Contests for Revenue Share

Ron Lavi **(Technion)**

19 November, 11:15-12:30, bld. 72, room 465

Abstract: In a framework of contest design, we study capacity-constrained competition and analyze its resulting revenue shares. Our analysis contrasts the near-symmetric case, where firms have similar supply sizes, and the extremely asymmetric case, where one large firm dominates the market. In particular, we show that while in the near-symmetric case simple contests provide near optimal equilibrium revenues for all sellers, in the asymmetric case a large firm can design more complicated contests that yield disproportionally low equilibrium revenues for her smaller opponents.

Contests for Revenue Share*

Uriel Feige[†] Ron Lavi[‡] Moshe Tennenholtz[§]

November 8, 2015

Abstract

In a framework of contest design, we study capacity-constrained competition and analyze its resulting revenue shares. Our analysis contrasts the near-symmetric case, where firms have similar supply sizes, and the extremely asymmetric case, where one large firm dominates the market. In particular, we show that while in the near-symmetric case simple contests provide near optimal equilibrium revenues for all sellers, in the asymmetric case a large firm can design more complicated contests that yield disproportionally low equilibrium revenues for her smaller opponents.

*This work was conducted at Microsoft Research. A previous version of this paper, titled "Competition Among Asymmetric Sellers With Fixed Supply", was presented at the 14th ACM Conference on Electronic Commerce (EC'13), and its abstract (14 lines) was included in the conference proceedings.

[†]Weizmann Institute of Science. Email: uriel.feige@weizmann.ac.il.

[‡]Technion – Israel Institute of Technology. Email: ronlavi@ie.technion.ac.il. This author is supported by a Marie-Curie IOF fellowship.

[§]Technion – Israel Institute of Technology. Email: moshet@ie.technion.ac.il.

1 Introduction

We study a setting closely related to the two frameworks of contest design and of Blotto games, where budget-constrained buyers need to partition their budgets among multiple simultaneous contests. Departing from most of the literature, we assume that each contest is designed endogenously, as part of a two-stage game: First, sellers simultaneously design their contests. Second, given the set of contests being offered, buyers simultaneously decide how to partition their budgets among the different contests. The goal of each contest designer is to obtain the largest possible revenue share, i.e., to maximize the sum of budgets assigned to her contest. Our main interest is in understanding what types of contests maximize their designer's revenue share.

The study of contests and Blotto games stems from numerous economic applications. Examples relevant to our work include situations of rent-seeking and lobbying (Tullock, 2001) where instead of treating the regulatory committee as a single entity we study the case that the committee is composed of several distinct members, multiple simultaneous R&D races each aiming to attract the most research efforts (generalizing a single R&D race as in Che and Gale (2003)), and many other situations of competing all-pay auctions. Our particular interest in this setting stems from observations regarding the market for Internet search advertising. This is now a multi-billion dollar market, where thousands of advertisers post hundreds of millions of ads every day. Quite interestingly, market shares differ significantly from revenue shares in this market.[1] In 2012, for example, Google had a market share of about 65% and a revenue share of about 75%, while Microsoft (Bing) had a market share of about 20% but a revenue share of less than 10% which is less than half of its market share.[2] We aim to study the connection between market shares and revenue shares in an abstract theoretical model inspired by such markets. We ask how can a firm leverage its market share in order to maximize its revenue share.

Our theoretical abstraction captures markets with three main properties. First, each seller has a fixed number of goods to sell. For example, increasing the market share of a search engine requires increasing the number of users of the search engine. This is a long-term technological effort, not a short-term strategic production decision. Thus, in our model, a firm's market share is exogenously fixed, and firms aim to maximize their revenue shares given their fixed market shares. Second, the sales tools that we consider are quite general, and can capture various methods like bundling and non-linear pricing, exclusive contracts, quantity discounts, etc., in addition to standard linear pricing. In our model, a seller has a large and realistic strategy space to choose from, and a large

[1] The "items" in this market are ad slots on users' search results pages. The "market share" of a certain search engine is its number of available ad slots divided by the total number of available ad slots of all search engines. Its "revenue share" is the budget spent on ads in this search engine relative to the budgets spent on all search ads. In other industries, what we refer to as a market share is also commonly termed unit market share or volume market share, and what we refer to as revenue share is also commonly termed dollar market share or value market share.

