\theta)$ and $x(\theta')$.

In the following theorem, we prove that our framework can implement all mechanisms that are incentive compatible, anonymous, and top interference free, as well as satisfying an additional technical axiom (one which seems to be satisfied for reasonable mechanisms). Hence, requiring EMI and MMI does not restrict the designer in ways that current standard designs such as VCG and GSP do not.

Theorem 2.2. *A mechanism is derived from our framework if and only if it satisfies IC, AM, TIF, and 2T.*

See Section 2.7 for a definition of the technical axiom (2T), a proof of the theorem, and a discussion of 2T.

2.5 Pricing Functions

In the preceding section we designed a general framework. In this section, we apply it to the desired special case of transitioning from a GSP auction to a VCG auction. To do so, we must decide how to price GSP bidders, since our framework is silent about how they should be charged. We make perhaps the simplest decision, to charge them the same amount a truthful bidder would be charged in the same slot, and show that this has several desirable properties.

We begin by discussing what happens when users can evolve and switch types: that is, they are either TF or GSP bidders, and in each case act as utility maximizers. When the hybrid mechanism is first put into use, all bidders are GSP bidders, but as time goes on some bidders will change to being truthful. Given perfect rationality, this means i's bid changes from $b(i)$ to $v(i)$. We can then show that if GSP bidders begin from the *lowest revenue* Symmetric Nash Equilibrium (SNE) then the revenue and allocations are unaltered, provided a particular hybrid pricing function is employed, irrespective of the order in which users transition from GSP to TF bids. (As SNE rank bidders in decreasing order of bid, in this section we assume that ad i is in slot i.) Formally,

Proposition 2.1. *For any GSP algorithm where GSP bidders bid as in the lowest revenue SNE and truthful bidders bid truthfully, the revenue, allocation and prices paid will be independent of the number and identity of TF and GSP bidders if and only if the payment function satisfies MMI and EMI and further satisfies*

$$p^{(i-1)}(\Pi^{(i)}) = \begin{cases} p^{(i)}(\Pi^{(i+1)}) + v(i)(f_{i-1} - f_i) & \textit{if } \theta_i = (TF, v(i)) \\ b(i)f_{i-1} & \textit{if } \theta_i = (GSP, b(i)) \end{cases} \tag{2.1}$$

Proof. Both directions of the proof follow almost directly from the definitions of a lowest SNE, EMI and MMI. In order for payments of TF and GSP bidders to be identical, the payment functions $p^{(i)}(\Pi^{(i+1)})$ must be the same regardless of whether i is TF or GSP, and independent of the mix of bidder types in $\Pi^{(i+1)}$. Consequently no GSP or TF bidder wants to change bid or position, since by the definition of an SNE

$$\left(v(i) - p^{(i)}(\Pi^{i+1})\right) f_i \geq \left(v(i) - p^{(j)}(\Pi^{j+1})\right) f_j \quad \text{for all } i, j. \tag{2.2}$$

By standard arguments about the lowest SNE (see, e.g., [Var07, RGKK13]) we in fact have that for all i,

$$b(i) = b(i+1) + v(i)(f_{i-1} - f_i). \tag{2.3}$$

Thus, by induction, the two conditions of (2.1) are in fact equal at the lowest SNE. This gives that the form is necessary and sufficient for prices to coincide, as pricing must be equivalent to the case $\theta_j = (GSP, b(j))$ for all $j \geq i$. As this outcome is equivalent to the outcome of a truthful auction, it follows that EMI and MMI are satisfied as well. □

In the statement of the proposition, we say that it applies to any GSP algorithm. By this we mean that GSP is, strictly speaking, just a payment rule. It can be applied to a variety of rank score allocation rules. As long as the one chosen admits an SNE, the proposition applies. The following corollary results from applying a sufficient condition [RGKK13] for this.

Corollary 2.1. *The result holds true for any GSP ranking function that uses a rank score of the form*

$$y(b,i) = \max\big(g(i)b - h(i), 0\big), \tag{2.4}$$

where g and h are arbitrary non-negative values that can depend on i.

The necessary and sufficient conditions only hold when we start from a *lowest* SNE. For example if we are in another SNE, then moving just one bidder i from GSP to TF will not change the position or prices paid by i or those below i, but potentially changes prices (and hence positions) of bidder(s) above i (since equality in (2.3) need no longer hold for

$i-1$). Hence, we want to construct price functions that satisfy EMI and MMI when other equilibria hold, and for more general non-truthful prices. Specifically, we shall consider two natural examples, where the pricing functions for ad i are the same for truthful and non-truthful i. First,

$$\mathbf{A}: \quad p^{(i-1)}(\Pi^{(i)}) = \max\left(p^{(i)}(\Pi^{(i+1)}) + v_{max}(\Pi^{(i)})(f_{i-1} - f_i), b_{max}(\Pi^{(i)})f_{i-1}\right) \quad (2.5)$$

where

$$v_{max}(\Pi^{(i)}) \stackrel{def}{=} \max\{v(\theta_j) : j \geq i \cap \theta_j = (TF, v(j))\} \quad (2.6)$$

$$b_{max}(\Pi^{(i)}) \stackrel{def}{=} \max\{b(\theta_j) : j \geq i \cap \theta_j = (GSP, b(j))\} \quad (2.7)$$

are the largest TF valuation and GSP bid at or below i, respectively, and

$$\mathbf{B}: \quad p^{(i-1)}(\Pi^{(i)}) = \max\Big(p^{(\arg\max v(\Pi^{(i)}))}(\Pi^{(\arg\max v(\Pi^{(i)})+1)}) + v_{max}(\Pi^{(i)})(f_{i-1} - f_{\arg\max v(\Pi^{(i)})}), b_{max}(\Pi^{(i)})f_{i-1}\Big) \quad (2.8)$$

where

$$\arg\max v(\Pi^{(i)}) \stackrel{def}{=} \arg\max_j \{v(j) : j \geq i \cap \theta_j = (TF, v(j)\} \quad (2.9)$$

is the identity of the largest TF ad at or below i.

Either of these hybrid auctions is consistent with Proposition 2.1, and so in fact they are identical in this case. It is easy to see, however, that they do differ in other scenarios. In some sense, we can think of these as the two extremes of reasonable prices. In any "reasonable" extension of GSP, an advertiser ought to pay at least the bid of a GSP bidder below him, and **B** is the lowest set of prices consistent with this, EMI, and MMI. At the other extreme, **A** charges the highest prices that are consistent with no GSP bidder paying more than his bid.

Observation 1. *When setting prices according to (2.5), GSP bidders pay exactly their bid when an indifferent TF bidder is put below them.*

Proof. Let b be the GSP bidder assigned to slot i and v be the value of the truthful bidder assigned to slot $i+1$ and indifferent between that and slot i. Then

$$f_i(v-b) = f_{i+1}v - p^{(i+1)}(\Pi^{(i+1)}).$$

Rewriting shows that the GSP bidder pays exactly his bid. □

Since the main concern with switching to VCG is the loss of revenue, we use **A** in our simulations. However, this does come at a cost in welfare relative to **B**, since it will tend to put higher TF bidders below lower GSP bidders more often.

2.6 Simulation Results

We saw in Proposition 2.1 that if bidders always play the lowest SNE, we can perfectly identify which bidders have adjusted to the new truthful auction, and that adjustment consists of instantly switching to the advertiser's true value, then there would be no effect on efficiency or revenue from switching to the hybrid auction. Of course, none of these assumptions are realistic. In this section, we discuss a variety of simulations that analyze the practical effects of the hybrid auction in more realistic scenarios.

2.6.1 Simulation Setup

We base our simulations on a *non-random* sample of Bing data on 3984 auctions. It is a filtered subset of a larger random sample that ensures the auctions are "interesting." In particular, we wanted thick auctions (with at least 12 participants), and with other properties such that techniques for inferring true values from GSP bids could give reasonable answers. The metrics have been normalized. Nevertheless, we believe the sample is representative enough to allow a meaningful exploration of our approach.

We restrict each auction to the top 12 participants, and only actually run an auction for the top three slots. In order to run our simulations we need to have an estimate of true valuation of GSP ads. One estimate of true valuations is to assume that GSP ads have played the minimum symmetric Nash equilibrium and invert their bids to their valuations. In this case, we would essentially be baking in the first of the assumptions from Proposition 2.1, so unsurprisingly the transition would happen without any changes as the allocation and payment of ads remain identical at each point of time.

Instead, we use the stochastic formulation from Pin and Key [PK11]. This approach derives the valuations under the hypothesis that each advertiser chooses her bid to maximize her expected net utility under the assumption that she faces a stationary bid distribution -the bids of other advertisers are from a distribution which does not change over time. In our calculations we assume that the CTRs are known, with the opposing bid distribution estimated from the opponents' empirical bid distributions.

In our simulations we run four different mechanisms.

- **GSP.** The first mechanism is GSP run on the original set of bids when no updates have happened. This represents the current state of the world and serves as a benchmark to which the other approaches can be compared.
- **VCG-V.** The second mechanism is VCG run on the final set of true valuations when all the ads have updated their bids. This represents the ideal end state when all bidders have transitioned to being truthful. It also serves as a sanity check on the reasonableness of our value estimation (i.e. it should display similar performance to GSP).
- **HYBRID.** The third mechanism is the one derived from our framework, using pricing rule **A** described in (2.5).
- **VCG-B.** Finally the last mechanism is VCG run on the current set of bids when some ads have updated their bids and some have not. This is the obvious alternative strategy for transitioning: simply transition directly to VCG and wait for bidders to catch up.

2.6.2 Perfectly Rational Bidders

The first set of simulations we run assumes that bidders are perfectly rational and that they know that any bid change will result in them being classified as truthful. Such a bidder would directly update his bid to his true valuation without bidding any other intermediate amounts since the mechanism is incentive compatible. In this simulation we try to measure the possible revenue loss caused by difference in time patterns of renewing bids for advertisers. In this simulation we assume that at each time step one randomly selected ad decides to change its bid from the GSP value to the true value.

Figure 2.2 shows the normalized average revenue, welfare, and click yield[4] for different mechanisms during the transition. The estimated revenue from ultimately running VCG (i.e. VCG-V) is close to GSP, which is consistent with the reasonableness of our value estimation procedure. Immediately switching to VCG (curve VCG-B) results in a significant revenue drop, which is steadily recovered as more advertisers update their bids. In contrast, there is a more modest revenue drop under the hybrid mechanism (since bidders are not always following the lowest SNE). In particular, revenue always dominates directly switching, substantially so in the initial time steps. In both the estimated welfare and click yield there are no significant differences between VCG-B and Hybrid auction. The observation that welfare and click yield do not differ much in the Hybrid auction and in the VCG-B strengthens the importance of the revenue improvement that the former has over the latter because it is not coming at the cost of other important metrics. Note also that, in the worst case, the drop in welfare is less than 1.5% and the drop in click yield is less than 0.3% from the optimal case. Thus, we focus on revenue in our subsequent simulations.

[4]Here click yield includes both bad clicks, the ones users do not stay in the clicked website for enough time, and good clicks.

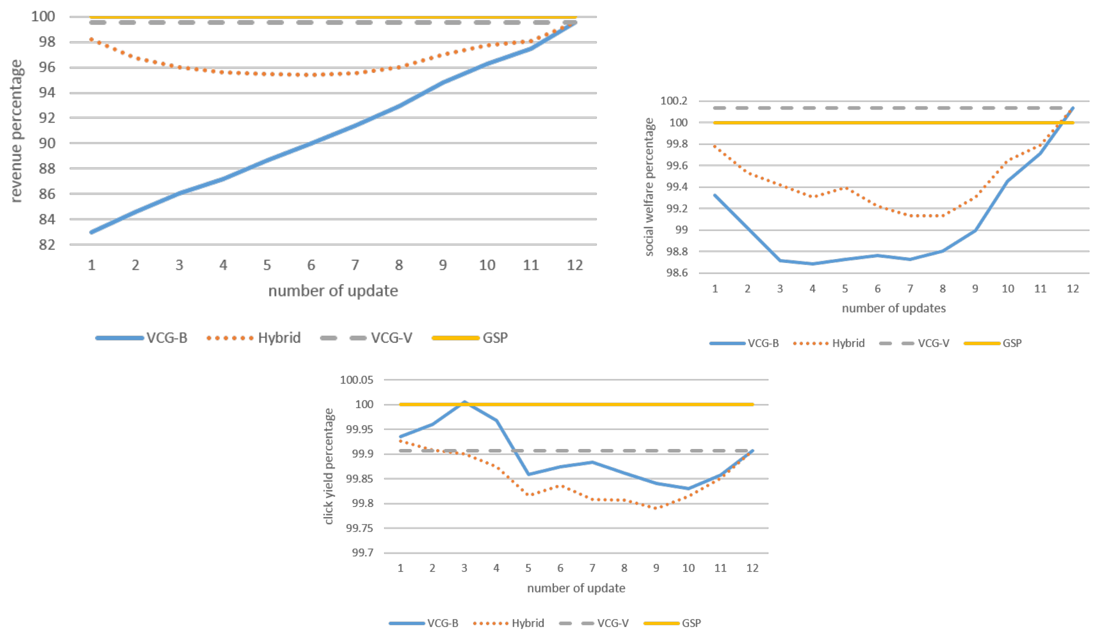

Figure 2.2: Perfectly Rational Bidders

2.6.3 Cautious Bidders

The second set of simulations relaxes the idea that bidders are willing to immediately jump to their true value, no matter how large a bid increase this implies. Instead we parameterize them with a triple (p, q, i). At each time step, bidders decide randomly whether to update their bid, doing so with probability p. If their consumer surplus decreased in the last step (i.e. because the bid changes of others changed their slot or increased their price) they update with a higher probability q. This allows us to model advertisers who are attentive only when needed. Finally, when they update they increase their bid by a percentage i until they reach their true value. Bidders are treated as truthful as soon as they change their bid.

Different parameterizations lead to somewhat different pictures, but all share the same general trends as in Figure 2.2. Note however, that relative to that figure their x-axis has been compressed, since it now takes significantly more than 12 rounds for all bidders to fully adapt. Figure 2.3, with parameters (0.3,0.6,0.1) shows that these cautious updates hurt the performance of the hybrid relative to a direct switch to VCG. There is still a benefit for the first 8 rounds, but then essentially all bidders are classified as truthful, so performance is the same as if we had switched directly. This points to the need for more subtle methods of determining how to treat bidders than just on the basis of whether or not they have updated their bids since the process began. Figure 2.4 shows that the benefits persist longer if we have bidders who are lax about updating (unless something bad happens) with parameters (0.1,0.9,0.1). Larger values of i (not shown) lead to more of the benefits of the hybrid approach being maintained, since the period when a bidder is not GSP but not yet truthful is shortened.

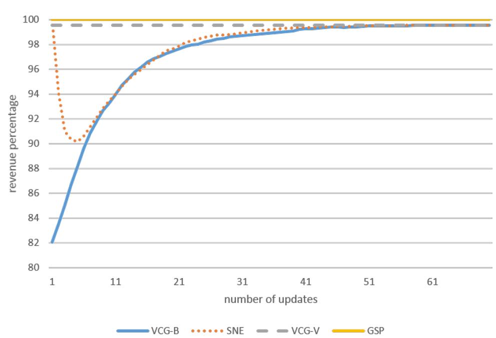

Figure 2.3: Cautious Bidders

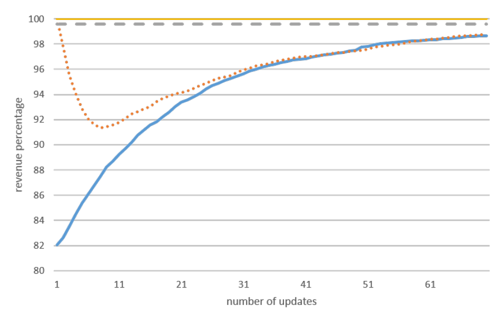

Figure 2.4: Cautious Bidders, Lax Updates

2.6.4 Identification Strategies

In the preceding simulations, bidders are homogeneous. Thus, in some sense trying to identify which bidders are truthful is meaningless: bidders differ only in the number of times they have increased their bid. In this simulation, we consider instead a heterogeneous population of bidders. Some update more frequently and in smaller increments, representing advertisers who use automated tools to optimize their campaigns, while others update only occasionally (and in practice there are many advertisers who go long periods of time between updates). Intuitively, we can likely tolerate simply treating the former as truthful immediately since they will rapidly adapt to the new setting, while we can update the latter as soon as they first change their bid, since that first change represents a large step towards their true value. Advertisers in between are more problematic.

Here we model each bidder's behavior by an update factor x where $x \in [0..1]$ is a real number. The bidder with update factor x updates her bid with probability of x and increases her bid by factor $1 + \frac{b}{x}$ where base increment percentage b is the parameter of the simulation. Note that the larger x is the less is the increment factor. At the start we assign the update factors to bidders by selecting a random number from $[0..1]$.

Our identification strategy is as follows. We update a bidder to a TF bidder with probability $IncPer/BasePer$ where $IncPer$ is the total percentage of increases that the bidder has made since the start of the transition and $BasePer$ is a parameter of the simulation. Moreover we update a bidder to a TF bidder with probability $IncNum/BaseNum$ where $IncNum$ is the total number of increases that the bidder has made since the start of the transition and $BaseNum$ is a parameter of the simulation. Therefore, we describe our simulations by a triple $(b, BasePer, BaseNum)$. Note that the smaller $BasePer$ or $BaseNum$ are, the more likely is to tag an ad as a TF bidder.

Figure 2.5 shows the performance of the hybrid auction in three different simulations: $(0.2, 0.1, 5)$, $(0.2, 0.2, 10)$, and $(0.2, 0.3, 30)$. In all the simulations the base increment percentage b is 0.2. In simulation $(0.2, 0.1, 5)$ values $BasePer$ and $BaseNum$ are the smallest among others, hence, we tend to tag a bidder as a TF ad more likely. As one expects, this reduces the overall revenue as it is more likely that a bidder who is not fully adapted (does not bid his true valuation) get tagged as a TF bidder. In simulation $(0.2, 0.3, 30)$ values $BasePer$ and $BaseNum$ are the largest among others, hence, we are less likely to tag a bidder as a TF ad and hence get more revenue. Although these simulations suggest that we should less likely tag bidders as TF, but we note that in order to give incentive to bidders to update at all, we cannot decrease this probability too much. Thus, perhaps the biggest message is that there is a tension between preserving revenue and giving bidders an incentive to update their bids.

2.7 Obtainable Mechanisms by the Payment Framework

In this section we prove Theorem 2.2. Before we begin, we observe that our requirements ensure that the mechanism allocates truthful ads in increasing order of value.

Observation 2. *If $\mathcal{M}$ satisfies IC, AM, and TIF then it allocates TF ads in increasing order.*

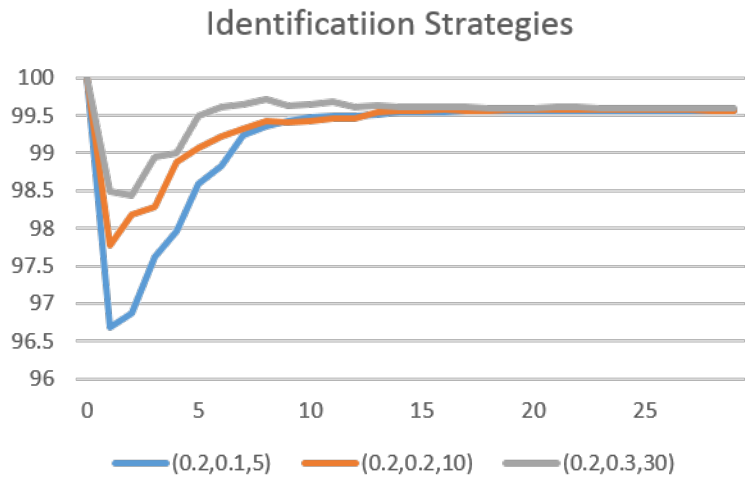

Figure 2.5: Identification Strategies

Proof. Suppose for contradiction that $v_1 < v_2$ but ad (TF, v_1) got a better slot than ad (TF, v_2). By IC, if we replace (TF, v_1) by another copy of (TF, v_2), its slot can only improve. By TIF, this means that the slot of the original copy is unchanged. But by AM, if the original ads had been permuted the ad with value v_1 could raise its bid to v_2 and receive a worse slot, contradicting IC. $\square$

We begin with sufficiency. Let $\mathcal{M}' = (x', p')$ be a mechanism with allocation function x' and payment function p' which satisfies IC, AM, TIF, and 2T (see below). We use $\mathcal{M}'$ to propose payment rule $\mathcal{P}$ such that mechanism $\mathcal{M}$ derived from applying $\mathcal{P}$ to TIFP framework is equivalent to $\mathcal{M}'$. In order to deal with technicalities of ties, we assume that in the case of ties $\mathcal{M}'$ use the same tie breaking rule as our framework.

Let $\Pi^{(k)} = (\pi_k, \ldots, \pi_n)$ for $k \in \{3, \ldots, n\}$ be the assignment of $\mathcal{M}'$ for positions k to n. We say that ad i is less than j with regard to $\Pi^{(k)}$ (show by $i \prec_{\Pi^{(k)}} j$), if there exist a type profile θ such that

- The allocation $x'(\theta)$ is the same as $\Pi^{(k)}$ for positions k to n.
- $i, j \in \theta$.
- Ad i is assigned to position $k-1$ and ad j is assigned to a position better than i $(x'(i) > x'(j))$.

Intuitively, $i \prec_{\Pi^{(k)}} j$ means that fixing $\Pi^{(k)}$ allocation x' prefers to assign i to position k over j. In the following lemma we prove that $\prec_{\Pi^{(k)}}$ is a total order over all the ads which can be assigned to position $k-1$ fixing $\Pi^{(k)}$. It turns out that TIF is almost, but not quite strong enough to prove this lemma. The difficulty is that it has no "bite" when applied to the case of $k = 3$ (i.e. the final 2 slots). That is, ad in slot 2 has no (non-fixed) ads below it (so the definition is vacuous), while an ad in slot 1 that changes type and remains in slot 1 forces the other ad to stay in slot 2 (again making the definition vacuous). Thus, we need a property that ensures the mechanism is well-behaved in this case.

Definition 2.7 (Two Transitive(2T))**.** *$\mathcal{M}$ is two transitive if for all choices of $\Pi^{(3)}$ the relation $\prec_{\Pi^{(3)}}$ is transitive.*

Lemma 2.1. *The relation $\prec_{\Pi^{(k)}}$ is antisymmetric (if $i \prec_{\Pi^{(k)}} j$ then $j \not\prec_{\Pi^{(k)}} i$), total, and transitive.*

Proof. We prove the lemma by contradiction. To show antisymmetry, let θ be a type profile for which $i \prec_{\Pi^{(k)}} j$ and θ' be a type profile for which $j \prec_{\Pi^{(k)}} i$. Consider a sequence of type profiles that are intermediate between θ' and θ in the sense that the transition from one profile to the next results from changing the type of a single ad from its value in θ' to its value in θ. Let $\theta' = \theta_1, \theta_2, \dots, \theta_{a-1}, \theta_a = \theta$ be the sequence of type profiles.

We show that this sequence maintains the following invariant: the ad in slot $k-1$ is not i and has already changed its type. Clearly this is true for θ', since j is in slot $k-1$ in $x(\theta')$ and its type will never change. Suppose it is true for θ_b, and let ad e be the ad that changes type between θ_b and θ_{b+1}. By our invariant, e is not in slot $k-1$ of $x(\theta_b)$. Thus, by TIF, if the ad in slot $k-1$ changes from $x(\theta_b)$ to $x(\theta_{b+1})$ it must be that in $x(\theta_{b+1})$ the ad in slot $k-1$ is in fact e, which is not i and has already changed its type. This shows that i is not is slot $k-1$ of $x(\theta)$, contradicting our assumption.

To see that the relation is total, take some θ and i and j. Create θ' by replacing all ads other than i and j shown in a slot above k in $x(\theta)$ with a copy of either i or j. Thus, by an inductive argument that shows this does not change the allocation below slot $k-1$, some copy of i or j must be allocated to slot $k-1$.

Finally, transitivity follows via a similar construction. If $k > 3$, simply transform θ to θ' by replacing all ads above slot k with copies of one of the relevant i, j, or ℓ. If $k = 3$, transitivity is by assumption (i.e. 2T). □

Now we are ready to specify how we build set of payment functions $\mathcal{P} = \{p^{(i)}\}_{i\in[n]}$ from $\mathcal{M}'$. Let $\Pi^{(k)}$ be a valid assignment of ads to positions k to n. By Lemma 2.1 we have a total ordering of candidate ads for position $k-1$, so set $p^{(k-1)}(\Pi^{(k)})$ to be the infimum of TF ads among those candidates.

Now we prove that the mechanism ($\mathcal{M}$) derived from our framework using payment rule $\mathcal{P}$ always gives the same allocation and payments as $\mathcal{M}'$ which finishes the proof of this sufficiency.

Lemma 2.2. *Mechanism $\mathcal{M} = (x, p)$ is the same as $\mathcal{M}' = (x', p')$.*

Proof. We need only verify that $\mathcal{M}$ always gives the same allocation as $\mathcal{M}'$ as the fact that the payments are the same (at least up to a constant) then follows via revenue equivalence. (The constant can be matched by changing p to p'' to include this constant shift.)

Now we prove by contradiction that the allocation functions x and x' are the same. Let θ be a type profile for which $x(\theta) \neq x(\theta')$. Let $\Pi^{(k)} = (\pi_k, \dots, \pi_n)$ be the largest common suffix of $x(\theta)$ and $x(\theta')$ and assume for now that $p^{(k-1)}(\Pi^{(k)})$ is finite. Let e be the ad assigned to position $k-1$ in $x(\theta)$ and e' be the ad assigned to position $k-1$ in $x'(\theta)$. Note that

$$e' \prec_{\Pi^{(k)}} e \tag{2.10}$$

since $x'(\theta)$ assigned positions $k-1$ to e'.

Because $x(\theta)$ assigns position $k-1$ to e as opposed to e' this means that

$$p^{(k-2)}(e, \pi_k, \dots, \pi_n) < p^{(k-2)}(e', \pi_k, \dots, \pi_n)$$

(Recall Lines **??** and 17 of algorithm AA). This means that there exists a TF ad x with valuation $\nabla^{(k-2,k-1)}\mathcal{P}((e, \pi_k, \dots, \pi_n)) < v(x) < \nabla^{(k-2,k-1)}\mathcal{P}((e', \pi_k, \dots, \pi_n))$. Now if we replace the rest of ads with TF bidders with valuation $v(x)$ then they appear before e but after e'. This contradicts with Equation 2.10 and the fact that $\prec_{\Pi^{(k)}}$ is a total order for any $\Pi^{(k)}$.