[2] Source: eMarketer, June 2013. Search advertising revenue share of leading search sites in the United States from 2011 to 2015. Available from http://www.statista.com/statistics/255863/search-ad-revenue-share-at-leading-search-sites-in-the-us/.

strategy space of the opponent to defend against. Third, we are interested to focus attention on the supply side and study asymmetries, even extreme, among the sellers, while assuming the simpler case of identical buyers on the demand side. An important property of the buyers that we do capture is their budget-constraints. This is a significant characteristic of Internet advertisers that influences most of their actions and considerations.

Previous literature studies revenue maximization (effort maximization) in a single contest, and various settings of multiple contests (Blotto games) with exogenously fixed contest success functions. The issue of competing contests that are endogenously designed to maximize revenue shares is rarely touched upon. Other related literature on capacity-constrained competition mostly focuses on standard linear pricing in contrast to contests that capture a much wider variety of sales tools. The recently expanding literature on antitrust economics does study various pricing tools like exclusivity contracts but focuses mainly on how to increase market shares endogenously. The issue of market shares versus revenue shares is not dealt with in this literature. Section 1.1 provides many references and a detailed discussion.

Our first result shows that a very simple, normative contest provides a strong safety level to each seller that uses it. In the terminology of Blotto games, this is the well-studied "lottery contest success function". This contest uses a simple uniform-price rule to clear its supply. Friedman (1958) shows that if all sellers happen to choose this contest, there is a unique pure NE in the resulting Blotto game. In this NE, the revenue share of each seller is exactly her market share. We generalize this result to incorporate the sellers' strategic point of view, and show the following result: If a specific seller, i, chooses this contest, then for any choice of contests of the other sellers, the revenue share of seller i in any pure or mixed NE of the resulting Blotto game will be at least her market share minus $\frac{1}{n}$, where n is the number of buyers. Thus, if the fraction $\frac{1}{n}$ is negligible relative to the market share of a certain seller, she can obtain a revenue share very close to her market share in all possible (pure as well as mixed) Nash equilibria outcomes. Various tools like exclusive contracts, quantity discounts, etc., cannot help her opponents in this case.

However, cases where the market share of a seller is comparable to $\frac{1}{n}$ should not be overlooked. This can happen either because a seller's market share is low, or because the number of buyers is low, or both. In particular, in our motivating scenario of online advertising, the number of buyers (advertisers) in a given market is usually quite small, as advertisers target very specific user features.[3] In such cases, the above result no longer provides a tight connection between market shares and revenue shares. For example, if there are two extremely asymmetric sellers, the small one has a market share below $\frac{1}{n}$ while the large one provides the rest of the supply, the above result only provides a meaningless non-positive lower bound for a revenue share. This is no accident as, in fact, we show examples where in similar cases an extremely large seller can obtain a revenue

[3]For example, there could be a market for advertisers who target users that searched "best pizza in NYC" and a different market for advertisers who target users that searched "SFO to JFK flights". These are two separate markets. The thousands of Internet advertisers target such specialized markets, resulting in very few advertisers per market.

share of 100%, effectively driving his much smaller competitor out of the market. Thus, our results demonstrate a sharp contrast between the case of near-symmetric supply sizes and the case of extremely-asymmetric supply sizes. In the former case, all sellers are able to obtain revenue shares almost the same as their market shares, in all Nash equilibria outcomes, using a simple contest structure and linear pricing. In the latter case, a large seller may decrease the revenue share of her significantly smaller opponent substantially below her market share, using various sales tools like exclusivity contracts.

What tools are most useful in our theoretical framework for a very large seller competing against a very small seller? Does the answer vary according to the extent of differences in supply sizes? How to quantify "very large" vs. "very small"? These questions turn out quite challenging and we leave the full characterization of the powers of a large seller, as the market share of her smaller opponent gradually increases from $\frac{1}{n}$ to $\frac{1}{2}$, for future research. In this paper we make a first step towards answering these questions by characterizing the case when the large seller has a market share of $1 - \frac{1}{n}$ and the small seller has a market share of $\frac{1}{n}$. We believe that the analysis of this case is quite revealing. To begin with, we prove that any monotone contest (where a bidder that increases his bid never receives a strictly smaller prize) used by the large seller cannot eliminate the possibility that the small seller will end up with a revenue share close to her market share in some Nash equilibrium outcome.

Despite this, our analysis of this case does identify a contest for the large seller that decreases the revenue share of the small seller significantly below her market share, in a strong sense: for every contest that the small seller uses, and for every pure or mixed Nash equilibrium of the resulting game. This contest incorporates a certain non-monotonicity, alongside exclusivity, to overcome the limitations of monotone contests. Broadly speaking, the non-monotonicity that we suggest is, conceptually, somewhat related to realistic situations where a large seller offers attractive deals to her rival's customers, by this tempting them to switch over. We further discuss this in Section 5.2.

The fact that such a non-monotonic rule is helpful may seem to contradict an intuitive rule of thumb, phrased by McAfee (2005) as follows: *"If you offer discounts to your rival's customers, it will cause your rival to fight to hold onto his customers, and he will do this by cutting prices. He will then take some of your customers away from you. In the end, you will get some of his customers, he will get some of yours, and you will both be selling at lower prices. If, on the other hand, you reward loyalty by offering a better deal to customers that have been with you for a while, you make your customers expensive to poach. Your rivals are discouraged from poaching them, and tend to respond in kind."* We view our analysis as suggesting that this argument (and in particular the optimality of the stability that it suggests) depends on the underlying details of the market more than what initially may seem.

The remainder of the paper is organized as follows. Related literature is discussed in Section 1.1. Section 2 formally lays out our model and terminology. Section 3 discusses the proportional allocation policy. Section 4 discusses the general use of exclusivity and its impact on revenue shares. Section 5 discusses the case of two extremely asymmetric sellers. Its shows how a very large seller may significantly decrease the revenue share of her significantly smaller opponent. Section 6 summarizes, and describes additional applications of our model. Some additional proofs are provided in an Appendix.

1.1 Related Literature

We give more details on four strands of related literature: Blotto games, capacity-constrained competition, antitrust economics, and single contest design.

The framework of Blotto games captures the exact setting that we study here: Colonels (buyers) have a fixed number of soldiers (buyers' budgets) which they need to partition among multiple battlefields (sellers). Borel (1921) first studied a Blotto game with three battlefields (sellers) and two colonels (buyers) that care about winning a majority of the battlefields, assuming an "auction contest success function" (auction CSF) where the winner of any specific battlefield is the colonel who sent the largest number of soldiers to this battlefield. Laslier and Picard (2002) generalize Borel's model to any number of battlefields. Motivated by the allocation of advertising expenditures across different marketing areas, Friedman (1958) studies a similar model, but with a lottery CSF, and assuming additive colonels' utility functions (sum of battlefields won). Our model mostly follows that of Friedman (1958), as our motivating example is an advertising market as well. In particular, our results demonstrate the importance of the lottery CSF and its advantages, as a result of the model rather than as an assumption in it. Snyder (1989) studies a class of CSFs that generalize the lottery CSF, describing properties of NE and sufficient conditions for its existence and uniqueness. His motivation comes from political campaigns that distribute advertising budgets among different legislative districts (the "markets"). Myerson (1993) studies a version of the game with a continuum of buyers. This early literature was focused mostly on the "demand side", having studied various aspects of the generals (which are the buyers) while treating battlefields (which are the sellers in our case) as fixed exogenous entities. We on the other hand focus our attention on the sellers, treating them as strategic entities. In particular, we study the effect that the contest success function of a specific battlefield has on the number of soldiers being sent to this battlefield.