Now we deal with the case where $p^{(k-1)}(\Pi^{(k)})$ is infinite. Intuitively, this is the case where only non-truthful ads can be shown before the suffix $\Pi^{(k)}$. Since prices are all infinite, we need a way for the algorithm to match the order that $\mathcal{M}'$ chooses. We do this by allowing prices of the form (∞, a), where a is the type of a non-TF ad. The total order $\prec_{\Pi^{(k)}}$ then gives a well defined notion of the lowest price as the one whose a is lowest according to that ordering. With this enlarged set of prices, the proof proceeds as before. □

Finally, the necessary part is easy to prove. Let $\mathcal{M}$ be a mechanism derived from using TIFP framework. The AM and TIF properties follow by the fact that in algorithm AA (Figure 1) when assigning the next TF ad, AA uses only its value and neither its index nor the value of higher TF ads. The IC property of $\mathcal{M}$ is the result of Theorem 2.1. 2T follows from the greedy nature of the allocation.

2.7.1 Discussion of Two Transitivity (2T)

Two transitivity is a technical assumption. The intuition is that we require the mechanism to be well-behaved when considering the top two slots, which TIF is not strong enough to enforce. If all non-TF ads are of the same nature, a sufficient condition is that $\mathcal{M}'$ is monotone for non-TF ads (MN). That is, if a non-TF ad raises its bid, it gets a (weakly) better slot.

Lemma 2.3. *IF $\mathcal{M}$ satisfies IC, AM, TIF, and MN and all non-TF ads have the same nature then it satisfies 2T*

Proof. By AM and IC/MN, $\prec_{\Pi(3)}$ is transitive if all 3 ads are TF or non-TF respectively. Thus, WLOG let 2 be TF and 1 be non-TF (replace IC by MN below if only 1 is TF). There are 3 cases.

Case 1: $\boldsymbol{N \prec V_1 \prec V_2}$. By IC, $N \prec V_2$ (otherwise a truthful ad V_1 could raise its bid and go from slot 1 to slot 2).

Case 2: $\boldsymbol{V_1 \prec N \prec V_2}$. By AM+IC, $V_1 \prec V_2$ (otherwise a truthful ad V_2 could raise its bid and go from slot 1 to slot 2 when facing N).

Case 3: $\boldsymbol{V_1 \prec V_2 \prec N}$. By IC, $V_1 \prec N$ (otherwise a truthful ad V_1 could raise its bid and go from slot 1 to slot 2). □

Such a nice sufficient condition is not obvious if there is more than 1 nature of non-TF ad. Non-degenerate examples still appear to satisfy 2T, but we do not know of a less technical way to explain the way in which they are non-degenerate. To see why, consider an example with 2 slots. If the bidders have at least one truthful ad, this becomes a form of second price auction. However, when there are two non-truthful ads of different natures, nothing obviously constrains the rule for determining the order in which they are shown in a way that corresponds to enforcing transitivity. This same example shows why we require MN in Lemma 2.3. Without it we would be equally at a loss in this situation.

2.8 Concluding Remarks

We have described a framework for a hybrid mechanism design where bidders may be either truthful or not, and compete for goods where there is an intrinsic ordering by worth amongst the goods. If we can identify each group in this dichotomy, then we have shown that the fundamental building block is a payment function, which maps the bids for lower value goods to the payment for a good, and which needs to satisfy two properties. The two properties, EMI and MMI relate solely to the bids of truthful agents, and place constraints on the discrete derivative of the payment function; the payment for non-TF bidders may be arbitrary.

We have given details of a bottom-up allocation procedure, which when used with an EMI and MMI payment function gives an Incentive Compatible (IC) mechanism, and hence gives incentive for bidders to change type to truthful. If in addition, the mechanism is Top Interference Free, TIF , then this characterization is both necessary and sufficient for IC anonymous mechanisms which satisfy an additional technical axiom. Any mechanism derived from a bottom-up procedure, such as standard GSP or VCG, are all examples of TIF mechanisms.

The motivation for the work of this chapter is sponsored search, where ad-slots are auctioned off. But as the previous paragraph suggests, our framework does not have to be tied to this setting. Our particular starting point was wanting to provide a pathway for migrating from one non-truthful mechanism (GSP) to a truthful mechanism whilst mitigating the revenue loss that would occur if there was a switch to a truthful mechanism but bidders did not immediately update (increase) their bids to their true valuation. We have shown empirically that revenue loss can indeed by mitigated if users do switch between being non-truthful with a GSP bid or truthful and bid their true valuation, but that the situation is more nuanced if users update partially (i.e. are not fully truthful).

The latter illustrates the difficulties involved in ascertaining a users type, and points to the need for further research into the best way to classify users, and the trade-offs involved in mis-classification. The example pricing functions we gave appear to work well for vanilla sponsored-search auctions, and this was our intended target. For use in richer settings (e.g. where there are richer ad variants), the pricing functions can be straightforwardly adapted for bottom-up or TIF settings, but would need additional work to adapt them to non-TIF settings.

CHAPTER 3

Deterministic Revenue Monotone Mechanism Design

3.1 Introduction

Fueled by the growth of Internet and advancements in online advertising techniques, today more and more online firms rely on advertising revenue for their business. Some of these firms include news agencies, media outlets, search engines, social and professional networks, etc. Much of this online advertising business is moving to what's called *programmatic* buying where an advertiser bids for each single impression, sometimes in real-time, depending on how he values the ad opportunity. This work is motivated by the need of a desired property in the auction mechanisms that are used in these bid-based advertising systems.

A standard mechanism for most auction scenarios is the famous Vickrey-Clarke-

Groves (VCG) mechanism. VCG is *incentive-compatible* (IC) and maximizes *social welfare* (see [NRTV07, Chapter 9] for a detailed explanation). Incentive-compatibility guarantees that the best response for each advertiser is to report its true valuation. This makes the mechanism transparent and removes the load from the advertisers to calculate the best response. *Social welfare* is the sum of the valuations of the winners. This value is treated as a proxy for how much all the participants gain from the transaction. What makes VCG mechanism versatile is that it reduces the mechanism design problem into an optimization problem for any scenario.

Even though this versatility of VCG mechanism makes it a popular choice mechanism, however, it doesn't satisfy an important property, namely, that of *revenue-monotonicity*. Revenue-monotonicity says that if one increases the bid values or add new bidders, the total revenue should not go down. To see that VCG is not revenue-monotone, consider a simple example of two items and three bidders (A, B, and C). Say bidder A wants only the first item, and has a bid of 2. Similarly bidder B wants only the second item, and has a bid of 2. Bidder C wants both the items or nothing, and has a bid of 2. Now if only bidders A and B participate in the auction, then VCG gives a revenue of 2, however, if all the three bidders participate, then the revenue goes down to 0.

This lack of revenue-monotonicity (which has been noted several times in the literature) is one of the serious practical drawbacks of the celebrated VCG mechanism. To think of it, an online firm that depends on advertising revenue puts significant resources in its sales efforts to attract more bidders as the general belief is that more bidders imply more competition which should lead to higher prices. Now to tell this firm that their revenue can go down if they get more bidders can be strategically very confusing for them. To see this from another perspective, say in a search engine firm, there is a team which makes a UI change that increases the click-through probability (CTR) of the search ads. These changes are thought of as good changes in the firm as they increase the effective bid of the bidders (the effective bid of a bidder in search advertising is a function of its cost-per-click bid and the CTR of its ad). Now if after making the change, the revenue goes down, what was supposed to be a good change may seem like a bad change. The point we are trying to make is that there are many teams in a firm, and for these teams to function properly, it is important that the auction mechanisms satisfy *revenue-monotonicity*.

In this chapter, with a focus on auctions arising in advertising scenarios, we seek to understand mechanisms that satisfy this additional property of revenue monotonicity (RM). It is well known that for various settings (including ours), no mechanism can satisfy both IC and RM properties while attaining optimal social welfare. In fact it is known that one cannot even hope to get Pareto-optimality in social welfare while attaining both IC and RM [RCLB11]. Thus to overcome this bottleneck and develop an understanding of RM mechanisms, we relax the requirement of attaining full social welfare, and define the notion of *price of revenue-monotonicity* (PoRM). Price of revenue-monotonicity of an IC and RM mechanism M is the ratio of optimal social welfare to the social welfare attained by the mechanism M. The goal is to design mechanisms that satisfy IC and RM properties and at the same time achieve low price of revenue-monotonicity.

We study two different advertising settings in this chapter. The first setting we study is the *image-text* auction. In image-text auction there is a special box designated for advertising in a publisher's website which can be filled by either k text-ads or a single image-ad. The second setting is the *video-pod* auction where an advertising break of a certain duration in a video content can be filled with multiple video ads of possibly different durations.

We note that revenue-monotonicity is an *across-instance* constraint as it requires total revenue to behave in a certain manner across different instances, where a single instance is defined by fixing the *type* of the buyers. Note that incentive-compatibility is also an across-instance constraint. A lot of research effort has gone into understanding incentive-compatibility, which has resulted in useful tools for designing incentive-compatible mechanisms. Surprisingly, hardly any work has gone into understanding and building tools for designing mechanisms which satisfy the desired property of revenue-monotonicity. We believe that understanding revenue-monotonicity will shed new fundamental insights into the design of mechanisms for many practical scenarios.

3.1.1 Related Work

Ausubel and Milgrom [AM02] show that VCG satisfies RM if bidders' valuations satisfy *bidder-submodularity*. Bidders' valuations satisfy bidders submodularity if and only if for any bidder i and any two sets of bidders S, S' with $S \subseteq S'$ we have $\text{WF}(S \cup \{i\}) - \text{WF}(S) \geq \text{WF}(S' \cup \{i\}) - \text{WF}(S')$, where $\text{WF}(S)$ is the maximum social welfare achievable using only S. Note that this is a general tool one can use to design revenue monotone mechanisms - restrict the range of the possible allocations such that we get bidder-submodularity when we run VCG on this range. However, we can show that this general tool is not so powerful by showing that for our auction scenarios, it is not possible to get a mechanism with PoRM better than $\Omega(k)$ by using the above tool.

Ausubel and Milgrom [AM02] also show that bidder-submodularity is guaranteed when the goods are substitutes, i.e., the valuation function of each bidder is submodular over the sets of goods. However, for many practical scenarios, including ours, the valuation function of the bidders is not submodular. Ausubel and Milgrom [AM02] design mechanisms which select allocations that are in the core of the exchange economy for combinatorial auctions. Here an allocation is in the core if there is no coalition of bidders and the seller to trade with each other in a way which is preferred by all the members of the coalition to the allocation. Day and Milgrom [DM08] show that core-selecting mechanisms that choose a core allocation which minimizes the seller's revenue satisfy RM given bidders follow so called *best-response truncation strategy*. Therefore the core selecting mechanism designed by [AM02] satisfies RM if the participants play such best-response strategy; although this mechanism is not incentive-compatible.

Rastegari *et al.* [RCLB11] prove that no mechanism for general combinatorial auctions which satisfies IC and RM can achieve weakly maximal social welfare. An allocation is weakly maximal if it cannot be modified to make at least one participant better off without hurting anyone else. In another work [RCLB09] they design a randomized mechanism for combinatorial auctions which achieves weak maximality and expected revenue monotonicity.

Another related work to this chapter is around the characterization of mechanisms that achieve the IC property. The classic result of Roberts [Rob79] states that affine maximizers are the only social choice functions that can be implemented using IC mechanisms when bidders have unrestricted quasi-linear valuations. Subsequent works study the restricted cases [Roc87, LMN03, BCL+06, SY05].

There is also an extensive body of research around designing mechanisms with good bounds on the revenue. Myerson [Mye81] designs a mechanism which achieves the optimal expected revenue in the single parameter Bayesian setting. Goldbert *et al.* consider optimizing revenue in prior-free settings (see [NRTV07] for a survey on this).

3.1.2 Results and Overview of Techniques

As mentioned earlier, we study two settings: 1) image-text auction, and 2) video-pod auction. Both these settings can be described using the following abstract model. Say there is a seller selling k identical items to n participants/buyers. Participant i wants either d_i items or nothing, and has a valuation of v_i if gets d_i items or 0 otherwise. Demand d_i is assumed to be public knowledge, and valuation v_i is assumed to be the private information of the participant i. We want to design a mechanism that is incentive-compatible, individually-rational (IR), revenue-monotone, and maximizes social welfare.

For the image-text auction, the demand $d_i \in \{1, k\}$, i.e., each participant wants either 1 item (text ads) or k items (image ad). For the video-pod auction, an item corresponds to a unit time interval (say one second), and the demand d_i could be any number between 1 and k, i.e., $d_i \in [k]$.

The first result of this chapter is the following theorem.

Theorem 3.1. *We design a deterministic mechanism for image-text auction (*MITA*) which satisfies Individual Rationality (IR), IC, and RM with* PoRM *of at most* $\sum_{i=1}^{k} \frac{1}{i} \simeq \ln(k)$*, i.e., the ratio of* MITA*'s welfare over the optimal welfare is at most* $\ln(k)$.

The proof of Theorem 3.1 appears in Section 3.3. We outline our mechanism over here: Let $v_1 \geq \ldots \geq v_{n_1}$ be the valuations of text-participants and V_1 be the maximum valuation of the image-participants. If $\max_{j \in [k]} j \cdot v_j$ is less than V_1, MITA gives all the items to the image-participant who has valuation V_1, otherwise MITA picks the highest j^* text-participants as the winners where j^* is the maximum number in $[k]$ such that $j^* \cdot v_{j^*} \geq V_1$. Note that the j that maximizes $j \cdot v_j$ might be less than the $j*$ which is the largest j such that $j \cdot v_j \geq V_1$. Also note that MITA sometimes picks less than k text ads as the winner (even if there are k or more text ads). VCG always picks the maximum number of text ads (if it decides to allocate the slot to text ads); this is one of the reasons why VCG fails to satisfy RM. When we allow lesser number of text ads to be declared as winners, intuitively, this increases the competition which boosts the revenue and thus helps in achieving RM. Although this comes with a loss in social welfare.

We can also show that the above mechanism achieves the optimal PoRM for the image-text auction by proving a matching lower bound. We show that a mechanism that satisfies IR, IC, RM, and two additional mild assumptions of Anonymity (AM) and Independence of Irrelevant Alternatives (IIA) cannot achieve a PoRM better than $\sum_{i=1}^{k} \frac{1}{i}$. Anonymity means that the auction mechanism doesn't depend on the identities of the participants (a formal definition appears in Section 3.5). IIA means that decreasing the bid of a loosing participant shouldn't hurt any winner. Note that our mechanism satisfy both AM and IIA as well. Formally, we prove the following theorem whose proof appears in Section 3.5.

Theorem 3.2. *There is no deterministic mechanism which satisfies IR, IC, RM, AM, and IIA and has* PoRM *less than* $\sum_{i=1}^{k} \frac{1}{i}$.

Finally we prove the following theorem for video-pod auctions.

Theorem 3.3. *We design a Mechanism for Video-pod Auction (*MVPA*) which satisfies IR, IC, and RM with* PoRM *of at most* $(\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)$.

We give the formal proof of Theorem 3.3 in Section 3.4, and outline the mechanism here. MVPA partitions the participants into $(\lfloor \log k \rfloor + 1)$ groups where each group $g \in [\log k]$ contains only the participants whose demands are in the range $[2^{g-1}, 2^g)$. MVPA selects winners only from one group. We round up the size of each participant in group g

to 2^g, thus we can have at most $\frac{k}{2^g}$ number of winners from the group g. Let $v_1^{(g)} \geq \ldots \geq v_p^{(g)}$ be the sorted valuations of all the participants in group g. We define the Max Possible Revenue of Group g ($\textsc{mprg}(g)$) to be

$$\textsc{mprg}(g) = \max_{j \in [k/2^g]} j \cdot v_j^{(g)}.$$

As the name of $\textsc{mprg}(g)$ suggests, its value captures the maximum revenue we can truthfully obtain from group g without violating revenue-monotonicity. Let g^* be the group with the highest $\textsc{mprg}$ value and group g' be the group whose $\textsc{mprg}$ is second highest. The set of winners are the first j participants from group g^* where j is the largest number in $[k/2^g]$ such that $j \cdot v_j^{(g^*)}$ is greater than or equal $\textsc{mprg}(g')$. We show that PoRM of $\textsc{mvpa}$ is $(\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)$.

3.2 Preliminaries

Let $N = \{1, \ldots, n\}$ be the set of all participants, and k be the number of identical items. We denote the type of participant i by $\theta_i = (d_i, v_i) \in [k] \times \mathbb{R}^+$, where d_i is the number of items participant i demands and v_i is her valuation for getting d_i items. Note that the valuation of player i for getting less than d_i items is 0. Now in the image-text auction, participants have demand of either 1 or k. In the video-pod auction participants can have arbitrary demands in $\{1, \ldots, k\}$. Lets denote the set of all possible types $[k] \times \mathbb{R}^+$ by Θ and the set of all type profiles of n participants by $\Theta^n = \underbrace{\Theta \times \ldots \times \Theta}_{n}$.

A deterministic mechanism $\mathcal{M}$ consists of an allocation rule $x : \Theta^n \to 2^n$ which maps each type profile to a subset of participants as the winners, and payment rule p :

$\Theta^n \to (\mathbb{R}^+)^n$ which maps each type profile to the payments of each participant.

Let $\theta = (\theta_1, \theta_2, \ldots, \theta_n) \in \Theta^n$ be a specific type profile. Also let $\mathcal{A}_\theta$ be the set of all feasible solutions, i.e.,

$$\mathcal{A}_\theta = \left\{ S \subseteq N \middle| \sum_{i \in S} d_i \leq k \right\}.$$

For each feasible solution $A \in \mathcal{A}_\theta$, the social welfare of A (denoted by $\textsc{wf}(A)$) is equal to $\sum_{\theta_i \in A} v_i$. To evaluate the social welfare of a mechanism $\mathcal{M}$ on a type profile θ, we compare the welfare of its solution to the optimal solution.

Definition 3.1. *The welfare ratio of mechanism $\mathcal{M} = (x, p)$ on type profile $\theta \in \Theta^n$ (denoted by $\textsc{wfr}(\mathcal{M}, \theta)$) is the following.*

$$\textsc{wfr}(\mathcal{M}, \theta) = \frac{\max_{A \in \mathcal{A}_\theta} \textsc{wf}(A)}{\textsc{wf}(x(\theta))}$$

To capture the worst-case loss in social welfare across all type profiles, we define the notion of *price of revenue-monotonicity*.

Definition 3.2. *The Price of Revenue Monotonicity of a mechanism $\mathcal{M}$ (denoted by* PoRM$(\mathcal{M})$*) is defined as follows:*

$$\text{PoRM}(\mathcal{M}) = \max_{\theta \in \Theta^n} \textsc{wfr}(\mathcal{M}, \theta)$$

The desired goal is to design mechanisms which have low PoRM value, where the best possible value is 1.

Note that since we are interested in mechanisms with bounded PoRM, we restrict ourselves to mechanisms that satisfy consumer sovereignty. Consumer sovereignty says that any participant can be a winner as long as he bids high enough.

Now we will define a weakly monotone allocation rule which is used in the characterization of deterministic IC mechanisms. Let function $x_i : \Theta^n \to \{0, 1\}$ be the restriction of function x to participant i. Here $x_i(.)$ is one if participant i is a winner and zero otherwise.

Definition 3.3. *We call allocation function x is weakly monotone if for any type profile $\theta \in \Theta^n$ and any participant $i \in [n]$ with demand d_i, function $x_i\left((d_i, v_i), \theta_{-i}\right)$ is a non-decreasing function in v_i.*

Note that if a deterministic mechanism $\mathcal{M}$ satisfies consumer sovereignty and has a weakly monotone allocation function then function $x_i((d_i, v_i), \theta_{-i})$ is a single step function. The value at which the function $x_i((d_i, v_i), \theta_{-i})$ jumps from zero to one, i.e. the smallest value at which the participant i becomes a winner, is called *critical value*.

Definition 3.4. *Let $\mathcal{M} = (x, p)$ be a deterministic mechanism that satisfy consumer sovereignty and has a weakly monotone allocation function, the critical value of participant i in type profile θ is $v_i^* = sup\{v_i | x_i\left((d_i, v_i), \theta_{-i}\right) = 0\}$.*

The following lemma characterizes deterministic IC mechanisms (first given by [Mye81]). We provide a proof sketch for the sake of completeness (for a complete proof see *e.g.* [NRTV07]).

Lemma 3.1. *Let $\mathcal{M} = (x, p)$ be a mechanism which satisfies IR. Mechanism $\mathcal{M}$ is truthful (IC) if and only if the followings hold.*

1. *x is weakly monotone.*

2. *If participant i is a winner then its payment is its critical value (v_i^*).*

Proof. First we prove that if $\mathcal{M}$ is truthful then it satisfies both conditions 1 and 2. We prove the first condition by contradiction. If x is not monotone then there exist participant i, type profile θ, and two values $v_i^{(1)} > v_i^{(2)}$ such that i wins in type profile $\left((d_i, v_i^{(2)}), \theta_{-i}\right)$ but loses in type profile $\left((d_i, v_i^{(1)}), \theta_{-i}\right)$. This makes incentive for participant i to lie for type profile $\left((d_i, v_i^{(1)}), \theta_{-i}\right)$ and announce its valuation as $v_i^{(2)}$.

Consider an arbitrary participant i who is a winner, now we prove that the payment of participant i is its critical value. Assume for contradiction that mechanism $\mathcal{M}$ charges participant i amount c_i where $c_i < v_i^*$ in a type profile $((d_i, v_i), \theta_{-i})$. In this case, if participant i had type $(d_i, \hat{v}_i)$ where $c_i < \hat{v}_i < v_i^*$ then i is not a winner in $((d_i, \hat{v}_i), \theta_{-i})$ as v_i^* is the critical value. Therefore, if the real type of participant i is $(d_i, \hat{v}_i)$, she has incentive to lie her type as (d_i, v_i), become a winner, and pay c_i. Hence, the payment cannot be less than v_i^*. Now suppose that there exists value v_i for which mechanism $\mathcal{M}$ charges i amount c_i which is more than v_i^*. In this case, if participant i had type $(d_i, \hat{v}_i)$ where $v_i^* < \hat{v}_i < c_i$ then i is still a winner (as v_i^* is the critical value) and pays at most $\hat{v}_i$ (as $\mathcal{M}$ satisfies IR). Therefore, she has an incentive to lie her type as $(d_i, \hat{v}_i)$, become a winner, and pay at most $\hat{v}_i$. Hence, the payment cannot be more than v_i^* for any winning valuation v_i.

For the other direction, it is easy to check that any IR mechanism that satisfies conditions 1 and 2 is truthful. □

3.3 Image-text Auctions

In this section we give our *Mechanism for Image-Text Auction* (MITA) which satisfies IR, IC, RM, and PoRM(MITA) $\leq \ln k$. Recall that in the image-text auction we have k identical items to sell and there are two groups of participants: the ones who want all the k items which we call *image-participants*; and the ones who want only one item which we refer to as *text-participants*. As a result there are also two possible types of outcome: MITA gives all the items to an image-participant; or it gives an item to each member of a subset of the text-participants.

We start with explaining why VCG fails to satisfy RM and how we address this issue in MITA. Consider the type profile where we have one image-participant with type $(k, 1)$ and one text-participant with type $(1, 1)$. In this case either of the participants can be the winner. The payment of the winner in VCG is her critical value which is one. However if we add one more text-participant with the same type $(1, 1)$, the two text-participants win and each of them pay zero. The reason for the payment drop is that VCG always selects k winners from the text-participants. This decreases the critical value of each text participant as the valuation of the other text-participants helps her to win against image-participants. In our mechanism we overcome this issue by not guaranteeing that the maximal number of text-participants can win an item. In other words, in our mechanism it is possible that less than k text-participants win an item even if there are more than k text-participants. This way, intuitively, even if the number of text-participants increase, it potentially creates more competition and hence increases the payments.

Let θ be an arbitrary type profile where there are n_1 text-participants with types $(1, v_1), \ldots, (1, v_{n_1})$ and n_2 image-participants with types $(k, V_1), \ldots, (k, V_{n_2})$. We define mechanism MITA $= (x^{\text{MITA}}, p^{\text{MITA}})$ by giving allocation function x^{MITA} which is weakly monotone. Given the allocation function, we obtain payment function p^{MITA} using the critical values defined in Lemma 3.1 which makes the mechanism truthful.

Allocation rule of MITA. Without loss of generality we assume that $v_1 \geq v_2 \geq \ldots \geq v_{n_1}$ and $V_1 \geq V_2 \geq \ldots \geq V_{n_2}$. Also, we assume that $n_1 \geq k$, if not, we add fake text-participants with value 0. For each $j \in [k]$ we consider value $j \cdot v_j$. Let candidate set C_θ contains all the values $j \in [k]$ such that $j \cdot v_j$ is greater than or equal to V_1, *i.e.*, $C_\theta = \{j \in [k] | j \cdot v_j \geq V_1\}$. If C_θ is empty, the image-participant with type (k, V_1) wins. If C_θ is non-empty then let j^* be the maximum member of C_θ, *i.e.*, $j^* = \max_{j \in C_\theta} j$. In this case the first j^* text-participants win.

Observation 3. *Allocation function x^{MITA} is weakly monotone.*

Proof. Recall from Definition 3.3, in order to prove that x^{MITA} is weakly monotone, we have to show that for any participant $i \in [n]$ with demand d_i, function $x_i\left((d_i, v_i), \theta_{-i}\right)$ is a non-decreasing function in v_i.

If i is an image-participant then i wins if its valuation is larger than $\max(W, \max_{j \in [k]} j \cdot v_j)$ where W is the largest valuation of the image-participants in θ_{-i}. Moreover, bidder i loses for any value smaller than or equal to $\max(W, \max_{j \in [k]} j \cdot v_j)$. Therefore x_i is weakly monotone.

If i is a text-participant then let $v'_1 \geq v'_2 \geq \ldots$ be the sorted valuations of the text-participants and V_1 be the largest valuation of image-participants in θ_{-i}. Let t be the smallest value such that there exist $j \in [k-1]$ where $v'_{j+1} \leq t \leq v'_j$ and $(j+1) \cdot t$ is greater than or equal to V_1. If the valuation of bidder i is larger than or equal to t then she wins since $(j+1) \cdot t \geq V_1$ otherwise she does not win since t is the smallest value for which there exist $j \in [k-1]$ such that $(j+1) \cdot t \geq V_1$. Therefore x_i is weakly monotone. $\square$

In the following lemma we obtain the critical value (or truthful payments) of the winners in x^{MITA} using Lemma 3.1. The lemma also gives an intuition to why we select j^* text-participants to win, which is the maximum j such that $j \cdot v_j \geq V_1$.