Recent literature on Blotto games continues to focus its attention on various properties and generalizations of the demand side (the colonels). For example, Weinstein (2012) and Roberson (2006) study colonels with asymmetric budgets (while we study here the more basic case of symmetric budgets), Kvasov (2007) and Roberson and Kvasov (2012) study colonels' utilities that are quasi-linear w.r.t. their budgets (while we study the more classic case of "use-it-or-lose-it" budgets that seems to be a better fit to advertising markets). Hart (2008) studies discrete (non-continuous) budgets, which is a property that we also partly address by studying bid rigidity. Many additional results and details are surveyed in Kovenock and Roberson (2012). Robson (2005) is the only other paper that we are aware of that studies (among other things) sellers/battlefields in a similar way to our interest here. Specifically, he studies comparative statics when prize values and battlefield CSFs change. However, he assumes only two buyers (while we allow any number of buyers) and a class of CSFs that does not capture many CSFs that prove useful in our setting (while we do not cast any restrictions on the possible CSFs). For example, exclusivity cannot be captured by any CSF in his class, let alone non-monotonicity. As mentioned above, these two ingredients turn up to be especially important in our analysis. More generally, our paper sheds light on several new contest success functions that has not been analyzed previously and that turn up useful in the context of revenue maximization.

A central aspect of our model is the capacity-constraints of the sellers, that have fixed supply sizes that cannot be changed. Earlier studies on competition between capacity-constrained firms usually adopt the Edgeworth-Bertrand framework (Tirole, 1988). Two relevant examples are Acemoglu, Bimpikis and Ozdaglar (2009), that study resulting market efficiency in a two-stage competition model where sellers choose their capacities in the first stage and then engage in a price competition in the second stage, and Dudey (1992), that studies a dynamic variant, where buyers arrive one after the other, and two competing sellers may adjust prices after each sale. He shows, quite surprisingly, that in the dynamic setting sellers earn positive profits in equilibrium, unlike Bertrand's original conclusion. Many additional generalizations have been studied, see for example Ghemawat and McGahan (1998), Johari, Weintraub and Van Roy (2010), and Martínez-de Albéniz and Talluri (2011). Our work significantly differs from this literature in two main aspects. First, on the sellers' side, this literature focuses on pure price competition while our approach captures a wide array of non-linear pricing aspects of the competition. The generality of our approach has the additional advantage of endogenously incorporating modeling issues like the one raised by Davidson and Deneckere (1986) regarding the importance of the rationing rule being assumed (when total demand exceeds capacity). While the explicit choice of the rationing rule was a subject for debate in previous literature, our model includes this choice endogenously. Second, on the buyers' side, we model the buyers as having budget-constraints (an unavoidable and important assumption given our motivation) while most of the literature on capacity-constrained competition does not usually discuss budgets.[4]

There is also a growing body of literature on competitive non-linear pricing and antitrust economics. This literature studies the usefulness (and fairness) of various pricing tools like discounts for purchases of larger volumes, exclusivity contracts, bundling, etc. Recent studies that analyze the use of such tools and their resulting welfare effects include for example Chone and Linnemer (2015); Calzolari and Denicolò (2013); Armstrong (2013), and Armstrong and Vickers (2010). Whinston

[4]Earlier capacity-constrained models like that of Levitan and Shubik (1972) do capture the possibility of budget constraints since they model the buyers' side using a simple aggregate demand curve. This only emphasizes the importance of linear prices to previous models (as a demand curve takes into account only the unit price offered by the competing firms), while we are interested in much more general selling policies.