Lemma 3.2. *If C_θ, where $C_\theta = \{j \in [k] | j \cdot v_j \geq V_1\}$, is empty then the first image-participant wins all the items with critical value $\max(V_2, \max_{j\in[k]} j \cdot v_j)$. If C_θ is not empty, the first j^* text-participants win the items where $j^* = max_{j \in C_\theta} j$ and all of them have critical value $\max(v_{k+1}, \frac{V_1}{j^*})$.*

Proof. We find the critical value (Definition 3.4) of a winner by showing that if she has any valuation larger than the critical value she wins and for any valuation less than the critical value she doesn't.

If C_θ is empty then the first image-participant (with type (k, V_1)) wins all the items. As long as V_1 is larger than $\max(V_2, \max_{j\in[k]} j \cdot v_j)$ participant (k, V_1) wins. If V_1 is less than $\max(V_2, \max_{j\in[k]} j \cdot v_j)$ then she loses to the image-participant (k, V_2) if $\max(V_2, \max_{j\in[k]} j \cdot v_j) = V_2$, or loses to the text-participants if $\max(V_2, \max_{j\in[k]} j \cdot v_j) =$

$\max_{j \in [k]} j \cdot v_j$. This means that the critical value of the first image-participant is $\max(V_2, \max_{j \in [k]} j \cdot v_j)$ if she is the winner.

If C_θ is non-empty then the first j^* text-participants win. Let $i \in [j^*]$ be an arbitrary winner. First we observe that for any valuation v'_i greater than or equal to $\max(v_{k+1}, \frac{V_1}{j^*})$, participant i remains as a winner in type profile $\theta' = ((1, v'_i), \theta_{-i})$. This is because for any such change in valuation of participant i number j^* remains in set $C_{\theta'}$. Moreover, this change does not add any new number j' to $C_{\theta'}$ such that $j' > j^*$ because the valuations of the text-participants with index greater than j^* are not changed in θ'.

In order to prove that for any valuation v'_i less than critical value $\max(v_{k+1}, \frac{V_1}{j^*})$, participant i is not a winner we consider two cases: (A) when the critical value is equal to $\frac{V_1}{j^*}$, and (B) when the critical value is equal to v_{k+1}.

Case (A): We prove this case by contradiction. Let v'_i be a valuation less than $\frac{V_1}{j^*}$ for which participant i is in the set of winners in type profile $\theta' = ((1, v'_i), \theta_{-i})$. Because v'_i is less than $\frac{V_1}{j^*}$, the number of winners which contains participant i cannot be less than or equal to j^* in type profile θ'. Let $j' \in [k]$ which is greater than j^* be the number of winners in θ'. This means that there are at least j' participants whose valuation is larger than $\frac{V_1}{j'}$ in θ'. Note that all the valuations in θ is the same as θ' except v_i which is decreased to v'_i, therefore, there are also at least j' participants whose valuation is larger than $\frac{V_1}{j'}$ in θ and hence j' is in set C_θ. This contradicts with the fact that j^* is the largest member of C_θ.

case (B): In case (B) we have $\max(v_{k+1}, \frac{V_1}{j^*}) = v_{k+1}$ which implies that $k \cdot v_{k+1}$

is larger than V_1 as $j^* \in [k]$. Therefore Case (B) can only happen when $j^* = k$. Now consider participant i decreases its valuation to value v'_i that is less than v_{k+1}, then it cannot be a winner as there are k other participants whose valuations are more than v'_i while we have only k items. $\square$

The payment function of MITA is set to the critical values of the winners as specified in Lemma 3.2 which by using Observation 3 and Lemma 3.1 implies MITA satisfies IC. Moreover, as the payments are always less than the participants' bid IR property of MITA follows. Finally in the following lemma we show that MITA is revenue monotone.

Lemma 3.3. *Let θ' be the type profile obtained by either increasing the valuation of a participant or adding a new participant to the type profile θ, then we have* $\text{REVENUE}(\text{MITA}, \theta') \geq \text{REVENUE}(\text{MITA}, \theta)$.

Proof. Let $v_1 \geq v_2 \geq \ldots$ be the valuations of text-participants and $V_1 \geq V_2 \geq \ldots$ be the valuations of image-participants in θ. Similarly let $v'_1 \geq v'_2 \geq \ldots$ be the valuations of text-participants and $V'_1 \geq V'_2 \geq \ldots$ be the valuations of image-participants in θ'. Note that for any i we have $v_i \leq v'_i$ and $V_i \leq V'_i$ as we have one more participant or a higher valuation in θ'. Let x be the new added participant or the participant which has higher valuation in θ'.

We prove this lemma by considering the value of $\text{REVENUE}(\text{MITA}, \theta)$ for the case when text-participants win and the case when an image-participant wins. If an image-participant wins then it means that $V_1 > \max_{j \in [k]} j \cdot v_j$ and she pays $\max(V_2, \max_{j \in [k]} j \cdot v_j)$ which is the total revenue.

If text-participants win then it means $V_1 \le \max_{j\in[k]} j \cdot v_j$ and there are j^* winners where each of them pays $\max(v_{k+1}, \frac{V_1}{j^*})$. If $\max(v_{k+1}, \frac{V_1}{j^*}) = \frac{V_1}{j^*}$, then the total revenue is V_1. If $\max(v_{k+1}, \frac{V_1}{j^*}) = v_{k+1}$, it implies that $k \cdot v_{k+1}$ is larger than V_1. Remember that $C_\theta = \{j \in [k] | j \cdot v_j \ge V_1\}$ and $j^* = \max_{j\in C_\theta} j$ therefore $j^* = k$ and hence the total payment of the winners is $k \cdot v_{k+1}$.

In summary the total revenue for type profile θ is the following.

$$\text{REVENUE}(\text{MITA}, \theta) =$$
$$\begin{cases} \max(V_2, \max_{j\in[k]} j \cdot v_j) & V_1 > \max_{j\in[k]} j \cdot v_j \ (A) \\ \max(V_1, k \cdot v_{k+1}) & V_1 \le \max_{j\in[k]} j \cdot v_j \ (B) \end{cases}$$

Similarly the total revenue for type profile θ' is the following.

$$\text{REVENUE}(\text{MITA}, \theta') =$$
$$\begin{cases} \max(V'_2, \max_{j\in[k]} j \cdot v'_j) & V'_1 > \max_{j\in[k]} j \cdot v'_j \ (A) \\ \max(V'_1, k \cdot v'_{k+1}) & V'_1 \le \max_{j\in[k]} j \cdot v'_j \ (B) \end{cases}$$

Note that because for any i we have $v_i \le v'_i$ and $V_i \le V'_i$ the following inequalities are straight forward.

$$V_1 \le V'_1 \tag{3.1}$$

$$V_2 \le V'_2 \tag{3.2}$$

$$\max_{j\in[k]} j \cdot v_j \le \max_{j\in[k]} j \cdot v'_j \tag{3.3}$$

$$k \cdot v_{k+1} \le k \cdot v'_{k+1} \tag{3.4}$$

If both $\textsc{revenue}(\textsc{mita}, \theta)$ and $\textsc{revenue}(\textsc{mita}, \theta')$ take their value from Case (A) then the proof of the lemma follows from Equations (3.2) and (3.3). Similarly if both $\textsc{revenue}(\textsc{mita}, \theta)$ and $\textsc{revenue}(\textsc{mita}, \theta')$ take their value from Case (B) then the proof of the lemma follows from Equations (3.1) and (3.4).

If $\textsc{revenue}(\textsc{mita}, \theta)$ takes its value from Case (A) and $\textsc{revenue}(\textsc{mita}, \theta')$ takes from Case (B) then it means that participant x is a text-participant which causes $\max_{j\in[k]} j \cdot v'_j$ to be larger than V'_1. The following proves the theorem for this case.

$$
\begin{aligned}
&\textsc{revenue}(\textsc{mita}, \theta) \\
&= \max(V_2, \max_{j\in[k]} j \cdot v_j) \\
&< V_1 && \textsc{revenue}(\textsc{mita}, \theta) \text{ takes its value from Case } (A) \\
&= V'_1 && \text{participant } x \text{ is a text-participant} \\
&\leq \max(V'_1, k \cdot v'_{k+1}) \\
&= \textsc{revenue}(\textsc{mita}, \theta')
\end{aligned}
$$

If $\textsc{revenue}(\textsc{mita}, \theta)$ takes its value from Case (B) and $\textsc{revenue}(\textsc{mita}, \theta')$ takes from Case (A) then it means that participant x is an image-participant. The following proves the theorem for this case.

$$\begin{aligned}
&\textsc{revenue}(\textsc{mita}, \theta) \\
&= \max(V_1, k \cdot v_{k+1}) \\
&< \max_{j \in [k]} j \cdot v_j && \text{\textsc{revenue}(\textsc{mita}, θ) takes its value from Case (B) and the fact that $v_k \geq v_{k+1}$} \\
&= \max_{j \in [k]} j \cdot v'_j && \text{x is an image-participant} \\
&\leq \max(V'_2, \max_{j \in [k]} j \cdot v'_j) \\
&= \textsc{revenue}(\textsc{mita}, \theta')
\end{aligned}$$

□

In the above we proved that MITA satisfies IR, IC, and RM. In the following theorem we bound the PoRM of MITA and finish this section.

Theorem 3.4. PoRM(MITA) $\leq \ln k$.

Proof. Let A be the set of winner(s) which realizes the maximum social welfare in type profile θ. If A contains only one image-participant with valuation V_1 then we also have $V_1 \geq \max_{j \in [k]} j \cdot v_j$. Mechanism MITA also selects an image-participant with the same valuation if $V_1 > \max_{j \in [k]} j \cdot v_j$ and hence PoRM(MITA) is 1. Otherwise we have $V_1 = \max_{j \in [k]} j \cdot v_j$ where MITA selects a set of text-participants which overall gives social welfare V_1 and hence again the PoRM(MITA) is 1.

Now we consider the case when A contains text-participants. By adding enough dummy participants with value zero, and without loss of generality, we assume that set A contains the first k text-participants with highest valuations $v_1 \geq v_2 \geq \ldots \geq v_k$. Mechanism MITA selects either the first j^* text-participants with highest valuations ($v_1 \geq v_2 \geq \ldots \geq v_{j^*}$) or selects an image-participant with valuation V_1. Remember that j^* is the greatest number in set $C_\theta = \{j | j \in [k] \wedge j \cdot v_j \geq V_1\}$ which implies the following.

$$\forall j' \in \{j^* + 1, \ldots, k\} \quad v_{j'} < \frac{V_1}{j'} \tag{3.5}$$

Note that if MITA selects an image-participant then Equation (3.5) holds for $j^* = 0$.

Now we consider the following two cases to prove the theorem.

If MITA selects an image-participant then we have the following.

$$\begin{aligned} \text{PoRM}(\text{MITA}) &= \frac{\sum_{j \in [k]} v_j}{V_1} \\ &\leq \frac{\sum_{j \in [k]} V_1/j}{V_1} && \text{Equation (3.5)} \\ &\leq \ln k \end{aligned}$$

If $\textsc{mita}$ selects the first j^* text-participants then we have the following.

$$\begin{aligned}
\textsc{PoRM}(\textsc{mita}) &= \frac{\sum_{j\in[k]} v_j}{\sum_{j\in[j^*]} v_j} \\
&\le \frac{\sum_{j\in[j^*]} v_j + \sum_{j=j^*+1}^{k} v_j}{\sum_{j\in[j^*]} v_j} \\
&\le \frac{\sum_{j\in[j^*]} v_j + \sum_{j=j^*+1}^{k} V_1/j}{\sum_{j\in[j^*]} v_j} \\
&\quad \text{Equation (3.5)} \\
&\le \frac{\sum_{j\in[j^*]} v_j + \sum_{j=j^*+1}^{k} (\sum_{j\in[j^*]} v_j)/j}{\sum_{j\in[j^*]} v_j} \\
&\quad \text{because } V_1 \le \sum_{j\in[j^*]} v_j \\
&\le \ln k
\end{aligned}$$

□

3.4 Video-pod Auctions

In this section we design a Mechanism for Video-Pod Auction (MVPA) which satisfies IR, IC, and RM whose PoRM is at most $(\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)$. Note that all the $\log$ functions are in base 2. Let $\theta = ((d_1, v_1), \ldots, (d_n, v_n)) \in \Theta^n$ be an arbitrary type profile of n participants. We define the allocation and payment function of MVPA for this type profile.

Mechanism MVPA partitions the participants into $\lfloor \log k \rfloor + 1$ groups $G^{(1)}, \ldots, G^{(\lfloor \log k \rfloor + 1)}$ where group $G^{(g)}$ contains all the participants whose demand is in the range $[2^{g-1}, 2^g)$. Mechanism MVPA selects winners only from one group $G^{(g)}$.

Definition 3.5. *Let* $M^{(g)}$ *be equal to* $\max(\lfloor \frac{k}{2^g} \rfloor, 1)$ *which is the maximum number of winners* MVPA *selects from group* $G^{(g)}$.

Note that we can select at least $\lfloor \frac{k}{2^g} \rfloor$ winners from $G^{(g)}$ since there are k items and the demand of each participant is at most 2^g. Moreover, from the last group $G^{(\lfloor \log k \rfloor + 1)}$ we can select at least one winner although $\lfloor \frac{k}{2^{(\lfloor \log k \rfloor + 1)}} \rfloor = 0$, since we assume the demand of all the participants are from the set $[k]$.

Let $(d_1^{(g)}, v_1^{(g)}), \ldots, (d_p^{(g)}, v_p^{(g)})$ be the types of all the participants in group g where $p = |G^{(g)}|$. Here by adding enough dummy participants, we assume p is always larger than $M^{(g)}$. Also, without loss of generality we assume $v_1^{(g)} \geq v_2^{(g)} \geq \ldots \geq v_p^{(g)}$. We define the Max Possible Revenue of Group g (MPRG(g)) to be the following.

$$\text{MPRG}(g) = \max_{j \in [M^{(g)}]} j \cdot v_j^{(g)}$$

As the name MPRG suggests, we will see that its value captures the maximum revenue can be truthfully obtained from group g. Let $G^{(g^*)}$ be a group with the maximum MPRG and $G^{(g')}$ be a group with the second maximum MPRG breaking the ties arbitrarily.

The set of winners selected by MVPA is

$$\left\{(d_1^{(g^*)}, v_1^{(g^*)}), \ldots, (d_j^{(g^*)}, v_j^{(g^*)})\right\}$$

where j is the largest number in $[M^{(g^*)}]$ for which $j \cdot v_j^{(g^*)}$ is larger than or equal to MPRG(g'). In other words, the number of winners (j) is the largest number in $[M^{(g^*)}]$ for which $j \cdot v_j^{(g^*)} \geq \text{MPRG}(g')$.

Now we use Lemma 3.1 to show that MVPA is truthful and obtain the payments of winners.

Observation 4. *Allocation function x^{MVPA} is weakly monotone.*

Proof. Note that MVPA sorts the participants according to their valuation and selects the first j participants. Therefore if any participant i increases its valuation it only helps her to enter the winning set. Hence, the observation follows. $\square$

In the rest of this section we drop the group identifier of $M^{(g^*)}$ and simply use M unless it is about another group.

In the following lemma we find the critical value of each winner i which is actually equal to its payment (p_i^{MVPA}).

Lemma 3.4. *Let set of winners $x^{\text{MVPA}}(\theta)$ contains the first j participants with highest valuations from $G^{(g^*)}$ and $v_{M+1}^{(g^*)}$ be the $(M+1)$th highest valuation in group $G^{(g^*)}$ which is zero if it does not exist. Then, the payment of participant i is the following.*

$$p_i^{\text{MVPA}}(\theta) = \begin{cases} \max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)}) & i \in x^{\text{MVPA}}(\theta) \\ 0 & i \notin x^{\text{MVPA}}(\theta) \end{cases}$$

Proof. If participant i is not a winner then its payment is zero. When participant i is a winner then we prove that its payment is equal to its critical value (Definition 3.4). In order to prove that value $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ is the critical value of participant i, we show that for any value larger than $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ participant i still wins and for any value less than it she loses.

Remember that $v_1^{(g^*)} \geq v_2^{(g^*)} \geq \ldots \geq v_p^{(g^*)}$ are the valuations of participants in group $G^{(g^*)}$ and $v_1^{(g^*)}, v_2^{(g^*)}, \ldots, v_j^{(g^*)}$ are the valuations of the winners. Because group $G^{(g^*)}$ is the group with the maximum MPRG, we have $v_j^{(g^*)} \geq \frac{\text{MPRG}(g')}{j}$. As there can be at most M winners from group $G^{(g^*)}$ we have $v_j^{(g^*)} \geq v_{M+1}^{(g^*)}$. Therefore we have

$$v_j^{(g^*)} \geq \max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)}). \tag{3.6}$$

Let participant i with type profile $(d_i^{(g^*)}, v_i^{(g^*)})$ be the ith winner in group g^* where $i \in [j]$. We show that for any valuation greater than or equal to $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ participant i remains in the winning set. Equation (3.6) implies that there are j participants in group $G^{(g^*)}$ whose valuations are larger than $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$. If we decrease the valuation of participant i to $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ we still have j participants in group $G^{(g^*)}$ with valuations at least $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$. Therefore, the value $\text{MPRG}(g^*)$ will be at least $\text{MPRG}(g')$ and group $G^{(g^*)}$ remains the winning group,‘ hence participant i remains in the winning set.

Now we prove that if the valuation of participant i is less than the $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$, she cannot be in the winning set. In order to prove this we consider two cases: (A) when $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ is equal to $v_{M+1}^{(g^*)}$, and (B) when $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ is equal to $\frac{\text{MPRG}(g')}{j}$.

Case (A): If $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)}) = v_{M+1}^{(g^*)}$ and the valuation of participant i is less than $v_{M+1}^{(g^*)}$ then it means that there are M participants who have valuations greater than the valuation of participant i. As there can be at most M winners from group $G^{(g^*)}$, participant i cannot be a winner.

Case (B): We prove this case by contradiction. Suppose $\max(\frac{\text{MPRG}(g')}{j}, v^{(g^*)}_{M+1}) = \frac{\text{MPRG}(g')}{j}$ and $\theta' = \Big((d_i^{(g^*)}, v_i^{(g^*)'}), \theta_{-i}\Big)$ be a type profile in which the valuation of participant i is less than $\frac{\text{MPRG}(g')}{j}$ while she is still winner. Because the valuation of participant i ($v_i^{(g^*)'}$) is less than $\frac{\text{MPRG}(g')}{j}$ and i is in the winning set, in order for $\text{MPRG}(g^*)$ to be larger than $\text{MPRG}(g')$, there has to be more than j winners . Let $j' > j$ be the number of winners in θ'. Having j' winners in θ' and in order for $G^{(g^*)}$ to be the group with the highest MPRG we conclude that there are j' participants with valuation greater than $\frac{\text{MPRG}(g')}{j'}$. Note that the only difference between θ and θ' is that the valuation of participant i is higher in θ. Therefore, there are also at least j' participants with valuation greater than $\frac{\text{MPRG}(g')}{j'}$ in θ. This contradicts with the way we select the number of winners (j) in θ which is the maximum number for which $j \cdot v_j^{(g^*)}$ is larger than $\text{MPRG}(g')$. □

The allocation function x^{MVPA} is weakly monotone (Observation 4) and the payments of the winners are their critical values (Lemma 3.4), therefore by Lemma 3.1 we conclude that MVPA satisfies IC.

In the rest of this section first we prove that MVPA satisfies RM and then bound its PoRM.

Proposition 3.1. *The total revenue of mechanism* MVPA *for type profile* θ *(*$\text{REVENUE}(\text{MVPA}, \theta)$*) is the following.*

$$\text{REVENUE}(\text{MVPA}, \theta) = \max(\text{MPRG}(g'), M \cdot v^{(g^*)}_{M+1})$$

where g' *is a group with the second highest* MPRG.

Proof. From Lemma 3.4 we know that there are j winners and each of them pay $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$. Therefore the sum of payments or the revenue of MVPA is $j \cdot \max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$. The proof of the proposition follows if we show that when $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ is equal to $v_{M+1}^{(g^*)}$ then the number of winners (j) is equal to M.

If $\max(\frac{\text{MPRG}(g')}{j}, v_{M+1}^{(g^*)})$ is equal to $v_{M+1}^{(g^*)}$ then as $v_{M+1}^{(g^*)} \leq v_M^{(g^*)}$, we have $M \cdot v_M^{(g^*)} \geq \frac{\text{MPRG}(g')}{j}$. Remember that j is the maximum number in the set $[M]$ for which $j \cdot v_j^{(g^*)}$ is larger than $\text{MPRG}(g')$. Therefore j is equal to M. □

Lemma 3.5. *Let θ' be the type profile obtained by either adding a new participant or increasing the valuation of a participant in θ. Then,*

$$\text{REVENUE}(\text{MVPA}, \theta') \geq \text{REVENUE}(\text{MVPA}, \theta).$$

Proof. Let x be the new added participant or the participant which has the increased valuation in θ'. Throughout the proof we show MPRG of each group g in type profile θ by $\text{MPRG}_\theta(g)$ and in type profile θ' by $\text{MPRG}_{\theta'}(g)$. Similarly, we show the jth highest valuation of the participants of group g by $v_j^{(g^*,\theta)}$ in type profile θ and by $v_j^{(g^*,\theta')}$ in type profile θ'.

As the jth highest valuation of the participants of each group can only increase by adding participant x, we conclude

$$\forall g, \forall j \quad v_j^{(g,\theta')} \geq v_j^{(g,\theta)}. \tag{3.7}$$

Remember that MPRG_θ of each group g is $\max_{j \in [M^{(g)}]} j \cdot v_j^{(g,\theta)}$ and using Equation (3.7) we get

$$\forall g \quad \text{MPRG}_{\theta'}(g) \geq \text{MPRG}_{\theta}(g). \tag{3.8}$$

In order to prove this lemma we consider two cases: (A) adding participant x does not change the winning group $G^{(g^*)}$, and (B) adding x changes the winning group.

Case (A): Let g'' be a group with the second highest MPRG in θ', it is possible that g' is equal to g''.

$$\begin{aligned}
\text{REVENUE}(\text{MVPA}, \theta') =& \max(\text{MPRG}_{\theta'}(g''), M \cdot v_{M+1}^{(g^*,\theta')}) \\
& \text{Proposition 3.1} \\
\geq& \max(\text{MPRG}_{\theta'}(g'), M \cdot v_{M+1}^{(g^*,\theta')}) \\
& \text{definition of } g'' \\
\geq& \max(\text{MPRG}_{\theta}(g'), M \cdot v_{M+1}^{(g^*,\theta)}) \\
& \text{Equations (3.7) and (3.8)} \\
=& \text{REVENUE}(\text{MVPA}, \theta)
\end{aligned}$$

Case (B): Let $G^{(g'')}$ be a group with the highest MPRG in θ'. We have

$$\text{MPRG}_{\theta}(g^*) \geq \text{MPRG}_{\theta}(g') \tag{3.9}$$

as g^* has the highest and g' has the second highest MPRG in θ.

$$\begin{aligned}
\text{MPRG}_{\theta}(g^*) &\geq M \cdot v_M^{(g^*,\theta)} && \text{As } \text{MPRG}_{\theta}(g^*) = \max_{j \in [M]} j \cdot v_j^{(g^*,\theta)} \\
&\geq M \cdot v_{M+1}^{(g^*,\theta)} && \text{As } v_M^{(g^*,\theta)} \geq v_{M+1}^{(g^*,\theta)}
\end{aligned} \tag{3.10}$$

Let $\hat{g}$ be the group with second highest MPRG in θ'. Because g^* is no longer the winning group in θ' it can be a candidate for the group with the second highest MPRG in θ' and hence we have the following.

$$\text{MPRG}_{\theta}(g^*) \leq \text{MPRG}_{\theta'}(g^*) \leq \text{MPRG}_{\theta'}(\hat{g}) \tag{3.11}$$

The following equations conclude the proof of this case.

$$\begin{aligned}
\text{REVENUE}(\text{MVPA}, \theta) &= \max(\text{MPRG}_{\theta}(g'), M \cdot v_{M+1}^{(g^*,\theta)}) \\
&\leq \text{MPRG}_{\theta}(g^*) \\
&\quad \text{by Equations (3.9) and (3.10)} \\
&\leq \text{MPRG}_{\theta'}(\hat{g}) \\
&\quad \text{by Equation (3.11)} \\
&\leq \max(\text{MPRG}_{\theta'}(\hat{g}), M^{(g'')} \cdot v_{M^{(g'')}+1}^{(g'',\theta)}) \\
&= \text{REVENUE}(\text{MVPA}, \theta')
\end{aligned}$$

□

The following lemma which bounds PoRM of MVPA finishes this section.

Theorem 3.5. $\text{PoRM}(\text{MITA}) \leq (\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)$

Proof. Let $\text{WF}(g)$ to be the maximum social welfare achievable if we select the winners only from group $G^{(g)}$. Let A be a set of winner(s) which realizes the maximum welfare in type profile θ. Note that as there are $\lfloor \log k \rfloor + 1$ groups, one group ($\hat{g}$) has a subset of participants from A whose social welfare is at least $\frac{\text{WF}(A)}{\lfloor \log k \rfloor + 1}$ and hence the following.

$$\text{WF}(\hat{g}) \geq \frac{\text{WF}(A)}{\lfloor \log k \rfloor + 1} \tag{3.12}$$

Now we prove the following claim about $\text{MPRG}(\hat{g})$.

Claim 3.1. $\text{MPRG}(\hat{g}) \geq \frac{\text{WF}(\hat{g})}{2 + \ln k}$

Proof. Let B be the set of participants from group $G^{(\hat{g})}$ which give the maximum social welfare. Because the demands of all the participants of $G^{(\hat{g})}$ is in range $[2^{\hat{g}-1}, 2^{\hat{g}})$, size of B is at most $\lfloor k/2^{\hat{g}-1} \rfloor$. Remember from Definition 3.5 that $M^{(\hat{g})} = \max(\lfloor k/2^{(\hat{g})} \rfloor, 1)$ is the maximum number of winners that MVPA potentially selects from group $G^{(\hat{g})}$. Therefore, we have $|B| \leq 2 \cdot M^{(\hat{g})} + 1$.

Throughout the proof, we drop the superscript from $M^{(\hat{g})}$ and simply refer to it as M.

Let $v_1 \geq v_2 \geq \ldots \geq v_{2 \cdot M+1}$ be the valuations of the participants in B; if B has less than $2 \cdot M + 1$ participants we add enough dummy participants with valuations zero. Remember that $\text{MPRG}(\hat{g}) = \max_{j \in [M]} j \cdot v_j^{(\hat{g})}$ where M is at least 1 (see Definition 3.5) which implies

$$v_i \leq \frac{\text{MPRG}(\hat{g})}{i} \quad \forall i \in [M] \tag{3.13}$$

The following equations conclude the proof of the claim.

$$\begin{aligned}
\textsc{wf}(\hat{g}) &= \sum_{i=1}^{2\cdot M+1} v_i \\
&= \sum_{i=1}^{M} v_i + \sum_{i=M+1}^{2\cdot M+1} v_i \\
&\leq \sum_{i=1}^{M} v_i + \sum_{i=M+1}^{2\cdot M+1} v_M \\
&\quad \text{replacing } v_i \text{ with } v_M \text{ for } i > M \\
&\leq \sum_{i=1}^{M} \frac{\text{MPRG}(\hat{g})}{i} + \sum_{i=M+1}^{2\cdot M+1} \frac{\text{MPRG}(\hat{g})}{M} \\
&\quad \text{by Equation (3.13)} \\
&\leq (2 + \ln k)\textsc{mprg}(\hat{g})
\end{aligned}$$

□

Remember $G^{(g^*)}$ is the group with maximum MPRG value. Let j be the number for which MPRG(g^*) is equal to $j \cdot v_j^{(g^*)}$. Allocation function x^{MVPA} selects the first j^* participants from group $G^{(g^*)}$ where j^* is the maximum number for which $j^* \cdot v_{j^*}^{(g^*)}$ is larger than MPRG(g'). Therefore we can conclude that $j \leq j^*$ and hence

$$\textsc{wf}(x^{\text{MVPA}}(\theta)) \geq \textsc{mprg}(g^*). \tag{3.14}$$

The following equations conclude the proof of the theorem.

$$\begin{aligned}
\text{WF}(x^{\text{MVPA}}(\theta)) &\geq \text{MPRG}(g^*) \\
&\quad \text{by Equation (3.14)} \\
&\geq \text{MPRG}(\hat{g}) \\
&\quad G^{(g^*)} \text{ has the highest MPRG} \\
&\geq \frac{\text{WF}(\hat{g})}{2 + \ln k} \\
&\quad \text{Observation 9} \\
&\geq \frac{\text{WF}(A)}{(\lfloor \log k \rfloor + 1) \cdot (2 + \ln k)} \\
&\quad \text{by Equation (3.12)}
\end{aligned}$$

□

3.5 Lower Bound

In this section we prove Theorem 3.2. As mentioned earlier we need two additional mild assumptions of *anonymity* and *independence of irrelevant alternatives* (which we define below) on the class of mechanisms for which we prove our lower bound.

Definition 3.6. *A mechanism ($\mathcal{M} = (x, p)$) is anonymous (AM) if the following holds: Suppose $\theta_1, \theta_2 \in \Theta^n$ are two type profiles which are permutations of each other (i.e. the set of type profiles are same just that the identities of participants to whom those types belongs are different). Say $\theta_2 = \pi(\theta_1)$. Also say $x(\theta_1) = S_1$ and $x(\theta_2) = S_2$. Then $S_2 = \pi(S_1)$.*

Definition 3.7. *Let $\theta \in \Theta^n$ be an arbitrary type profile and $i \in N$ be an arbitrary participant with type $\theta_i = (d_i, v_i)$. A mechanism ($\mathcal{M} = (x, p)$) satisfies Independence of Irrelevant Alternatives (IIA) if we decrease the bid of a losing participant, say participant i, to $\hat{v}_i < v_i$ then the new set of winners is a super set of the previous one, i.e., $x(\theta) \subseteq x((d_i, \hat{v}_i), \theta_{-i})$. In other words, decreasing the bid of a losing participant does not hurt any winner.*

The proof outline of Theorem 3.2 is the following. Let $\mathcal{M}^* = (x^*, p^*)$ be a mechanism which satisfies all the five properties and has the optimal PoRM *opt* (*i.e.*, *opt* $=$ PoRM$(\mathcal{M}^*)$). We study the behavior of $\mathcal{M}^*$ in a few type profiles. Let ϵ be an arbitrary small positive real value. First we show that when there are only two participants with types $(k, 1)$ and $(k, 1+\epsilon)$, $\mathcal{M}^*$ gives all the k items to the participant with type $(k, 1+\epsilon)$. The revenue of $\mathcal{M}^*$ from these two participants is 1. Then, we add k more participants to create type profile $\theta = \big((1, 1-\epsilon), (1, \frac{1}{2}-\epsilon), \ldots, (1, \frac{1}{k}-\epsilon), (k, 1), (k, 1+\epsilon)\big)$. The RM property requires $\mathcal{M}^*$ to make at least the same revenue for θ. From this constraint we are able to show that $\mathcal{M}^*$ assigns all the items to participant $k+2$ with type $(k, 1+\epsilon)$ and hence gets social welfare $1+\epsilon$. Note that the maximum social welfare happens when the set of winners is $\{1, \ldots, k\}$ which implies WFR$(\mathcal{M}^*, \theta) \geq \sum_{i=1}^{k} \frac{1}{i} - k \cdot \epsilon$ (see Definition 3.1). Because PoRM$(\mathcal{M}^*) \geq$ WFR$(\mathcal{M}^*, \theta)$ for any $\theta \in \Theta^n$ we conclude that *opt* $\geq \sum_{i=1}^{k} \frac{1}{i}$.

First we study the behavior of $\mathcal{M}^*$ when we have only two participants with types $(k, 1)$ and $(k, 1+\epsilon)$.

Lemma 3.6. *Mechanism $\mathcal{M}^*$ in type profile $((k,1),(k,1+\epsilon))$ gives all k items to the second participant and make one unit of revenue, i.e., $x^*((k,1),(k,1+\epsilon)) = \{2\}$ and $p^*((k,1),(k,1+\epsilon)) = (0,1)$.*

Proof. First we study type profile $((k,v_1),(k,v_2))$ for general values $v_1, v_2 \in \mathbb{R}^+$ where $v_1 < v_2$. We prove that $\mathcal{M}^*$ gives all the items to the second participant.

Claim 3.2. $x^*((k,v_1),(k,v_2)) = \{2\}$ *for any* $v_1, v_2 \in \mathbb{R}^+$ *where* $v_1 < v_2$.

Proof. First note that M^* has to have a winner for this type profile because otherwise its social welfare will be zero while the maximum social welfare is v_2. This makes the social welfare ratio of $\mathcal{M}^*$ to be undefined.

Now we prove that if $x^*((k,v_1),(k,v_2)) = \{1\}$ then $\mathcal{M}^*$ either violates IC or AM. Lets call type profile $((k,v_1),(k,v_2))$ by $\theta^{(1)}$ and suppose for the sake of contradiction $x^*(\theta^{(1)}) = \{1\}$. From Lemma 3.1 we know that if participant 1 increases his bid to v_2 she still wins, hence $x^*(\theta^{(2)}) = \{1\}$ where $\theta^{(2)} = ((k,v_2),(k,v_2))$. Now if in type profile $\theta^{(2)}$ participant 2 decrease his bid to v_1, again from Lemma 3.1 we conclude that she cannot win, *i.e.*, $x^*(\theta^{(3)}) = \{1\}$ where $\theta^{(3)} = ((k,v_2),(k,v_1))$. Type profile $\theta^{(1)}$ is $\theta^{(3)}$ with participant 1 swapped with participant 2 but in both of them the first participant wins which contradicts with AM. $\square$

Claim 3.2 directly proves that the winner in type profile $((k,1),(k,1+\epsilon))$ is the second participant. The only thing remains is to show that her payment (p_2) is 1. Note that payment p_2 cannot be less than one because otherwise by Lemma 3.1 participant 2 wins all the items in type profile $((k, 1), (k, p_2))$ which contradicts with Claim 3.2. Payment p_2 cannot be larger than one because otherwise for any value $1 < v_2 < p_2$ participant 2 wins all the items in type profile $((k, 1), (k, v_2))$. This contradicts with Lemma 3.1 which states that the payment p_2 is the smallest value for which participant 2 wins the items. □

Now we add k more participants each of which wants only one item. In the following lemma we prove that RM forces $\mathcal{M}^*$ to assign all of the items to one of the participants who want all the items.

Lemma 3.7. *For the set of $k + 2$ participants with type profile $\theta^{(0)} = \left((1, 1 - \epsilon), (1, \frac{1}{2} - \epsilon), \ldots, (1, \frac{1}{k} - \epsilon), (k, 1), (k, 1 + \epsilon)\right)$, mechanism $\mathcal{M}^*$ assigns all the k items to either participant $k + 1$ or participant $k + 2$, i.e., $x^*\left(\theta^{(0)}\right) = \{k + 1\}$ or $x^*\left(\theta^{(0)}\right) = \{k + 2\}$.*

Proof. We prove the lemma by contradiction that if $\mathcal{M}^*$ assigns the items to a subset of the first k participants it satisfy be RM. We consider a class of k type profiles $(\theta^{(1)}, \ldots, \theta^{(k)})$ where $\theta^{(i)}$ is built from $\theta^{(i-1)}$. The only possible difference between $\theta^{(i)}$ and $\theta^{(i-1)}$ is in the valuation of participant i. If participant i is a winner in $\theta^{(i-1)}$, then we obtain $\theta^{(i)}$ by increasing the valuation of the ith participant from $\frac{1}{i} - \epsilon$ to $1 - \epsilon$. Note that the payment of participant i in $\theta^{(i-1)}$ is at most her valuation which is $\frac{1}{i} - \epsilon$ and in $\theta^{(i)}$ it remains the same by Lemma 3.1. If participant i is not a winner in $\theta^{(i-1)}$ then we obtain $\theta^{(i)}$ by decreasing his valuation to zero. Note that by IIA, no winner turns to a loser in $\theta^{(i)}$.

Let $j \in \{1, \ldots, k\}$ be the largest number for which participant j is a winner in

$\theta^{(j-1)}$ and we increase his valuation to $1-\epsilon$ in $\theta^{(j)}$. Note that at the start in type profile $\theta^{(0)}$ the set of winners is a non-empty subset of $\{1,\ldots,k\}$. Therefore there is at least one such j for which participant j is a winner in $\theta^{(j)}$ since decreasing the non-winners valuation does not reduce the size of the winners.

Now we prove that there is no winner in the set of participants $\{j+1,\ldots,k\}$ in type profile $\theta^{(j)}$. Assume otherwise and let $p \in \{j+1,\ldots,k\}$ be the smallest number for which participant p is a winner in $\theta^{(j)}$. Note that when we decrease the valuation of each participant $j < p' < p$ to zero to obtain $\theta^{(p')}$, participant p remains as a winner in all of them by IIA. Therefore, participant p is a winner in type profile $\theta^{(p-1)}$ and we increase his valuation in $\theta^{(p)}$ which contradicts with the fact that j is the largest number for which participant j is a winner in $\theta^{(j-1)}$.

The payment of participant j in $\theta^{(j-1)}$ is at most its valuation which is $\frac{1}{j}-\epsilon$. When we increase his bid to $1-\epsilon$ in type profile $\theta^{(j)}$ its payment remains the same by Lemma 3.1. Note that by construction of $\theta^{(j)}$ the valuation of all participants in $\{1,\ldots,j\}$ is either zero or $1-\epsilon$. If the valuation of them is $1-\epsilon$ and they are winner, by AM their payment is $\frac{1}{j}-\epsilon$. Therefore the total payments or revenue of $\mathcal{M}^*$ in $\theta^{(j)}$ is at most $j\cdot(\frac{1}{j}-\epsilon) = 1-j\cdot\epsilon$ since there is no other winner in set of participants $\{j+1,\ldots,k\}$ in type profile $\theta^{(j)}$.

Note that type profile $\theta^{(j)}$ is obtained from type profile $((k,1),(k,1+\epsilon))$ by adding k more participants. However the revenue of $\theta^{(j)}$ is $1-j\cdot\epsilon$ that is strictly less than 1 which is the revenue of $((k,1),(k,1+\epsilon))$ by Lemma 3.6. This contradicts with the RM property of $\mathcal{M}^*$, hence $\mathcal{M}^*$ has to assign the items to either participant $k+1$ or $k+2$. $\square$

Now we show how from Lemma 3.7 we can derive Theorem 3.2. Note that the maximum welfare for type profile $\theta^{(0)} = \big((1, 1-\epsilon), (1, \frac{1}{2}-\epsilon), \ldots, (1, \frac{1}{k}-\epsilon), (k, 1), (k, 1+\epsilon)\big)$ realized when we give one item to each of the first k participants for which we get the total social welfare $\sum_{i=1}^{k} \frac{1}{i} - k \cdot \epsilon$, *i.e.*, the nominator of Definition 3.1 for this type profile is $\sum_{i=1}^{k} \frac{1}{i} - k \cdot \epsilon$. The denominator of Definition 3.1 is at most $1+\epsilon$ by Lemma 3.7. Therefore the ratio of the welfare for this type profile is at least $\frac{\sum_{i=1}^{k} 1/i - k \cdot \epsilon}{1+\epsilon}$. Because *opt* is the maximum ratio over all type profiles (see Definition 3.2) we have *opt* $\geq \frac{\sum_{i=1}^{k} 1/i - k \cdot \epsilon}{1+\epsilon}$ which results in *opt* $\geq \sum_{i=1}^{k} \frac{1}{i} - \epsilon'$ where $\epsilon' = \frac{\epsilon(k - \sum_{i=1}^{k} 1/i)}{1+\epsilon}$.

Note that the value ϵ' can be made arbitrarily small by selecting a sufficiently small value for ϵ. Therefore we prove that for any positive small real value ϵ' we have *opt* $\geq \sum_{i=1}^{k} \frac{1}{i} - \epsilon'$ which implies Theorem 3.2.

CHAPTER 4

Randomized Revenue Monotone Mechanism Design

4.1 Introduction

Many Internet firms including search engines, social networks, and online publishers rely on online advertising revenue for their business; thus, making online advertising an essential part of the Internet. Online advertising consists of showing a few ads to a user when she accesses a web-page from a publisher's domain. The advertising can happen in different formats such as text-ads, image-ads, video-ads, or a hybrid of them.

A key component in online advertising is a mechanism which selects and prices the set of winning ads. In this chapter we study the design of mechanisms for Combinatorial Auction with Identical Items (CAII). In CAII we want to sell k identical items to a group of bidders; each demand a number of items from $\{1, \ldots, k\}$ and has a single-parameter valuation for obtaining them. Although CAII is a well-motivated model on its own, we note that a few important advertising scenarios such as image-text and video-pod auctions can be modeled by CAII. In image-text auction we want to fill an advertising box on a publisher's web-page with either one image-ad or k text-ads. We note that a large portion of Google AdSense's revenue is from this auction. Image-text auction is a special case of CAII where participants either demand only one item (text-ads) or all k items (image-ads). In video-pod auction there is an advertising break of k seconds which should be filled with video-ads each with certain duration and valuation.

When designing a mechanism, typically one focusses on attaining incentive-compatibility, and maximizing social welfare and/or revenue. In Chapter 3, we argue that the mechanisms for online advertising should satisfy an additional property of revenue-monotonicity. We bring the discussions of the previous chapter in this chapter as well to keep it independent. Revenue-monotonicity is a natural property which states that the revenue of a mechanism should not decrease as the number of bidders increase or if the bidders increase their bids. The motivation is that any online firm typically has a large sales team to attract more bidders on their inventory or they invest in new technologies to make bids more attractive. The typical reasoning is that more bidders (or higher bids) lead to more competition which should lead to higher prices. However, lack of revenue-monotonicity of a mechanism is conflicting with this intuitive and natural reasoning process, and can create significant confusion from a strategic decision-making point of view.

Even though Revenue Monotonicity (RM) seems very natural, we note that majority of the well-known mechanisms do not satisfy this property [RCLB09, RCLB11]. For example the famous Vickrey-Clarke-Groves (VCG) mechanism fails to satisfy RM as adding one more bidder might decreases the revenue to zero. To see this, consider two identical items to be sold to two bidders. One wants one item with a bid 2, and the other one wants both items with a bid 2. In this case the revenue of VCG mechanism is 2 (for a proof, see for instance [NRTV07, Chapter 9]). Now suppose we add one more bidder who wants one item with a bid of 2. In this case the revenue of VCG goes down to 0!

It is known that if we require mechanisms to satisfy both RM and IC, not only the mechanism cannot get the maximum social welfare but it can also not achieve Pareto-optimality in social welfare [RCLB11]. Remember that in Chapter 3 we introduced the notion of *Price of Revenue Monotonicity* (PoRM) to capture the loss in social welfare for RM mechanisms. Here a mechanism has PoRM of α if its social welfare is at least $\frac{1}{\alpha}$ fraction of the maximum social welfare in any type profile of participants. In Chapter 3, we saw that under a mild condition, the PoRM of any deterministic mechanism for the CAII problem is at least $\ln(k)$, *i.e.*, no deterministic mechanism can obtain more than $\frac{1}{\ln(k)}$ fraction of the maximum social welfare. In fact this impossibility result holds even for the case when participants demand either all the items or only one item.

This work is motivated by the desire to design better mechanisms for CAII. However, the above impossibility result of Chapter 3 is a bottleneck towards this goal. To overcome this, in this chapter, we resort to randomized mechanisms. We say a randomized mechanism satisfies RM if it satisfies RM in expectation[1]. Similarly, a randomized

[1] Since in a typical online advertising setting, there is a large number of auctions being run everyday, we

mechanism has PoRM of α if its expected social welfare is not less than $\frac{1}{\alpha}$ fraction of the maximum social welfare. We significantly improve the performance by designing a randomized mechanism with a constant PoRM. In particular, our randomized mechanism achieves a PoRM of 3.

Finally, we study Multi-group Combinatorial Auction with Identical Items (MCAII) that generalizes CAII. In MCAII bidders are partitioned into multiple groups and the set of winners has to be only from one group. The motivation is that the publisher sometimes require the ads to be of same format or size for a given ad slot. We design a randomized mechanism for MCAII that satisfies IC and RM with PoRM $O(\log k)$. An easy corollary of Chapter 3 gives a deterministic mechanism with a PoRM $O(\log^2 k)$. We give evidence that this factor for randomized mechanisms cannot be improved.

4.2 Related Works

[RCLB11] show that for combinatorial auctions, no deterministic mechanism that satisfies RM and IC can get weak maximality. A mechanism is Weakly Maximal (WM) if it chooses an allocation which cannot be augmented to make a losing participant a winner without hurting a winning participant. [RCLB09] study randomized mechanisms for combinatorial auctions which satisfy RM and IC. Note that a simple mechanism which chooses a maximal allocation uniformly at random ignoring the valuations of bidders satisfies RM, IC, and WM. [RCLB09] add another constraint that a mechanism has to also

get sharp concentration bounds.

satisfy Consumer Sovereignty (CS) which means that if a bidder increases her bid high enough, she can win her desired items. Now a new issue is that there is no randomized mechanism which satisfies RM, IC, WM, and CS [RCLB09]. In order to avoid this issue they relax CS constraint as follows. For each participant i there has to be λ different valuations $v_1 > v_2 > ... > v_\lambda$ such that for $j \in \{1, \ldots, \lambda\}$, we have $w_i(v_j) > w_i(v_{j+1}) + \sigma$ where w_i is the probability of winning for participant i and $\sigma > 0$. Roughly speaking relaxed CS constraint means that if participant i increases her bid from zero to infinity she sees at least λ jumps of length σ in her winning probability. The idea of their mechanism is that for each participant i they find λ constant values $c_{i,1} > c_{i,2} > ... > c_{i,\lambda}$ such that regardless of valuations of the other bidders; if the bid of bidder i is between $c_{i,j}$ and $c_{i,j+1}$ then her winning probability is at least $j * \sigma$. In order to find the constants for each participant they solve a LP whose constraints force RM, IC, Relaxed CS, and WM. As you may notice although this mechanism achieves WM, RM and relaxed CM, but can do very poorly in terms of PoRM. For example suppose you have n participants and each of them wants all items. The valuation of each participant i is bigger than its highest constant $c_{i,1}$. In this case all the participants can win with probability at most $1/n$. Now suppose that the valuation of one of the participants is infinity. She still wins with probability $1/n$ which shows that the PoRM of their mechanism is at least n.

[DRS09] show that VCG is revenue monotone if and only if the feasible subsets of winners form a matroid. [AM02] show that if valuations of bidders satisfy *bidder-submodularity* then VCG satisfies RM. Here valuations satisfy bidders submodularity if and only if for any bidder i and any two sets of bidders S, S' with $S \subseteq S'$ we have $\text{WF}(S \cup \{i\}) - \text{WF}(S) \geq \text{WF}(S' \cup \{i\}) - \text{WF}(S')$, where $\text{WF}(S)$ is the maximum social welfare achievable using only bidders in S. Note that we can restrict the set of possible allocations in a way such that bidder-submodularity holds. Then we can use VCG on this restricted set of allocations and hence achieve RM. However we can show that it is not possible to get a mechanism with PoRM better than $\Omega(k)$ by using the mentioned tool.

[AM02] design a mechanism which is in the core of the exchange economy for combinatorial auctions. A mechanism is in the core if there is no subset of participants including the seller which can collude and trade among each other such that all of them benefit more than the result of the mechanism. [DM08] show that a core-selecting mechanism which selects an allocation that minimizes the seller's revenue satisfies RM given bidders follow so called *best-response truncation strategy*. Therefore, the mechanism of [AM02] satisfies RM if it selects an allocation that minimizes the seller's revenue and the participants follow best-response strategy, however, this mechanism does not satisfy IC.

Another line of related works is around characterizing incentive compatible mechanisms. The classic result of [Rob79] tells that affine maximizers are the only social choice functions which can be implemented using mechanisms that satisfy IC when bidders have unrestricted quasi-linear valuations. Subsequent works study some restricted cases, see *e.g.* [Roc87, LMN03, BCL+06, SY05].

There is also a large body of research around designing mechanisms with good bounds on the revenue. In the single parameter Bayesian setting [Mye81] designs a mechanism which achieves the optimal expected revenue. [GHW01, GHK+06] consider optimizing revenue in prior-free settings (see *e.g.* [NRTV07] for a survey on this).

4.3 Results and Overview of Techniques

To give intuition about our approach, we first start with ideas that will not work but are potentially good candidates. To keep the explanation easier let us first focus on deterministic mechanisms. Note that the payment of each participant in a deterministic mechanism which satisfies IC is her critical value, *i.e.*, the minimum valuation for which she still remains a winner. Assume that all participants demand only one item. In this case we can simply give all the items to the highest k bidders, which sets the critical value (the payment) of each winner to the valuation of the $(k+1)$th highest bidder. If we add one more participant the valuation of the $(k+1)$th highest bidder increases, therefore, the payment of each winner increases and hence the mechanism satisfies RM.

Now assume we have two types of bidders: A bidder of type A who demand all k items, and a bidder of type B who demands a single item. This scenario is equivalent to the image-text auction for which there is a lower-bound of $\ln(k)$ for the PoRM of deterministic mechanisms (see Chapter 3). However using randomization we can simply get a PoRM of 2. Flip a coin and with probability half give all items to the highest type A bidder and with probability half give k items to the k highest bidders of type B. Here, the expected social welfare is at least half of the maximum social welfare. Note that when the coin flip selects bidders of type A the auction simply transforms to the second price auction of selling one package of items which has RM. When it selects bidders of type B the auction transforms to the case when all bidders demand one item which we explained earlier and has RM. Therefore, the expected revenue is monotone and hence the mechanism satisfies RM. Expanding the above idea we can partition the bidders into $\log(k)$ groups such that the bidders of each group $i \in \log(k)$ has demand in $[2^i, 2^{i+1})$. Then, we randomly select one group and choose the winners from the selected group. However, this partitioning approach does not lead to a PoRM better than $\log(k)$.

As a second approach instead of partitioning the bidders and sort them by their valuation, we can sort them according to their Price Per Item (PPI) which is the valuation of a participant divided by the number of items she demands. Now consider a simple greedy algorithm as follows. Start from the top of the sorted list of bidders and at each step do the following. If the number of remaining items is enough to serve the current bidder give the items to the bidder and proceed; otherwise stop. Let us call the bidder at which the greedy algorithm stops the *runner-up bidder*. Note that the runner-up bidder has the largest PPI among the loser bidders and let p be her PPI. If each of the winner participant had PPI less than p then she could not win. Therefore, the critical value of each winner participant is her demand multiplied by p. Although value p increases if we add more bidders, the number of items sold might decrease. For example consider the case when the bidder with the highest PPI demands all k items. In this scenario she wins all items and pays k multiplied by the PPI of the runner-up bidder. Now if we add one more bidder whose PPI is more than the highest bidder but demands only one item; the new bidder wins and we sell only one item. This potentially decreases the revenue of the greedy mechanism.

For our mechanism we use a combination of the above ideas and an extra technique. We partition the bidders into two groups: high-demand bidders who demand more than $k/2$ items, and low-demand bidders who demand less than or equal to $k/2$ items. With probability $1/3$ the winner is a high-demand bidder with the largest valuation. Similar to the partitioning approach the critical value of the winner is the second largest valuation of the high-demand bidders which can only increase if we add more bidders. With probability $2/3$ we do the following with the low-demand bidders. First we run the greedy algorithm over the low-demand bidders and find the runner-up bidder. The important observation here is that because there is no high-demand bidder, sum of winners' demands (A) is larger than $k/2$. Therefore we are sure that we sell at least $k/2$ items where the price of each item is the PPI of the runner-up bidder. Now we select each winner of the greedy algorithm with probability $\frac{k/2}{A}$ as the true winner of our mechanism. This random selection makes sure that the expected number of sold items is exactly $k/2$. The exact number $k/2$ is important since the expected revenue of the mechanism is $k/2$ multiplied by the PPI of the runner-up bidder. Therefore as the PPI of the runner-up bidder increases if we add more bidders the expected revenue is monotone.

Now we explain ideas used to design our mechanism for MCAII. We first note that as a corollary of the result of Chapter 3, we get a deterministic mechanism with a PoRM of $\log^2(n)$. In our mechanism, we assign a value to each group and use it as the criterion in order to select the winner group. Note that a simple value that can be assigned to each group is the maximum social welfare obtainable by the group. However, this way we cannot guarantee RM. Because suppose participant $i^{(g)}$ of group $G^{(g)}$ increases her bid high enough which guarantees that $G^{(g)}$ wins against all other groups no matter what are the valuations of the other participants of $G^{(g)}$. Therefore, the critical values of the other members of $G^{(g)}$ decreases as $i^{(g)}$ increases her bid and hence can decrease the revenue of the mechanism.

We refer to our assigned value to each group as the Maximum Possible Revenue of the Group (MPRG). As name MPRG suggests, it shows the maximum revenue we can obtain from each group without the fear of violating RM. For each $j \in \{1, \ldots, k\}$ and group $G^{(g)}$, let $u_j^{(g)}$ be the maximum price can be set for a single item so that we can sell at least j items to low-demand bidders of group $G^{(g)}$. More formally, $u_j^{(g)}$ is the maximum value where the sum of demands of low-demand bidders whose PPI is larger than $u_j^{(g)}$ in group $G^{(g)}$ is at least j. The MPRG of group $G^{(g)}$ is $\max(V^{(g)}, \max_{j\in\{1,\ldots,k/2\}} j \cdot u_j^{(g)})$ where $V^{(g)}$ is the highest valuation of high-demand bidders. Intuitively, MPRG either sells items to high-demand bidders and obtains revenue of at most $V^{(g)}$ or sells items to low-demand bidders in which we can sell a number of items between 1 and $k/2$. We select a group with the highest MPRG and choose the winners from this group. We are able to show that we can obtain a revenue of at least the second highest MPRG. We prove that our mechanism satisfies RM by showing that the second highest MPRG increases if we add more bidders.

We show that the MPRG of each group is at least $1/\ln(k)$ fraction of the maximum social welfare obtainable by the group. Therefore, as we select the winning group using the MPRGs of groups, the PoRM of our mechanism is $O(\ln(k))$. We provide evidence that indeed the MPRG of each group is the closest value to its social welfare that can be safely used for selecting the winning group without violating RM. Moreover, any randomization over the groups for selecting the winning one according to MPRG cannot improve the PoRM factor.

4.4 Preliminary

Let assume we have a set of n bidders $\{1, \ldots, n\}$ and a set of k identical items. Let type profile θ be a vector containing the type of each bidder i which we show by θ_i. Here θ_i is pair $(d_i, v_i) \in [k] \times \mathbb{R}^+$ where d_i is the number of items she demands and v_i shows her valuation for getting d_i items. Here we assume the demands are publicly known because in our scenario they represent the length of video-ads stored in database while the valuations are private to bidders.

Note that having higher valuation does not necessarily mean that the bidder is more desirable to the seller as she might have a large demand. We define Price Per Item (PPI) of bidder i to be $\frac{v_i}{d_i}$ which we use in our mechanism to compare bidders.

We show a randomized mechanism ($\mathcal{M}$) by pair (w, p) where $w_i(\theta)$ shows the winning probability of bidder i in type profile θ and $p_i(\theta)$ is her expected payment.

We use the following Theorem in this chapter frequently which is a well-known characteristic of the truthful randomized mechanisms in the single parameter model (see *e.g.* [NRTV07]).

Theorem 4.1. *Randomized mechanism $\mathcal{M} = (w, p)$ is truthful if and only if for any type profile θ and any bidder i with type (d_i, v_i) the followings hold.*

1. *Function $w_i\left((d_i, v_i), \theta_{-i}\right)$ is weakly monotone in v_i.*

2. $p_i(\theta) = v_i \cdot w_i(\theta) - \int_0^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt$

4.5 Combinatorial Auction with Identical Items

We build a randomized mechanism ($\mathcal{M} = (w, p)$) satisfying revenue monotonicity and incentive compatibility such that PoRM($\mathcal{M}$) is equal to 3.

We call a bidder *high-demand* bidder if her demand is greater than $\lfloor k/2 \rfloor$ otherwise we refer it as *low-demand* bidder. Mechanism $\mathcal{M}$ partitions the bidders into two groups of low-demand and high-demand bidders and with probability $1/3$ selects the winning set from the high-demand bidders and with probability $2/3$ from the low-demand bidders.

We will see that mechanism $\mathcal{M}$ favors high-demand bidders with larger valuations and favors low-demand-bidders with larger PPIs while breaking the ties by the index number of the bidders.

Definition 4.1. *We call low-demand bidder l_1 is* more valuable *than low-demand bidder l_2 and show it by ($l_1 \succ l_2$) if* $\text{PPI}_{l_1} > \text{PPI}_{l_2} \vee (\text{PPI}_{l_1} = \text{PPI}_{l_2} \wedge l_1 < l_2)$. *Similarly we call high-demand bidder h_1 is* more valuable *than high-demand bidder h_2 and show it by ($h_1 \succ h_2$) if* $v_{h_1} > v_{h_2} \vee (v_{h_1} = v_{h_2} \wedge h_1 < h_2)$.

Let's assume that there are ℓ low-demand bidders and h high-demand bidders. By adding some dummy bidders with demand 1 and valuation zero we assume that the sum of demands of low-demand bidders is always greater than k. Without loss of generality we assume that the first ℓ bidders are low-demand bidders and $i \succ i+1$ for any $i \in [\ell-1]$ (the PPIs of the low-demand bidders decreases by their index) and the remaining h bidders are high-demand-bidders while $i \succ i+1$ for any $i \in \{\ell+1, \dots n-1\}$ (the valuations of the high-demand bidders decreases by their index).

Definition 4.2. *We call low-demand bidder r the runner-up bidder if r is the smallest value in set $[\ell]$ for which $\sum_{i=1}^{r} d_i \geq k$.*

Later we will see that the runner-up bidder is the bidder with the largest PPI and smallest index number who has zero probability of winning. We simply refer to the runner-up bidder as r.

We define A to be $\sum_{i=1}^{r-1} d_i$ which is the sum of demands of low-demand bidders that have PPIs greater than or equal to that of r and have positive probability of winning (see Figure 4.1).

Observation 5. *We have $\lceil k/2 \rceil \leq A < k$.*

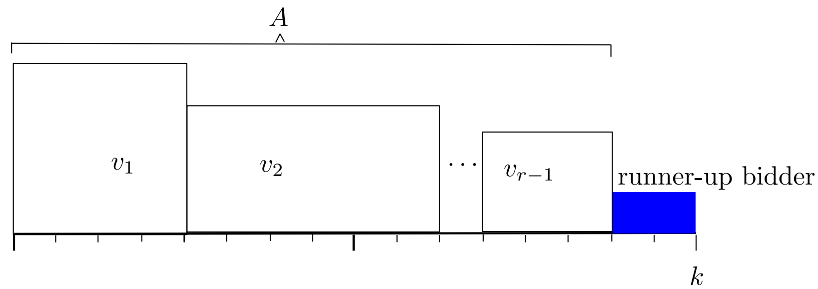

Figure 4.1: Each rectangle corresponds to a bidder where the height, width, and area represent PPI, demand, and valuation of the bidder respectively. The dark rectangle corresponds to the runner-up bidder whose demand crosses the value k.

Proof. Inequality $A < k$ is the direct result of the way we select runner-up bidder r. Inequality $\lceil \frac{k}{2} \rceil \leq A$ follows from the fact that $\sum_i^r d_i \geq k$ and the demand of the runner-up bidder is less than or equal to $\lfloor \frac{k}{2} \rfloor$ by definition of low-demand bidders. $\square$

Now we are ready to precisely define how $\mathcal{M}$ selects and charges the set of winners. With probability $1/3$ $\mathcal{M}$ selects the most valuable high-demand bidder (which is the high-demand bidder with largest valuation breaking the ties by index). Therefore, the winning bidder in this case is $\ell + 1$ and she pays $v_{\ell+2}$ which is the second highest valuation among high-demand bidders. Therefore her expected payment is $p^{\mathcal{M}}_{\ell+1}(\theta) = \frac{v_{\ell+2}}{3}$.

If we did not select the largest high-demand bidder then mechanism $\mathcal{M}$ uniformly at random selects the winner set from the first $r - 1$ low-demand bidders where the probability of selecting each bidder $i \in [r-1]$ is $\frac{\lceil k/2 \rceil}{A}$. In this case if bidder i gets selected she has to pay $d_i \cdot \text{PPI}_r$. Therefore, her expected payment is $p^{\mathcal{M}}_i(\theta) = \frac{2\lceil k/2 \rceil}{3A} \cdot d_i \cdot \text{PPI}_r$ since with probability $2/3$ we select low-demand-bidders and with probability $\frac{\lceil k/2 \rceil}{A}$ bidder i gets selected. The probability $\frac{\lceil k/2 \rceil}{A}$ is selected in a way such that if the low-demand bidders win, the expected number of allocated items is $\lceil k/2 \rceil$ since the sum of demands of the first $r-1$ low-demand bidders is A and each of them gets selected with probability $\frac{\lceil k/2 \rceil}{A}$.

In summary the expected payments of the bidders in mechanism $\mathcal{M}$ is the following.

$$p_i^{\mathcal{M}}(\theta) = \begin{cases} 0 & \ell + 1 < i \\ \frac{v_{\ell+2}}{3} & i = \ell + 1 \\ 0 & r \leq i \leq \ell \\ \frac{2\lceil k/2 \rceil}{3A} \cdot d_i \cdot \text{PPI}_r & 1 \leq i < r \end{cases} \tag{4.1}$$

In the following first we prove that the allocation function of $\mathcal{M}$ is monotone and then we show that the unique expected payments of the winners calculated using Lemma 5.2 is equal to the expected payment of mechanism $\mathcal{M}$ which proves that $\mathcal{M}$ truthful.

Observation 6. *$w_i\left((d_i, v_i), \theta_{-i}\right)$ is monotone in v_i.*

Proof. If bidder i is a high-demand bidder then clearly increasing her bid just increases her chance to be the high-demand bidder with the largest valuation and hence win with probability $1/3$. If bidder i is a low-demand bidder then increasing her bid just increases her PPI and hence can help her to go over the PPI of the runner-up bidder and win with probability $\frac{2\lceil \frac{k}{2} \rceil}{3 \cdot A}$. □

The following lemma shows the expected payment of each winner.

Lemma 4.1. *The truthful expected payment of bidder i ($p_i(\theta)$) calculated by item 2 of Lemma 5.2 is the following.*

$$p_i(\theta) = \begin{cases} 0 & \ell + 1 < i \\ \frac{v_{\ell+2}}{3} & i = \ell + 1 \\ 0 & r \le i \le \ell \\ \frac{2\lceil k/2 \rceil}{3A} \cdot d_i \cdot \text{PPI}_r & 1 \le i < r \end{cases} \tag{4.2}$$

Proof. Remember that the first ℓ bidders are low-demand bidders which have non-decreasing PPIs, r is the low-demand runner-up bidder, and finally among high-demand bidders, bidder $\ell + 1$ has the largest valuation and bidder $\ell + 2$ has the second largest valuation.

The probability of winning for bidder i when $\ell + 1 < i$ is zero since she is a high-demand bidder who either does not have the highest valuation or has the highest valuation but has larger index number (see Definition 4.1). Because function w_i is monotone we conclude that $w_i\left((d_i, t), \theta_{-i}\right)$ is equal to zero for any $t \le v_i$. Hence by calculating the formula in item 2 of Lemma 5.2 we get $p_i(\theta) = 0$.

We calculate the truthful expected payment of bidder $\ell + 1$ by using the formula in item 2 of Lemma 5.2.

$$\begin{aligned}
p_{\ell+1}(\theta) &= v_{\ell+1} \cdot w_{\ell+1}(\theta) - \int_0^{v_{\ell+1}} w_{\ell+1}\left((d_{\ell+1}, t), \theta_{-\ell+1}\right) dt \\
&= \frac{1}{3} v_{\ell+1} - \int_0^{v_{\ell+1}} w_{\ell+1}\left((d_{\ell+1}, t), \theta_{-\ell+1}\right) dt \\
&= \frac{1}{3} v_{\ell+1} - \int_0^{v_{\ell+2}} w_{\ell+1}\left((d_{\ell+1}, t), \theta_{-\ell+1}\right) dt - \int_{v_{\ell+2}}^{v_{\ell+1}} w_{\ell+1}\left((d_{\ell+1}, t), \theta_{-\ell+1}\right) dt \\
&= \frac{1}{3} v_{\ell+1} - \int_{v_{\ell+2}}^{v_{\ell+1}} w_{\ell+1}\left((d_{\ell+1}, t), \theta_{-\ell+1}\right) dt \\
&= \frac{1}{3} v_{\ell+1} - \frac{1}{3}(v_{\ell+1} - v_{\ell+2}) \\
&= \frac{v_{\ell+2}}{3}
\end{aligned}$$

The first equality is item 2 of Lemma 5.2, the second equality follows from the fact that probability of winning for bidder $\ell + 1$ ($w_{\ell+1}(\theta)$) is $1/3$, the third one is breaking the domain of integration, the forth and fifth equalities are followed by noting that probability of winning for bidder $\ell + 1$ is zero if his valuation is less than $v_{\ell+2}$ and is $1/3$ if his valuation is greater than or equal to $v_{\ell+2}$.

The probability of winning for bidder i when $r \le i \le \ell$ is zero since she is a low-demand bidder which has PPI less than or equal to PPI_r. Because function w_i is monotone we conclude that $w_i\left((d_i, t), \theta_{-i}\right)$ is equal to zero for any $t \le v_i$. Hence by calculating the formula in item 2 of Lemma 5.2 we get $p_i(\theta) = 0$.

The only part remaining is to show that $p_i(\theta) = \frac{2\lceil k/2 \rceil}{3A} \cdot d_i \cdot \text{PPI}_r$ for $1 \le i < r$ using item 2 of Lemma 5.2. In order to calculate $\int_0^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt$ we consider the curve of allocation function $w_i\left((d_i, t), \theta_{-i}\right)$ when t increases from zero to v_i (see Figure 4.2).

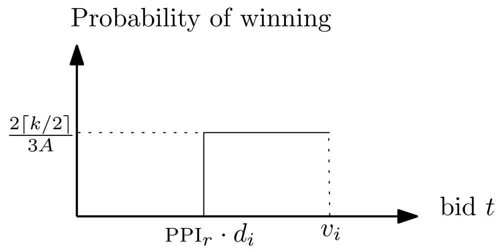

Figure 4.2: The horizontal axis represents the bid of bidder i and the vertical axis shows the probability of bidder i winning. As bidder i increases her bid; at the point when her PPI is equal to the PPI of the runner-up bidder she gets allocated with probability $\frac{2\lceil k/2\rceil}{3A}$.

Observation 7. *For any bidder* $1 \leq i < r$ *allocation function* $w_i((d_i, t), \theta_{-i})$ *is equal to zero when* $t < d_i \cdot \text{PPI}_r$ *and is equal to* $\frac{2\lceil k/2\rceil}{3A}$ *when* $t \geq d_i \cdot \text{PPI}_r$.

Proof. Remember the runner-up bidder r has the smallest index for which $\sum_{j=1}^{r} d_j \geq k$. Mechanism $\mathcal{M}$ allocates all the low-demand bidders which have index less than r (have PPIs greater than or equal to the runner-up bidder) with probability $\frac{2\lceil k/2\rceil}{3A}$. Therefore as far as $t \geq d_i \cdot \text{PPI}_r$ bidder i is more valuable than the runner-up bidder (see Definition 4.1) and wins with probability $\frac{2\lceil k/2\rceil}{3A}$ in type profile $((d_i, t), \theta_{-i})$.

Now assume that $t < d_i \cdot \text{PPI}_r$ and $\theta' = ((d_i, t), \theta_{-i})$. Our objective is to show that the probability of bidder i winning is zero for type profile θ' and hence finish the proof of the observation. Note that $\sum_{j=1}^{r} d_j \geq k$ and in the new type profile θ' bidder i has PPI less than the PPIs of all bidders $j \in [r]$ where $j \neq i$ since $t < d_i \cdot \text{PPI}_r$. In other words, bidder i is the least valuable bidder in $[r]$ (see Definition 4.1) while $\sum_{j=1}^{r} d_j \geq k$. Therefore, bidder i is either the runner-up bidder in θ' or has PPI less than the runner-up bidder (see

Definition 4.2). Hence has zero probability of winning. □

The following equalities shows the expected payment of bidder i for $1 \le i < r$.

$$
\begin{aligned}
p_i(\theta) &= v_i \cdot w_i(\theta) - \int_0^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt \\
&= \frac{2\lceil k/2 \rceil}{3A} v_i - \int_0^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt \\
&= \frac{2\lceil k/2 \rceil}{3A} v_i - \int_0^{d_i \cdot \mathrm{PPI}_r} w_i\left((d_i, t), \theta_{-i}\right) dt - \int_{d_i \cdot \mathrm{PPI}_r}^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt \\
&= \frac{2\lceil k/2 \rceil}{3A} v_i - \int_{d_i \cdot \mathrm{PPI}_r}^{v_i} w_i\left((d_i, t), \theta_{-i}\right) dt \\
&= \frac{2\lceil k/2 \rceil}{3A} v_i - \frac{2\lceil k/2 \rceil}{3A} (v_i - d_i \cdot \mathrm{PPI}_r) \\
&= \frac{2\lceil k/2 \rceil}{3A} (d_i \cdot \mathrm{PPI}_r)
\end{aligned}
$$

The first equality is item 2 of Lemma 5.2, the second equality follows from the fact that probability of winning for bidder i ($w_i(\theta)$) is $\frac{2\lceil k/2 \rceil}{3A}$, the third one is breaking the domain of integration, the forth and fifth equalities are followed from Observation 7. □

Let $\mathrm{REVENUE}(\mathcal{M}, \theta)$ denotes the expected revenue of mechanism $\mathcal{M}$ in type profile θ. We prove the following.

$$
\begin{aligned}
\text{REVENUE}(\mathcal{M},\theta) &= \sum_{i=1}^{n} p_i(\theta) && \text{definiton of REVENUE} \\
&= \sum_{i=1}^{r-1} p_i(\theta) + \sum_{i=r}^{\ell} p_i(\theta) + p_{\ell+1}(\theta) + \sum_{i=\ell+2}^{n} p_i(\theta) \\
&= \sum_{i=1}^{r-1} \frac{2\lceil k/2\rceil}{3A}(d_i \cdot \text{PPI}_r) + \frac{v_{\ell+2}}{3} && \text{Lemma 4.1} \\
&= \frac{2\lceil k/2\rceil}{3A} \cdot \text{PPI}_r \cdot \sum_{i=1}^{r-1} d_i + \frac{v_{\ell+2}}{3} \\
&= \frac{2\lceil k/2\rceil}{3} \cdot \text{PPI}_r + \frac{v_{\ell+2}}{3} && \text{as } A = \sum_{i=1}^{r-1} d_i
\end{aligned}
\tag{4.3}
$$

The following lemma proves that $\mathcal{M}$ is revenue monotone.

Lemma 4.2. *The expected revenue of mechanism $\mathcal{M}$ does not decrease if we add one more bidder or a bidder increases her bid.*

Proof. The expected revenue of mechanism $\mathcal{M}$ is $\frac{2\lceil k/2\rceil}{3} \cdot \text{PPI}_r + \frac{v_{\ell+2}}{3}$ by Equation (4.3). Remember that $v_{\ell+2}$ is the high-demand bidder with the second highest valuation, therefore, if we add one more bidder or a bidder increases her bid this value does not decrease. On the other hand PPI_r is the PPI of the runner-up bidder (see Definition 4.2) which is the PPI of the first low-demand bidder that crosses value k (see Figure 4.1) where the bidders are sorted according to their PPIs. The proof of the lemma follows by the fact that PPI_r also does not decrease as we add one more bidder or a bidder increases her bid. □

In the following lemma we show $\text{PoRM}(\mathcal{M}) = 3$ and finish this section.

Lemma 4.3. *The Price of Revenue Monotoncity (*PoRM*) of $\mathcal{M}$ is* 3*.*

Proof. We prove the lemma by showing that the expected social welfare of $\mathcal{M}$ in type profile θ is at least $\frac{1}{3}$ of the maximum social welfare ($\text{WF}(\theta)$). Let $\mathcal{S}$ be an arbitrary subset of bidders which VCG selects and realizes the maximum social welfare $\text{WF}(\theta)$. Let L be the sum of valuations of low-demand bidders in $\mathcal{S}$ and H be the sum of valuations of high-demand bidders in $\mathcal{S}$. Therefore, we have:

$$\text{WF}(\theta) = L + H \tag{4.4}$$

There can be at most one high-demand bidders in $\mathcal{S}$ since they have demand more than $\lfloor k/2 \rfloor$. As $v_{\ell+1}$ is the high-demand bidder with the largest valuation we have the following.

$$v_{\ell+1} \geq H \tag{4.5}$$

We also have

$$\sum_{i=1}^{r-1} v_i \geq L \cdot \frac{A}{k} \tag{4.6}$$

since $A = \sum_{i=1}^{r-1} d_i$ and the first $r-1$ bidders have larger PPIs than the rest as they are sorted non-increasingly according to their PPIs.

Remember that $\mathcal{M}$ selects bidder $\ell + 1$ with probability $1/3$ or selects each of the first $r - 1$ low-demand bidders with probability $\frac{2\lceil k/2 \rceil}{3A}$. The following equalities finishes the proof of the lemma.

$$\begin{aligned}
E[\text{WF}(\mathcal{M},\theta)] &= \sum_{i=1}^{r-1} \frac{2\lceil k/2\rceil}{3A} \cdot v_i + \frac{1}{3} v_{\ell+1} && \text{definition of expected social welfare} \\
&= \frac{2\lceil k/2\rceil}{3A} \cdot \sum_{i=1}^{r-1} v_i + \frac{1}{3} v_{\ell+1} \\
&\geq \frac{2\lceil k/2\rceil}{3A} \cdot L \cdot \frac{A}{K} + \frac{1}{3} H && \text{by Equation (4.5) and Equation (4.6)} \\
&\geq \frac{1}{3} L + \frac{1}{3} H && \text{algebra} \\
&= \frac{1}{3} \text{WF}(\theta) && \text{by Equation (4.4)}
\end{aligned}$$

□

4.6 Multigroup Combinatorial Auction with Identical Items

In this section we describe our Mechanism for Multigroup Combinatorial Auction with identical items (MMCA). We prove that MMCA satisfies IC and RM while has PoRM of at most $O(\log k)$.

In the following we define required notations to be used throughout this section. Let's assume we have m groups $G^{(1)}, \ldots, G^{(m)}$. We always show the group index of any variable within parenthesis in superscript. For group $G^{(g)}$ let $\left(D^{(g)}, V^{(g)}\right)$ be the type of a high-demand bidder with the highest valuation, $\left(d_1^{(g)}, v_1^{(g)}\right), \left(d_2^{(g)}, v_2^{(g)}\right), \ldots, \left(d_{n^{(g)}}^{(g)}, v_{n^{(g)}}^{(g)}\right)$ be the types of low-demand bidders, and $\text{PPI}_i^{(g)}$ be the PPI of ith low-demand bidder. Without loss of generality, for each group $G^{(g)}$ we assume that the number of low-demand bidders $n^{(g)}$ is larger than k and

$\text{PPI}_1^{(g)} \geq \text{PPI}_2^{(g)} \geq \ldots \geq \text{PPI}_{n^{(g)}}^{(g)}$. The following defines the maximum price per item with which we can sell at least j items to the low-demand bidders of group $G^{(g)}$.

Definition 4.3 (jth-item value)**.** *For each $j \in [k]$ we define* jth-item value *($u_j^{(g)}$) of group $G^{(g)}$ to be equal to the valuation of ith low-demand bidder $v_i^{(g)}$ where i is the minimum number for which $\sum_{t=1}^{i} d_t^{(g)}$ is greater than or equal to j (see Figure 4.3).*

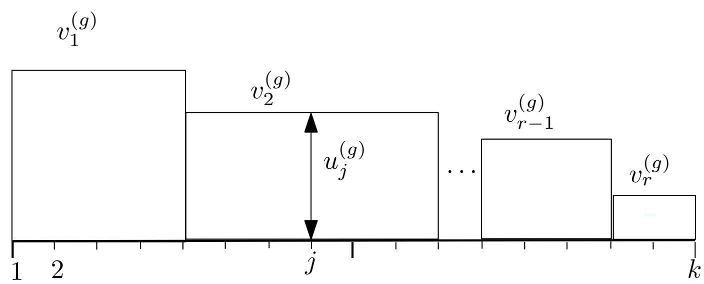

Figure 4.3: Each rectangle corresponds to a low-demand bidder of group $G^{(g)}$ where the height, width, and area represent PPI, demand, and valuation of the bidder respectively. Here bidders are sorted according to their PPIs and the valuation of the bidder who crosses item j is the jth item-value.

We are interested in assigning a value to each group which represents how much revenue can be obtained if we give the items to bidders of the group. Then, we give the items to bidders of a group with the highest assigned value.

Definition 4.4 ($\text{MPRG}^{(g)}$)**.** *The Maximum Possible Revenue of Group $G^{(g)}$ ($\text{MPRG}^{(g)}$) is equal to*

$$\max(V^{(g)}, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j^{(g)}).$$

The intuition for MPRG$^{(g)}$ is the following. The value $j \cdot u_j^{(g)}$ is the maximum revenue we can obtain if we sell exactly j items to low-demand bidders of group $G^{(g)}$, see Definition 4.3. Note that number j is taken from the set $[\lceil k/2 \rceil]$ meaning that we consider selling at most $\lceil k/2 \rceil$ items. This is because low-demand bidders have different demands from the range $[1..\lceil k/2 \rceil]$, therefore, we can guarantee selling at most $\lceil k/2 \rceil$ items without overselling the k items. In fact with randomization we make sure that we sell exactly $\lceil k/2 \rceil$ items in expectation. Finally we take the maximum of the highest valuation of high-demand bidders ($V^{(g)}$) and value $\max_{j \in [\lceil k/2 \rceil]} j \cdot u_j^{(g)}$.

The rest of this section is organized as follow. First, we describe the allocation function of MMCA. Second, we use Lemma 4.1 to derive the expected payments of the winners which determines the revenue of MMCA. Third, we show that the revenue does not decrease if we add one more bidder or a bidder increases her bid. Finally, we show that PoRM of MMCA is at most $O(\log k)$.

Let G^{g^*} be the group with highest MPRG, $g^* = \arg\max_g \text{MPRG}^{(g)}$, and $G^{\hat{g}}$ be the group with the second highest MPRG. Mechanism MMCA selects the winners from group $G^{(g^*)}$. Let R be equal to MPRG$^{\hat{g}}$. We think of R as a reserved value such that we must obtain at least R revenue from group $G^{(g^*)}$. In the rest of this section all the discussions are about group $G^{(g^*)}$ unless mentioned otherwise, henceforth, we drop the group identifiers from variables.

Similar to Section 4.4, in the following, we define *runner-up* bidder to be the low-demand bidder with highest PPI which cannot be a winner if we sort bidders by their PPIs (see Figure 4.4).

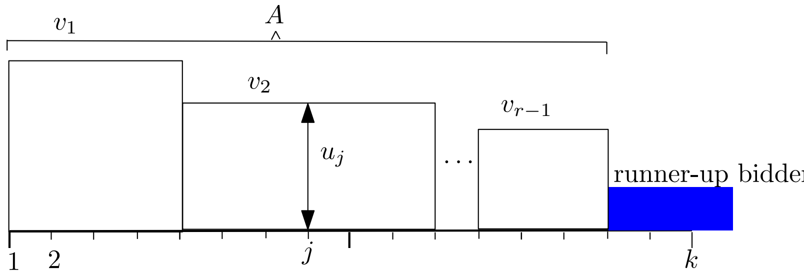

Figure 4.4: Each rectangle corresponds to a low-demand bidder of group $G^{(g^*)}$ where the height, width, and area represent PPI, demand, and valuation of the bidder respectively. Here bidders are sorted according to their PPIs and the low-demand bidder whose demand crosses kth item is the runner-up bidder.

Definition 4.5 (runner-up bidder)**.** *We call low-demand bidder r the runner-up bidder if r is the smallest number in set $[n]$ for which*

$$\sum_{i=1}^{r} d_i \geq k.$$

Remember that R is equal to the second largest MPRG of groups, V is the largest valuation of high-demand bidders of the winning group $G^{(g^*)}$, and u_j is largest price of an item for which we can sell at least j items to low-demand bidders of $G^{(g^*)}$ (see Definition 4.3). Note that as MPRG $= \max(V, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j)$ of $G^{(g^*)}$ is larger than R we have either: (1) $V > \max(R, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j)$ or (2) $\max_{j \in [\lceil k/2 \rceil]} j \cdot u_j \geq \max(R, V)$.

Case (1) is easy, if $V > \max(R, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j)$ then the set of winners contains only a high-demand bidder with valuation V breaking the ties arbitrarily. In Case(2) let j^* be the largest number from $[\lceil k/2 \rceil]$ such that $j^* \cdot u_{j^*} \geq \max(R, V)$. If j^* is less than $\lceil k/2 \rceil$

then the set of winners is all low-demand bidders whose PPI is greater than or equal to u_{j^*}. Otherwise if j^* is equal to $\lceil k/2 \rceil$ then we need to include all the low-demand bidders whose PPI is larger than or equal to $\max(R, V)/\lceil k/2 \rceil$ since, roughly speaking, their PPI is high enough to win against both high-demand bidders and the group with second highest MPRG. Note that the sum of demands of such bidders might exceed $\lceil k/2 \rceil$, therefore, we need randomization to guarantee selling exactly $\lceil k/2 \rceil$ items in expectation.

Definition 4.6 (**a** and **A**). *If $\lceil k/2 \rceil \cdot u_{\lceil k/2 \rceil} \geq \max(R, V)$, we define a to be the largest number in $[r-1]$ such that* $\text{PPI}_a \geq \max(R, V)/\lceil k/2 \rceil$, *i.e., number a is the smallest index in the set of all low-demand bidders that have index greater than the runner-up bidder and have* PPI *larger than or equal to $u_{\lceil k/2 \rceil}$. We also define number A to be the sum of demands of the first a low-demand bidders, i.e., $A = \sum_{i=1}^{p} d_i^{(g)}$.*

Now we are ready to formally define the allocation function of MMCA.

Definition 4.7 (Allocation Function of MMCA).

1. *If $V > \max(R, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j)$ then the set of winners contains only a high-demand bidder with valuation V breaking the ties arbitrarily.*
2. *If j^* is less than $\lceil k/2 \rceil$ then the set of winners is all low-demand bidders whose* PPI *is greater than or equal to u_{j^*}.*
3. *If j^* is equal to $\lceil k/2 \rceil$ then each of the first a low-demand bidder wins with probability $\frac{\lceil k/2 \rceil}{A}$ independently.*

In the following lemma we calculate the critical values of winners using Lemma 5.2.

Lemma 4.4. *The critical values of winners are the following considering different conditions of Definition 4.7.*

1. *If item 1 happens then the critical value of the winner is* $\max(R, \max_{j\in[\lceil k/2\rceil]} j \cdot u_j, V_2)$ *where* V_2 *is the second highest valuation of high-demand bidders.*

2. *If item 2 happens then the critical value of each winner* i *is* $d_i \cdot (\max(R, V)/j^*)$.

3. *If item 3 happens then the critical value of each winner* i *is* $\frac{\lceil k/2\rceil}{A} \cdot d_i \cdot \max(\text{PPI}_r, (\max(R, V)/\lceil k/2\rceil))$.

Proof. We consider all the three conditions separately.

item 1: We show that if the valuation V of the winner goes below $\max(R, \max_{j\in[\lceil k/2\rceil]} j \cdot u_j, V_2)$ then player V cannot be a winner. If the value of $\max(R, \max_{j\in[\lceil k/2\rceil]} j \cdot u_j, V_2)$ is equal to R then decreasing V to a value less than R causes a group with the second highest MPRG win and hence changes player V to a loser participant. If it is equal to $\max_{j\in[\lceil k/2\rceil]} j \cdot u_j$ then decreasing V to a value less than $\max_{j\in[\lceil k/2\rceil]} j \cdot u_j$ causes low-demand bidders win and hence changes player V to a loser participant. If it is equal to V_2 then decreasing V to a value less than V_2 causes the high-demand bidder with valuation V_2 win and hence changes player V to a loser participant.

item 2: Remember that j^* is the largest number from $[\lceil k/2\rceil]$ such that $j^* \cdot u_{j^*} \geq \max(R, V)$. Therefore if the valuation of any winner i goes below $d_i \cdot (\max(R, V)/j^*)$ then there exist no j in $[\lceil k/2\rceil]$ such that $j \cdot u_j$ is larger than $\max(R, V)$. Hence the winning set changes to either another group if $\max(R, V) = R$ or to high-demand bidders if $\max(R, V) = V$.

item 3: In this case we show that if the valuation of winner i goes below $d_i \cdot \max(\text{PPI}_r, (\max(R, V)/\lceil k/2 \rceil))$ then she has zero probability of winning and if it is larger than or equal to $d_i \cdot \max(\text{PPI}_r, (\max(R, V)/\lceil k/2 \rceil))$ then she has $\frac{\lceil k/2 \rceil}{A}$ probability of winning. If the valuation of i is more than $d_i \cdot \max(\text{PPI}_r, (\max(R, V)/\lceil k/2 \rceil))$ then by the way we define the allocation function (see Definition 4.7) he has probability $\frac{\lceil k/2 \rceil}{A}$ of winning otherwise her PPI is less than $\max(\text{PPI}_r, (\max(R, V)/\lceil k/2 \rceil))$. If PPI of i is less than PPI_r (the PPI of the runner-up bidder) then she cannot be a winner since the sum of demands of participants who have higher PPI than her is larger than k. If PPI of i is less than $\max(R, V)/\lceil k/2 \rceil$ then she cannot win because of the way we select winners in the allocation function (item 3 of Definition 4.7). □

Now we prove that MMCA satisfies RM.

Lemma 4.5. *If we add one more bidder or a bidder increases her bid the revenue of* MMCA *does not decrease.*

Proof. Let x be the new participant or the participant who has increased her bid. Let θ be the type profile before adding x and θ' be the type profile after adding x. We need to prove that $\text{REVENUE}(\text{MMCA}, \theta') \geq \text{REVENUE}(\text{MMCA}, \theta)$. First we prove that the revenue of MMCA is between the MPRG of the highest group and the second highest group.

Observation 8. *We have* $R \leq \text{REVENUE}(\text{MMCA}, \theta) \leq \text{MPRG}^{g^*}$.

Proof. If item 1 of Definition 4.7 happens then the revenue of MMCA is

$\max(R, \max_{j \in [\lceil k/2 \rceil]} j \cdot u_j, V_2)$ by Lemma 4.4. Note that in this case the revenue is less than V and more than R and hence the proof of the observation follows.

If item 2 of Definition 4.7 happens then the revenue of MMCA is $\max(R, V)$ by Lemma 4.4 since we sell j^* items. Note that in this case the revenue is less than $\max_{j \in [\lceil k/2 \rceil]} j \cdot u_j$ and more than R and hence the proof of the observation follows.

If item 3 of Definition 4.7 happens then the revenue of MMCA is $\max(\lceil k/2 \rceil \cdot \text{PPI}_r, R, V)$ by Lemma 4.4 since we sell $\lceil k/2 \rceil$ items in expectation. Note that in this case the revenue is less than $\max_{j \in [\lceil k/2 \rceil]} j \cdot u_j$ and more than R and hence the proof of the observation follows. □

By Observation 8 we know that the revenue of MMCA is between MPRG of the highest group and MPRG of the second highest group. Therefore if we add participant x and the group with highest MPRG changes then it means that the revenue in θ' is now $\text{MPRG}^{(g^*)}$ and hence increases.

Now we assume that after adding x the winning group does not change. Here we can check all the three cases that can happen in allocation function of MMCA (Definition 4.7) for new type profile θ' and see that adding x can only increase the revenue. □

Now we prove that PoRM of MMCA is at most $O(\ln k)$.

Lemma 4.6. $\text{PoRM}(\text{MMCA}) \leq (2 + \ln(k))$

Proof. First we show that for any group g MPRG of g is at least $\frac{1}{2+\ln(k)}$ fraction of the maximum social welfare obtainable by group g.

Observation 9. *For any group $G^{(g)}$ we have* $\text{MPRG}^{(g)} \geq \frac{1}{(2+\ln(k))} \cdot \text{WF}(G^{(g)})$.

Proof. Let set S contain bidders of $G^{(g)}$ which obtain the maximum social welfare and T be equal to $\text{MPRG}^{(g)}$. Note that set S can contain at most one high-demand bidder since their demand is larger than $\lceil k/2 \rceil$. Moreover, the valuation of high-demand bidder of S cannot be more than T because otherwise MPRG of $G^{(g)}$ will be larger than T. Therefore the total social welfare of S from high-demand bidders is at most T.

Now let us sort the low-demand bidders of S by their PPI and define value $u_{S,j}$ to be the maximum price per item with which we can sell at least j items to the low-demand bidders of S, similar to Definition 4.3. Note that for any $j \in [\lceil k/2 \rceil]$ we have $u_{S,j} \leq \frac{T}{j}$ because otherwise MPRG of $G^{(g)}$ will be larger than T. Moreover for any $j \in \{\lceil k/2 \rceil + 1, \ldots, k\}$ we have $u_{S,j} \leq \frac{T}{\lceil k/2 \rceil}$ because we sort by PPI. Therefore the social welfare of S from low-demand bidders is at most $\sum_{j=1}^{\lceil k/2 \rceil} \frac{T}{j} + \lceil k/2 \rceil \cdot \frac{T}{\lceil k/2 \rceil}$ which is at most $(1 + \ln(k)) \cdot T$.

Set S can get social welfare of at most T from the high-demand bidders and $(1 + \ln(k)) \cdot T$ from low-demand bidders, therefore, the proof of the observation follows. □

Note that by the way we define the allocation function of MMCA it always selects a set of winners such that their expected social welfare is at least the MPRG of the winning group, see Definition 4.7. Since MMCA selects a group with the highest MPRG as the winning group by Observation 9 we know that the maximum social welfare cannot be more than $(2 + \ln(k))$ times the highest MPRG and hence the proof of the lemma follows.

□

Now we provide evidence that no randomized algorithm can obtain PoRM better than $\Omega(\ln(k))$. Note that an optimum randomized mechanism ($\mathcal{M}^*$) has to first give a probability distribution to groups and pick the winning group according to that distribution. Let's assume p^* be the function which maps each group g to its probability of winning. Now we argue that p^* cannot be dependent to the social welfare of g otherwise $\mathcal{M}^*$ will not satisfy RM. Because, in this case a bidder can increase her bid high enough so that probability $p^*(g)$ goes to its upper-limit which decrease the critical value of the other bidders and hence break RM. We further argue that p^* should be dependent to one factor of each group which is the same across the group. For example if it is dependent to two factors then increasing the first factor might remove the load from the second factor and hence the second factor is free to go down without changing the probability. If p^* is dependent to only one factor then the best way to make it as close as possible to social welfare is to define $j \cdot u_j^{(g)}$ (see Definition 4.3) for each group similar to our MPRG value assigned to each group. Therefore, we have evidence that function p^* can be dependent to a value which can be off from social welfare by factor $\ln(k)$.

Now consider the simple case of image-text auction where we have two groups: text-ads and image-ads. We want to assign the winning probabilities to each group. Now suppose we are given two values a for image-ads and b for text-ads, and also we know that social welfare of text-ads is either b or $\ln(k) \cdot b$. If we pick each group with probability $1/2$ and $1/2$, the expected social welfare is $1/2$ approximation to its max value. The question is can we do better by having a more clever randomization.

We prove that no other randomization can give us a factor better than $1/2 + 1/(2\sqrt{\ln(k)})$. Suppose the value of a is equal to X and the value of b is equal to $X/\sqrt{\ln k}$ where the social welfare of text-ads can be $X \cdot \sqrt{\ln(k)}$. Now suppose $\mathcal{M}^*$ gives the items to image-ads with probability F and to text-ads with probability $1 - F$. If the social welfare of text-ads is $X/\sqrt{\ln(k)}$, we get PORM of $F + (1 - F)/\sqrt{\ln(k)}$ and if the social welfare of text-ads is $X \cdot \sqrt{\ln(k)}$, we get PORM of $(1 - F) + F/\sqrt{\ln(k)}$. If we want to maximize the minimum of the two PORMs we have to set F to $1/2$ which gives PORM of $1/2 + 1/(2\sqrt{\ln(k)})$. Note that if the number of groups increases to m then this factor changes to $1/m + 1/(m\sqrt{\ln(k)})$, therefore, the best way is to give probability one to the group with the best assigned value and lose factor $\ln(k)$.

4.7 Discussion and Future Works

Note that our mechanism satisfies the desired IC property and gets a good portion of the maximum social welfare. However, one may worry that by enforcing revenue monotonicity we may lose in total revenue since it is an additional constraint. To address this issue, firstly, we emphasize that lack of revenue-monotonicity can lead to loss of revenue and efficiency from second-order effects that could be much higher than leaving some revenue to attain this property. Most teams in a technology firm treat auction as a black box without knowledge of its internal procedures. For example, assume a team designs a new innovative ad format that increases the click-through-rate of certain ads. Without revenue-monotonicity, such changes can actually lower the revenue. The problem now is that this can result in this innovative change not getting into production system as it might be seen as a *bad* change. This can be a real problem in practice, and as auction designers, we need to address this problem. Secondly, our ultimate goal is indeed to design auctions that satisfy good properties in different dimensions including total revenue. However, one must realize that almost nothing is known theoretically about this important property of *revenue-monotonicity*. Therefore, to fully understand it formally, in our model, we isolate and just focus on RM property while maximizing social welfare. This work can be a starting point for the community to build more results around RM property, and hopefully a general framework to design RM mechanisms emerges. Once we have the state of the art understanding of RM, the next goal is to understand more objectives simultaneously.

CHAPTER 5

Core-competitive Auctions

5.1 Introduction

The VCG mechanism is a powerful mechanism that achieves an efficient outcome in an incentive compatible manner for a variety of scenarios. The simplicity of the VCG mechanism raised our hopes of wide application of this elegant theory in practice. However, it has been noted in the recent past that the applicability of VCG auction beyond the simple case of multiple homogeneous goods has remained limited. Ausubel and Milgrom [AM02] offer an explanation of why VCG in its purest form is often unsuitable to be used in practice. They write:

> [...] higher revenues also improve efficiency, since auction revenues can displace distortionary tax revenues. [...] Probably the most important disadvantage of the Vickrey auction is that the revenues it yields can be very low or zero, even when the items being sold are quite valuable.

To illustrate this point of low or zero revenue, consider the following example from spectrum auctions (taken from [AM02, AM06] which is similar to that of Chapters 3 and 4): consider 3 bidders who are participating in an auction for two spectrum licenses: the first bidder is willing to pay 2 billion for the package of 2 licenses while each of the other two bidders is willing to pay 2 billion for any individual license. The VCG outcome allocates to the second and third bidder, and charges a payment of zero to each of them. This is because the externality each winning bidder imposes on the rest of the bidders is zero. Note that, one can hardly blame the lack of revenue to the absence of competition; if one were to treat it as a market equilibrium problem and compute market clearing prices (say by means of a tatonnement procedure), the revenue would be non-trivial.

Thus, one natural question to ask is, for an auction outcome, how to formally say that it achieves *a competitive revenue*? To answer this, [AM02] introduced the notion of a *core* outcome in an auction setting. The notion of core is a fundamental and well-known notion in cooperative game theory and represents a way to share the utility produced by a group of players in a manner that no sub-group of players would want to deviate. In an auction setting, a set of winning buyers and their payments are said to be a core outcome if no sub-group of losing bidders can propose to the auctioneer (seller) an alternative higher-revenue outcome. For example, in the license example, the outcome implemented by VCG is not in the core since the first bidder (who wanted to purchase two licenses) could negotiate with the auctioneer that the licenses should be allocated to him for any price larger than zero. On the other hand, the outcome which allocates one license each to players 2 and 3 and charges each of them 1 billion is in the core, since in this case there is no alternative outcome that the first player can propose to the auctioneer which would be beneficial for both.

It is noteworthy that when the goods are *substitute*, the VCG outcome is a core outcome, and VCG revenue equals the core-outcome with the minimum revenue (the set of core outcomes is not unique) [AM02]. However, if the goods are not substitutes, the VCG outcome may lie outside the core. In fact, as shown in the above example, VCG revenue can be arbitrarily lower when compared to the minimum-revenue core outcome.

So can one design incentive-compatible auctions whose outcome is always in the core? Unfortunately, one can show that it is impossible to design an auction that (a) achieves a core outcome, and (b) has truth-telling as a dominant strategy equilibrium. So we must either relax (a) or (b). In Ausubel and Milgrom [AM02], the authors relax (b), and give a family of ascending package auctions (called *core-selecting* auctions) which are not truthful but whose equilibrium outcome is a core outcome. These auctions have been extremely successful in practice – variations of these were used in spectrum-license auctions in the United Kingdom, Netherlands, Denmark, Portugal, and Austria, and in the auction of landing-slot rights in the three New York City airports. See [DC12] for a complete discussion.

The focus of this chapter is on applications in Internet ad auctions (we will call them ad auctions from now on). There are several ad auction scenarios which are modeled as goods with complementarities. As a case study for our work, we use a very common scenario in ad auctions which has complementarities, namely that of Text-and-Image ad auction introduced in Chapters 3 and 4. In a Text-and-Image ad auction scenario an ad slot on a page can either accommodate k text ads (which are the traditional ads displayed next to search results) or one large image-ad. Notice that the example by Ausubel and Milgrom can be reproduced exactly in this setting by setting $k = 2$.

What auction should we use for the Text-and-Image setting? The core-selecting auction of [AM02] is not a good choice for this setting as the ascending package auctions are interactive procedures in which bidders submit a sequence of bids after provisional allocations and prices for the previous phase are revealed; such designs often result in long and time-consuming procedures which are justified for one-time spectrum auctions but unsuitable for Internet advertisement[1]. Moreover, because of the fast-paced nature of online advertisement, one cannot expect bidders to reach an equilibrium outcome for each individual ad auction if the underlying auction is not a truthful one.

In this chapter we investigate whether it is possible to design direct-revelation incentive-compatible auctions whose revenue is competitive against a core outcome (we call such auctions *core-competitive* auctions). In core-competitive auction design, we seek to relax (a) instead of (b) above. More precisely, we define *core revenue benchmark* as the smallest revenue among all the core-outcomes. We say that an auction is α-core-competitive if its revenue is at least an $1/\alpha$ fraction of the core revenue benchmark.

[1] One can eliminate the interactive aspect of package bidding auction by using a *proxy agent*, as Ausubel and Milgrom discuss in Section 3.4 of [AM02]. While this technique eliminates the communication burden, it is not enough to achieve incentive-compatibility.

We formally define the notion of core-competitiveness in Section 5.2, and later we focus on the design of core-competitive auctions for the Text-and-Image setting. We give a randomized universally-truthful mechanism which is $O(\ln \ln k)$-core-competitive, where k is the number of slots. We also give a lower bound showing that this factor is tight. We note that in ad auction settings, there are several repeated auctions with each auction generating only a small revenue. For such settings, a seller care about the overall performance and therefore randomized auctions are perfectly fine from a practical auction design perspective. We also study deterministic auctions since for some settings randomization may not be desired; for instance, for one time auctions like spectrum auctions. We give a deterministic mechanism which is $O(\sqrt{\ln(k)})-$core-competitive, and again show that this factor is tight for deterministic mechanisms.

Finally, to the best of our knowledge, the notion of core-competitiveness has not been studied before. It is our belief that developing tools and techniques for designing core-competitive auctions, and understanding the possibilities and limitations of such auctions, will be very useful from a practical auction design perspective.

5.1.1 Related Work

The line of inquiry that seeks to design package auctions that implement core outcomes in equilibrium was started by Ausubel and Milgrom [AM02]. This line has been further developed in [DM08,AB10,DC12,EK10,GL09,Lam10]. The authors design an iterative procedure that asks bidders in each round for packages they want to bid on as well as bid values for each of those packages. In each round a set of provisionally winning bids are identified. This proceeds until no further bids are issued in a given round. Our work differs from this line of work in the sense that we require incentive-compatibility; in the core-selecting package auctions literature, the focus is on implementing core outcomes *in equilibrium*.

Another stream of related work is the design of incentive compatible auctions that tries to optimize for revenue in a prior-free setting. This research direction was initiated in [GHW01, FGHK02, GH03] and resulted in a sequence of followup results which are too large to survey here. We refer to Hartline's book [Har13] for a comprehensive discussion. The first successful results gave auctions for the digital goods that approximate the $\mathcal{F}^2$ revenue benchmark, the maximum revenue one can extract from at least two players using fixed prices. More modern versions of this result [HY11, HH13, DHH13] compare against the envy-free benchmark (how much revenue it is possible to extract from an outcome where any two agents wouldn't like to swap places). This resulted in success stories for a large class of environments such as multi-units, matroids and permutation environments. The work of this chapter differs from the above line of work as we consider environments with complementarities, while the envy-free revenue literature mostly focused on environments with substitutes. In Section 5.2.4 we discuss in detail the relation between the envy-free benchmark and the core-revenue benchmark and we argue that the core-revenue benchmark captures some of the *no-envy* notions.

Closer to our line of inquiry is the work of [MV07] and [AH06]. In [MV07], they design revenue extraction mechanisms for general combinatorial auctions where their benchmark is the maximum social welfare extractable from all except one player (the one with the top bid). They use randomization to obtain a mechanism with $O(\log n)$ approximation factor. They also give a matching lower bound of $\Omega(\log n)$ for randomized mechanisms, and for deterministic mechanisms they give a lower bound of $\Omega(n)$. [AH06] study knapsack auction where there are k identical items and each bidder demands a certain number of them. Their benchmark is a version of envy-free pricing where a bidder has to pay at least as much as the bidders with lower demands [2]. They get an approximation ratio of $\alpha \cdot \mathit{OPT} - \lambda O(\log \log \log n)$ where *OPT* is the optimal envy-free revenue, α is a constant number and λ equals to the highest valuation of any bidder. Although their approach is useful when λ is much smaller than *OPT*; it performs poorly when λ is close to *OPT* which can be the case in the Image-and-Text auction.

We note that the revenue benchmarks of both the above papers are stronger than the core-benchmark. Thus, one might wonder if the mechanisms proposed in [MV07] and [AH06] perform better against the core benchmark? However, one can show that mechanisms given in both the above papers perform worse than our mechanism when compared to the core benchmark. The mechanism of [AH06] can perform arbitrarily bad compared to the core benchmark, and the mechanism of [MV07] still gets only $O(\log n)$ using randomization when compared to the core benchmark[3]. In some sense, this sug-

[2]This is also called monotone benchmark, see also [LR12] and [BKK+13] for its definition on digital goods auction.

[3]We refer to Section 5.2.4 for further discussion on this.

gests that a too strong benchmark that leads to large lower bounds in approximation ratio impedes the design of a good revenue-maximizing mechanism. We believe that the core benchmark is a more fundamental benchmark (as argued in series of papers starting with the work of [AM06]), and as our work shows, it looks amenable to a good multiplicative approximation ratio.

Finally, while we focus on the *forward* setting (i.e. an auctioneer selling goods to various buyers), there is a very extensive literature on the procurement (reverse auction) version of this problem (i.e. a buyer purchasing goods from various sellers). In this line of work, the goal is to design procurement auctions where the total amount paid by the buyer approximates a certain *frugality benchmark*. This line of work was initiated in [AT07] in which the frugality benchmark is defined as the best solution after the agents in the optimal solution are removed. A more sophisticated frugality benchmark was introduced in [KK05]. Their benchmark can be seen as the counterpart of the core-revenue benchmark in procurement settings. Frugality in the procurement setting is also a topic which is too broad to be completely covered here, but we mention a few recent papers on the topic: [KSM10, CEGP09, EGG07, IKNS10].

5.2 Preliminaries

5.2.1 Core Outcomes

We consider set $N = \{1, \ldots, n\}$ of single-parameter agents with value v_i for being allocated and value zero otherwise. The set of feasible allocations is specified by an *envi-*

ronment, which is a collection of subsets of players that can be simultaneously allocated $F \subseteq 2^N$. We say that an environment is *downward-closed* if every subset of a feasible set is also feasible, i.e., $X \in F$ and $Y \subseteq X$ imply $Y \in F$.

An outcome in such environment is a pair (X, p) where $X \in F$ corresponds to the selected set of players and $p \in \mathcal{R}^N$ is a vector of (possibly negative) payments. Players have *quasi-linear* utility functions, i.e., $u_i(X, p) = v_i - p_i$ if $i \in X$ and $u_i(X, p) = -p_i$ otherwise. We also define the utility of the auctioneer as its revenue $u_0(X, p) = \sum_{i=1}^n p_i$.

Throughout this chapter, given a vector $v \in \mathcal{R}^N$ and $S \subseteq N$, we define $v(S) := \sum_{i \in S} v_i$.

We can associate with the single parameter setting described above a *coalition value function* $w : 2^{\bar{N}} \to \mathcal{R}_+$ (where $\bar{N} = \{0\} \cup N$) given by:

$$w(S) = \begin{cases} \max_{X \in F, X \subseteq S, p \in \mathcal{R}_+^N} \sum_{i \in S} u_i(X, p) & 0 \in S \\ 0 & 0 \notin S \end{cases}$$

for every $S \subseteq \bar{N}$. The pair $(\bar{N}, w)$ defines a *cooperative game* with transferable utility. The coalition value of a set corresponds to the total utility that can be obtained by a certain set by defecting from the rest of the agents. Clearly, a coalition that doesn't contain the auctioneer can't obtain any value. A coalition containing the auctioneer can obtain utility equal to $\max_{X \in F, X \subseteq S, p \in \mathcal{R}_+^N} \sum_{i \in S} u_i(X, p) = \max_{X \in F, X \subseteq S} v(X)$.

An imputation of utilities for a coalition $S \subseteq \bar{N}$ corresponds to a vector of utilities $(u_i)_{i \in S}$ specifying how the coalition value is split between the agents, in other words, a vector $u_i \geq 0, \forall i \in S$ and $\sum_{i \in S} u_i \leq w(S)$. We say that an imputation of utilities for $\bar{N}$ is in the *core* if no coalition can defect and produce an imputation of utilities that is better for all agents in the coalition. Formally:

Definition 5.1 (core)**.** *Given a cooperative game $(\bar{N}, w)$ we define the core as the following set of utility imputations:*

$$\text{Core}(F, v) = \left\{ u \in \mathcal{R}_+^{\bar{N}}; \sum_{i=0}^{n} u_i = w(\bar{N}) \textit{ and } w(S) \leq \sum_{i \in S} u_i, \forall S \subseteq \bar{N} \right\}$$

Notice that $w(S) \leq \sum_{i \in S} u_i$ is a necessary and sufficient condition for S not wanting to defect. We say now that an outcome (X, p) is in the core if the utilities produced are in $\text{Core}(F, v)$. Precisely:

Definition 5.2 (core outcomes)**.** *Given a single parameter setting F and valuation profile v, an outcome (X, p) is in the core if the vector of utilities is in* $\text{Core}(F, v)$.

The following are important properties of core outcomes:

1. A core outcome is also a social welfare maximizing outcome, since $\sum_{i \in X} v_i = \sum_{i=0}^{n} u_i = w(\bar{N}) = \max_{X^* \in F} v(X^*)$;

2. The core is always non-empty, since the following allocation is always in the core: (X^*, p) where X^* maximizes $v(X)$ and $p_i = v_i$ or $i \in X^*$ and $p_i = 0$ otherwise;

3. Given a utility imputation $u \in \text{Core}(F, v)$, there is a core outcome that realizes this vector: select a set $X^* \in F$ maximizing $\sum_{i \in X^*} v_i$ and allocate to X and charge prices $p_i = v_i - u_i$ for $i \in X^*$ and $p_i = 0$ otherwise. The outcome clearly realizes

utilities for $i \in X^*$. For $i \notin X^*$, notice that $w(\bar{N}) = \sum_{i=0}^{n} u_i = (u_0 + \sum_{i \in X^*} u_i + \sum_{i \in N \setminus X^*} u_i) \geq w(\bar{N}) + \sum_{i \in N \setminus X^*} u_i$. So for all $i \in N \setminus X^*$, $u_i = 0$;

4. If the environment F is downward-closed, then for every $u \in \text{Core}(F, v)$ there is an outcome with non-negative payments that realizes it. The construction is the same as in the previous item. Note that if F is downward closed, $X^* \setminus i \in F$ for every $i \in X^*$, therefore: $v(X^*) = u_0 + u(X^*) \geq u_i + v(X^* \setminus i)$ so $u_i \leq v_i$ and hence $p_i = v_i - u_i \geq 0$.

The previous observations allow us to rephrase Definition 5.2 in a more direct way. Notice that in the following definition, $v(S \setminus X) \leq p(X \setminus S)$ is a simple rephrasing of the $w(S) \leq \sum_{i \in S} u_i$ condition.

Definition 5.3 (core outcomes - rephrased)**.** *Given a single parameter setting F and valuation profile v, an outcome (X, p) is in the core if $p_i \leq v_i$ for all $i \in N$ and for all $S \in F$,*

$$v(S \setminus X) \leq p(X \setminus S)$$

Definition 5.3 allows for a natural interpretation of the core in auction settings. If an outcome is not in the core, then there is a set S with $v(S \setminus X) > p(X \setminus S)$, which means that agents in $S \setminus X$ could come to the auctioneer and offer him to evict agents $X \setminus S$ and allocate to them instead, since they are able to collectively pay the auctioneer more than the revenue he is getting from $X \setminus S$. This characterizes core outcomes as outcomes for which no negotiation is possible between the auctioneer and losing coalitions

5.2.2 Core-revenue Benchmark

The discussion after Definition 5.3 shows that whenever an outcome is not in the core, the auctioneer can potentially raise his revenue by negotiating with losing coalitions. This suggests that the revenue of the core might be a natural benchmark against which to compare. We define as follows:

Definition 5.4. *Given a single parameter setting F and a valuation profile v, we define the core revenue benchmark as:*

$$\text{CoreRev}(F, v) := \min\{u_0 | u \in \text{Core}(F, v)\}.$$

Consider for example the case of multi-unit auctions, which can be modeled by $F = \{X \subseteq N; |X| \leq k\}$ for some fixed constant $k < n$ and agents sorted such that $v_1 > v_2 > \ldots > v_n$. It is straightforward from Definition 5.3 that an outcome is in the core iff it allocates to $X = \{1, \ldots, k\}$ and if $p_i \geq v_{k+1}$ for $i \in X$. Notice that the revenue from core outcomes range from $k \cdot v_{k+1}$ all the way to $\sum_{i=1}^{k} v_k$. The core benchmark corresponds to the minimum revenue of a core outcome, so for multi-unit auctions $\text{CoreRev}(F, v) = k \cdot v_{k+1}$.

It is not a coincidence that this is the same revenue as the VCG auction. In fact, it is a well-known fact that the core revenue is always at least the VCG revenue. This holds with equality when F is a matroid. For an in-depth discussion on the relation between the VCG mechanism and the core we refer the reader to Ausubel and Milgrom [AM02] and Day and Milgrom [DM08].

Lemma 5.1 ([AM02]). *For any environment F and any valuation profile v, the price paid by any agents in a core outcome is at least his VCG price. This implies in particular that the core revenue benchmark is at least the revenue of the VCG mechanism:*

$$\mathrm{CoreRev}(F, v) \geq \mathrm{VcgRev}(F, v) := \sum_{i \in X^*} [v(X^*_{-i}) - v(X^*) + v_i]$$

where $X^* = argmax_{X \in F} v(X)$ and $X^*_{-i} = argmax_{X \in F, i \notin F} v(X)$. Moreover, if F is a matroid, the the above expression holds with equality.

Proof. If (X, p) is a core outcome, by the condition in Definition 5.3, $v(X^*_{-i} \setminus X^*) \leq p(X^* \setminus X^*_{-i})$, which can be re-written as: $v(X^*_{-i}) - v(X^*) \leq -[v(X^* \setminus X^*_{-i}) - p(X^* \setminus X^*_{-i})] \leq v_i - p_i$. So $p_i \geq v(X^*_{-i}) - v(X^*) + v_i$ which is the revenue that the VCG mechanism extracts from player i.

If F is a matroid, then for each $i \in X^*$, X^*_{-i} is of the form $X^*_{-i} = X^* \cup j \setminus i$ and therefore the VCG payments are given by $p_i = \max\{v_j; j \notin X^*; X^* \cup j \setminus i \in F\}$. Now, we show that the VCG outcome is in the core: for any matroid basis $S \in F$, there is a one-to-one mapping between $\sigma : S \setminus X \to X \setminus S$ such that for $i \in S$ with $v_i > 0$, $X \cup i \setminus \sigma(i) \in F$, therefore, $p_{\sigma(i)} \leq v_i$. Summing this inequality for all $i \in S$ we obtain the core condition in Definition 5.3. □

The previous lemma says that when there is *substitutability* among agents, the core revenue benchmark is exactly the VCG revenue. When there are complementarities, however, the core revenue benchmark can be arbitrarly higher than the VCG revenue. Consider for example the famous example of [AM02, AM06] in which there are 3 players and 2 items: the first player has a valuation of 1 for the first item, the second player has a valuation of 1 for the second item and the third player has a valuation of 1 for getting both items. This example can be translated to our setting by taking the environment to be $F = \{\emptyset, \{1\}, \{2\}, \{3\}, \{1, 2\}\}$. The VCG auction allocates $X = \{1, 2\}$ and charges zero payments. So, $\mathrm{VcgRev}(F, v) = 0$. The core revenue, however, is equal to one ($\mathrm{CoreRev}(F, v) = 1$) since by taking the condition in Definition 5.3 with $X = \{1, 2\}$ and $S = \{3\}$, we get: $p_1 + p_2 \geq v_3 = 1$.

5.2.3 Core Competitive Auctions

Our goal in this paper is to be able to truthfully extract revenue that is competitive with the core-revenue benchmark. An auction for the single parameter setting consists of two mappings: (i) *allocation function*, that maps a profile of valuation functions to a distribution over allocations $x : \mathcal{R}_+^n \to \Delta(F)$, where $\Delta(F)$ denotes the set of probability distributions over F; (ii) *payment function*, that maps a profile of valuation functions to the expected payment of each agent: $p : \mathcal{R}_+^n \to \mathcal{R}_+^N$.

We abuse notation and define the maps $x_i : \mathcal{R}_+^N \to [0, 1]$ as the probability of winning for player i, i.e., $x_i(v) = p[i \in X(v)]$. A mechanism is said to to be *individually rational* if for all profiles v, $u_i(v) = v_i x_i(v) - p_i(v) \geq 0$. A mechanism is said to be *incentive-compatible* (a.k.a. *truthful*) if agents maximize their utility by reporting their true value. In other words:

$$v_i x_i(v) - p_i(v) \geq v_i x_i(v_i', v_{-i}) - p_i(v_i', v_{-i}) \quad \forall v_i'$$

The following lemma due to Myerson [Mye81] gives necessary and sufficient conditions for an auction to be individually rational and incentive compatible:

Lemma 5.2 ([Mye81]). *A mechanism defined by maps x and p is individually rational and incentive compatible if: (i) for every i and fixed valuations v_{-i} for other players, $v_i \mapsto x_i(v_i, v_{-i})$ is monotone non-decreasing; (ii) the payment function is such that $p_i(v_i) = v_i x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(u, v_{-i}) du$.*

Our goal in this chapter is to study auctions whose revenue is competitive with the core-revenue benchmark.

Definition 5.5 (core competitive auctions). *We say that an auction defined by x, p is α-core competitive if for every profile of valuation functions $v \in \mathcal{R}_+^N$,*

$$\sum_i p_i(v) \geq \alpha^{-1} \cdot \mathrm{CoreRev}(F, v).$$

5.2.4 Comparison with Other Benchmarks

A natural question at this point is how does the core benchmark compare with other revenue benchmarks. Perhaps one of the closest benchmarks in this spirit is the envy-free benchmark, which corresponds to the minimum revenue of an allocation for which an agent would not want to trade positions with a different agent. This benchmark has been successfully used in various papers ([GHK+05, HY11, HH13, DHH13] to cite a few) to design approximately-optimal revenue-extracting mechanisms. This benchmark, however, is very appropriate for *symmetric* settings, i.e., a setting in which whenever an allocation is feasible, a similar allocation with the names of agents permuted is also feasible. For asymmetric settings, however, it is not clear what the envy-freedom condition means since some agents can't be simply replaced by others. In an ad auction where ads can be either texts (occupying one slot) or images (occupying multiple slots), it is not clear how to define what the envy of an image for a text means, since the image is not able to replace a single text.

On the other hand, however, the core-revenue benchmark captures some notion of "envy", which is made explicit in Definition 5.3. One can think of the inequality in the defintion as the "envy" of an allocated image for a group of allocated text ads. Or more generaly, as the "envy" of a set of losing players for a set of winning players that they can replace. What the core benchmark doesn't capture, however, is the "envy" from one allocated agent for another allocated agents. For this reason, for symmetric settings, the envy free benchmark can be arbitrarily higher than the core-revenue benchmark, which boils down to the VCG revenue, as discussed in Section 5.2.1.

Another important benchmark against which to compare is the one introduced by Micali and Valiant [MV07]. Given any feasiblity set, the authors define as the maximum social welfare obtainable after the largest valued agent is excluded. Formally:

$$\mathrm{MV}(F, v) = \max_{X \in F, i^* \notin X} v(X)$$

where i^* is the agent with largest value[4].

[4]The benchmark of [MV07] is defined for a generic multi-parameter setting. For the exposition, we specialize it for the single-parameter setting we are studying.

Lemma 5.3. *For any environment F and any valuation profile v, the core revenue benchmark is dominated by the Micali-Valiant benchmark:*

$$\mathrm{MV}(F, v) \geq \mathrm{CoreRev}(F, v).$$

Proof. Let (X, p) be the outcome of the VCG auction. Now, define p' such that $p'_i = v_i$ if $i \in X \setminus i^*$, $p'_{i^*} = p_{i^*}$ and $p'_i = 0$ otherwise. First we show that (X, p') is in the core. Notice that if $i^* \notin X$, then $p'(X \setminus S) = v(X \setminus S) \geq v(S \setminus X)$ so clearly (X, p') is in the core. If $i^* \in X$, then $p'_{i^*} = v(X_{-i^*}) - v(X \setminus i^*)$ where X_{-i^*} is the allocation with $X \in F, i^* \notin X$ maximizing $v(\cdot)$. Therefore $p(X) = v(X \setminus i^*) + p'_{i^*} = v(X_{-i^*})$ therefore, $p(X \setminus S) = v(X_{-i^*}) - p(X \cap S) \geq v(S) - v(X \setminus S) = v(S \setminus X)$. Finally, notice that $\mathrm{CoreRev}(F, v) \leq p'(X) = \mathrm{MV}(F, v)$. □

Micali and Valiant [MV07] give an individually rational and incentive compatible randomized mechanism whose revenue is an $O(\log n)$ approximation of $\mathrm{MV}(F, v)$ and that such approximation factor is tight. This directly translates in a same factor approximation for the core-revenue benchmark. They also show that no deterministic auction can approximate $\mathrm{MV}(F, v)$ by a factor better then $\Omega(n)$.

One reason for which it is hard to improve the MV-benchmark even for very simple settings, MV is too stringent: for example, for the digital goods setting $F = 2^N$, $\mathrm{MV}(F, v) = \sum_i v_i - \max_i v_i$. Indeed, both lower bounds in [MV07] are given for the digital goods setting. For this setting, the core-revenue benchmark is zero, since there is no natural competition among the agents.

We believe that the core revenue benchmark provides a more achievable goal and therefore a more likely avenue for improvement for particular settings. For the Text-and-Image setting, for example, the lower bounds of [MV07] imply that no mechanism can approximate the MV-benchmark by a better factor then $\Omega(\log k)$ for randomized mechanisms and $\Omega(k)$ for deterministic mechanisms. For the $\mathrm{CoreRev}$-benchmark, however, we are able to obtain $O(\ln \ln k)$ and $O(\sqrt{\ln(k)})$ respectively.

The core revenue also has the important property of disentangling the problems of achieving high revenue for setting with substitutes and for settings with complements, since the former becomes trivial under the $\mathrm{CoreRev}$-benchmark while the latter is quite challenging. Under the MV-benchmark, both substitutes and complements are challenging.

5.3 $O(\sqrt{\ln(k)})$-core-competitive Auction for Text-and-Image Setting

5.3.1 Text-and-Image Setting

Consider k advertisement slots and n bidders. Each bidder either corresponds to a text ad, which demands one slot to be displayed, or an image, which demands all k slots. It is public information that whether each bidder is a text or an image. Each bidder's value for being displayed is given by v_i^T for text ads and v_i^I for image ads. The values are private information of the bidders.

Let n^T and n^I be the number of text and image ads respectively. We assume w.l.o.g. that $n^T \geq k + 1$ and $n^I \geq 2$ (adding a few extra bidders with value zero if necessary). We also assume that the indices of the players are sorted such that valuations of text ads are $v_1^T \geq v_2^T \geq \ldots \geq v_{n^T}^T$ and valuations of image ads are $v_1^I \geq v_2^I \geq \ldots \geq v_{n^I}^I$. For convenience, we define the *maximum extractable revenue of text ads* as:

$$\Phi^T := \max_{j \in \{1..k\}} j \cdot v_j^T$$

We will also denote the k-th harmonic partial sum by $\mathcal{H}_k = \sum_{j=1}^k \frac{1}{j} = O(\ln k)$. It is a well known fact that

$$\Phi^T \geq \frac{1}{\mathcal{H}_k} \sum_{j=1}^k v_j^T, \tag{5.1}$$

since $j \cdot v_j^T \leq \Phi^T$ for all j, so $\frac{1}{j}\Phi^T \geq v_j^T$. We finish the argument by summing the previous inequality for all $j = 1..k$.

5.3.2 A Deterministic Core-competitive Auction

We start by presenting a $O(\sqrt{\ln(k)})$-core competitive deterministic auction. We will use this mechanism as a building block for the more complicated randomized mechanism given in Section 5.4. As a first step, we provide a characterization of the core-revenue in that setting:

Lemma 5.4. *Given a Text-and-Image setting, if the highest value feasible set consists of text ads ($\sum_{i=1}^k v_i^T \geq v_1^I$) then* $\mathrm{CoreRev}(F, v) = \max\{k v_{k+1}^T, v_1^I\}$. *If the highest value feasible set consists of an image ad, then* $\mathrm{CoreRev}(F, v) = \max\{v_2^I, \sum_{i=1}^k v_i^T\}$.

Proof. Note that in Text-and-Image setting the winner set cannot contain both text and image ads. Now consider the special case where $\sum_{i=1}^k v_i^T = v_1^I$. In this case no matter from which group is the winning set, the sum of payments has to be at least $\sum_{i=1}^k v_i^T = v_1^I$.

Because if the sum of payments is less, then the non-winning group can offer more to the auctioneer and all of them benefit more.

Now consider the case where $\sum_{i=1}^{k} v_i^T > v_1^I$. In this case the winners are the first k text ads with sum of valuations $\sum_{i=1}^{k} v_i^T$. In order to be a core outcome, the sum of payments of the winners has to be more than valuations of image ads and hence more than v_1^I. The payment of each winner also has to be more than the valuation of the highest text ad who is not in the winning set which is v_{k+1}^T. Therefore, the sum of payments of the winners has to be more than $k \cdot v_{k+1}^T$. We conclude for this case that $\mathrm{CoreRev}(F, v) = \max\{kv_{k+1}^T, v_1^T\}$.

Now consider the case where $\sum_{i=1}^{k} v_i^T < v_1^I$. In this case, the winner is an image ad with value v_1^I. In order to be a core outcome, the payment of the winner has to be at least the value of the second best image ad which is v_2^I. The payment of the winner also has to be more than the sum of valuations of the highest k text ads which is $\sum_{i=1}^{k} v_i^T$. We conclude that $\mathrm{CoreRev}(F, v) = \max\{v_2^I, \sum_{i=1}^{k} v_i^T\}$. □

Recall the example by Ausubel and Milgrom discussed in the introduction: if we have two text ads and one image ad all with value 1, the text ads are selected and their payment is zero. The reason for that is that if any text ad decreases his value all the way to $\epsilon > 0$, the text ads are still selected. One way to get around this problem is picking the allocated set in such a way that a decrease in value for any given text significantly decreases the likelihood of the entire set being picked.

A natural way to do so is to allocate to the set which has the potential of generating the largest revenue. One proxy for that is the maximum extractable revenue Φ^T which corresponds to the maximum revenue you can extract by setting a uniform price. This motivates the mechanism that allocates to the highest value image ad if $v_1^I \geq \Phi^T$ and otherwise allocates to the j highest text ads where j is the maximum index such that $jv_j^T \leq v_1^I$. Here the payments are according to critical prices.

In the Ausubel and Milgrom example, for instance, the text ads are still allocated but their threshold is now $\frac{1}{2}$, so their total revenue is 1. This mechanism is clearly truthful since the allocation is monotone and its revenue is clearly an improvement over VCG. The gap between its revenue and the core-revenue benchmark can be as bad as $O(\ln k)$. Consider the following example: one image ad with value $\mathcal{H}_k$ and k text ads with value $1/i$ for $i = 1, \ldots, k$. The core-benchmark is $\mathcal{H}_k$ but the revenue of the mechanism is only $\Phi^T = 1$.

A way to improve this mechanism is to increase the weight attributed to the text ads by a factor of $\sqrt{\ln(k)}$. Now, we are ready to define our mechanism:

Allocation rule: If $v_1^I \geq \Phi^T \cdot \sqrt{\ln(k)}$ then allocate to the highest value image ad. Otherwise, allocate to the j text ads with largest values where j is the largest $j \leq k$ such that $j \cdot v_j^T \geq v_1^I / \sqrt{\ln(k)}$.

Pricing rule: Allocated bidders are charged according to critical values.

Lemma 5.5. *In the deterministic Text-and-Image mechanism, if the first image ad wins, her critical value is* $\max\{v_2^I, \Phi^T \cdot \sqrt{\ln(k)}\}$*. If a set of* j *text ads win, their critical value* *is* $\max\{v^T_{k+1}, v^I_1/(j \cdot \sqrt{\ln(k)})\}$.

Proof. Recall that the critical value of each winner is the minimum bid for which she remains a winner fixing the other bidders' bids.

The case where the winner is an image ad is easy to proof. Note that in this case the winner has value v^I_1. The minimum bid in order to remain the winner has to be at least the value of the second highest image ad which is v^I_2 and has to be larger than $\Phi^T \cdot \sqrt{\ln(k)}$ to win against text ads. Therefore, the critical value of the winner is $\max\{v^I_2, \Phi^T \cdot \sqrt{\ln(k)}\}$.

Now we consider the case where the winners are the j highest text ads. If $\max\{v^T_{k+1}, v^I_1/(j \cdot \sqrt{\ln(k)})\}$ is equal to v^T_{k+1}, we have $v^I_1/\sqrt{\ln(k)} \leq k \cdot v^T_{k+1}$ hence $j = k$ by the way we select j. Hence, the first k text ads win. Moreover, the winners' payments has to be at least v^T_{k+1} in order to be in the first k text ads, therefore, the critical value of the winners is $\max\{v^T_{k+1}, v^I_1/(j \cdot \sqrt{\ln(k)})\} = v^T_{k+1}$.

If $\max\{v^T_{k+1}, v^I_1/(j \cdot \sqrt{\ln(k)})\}$ is equal to $v^I_1/(j \cdot \sqrt{\ln(k)})$, we prove by contradiction that the critical value of winners is $v^I_1/(j \cdot \sqrt{\ln(k)})$. Lets assume that there exist value v' ($v' < v^I_1/(j \cdot \sqrt{\ln(k)})$) such that if a winner (W) bids v', she remains in the winning set and hence v' is her critical value. Let j' be the number of winners when W bids v'. We know that the value of Φ^T is at most $j' \cdot v'$ since W is in the winning set. The value of Φ^T has to be greater than $v^I_1/\sqrt{\ln(k)}$ in order for text ads to win against image ads. Therefore, we have $j' \cdot v' \geq v^I_1/\sqrt{\ln(k)}$. On the other hand we have $v' < v^I_1/(j \cdot \sqrt{\ln(k)})$ which implies $j \cdot v' < v^I_1/\sqrt{\ln(k)}$. Hence we conclude that $j' > j$ which contradicts with the fact that j is the largest number such that $j \leq k$ and $j \cdot v_j$ is larger than or equal to

$\frac{v_1^I}{\sqrt{\ln(k)}}$. □

Using the previous two lemmas we prove the following theorem and finish this section.

Theorem 5.1. *The deterministic Text-and-Image mechanism is $O(\sqrt{\ln(k)})$-core competitive.*

Proof. We prove the theorem by considering two cases: Case (i) when the first image ad wins and Case (ii) when the first j text ads win.

In Case (i) the winner is the image ad with value v_1^I and his payment $\max\{v_2^I, \Phi^T \cdot \sqrt{\ln(k)}\}$ by Lemma 5.5 is the revenue of our deterministic Text-and-Image mechanism. The value of $\mathrm{CoreRev}$ in this case is $\max\{v_2^I, \sum_{i=1}^{k} v_i^T\}$ by Lemma 5.4. Therefore, using Equation 5.1 we conclude that the revenue of our deterministic Text-and-Image mechanism is at least $\sqrt{\ln(k)}$ fraction of $\mathrm{CoreRev}$.

In Case (ii) the winners are the first j text ads. By Lemma 5.5 we know that their critical value is $\max\{v_{k+1}^T, v_1^I/(j \cdot \sqrt{\ln(k)})\}$. If their critical value is equal to $v_1^I/(j \cdot \sqrt{\ln(k)})$ then the total revenue of the mechanism is $v_1^I/\sqrt{\ln(k)}$. If their critical value is equal to v_{k+1}^T then it means that $v_{k+1}^T \geq v_1^I/(j \cdot \sqrt{\ln(k)})$, hence j is equal to k since j is the largest $j \leq k$ such that $j \cdot v_j^T \geq v_1^I/\sqrt{\ln(k)}$. Therefore, the total revenue is $k \cdot v_{k+1}^T$. As a result the total revenue in Case (ii) is $\max\{v_1^I/\sqrt{\ln(k)}, k \cdot v_{k+1}^T\}$. The value of $\mathrm{CoreRev}$ in this case is $\max\{kv_{k+1}^T, v_1^I\}$ by Lemma 5.4. Therefore the revenue of our deterministic Text-and-Image mechanism is at least $\sqrt{\ln(k)}$ fraction of $\mathrm{CoreRev}$. □

5.3.3 A $O(\sqrt{\ln(k)})$ Lower Bound for Deterministic Mechanisms

Now we show that $O(\sqrt{\ln(k)})$ is necessary for deterministic core-competitive mechanisms. Formally, we show that no mechanism that is anonymous and satisfies independence of irrelevant alternatives can provide an approximation ratio better then $O(\sqrt{\ln(k)})$. A word of caution: while anonymity and independence of irrelevant alternatives are commonly used assumptions in lower bounds for deterministic mechanisms [ADL12], they are not completely innocuous as shown by [AFG+11].

Definition 5.6. *A mechanism ($\mathcal{M} = (x,p)$) is anonymous if the following holds. Let v and v' be two valuation profiles that are permutations of each other (i.e. the set of valuations are the same but the identities of bidders are permuted). Say $v = \text{permutation}(v')$. If $x(v) = S_1$ and $x(v') = S'$, then $S' = \text{permutation}(S)$.*

Definition 5.7. *A mechanism ($\mathcal{M} = (x,p)$) satisfies independence of irrelevant alternatives if we decrease the bid of a losing participant, it does not hurt any winner. More formally, for every valuation profile v and loser participant $i \notin x(v)$, if we decrease the value of i from v_i^T to $\hat{v}_i^T < v_i$ then $x(v) \subseteq x(\hat{v}_i^T, v_{-i})$.*

Theorem 5.2. *Let M^* be a deterministic mechanism with optimum core competitive factor satisfying anonymity and independence of* *irrelevant alternatives. Then there exist a valuation profile for which revenue of M^* is at most $\sqrt{\ln(k)}$ of* $\mathrm{CoreRev}$.

Proof. Let valuation profile v consists of k text ads $\{v_1^T, v_2^T, \ldots v_k^T\}$ where value v_i^T is equal to $1/i$ and 2 image ads $\{v_1^I, v_2^I\}$ both with value $\sqrt{\ln(k)}$. Now we consider two cases:

Case (i) M^* allocates to an image ad. Note that the revenue of $\mathcal{M}^*$ is the payment of the winner and is at most $\sqrt{\ln(k)}$. Now, lets increase the valuation of the winner to $\ln(k)$ and build a new valuation profile v'. Note that by Lemma 5.2 the winner and his payment in v' remains the same as in v. Therefore, the revenue of v' is $\sqrt{\ln(k)}$ while its $\mathrm{CoreRev}$ by Lemma 5.4 is $\ln(k)$.

Case (ii) M^* allocates to a set of text ads . We build a group of k valuation profiles $v^{(1)}, \ldots, v^{(k)}$ and show that in at least one of them the difference between $\mathrm{CoreRev}$ and revenue of $\mathcal{M}^*$ is $\sqrt{\ln(k)}$. valuation profile $v^{(1)}$ is the same as v and we build $v^{(i+1)}$ from $v^{(i)}$ by the following procedure. If text ad v_{i+1}^T is a winner in $v^{(i)}$ then we obtain $v^{(i+1)}$ by increasing value of v_{i+1}^T to one in $v^{(i)}$. Otherwise, if text ad v_{i+1}^T is a loser in $v^{(i)}$ then we obtain $v^{(i+1)}$ by decreasing value of v_{i+1}^T to zero in $v^{(i)}$.

Let j be the largest number such that $j \leq k$ and text ad v_j^T is a winner in $v^{(j)}$. Now we claim that every text ad j' where $j' > j$ is a loser in $v^{(j)}$. Otherwise, if such j' exist then j' will also be a winner in $v^{(j')}$ since by independence of irrelevant alternative j' remains a winner in all valuation profiles $v^{(\ell)}$ for $j < \ell < j'$. This contradicts with the fact that j is the largest number. Therefore, we know that in valuation profile $v^{(j)}$ all the winners are between 1 and j and hence we have at most j winners. Note that $v^{(j)}$ is obtained from $v^{(j-1)}$ by increasing the value of v_j^T from $1/j$ to 1 and by Lemma 5.2 his payment is at most $1/j$. Also, all the winners in $v^{(j)}$ have valuation 1, so we claim that all the winners should pay the same amount. Before proving the claim, we show that this is enough to finishes the proof in this case. Mechanism $\mathcal{M}^*$ at valuation profile $v^{(j)}$ has at most j winners each paying at most $1/j$, therefore, the revenue of $\mathcal{M}^*$ is 1 while $\mathrm{CoreRev}$ of $v^{(j)}$ is $\sqrt{\ln(k)}$ (by Lemma 5.4).

We finish this section by proving the claim that the payments of winners of $\mathcal{M}^*$ at valuation profile $v^{(j)}$ are all the same. Assume otherwise and let a and b be two text ads in the valuation profile v where both are winners but they pay different amounts. w.l.o.g. assume $p_a < p_b$. Lets pick value x such that $p_a < x < p_b$ and x be different than all the valuations in $v^{(j)}$. Note that such x exists since there are finite number of bidders in $v^{(j)}$ but infinitely many numbers in range (p_a, p_b). Now if we decrease the valuation of bidder v_a^T from 1 to x and obtain valuation profile A she remains a winner by Lemma 5.2. If we decrease the valuation of bidder v_b^T from 1 to x and obtain valuation profile B she does not remain a winner by Lemma 5.2. Note that the single bidder in A with valuation x is a winner but the single bidder with valuation x in B is not a winner while A and B are permutations of each other. This contradicts with anonymity (see Definition 5.7) of $\mathcal{M}^*$. □

5.4 A Randomized $O(\ln(\ln(k)))$-core Competitive mechanism

In this section we improve the $O(\sqrt{\ln(k)})$-core competitive mechanism presented in the last section with the use of randomization. Recall that in the deterministic mechanism we decide on allocating to text or image ads based on the ratio v_1^I/Φ^T being above or below $\sqrt{\ln(k)}$. If we allow randomness, we can decide a threshold as a random function of this ratio. Optimizing the revenue as a function of this distribution, we obtain the following mechanism:

Allocation rule: Consider the ration $\psi = v_1^I/\Phi^T$:

- ⋆ if $\psi \leq 2$ allocate the items to the j largest text ads, where j is the largest number such that $jv_j^T \geq v_1^I/2$.
- ⋆ if $2 < \psi$, allocate to the highest valued image ad with probability $\min\{1, \ln(\psi)/\ln(\ln k)\}$. With the remaining probability, leave the items unallocated.

Pricing rule: Allocated bidders are charged according to Myerson's integral.

In the following lemma we calculate the critical values of winners and total revenue of our randomized mechanism.

Lemma 5.6. *The revenue of our mechanism is the following.*

$$\sum_i p_i(v) = \begin{cases} \max\{k \cdot v^T_{k+1}, v^I_1/2\} & \textit{case (i): } \psi < 2 \\ (v^I_1 + 2\Phi^T \ln(2) - 2\Phi^T)/\ln(\ln(k)) & \textit{case (ii): } 2 \leq \psi \leq \ln(k) \\ (\ln(k) \cdot \Phi^T + 2\Phi^T \ln(2) - 2\Phi^T)/\ln(\ln(k)) & \textit{case (iii): } \psi > \ln(k) \end{cases}$$

Proof. We consider three cases:

Case (i). In this case we have j text winners. We prove that the critical value of each of them is at least $v^I_1/(2j)$. Suppose not and assume that the critical value of text ad A is v'_A where $v'_A < v^I_1/(2j)$. This means that when A bids v'_A she still remains a winner. Therefore, there exists a number j' such that

$$j' \cdot v'_A > v^I_1/2 \tag{5.2}$$

in order for text ads to win against image ads. Using $v'_A < v^I_1/(2j)$ and Equation 5.2 we conclude that $j' > j$ which contradicts with the fact that j is the largest number that $jv^T_j > v^I_1/2$. Therefore the critical value of each of the j text winners is at least $v^I_1/(2j)$. Moreover if $kv^T_k > v^I_1/2$ then each winner's critical value must be more than v^T_{k+1} in order to be in the winning set. Therefore, the critical value of the winners is equals to $\max\{v^T_{k+1}, v^I_1/(2j)\}$ and the total revenue in this case is $\max\{k \cdot v^T_{k+1}, v^I_1/2\}$.

Case (ii). In this case the image ad with largest valuation v^I_1 wins and his expected payment is the expected total revenue of our mechanism.

$$\begin{aligned}
p(v_1^I) &= v_1^I x_1^I(v_1^I, v_{-1}) - \int_{2\Phi^T}^{v_1^I} x_1^I(u, v_{-1})du && \text{Lemma 5.2} \\
&= v_1^I \ln(v_1^I/\Phi^T)/\ln(\ln(k)) - \int_{2\Phi^T}^{v_1^I} \ln(u/\Phi^T)/\ln\ln(k)du && \text{replacing} x_1^I \\
&= v_1^I \ln(v_1^I/\Phi^T)/\ln(\ln(k)) - \left[(u \cdot \ln(u/\Phi^T) - u)/\ln(\ln(k))\right]_{2\Phi^T}^{v_1^I} && \text{solving the integral} \\
&= (v_1^I + 2\Phi^T \ln(2) - 2\Phi^T)/\ln(\ln(k))
\end{aligned}$$

Case (iii). Note that if v_1^I is larger than $\ln(k) \cdot \Phi^T$ then her probability of winning is one. Therefore, his payment will be the same as when her valuation is $\ln(k) \cdot \Phi^T$. Therefore, using case (ii) the payment v_1^I in this case is $(\ln(k) \cdot \Phi^T + 2\Phi^T \ln(2) - 2\Phi^T)/\ln(\ln(k))$. □

Theorem 5.3. *Core competitive factor of randomized Image-and-Text mechanism is* $\max\{2, 1.43 \cdot \ln(\ln(k))\}$.

Proof. We prove the theorem by considering three cases similar to Lemma 5.6:

Case (i) : $\psi < 2$. By Lemma 5.4 we know that if $\sum_1^k v_i^T \geq v_1^I$ then $\mathrm{CoreRev}$ is equal to $\max\{kv_{k+1}^T, v_1^I\}$. As the revenue of our mechanism in this case is $\max\{k \cdot v_{k+1}^T, v_1^I/2\}$ (by Lemma 5.6) the proof of the lemma follows. If $\sum_1^k v_i^T < v_1^I$ then by Lemma 5.4 we know that $\mathrm{CoreRev} = \max\{v_2^I, \sum_{i=1}^k v_i^T\}$ which is at most

v_1^I. Therefore, the core competitive factor for this case is 2 and the proof of the lemma follows.

Case (ii) : $2 \leq \psi \leq \ln(k)$. By Lemma 5.4 we know that if $\sum_1^k v_i^T \geq v_1^I$ then $\mathrm{CoreRev}$ is equal to $\max\{kv_{k+1}^T, v_1^I\}$. As $\Phi^T \geq kv_{k+1}^T$ and $v_1^I \geq 2\Phi^T$, we conclude that $\mathrm{CoreRev}$ is at most v_1^I. If $\sum_1^k v_i^T < v_1^I$ then by Lemma 5.4 we know that $\mathrm{CoreRev} = \max\{v_2^I, \sum_{i=1}^k v_i^T\}$ which is at most v_1^I. Therefore, in this case $\mathrm{CoreRev}$ is at most v_1^I. The revenue of our mechanism in this case is $(v_1^I + 2\Phi^T \ln(2) - 2\Phi^T)/\ln(\ln(k)) \simeq (v_1^I - 0.61\Phi^T)/\ln(\ln(k))$ (by Lemma 5.6). As $v_1^I \geq 2\Phi^T$, the revenue of our mechanism is at least $(v_1^I - 0.61v_1^I/2)/\ln(\ln(k)) = 0.695v_1^I/\ln(\ln(k))$, hence it is at least $0.695/\ln(\ln(k))$ fraction of $\mathrm{CoreRev}$ (*i.e.* $1.43 \cdot \ln(\ln(k))$-core competitive) and the proof of the lemma follows.

Case (iii): $\psi > \ln(k)$. In this case we have $v_1^I > \ln(k) \cdot \Phi^T$ which by Equation 5.1 implies $v_1^I \geq \sum_1^k v_i^T$. Hence the $\mathrm{CoreRev}$ in this case is $\max\{v_2^I, \sum_{i=1}^k v_i^T\}$ which is at most v_1^I. The rest of the proof is similar to case (ii) and the core competitive factor for this case is at least $1.43 \cdot \ln(\ln(k))$. □

5.5 A Lower Bound for Revenue of Randomized Mechanisms in Image-and-Text setting

In this section we prove lower bound of $\Omega(\ln(\ln(k)))$ for core-competitive factor of randomized mechanisms. The structure of the proof is as follows. Let assume $\mathcal{R}^* = (x^*, p^*)$ to be a truthful randomized mechanism (satisfying conditions of Lemma 5.2) with optimum core-competitive factor. We derive a distribution over valuation profiles for the Text-and-Image setting such that the expected revenue of $\mathcal{R}^*$ is at most 2 and the expected value of $\mathrm{CoreRev}$ is $\Omega(\ln(\ln(k)))$. Therefore, we conclude that for at least one of the valuation profiles in the support of α, R^* yields a revenue that is smaller than core revenue by factor $\Omega(\ln(\ln(k)))$.

A distribution over valuation profiles. Given k text ads and one image ad, define a distribution $\mathcal{D}$ over valuation profiles as the following. The value of each text ad is taking iid from the set $\{1, \frac{1}{2}, \dots, \frac{1}{k}\}$, each element has probability $\frac{1}{k}$. The value of the image ad is taken from set $\{H, \frac{H}{2}, \dots, \frac{H}{H}\}$ were each element has probability $\frac{1}{H}$, where $H = \lceil \mathcal{H}_k \rceil$.

In the following lemma we prove that the expected revenue of $\mathcal{R}^*$ is at most 2.

Lemma 5.7. *The expected revenue of $\mathcal{R}^*$ for α is at most* 2.

Proof. From the perspective of any given player, a randomized mechanism can be seen as a random threshold being offered to i as a function of v_{-i}. So the revenue that can be extracted from each agent i in expectation, is the revenue that can be extracted from i by using a random threshold, which is the maximum revenue that can be obtained from any given player by a fixed threshold (since the revenue from a random threshold is the expectation of revenue that can be obtained from a fixed threshold)[5].

It is simple to see that under $\mathcal{D}$ the best revenue that can be obtained by a single threshold from any given text ad is $1/k$ and the revenue that can be obtained from an image is 1. So, the total revenue is at most $k \cdot \frac{1}{k} + 1 = 2$. □

Lemma 5.8. *The expected value of the core revenue benchmark is doubly-logarithmic:* $\mathcal{E}_{v \sim \mathcal{D}} \text{CoreRev}(v) \geq \Omega(\ln(\ln(k)))$.

Proof. Throughout this proof, let v be a random variable drawn from $\mathcal{D}$. For any given text ad, $\mathcal{E}[v_i^T] = \mathcal{H}_k/k$. Now, we bounds its variance by:

$$\mathbf{Var}[v_i^T] = \mathcal{E}[(v_i^T)^2] - \mathcal{E}[v_i^T]^2 \leq \mathcal{E}[(v_i^T)^2] = \frac{1}{k}\sum_{j=1}^{k}\frac{1}{j^2} \leq \frac{\pi^2}{6 \cdot k} \leq \frac{2}{k}.$$

[5] Here is a simple mathematical derivation of those arguments for differentiable allocation function $x(v)$ (since monotone functions are almost-everywhere differentiable, the same argument can be easily extended just by performing the equivalent calculations on discontinuities) given an allocation $x(x)$, let $\hat{x}(v_i) = \mathcal{E}_{v_{-i}} x(v_i, v_{-i})$, then the expected revenue that can be extracted from agent i with distribution F is given by $p_i = \mathcal{E}_{v_i}[\int_0^{v_i} u \cdot \partial\hat{x}(u)du = \int_0^{\infty}\int_0^{v_i} u \cdot \partial\hat{x}(u)dudF(v)$. Inverting the order of the integration we get: $p_i = \int_0^{\infty}\int_u^{\infty} u \cdot \partial\hat{x}(u)dF(v)du = \int_0^{\infty} u \cdot \partial\hat{x}(u)(1 - F(u))du \leq \max_u[u \cdot (1 - F(u))] \cdot \int_0^{\infty} \partial\hat{x}(u)du \leq \max_u[u \cdot (1 - F(u))]$, which is the maximum revenue obtained from a single threshold.

Therefore, $\mathcal{E}[\sum_i v_i^T] = \lceil \mathcal{H}_k \rceil$ and $\mathbf{Var}[\sum_i v_i^T] \leq 2$. By Chebyshev's inequality

$$p\left(\left|\sum_i v_i^T - \mathcal{H}_k\right| \geq 2\right) \leq \frac{1}{2}.$$

By Lemma 5.4 we know that the $\mathrm{CoreRev}(v) = \min\{\sum_i v_i^T, v^I\}$. Now, we are ready to lower bound the core revenue benchmark:

$$\begin{aligned}
\mathcal{E}[\mathrm{CoreRev}(v)] &= \mathcal{E}\left[\min\left(\textstyle\sum_i v_i^T, v^I\right)\right] \\
&\geq \frac{1}{2} \cdot \mathcal{E}\left[\min\left(H-2, v^I\right)\right] && \text{by Chebyshev's inequality} \\
&\geq \frac{1}{2} \cdot \frac{1}{H} \sum_{i=1}^{H} \min\left(H-2, \frac{H}{i}\right) && \text{replacing } v^I \\
&= \Omega(\log H)
\end{aligned}$$

Since $H = O(\log k)$ we get that $\mathcal{E}[\mathrm{CoreRev}(v)] \geq \Omega(\ln(\ln(k)))$. □

Theorem 5.4. *The core-competitive factor of $\mathcal{R}^*$ is at least $\Omega(\ln(\ln(k)))$.*

Proof. Since $\mathcal{E}[\mathrm{CoreRev}(v)] = \Omega(\ln(\ln(k)))$ and $\mathcal{E}[\sum_i p_i(v)] = O(1)$, it follows from the probabilistic methods that there must be at least one valuation profile for which $\mathrm{CoreRev}(v) \geq \Omega(\ln(\ln(k))) \cdot \mathcal{E}[\sum_i p_i(v)]$. □

Note on inefficient allocations:

The auctions described in this chapter implement outcomes that are often not socially optimal. Moreover, even when more then one socially optimal allocation is available, the mechanism might allocate to an agent that is part of no efficient allocation. This is unlike, for example, the Micali-Valiant mechanism [MV07] which always allocates to a (random) subset of the agents allocated by the VCG mechanism. Next we show that sometimes allocating to agents which are not allocated in any efficient outcome is necessary in order to get core-competitiveness better then $O(\ln k)$.

Theorem 5.5. *Any mechanism for the Text-and-Image setting that only allocates for a subset of the agents selected by the VCG mechanism has $\Omega(\ln k)$ core competitive hardness.*

Proof. Consider k text ads with v_i^T drawn from the same distribution used for the previous lower bound and one image ad with $v_1^I = \mathcal{H}_k/2$. Using the expectation and variance of $\sum_i v_i^T$ computed earlier in this section, we know by Chebyshev's inequality that $\Pr(|\sum|_i v_i^T - \mathcal{H}_k| > \mathcal{H}_k/s) \leq \Omega(1/\mathcal{H}_k^2)$. So the image ad is allocated with probability $O(1/\mathcal{H}_k^2)$. Since the revenue obtained from any given text ad in expectation is at most $1/k$ (by Lemma 5.7), the total revenue is at most $k \cdot \frac{1}{k} + O(\frac{1}{\mathcal{H}_k^2}) \cdot \mathcal{H}_k/2 = O(1)$. The expected core revenue benchmark, however, is at least $(1 - \frac{1}{\mathcal{H}_k^2}) \cdot \mathcal{H}_k/2 = \Omega(\ln k)$. □

Bibliography

[AB10] Lawrence M Ausubel and Oleg V Baranov. Core-selecting auctions with incomplete information. *University of Maryland*, 2010. 121

[ABHM10] I. Ashlagi, M. Braverman, A. Hassidim, and D. Monderer. Monotonicity and implementability. *Econometrica*, 78(5):1749—1772, 2010. 12

[ADL12] Itai Ashlagi, Shahar Dobzinski, and Ron Lavi. Optimal lower bounds for anonymous scheduling mechanisms. *Math. Oper. Res.*, 37(2):244–258, 2012. 140

[AFG$^+$11] Gagan Aggarwal, Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, Nicole Immorlica, and Madhu Sudan. Derandomization of auctions. *Games and Economic Behavior*, 72(1):1–11, 2011. 140

[AFM07] Gagan Aggarwal, Jon Feldman, and S Muthukrishnan. Bidding to the top: Vcg and equilibria of position-based auctions. In *Approximation and Online Algorithms*, pages 15–28. Springer, 2007. 14

[AGM06] Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7. ACM, 2006. 14

[AH06] Gagan Aggarwal and Jason D Hartline. Knapsack auctions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1083–1092. Society for Industrial and Applied Mathematics, 2006. 122, 123

[AK08] A. Archer and R. Kleinberg. Truthful germs are contagious: a local to global characterization of truthfulness. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 21—30, 2008. 12

[AM02] Lawrence M Ausubel and Paul Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1–42, 2002. 3, 50, 51, 88, 89, 117, 118, 119, 120, 121, 128, 129

[AM06] Lawrence M Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. *Combinatorial auctions*, 17:22–26, 2006. 3, 118, 124, 129

[AMPP09] Gagan Aggarwal, S Muthukrishnan, Dávid Pál, and Martin Pál. General auction mechanism for search advertising. In *Proceedings of the 18th international conference on World wide web (WWW)*, pages 241–250. ACM, 2009. 15

[AMT07] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 279–287. ACM, 2007. 14

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

[AT07] Aaron Archer and Éva Tardos. Frugal path mechanisms. *ACM Transactions on Algorithms (TALG)*, 3(1):3, 2007. 124

[BCK+14] Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, and David Kurokawa. Optimising trade-offs among stakeholders in ad auctions. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 75–92, 2014. 9

[BCK+15] Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, and Mohammad Reza Khani. Mechanism design for mixed ads. To be submitted, 2015. 3

[BCL+06] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu'alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006. 52, 89

[BKK+13] Sayan Bhattacharya, Elias Koutsoupias, Janardhan Kulkarni, Stefano Leonardi, Tim Roughgarden, and Xiaoming Xu. Near-optimal multi-unit auctions with ordered bidders. pages 91–102, 2013. 123

[CDE+07] Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 262–271. ACM, 2007. 15

[CEGP09] Ning Chen, Edith Elkind, Nick Gravin, and Fedor Petrov. Frugal mechanism design via spectral techniques. *arXiv preprint arXiv:0912.3403*, 2009. 124

[CKKK11] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 81–90. ACM, 2011. 14

[DC12] Robert W Day and Peter Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, 2012. 3, 119, 121

[DHH13] Nikhil R. Devanur, Bach Q. Ha, and Jason D. Hartline. Prior-free auctions for budgeted agents. In *ACM Conference on Electronic Commerce*, pages 287–304, 2013. 122, 131

[DM08] Robert Day and Paul Milgrom. Core-selecting package auctions. *International Journal of game Theory*, 36(3-4):393–407, 2008. 3, 51, 89, 121, 128

[DRS09] Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 243–252. ACM, 2009. 88

[EGG07] Edith Elkind, Leslie Ann Goldberg, and Paul W Goldberg. Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 336–345. ACM, 2007. 124

[EK10] Aytek Erdil and Paul Klemperer. A new payment rule for core-selecting package auctions. *Journal of the European Economic Association*, 8(2-3):537–547, 2010. 121

[EOS05] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords. 2005. 8, 14

[ES10] Benjamin Edelman and Michael Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *The American Economic Review*, pages 597–602, 2010. 14

[FGHK02] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 72–81, 2002. 122

[FK14] Rafael Frongillo and Ian Kash. General truthfulness characterizations via convex analysis. In *Web and Internet Economics*, volume 8877 of *Lecture Notes in Computer Science*, pages 354–370. 2014. 12

[GH03] Andrew V. Goldberg and Jason D. Hartline. Competitiveness via consensus. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 215–222, 2003. 122

[GHK+05] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1164–1173, 2005. 131

[GHK+06] Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006. 90

[GHK14] Gagan Goel, Mohammad T. Hajiaghayi, and Mohammad Reza Khani. Randomized revenue monotone mechanisms for online advertising. In *Conference on Web and Internet Economics (WINE)*, 2014. 5

[GHW01] Andrew V Goldberg, Jason D Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Symposium on Discrete Algorithms*, pages 735–744, 2001. 90, 122

[GK14] Gagan Goel and Mohammad Reza Khani. Revenue monotone mechanisms for online advertising. In *International World Wide Web Conference (WWW)*, 2014. 4

[GKL15] Gagan Goel, Mohammad Reza Khani, and Renato Paes Leme. Randomized revenue monotone mechanisms for online advertising. In *Conference on Economics (WINE)*, 2015. 7

[GL09] Jacob K Goeree and Yuanchuan Lien. On the impossibility of core-selecting auctions. *Institute for Empirical Research in Economics University of Zurich Working Paper*, (452), 2009. 121

[GS14] Renato Gomes and Kane Sweeney. Bayes–nash equilibria of the generalized second-price auction. *Games and Economic Behavior*, 86:421–437, 2014. 15

[Har13] Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 2013. 122

[HH13] Bach Q. Ha and Jason D. Hartline. Mechanism design via consensus estimates, cross checking, and profit extraction. *ACM Trans. Economics and Comput.*, 1(2):8, 2013. 122, 131

[HY11] Jason D. Hartline and Qiqi Yan. Envy, truth, and profit. In *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*, pages 243–252, 2011. 122, 131

[IKNS10] Nicole Immorlica, David R Karger, Evdokia Nikolova, and Rahul Sami. First-price procurement auctions, 2010. 124

[IP15] IAB and PWC. IAB internet advertising revenue report, 2015. 1

[KK05] Anna R Karlin and David Kempe. Beyond vcg: Frugality of truthful mechanisms. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 615–624. IEEE, 2005. 124

[KSM10] David Kempe, Mahyar Salek, and Cristopher Moore. Frugal and truthful auctions for vertex covers, flows and cuts. In *FOCS*, volume 10, pages 472–482, 2010. 124

[Lam10] Laurent Lamy. Core-selecting package auctions: a comment on revenue-monotonicity. *International Journal of Game Theory*, 39(3):503–510, 2010. 121

[LMN03] Ron Lavi, Ahuva Mu'Alem, and Noam Nisan. Towards a characterization of truthful combinatorial auctions. In *Foundations of Computer Science*, pages 574–583, 2003. 52, 89

[LPLT12] Brendan Lucier, Renato Paes Leme, and Eva Tardos. On revenue in the generalized second price auction. In *Proceedings of the 21st international conference on World Wide Web*, pages 361–370. ACM, 2012. 14

[LR12] Stefano Leonardi and Tim Roughgarden. Prior-free auctions with ordered bidders. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, pages 427–434, 2012. 123

[LT09] R Paes Leme and Éva Tardos. Sponsored search equilibria for conservative bidders. In *Fifth Workshop on Ad Auctions*, 2009. 14

[MV07] Silvio Micali and Paul Valiant. Collusion-resilient revenue in combinatorial auctions. Technical Report MIT-CSAIL-TR-2007-052, MIT-CSAIL, November 2007. 122, 123, 132, 133, 134, 150

[Mye81] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 10, 52, 57, 90, 131

[NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007. 9, 48, 52, 57, 86, 90, 95

[PK11] Furcy Pin and Peter Key. Stochastic variability in sponsored search auctions: observations and models. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 61–70. ACM, 2011. 15, 32

[RCLB09] Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Stepwise randomized combinatorial auctions achieve revenue monotonicity. In *Symposium on Discrete Algorithms*, pages 738–747, 2009. 51, 86, 87, 88

[RCLB11] Baharak Rastegari, Anne Condon, and Kevin Leyton-Brown. Revenue monotonicity in deterministic, dominant-strategy combinatorial auctions. *Artificial Intelligence*, 175(2):441–456, 2011. 49, 51, 86, 87

[RGKK13] Ben Roberts, Dinan Gunawardena, Ian A. Kash, and Peter Key. Ranking and tradeoffs in sponsored search auctions. In *Proc. of the Fourteenth ACM Conf. on Electronic Commerce*, EC '13, pages 751–766, 2013. 14, 29

[Rob79] Kevin Roberts. The characterization of implementable choice rules. *Aggregation and Revelation of Preferences*, 12(2):321–348, 1979. 52, 89

[Roc87] Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200, 1987. 52, 89

[SY05] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Electronic Commerce*, pages 286–293, 2005. 52, 89

[TLB13] David R.M. Thompson and Kevin Leyton-Brown. Revenue optimization in the generalized second-price auction. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 837–852, New York, NY, USA, 2013. ACM. 15

[Var07] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007. 8, 14, 29

[VH14] Hal R. Varian and C. Harris. The VCG auction in theory and practice. *The American Economic Review*, 104(5):442–445, 2014. 9

[Vor09] Yevgeniy Vorobeychik. Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 583–590. AUAI Press, 2009. 15