(2008) gives a survey to earlier literature on antitrust economics, including earlier studies on exclusive contracts. The analysis in our paper gives rise to similar tools and pricing techniques, but our setting differs significantly. In particular, we focus on a setting where firms compete for revenue shares while having fixed market shares. In contrast, the literature on competitive non-linear pricing mostly analyzes how firms increase market shares endogenously, i.e., how pricing decisions enable firms to increase production on the expense of their opponents' production. The issue of market shares versus revenue shares is not dealt with in these models. These differences in settings are important, and lead to qualitatively different conclusions. For example, O'Brien and Shaffer (1997) and Bernheim and Whinston (1998) show that, with complete information, firms have no incentive to offer exclusivity or market-share discounts. Indeed, in light of these results, all the previously-mentioned recent studies assume incomplete information. In contrast, we do not need to introduce incomplete information in our model since, as we show, the criticality and usefullness of exclusivity and other non-linear pricing methods become apparent in our setting even with complete information. More broadly, the entire context of our analysis is completely different from these previous studies on antitrust economics, since they focus on the question of how to increase the firm's market share while we study how to increase the firm's revenue share, given exogenous market shares.

In contrast to the case of revenue maximization in multiple contests which is our focus here, the case of revenue maximization in a single contest has been studied quite extensively under various assumptions in the literature for all-pay auctions, rent-seeking competitions, and more abstract contest design. For example, Siegel (2009) characterizes NE outcomes including player efforts in a large class of "generic" contests; Moldovanu and Sela (2006) compare total effort (total revenue) extracted by a static (one stage) contest versus a dynamic (multiple stages) contest; Che and Gale (1998) analyze a political contest where caps on maximal bids are present; Clark and Riis (1998) analyze a contest that allocates multiple prizes either sequentially or simultaneously; Barut and Kovenock (1998) characterize revenue generated in Nash equilibria of a contest with multiple prizes. Michaels (1988) explicitly studies how to choose a CSF from the family of Tullock CSFs in order to maximize the revenue of the contest designer, and mentions a motivating situation where "a monopoly politician sets the parameters of the rent-seeking operation in such a way that he maximizes his residual income from the expenditures made by seekers", which is a motivation very similar to ours. However the contests that he analyzes are limited to the class of Tullock contests, while we analyze arbitrary contests. See Konrad (2009) and Dechenaux, Kovenock and Sheremeta (2014) for elaborate literature surveys for the case of a single contest.

2 Setup

There are m sellers and n buyers. There is only one type of good, and this good is divisible.[5] Each seller has a fixed supply of the good. Since the good is divisible, by scaling (changing units) we may assume without loss of generality that the total amount of the good is n. For convenience we shall refer to this amount as being composed of n identical *items*, with the understanding that each item is divisible as well. The number of items of seller j is denoted by $q_j n$, where $q_j > 0$ and $\sum_{j=1}^{m} q_j = 1$. Buyers have budgets, and in this work we consider the case of symmetric buyers in which the budgets of all buyers are equal, say, to B. Without loss of generality, we assume that $B = 1$. Throughout, we assume complete and perfect information. For convenience, we refer to sellers as feminine and to buyers as masculine.

The action of a buyer is to distribute his budget among the sellers. We refer to this as placing bids, where b_{ij} denotes the bid of buyer i at seller j, with $b_{ij} \geq 0$, and $\sum_j b_{ij} \leq 1$. The buyer's objective is to maximize the total number of items that he receives (this number need not be an integer, because items are divisible), or in situations in which randomization is involved, to maximize the expected number of items received. The buyer does not derive any utility from unspent budget. The objective of a seller is to maximize her (expected) revenue – the total amount of bids that she received. A seller does not derive any utility from left-over items. This setup gives rise to the following extensive-form game with $m + n$ players and two consecutive steps:

1. First, sellers simultaneously design contests. Each seller designs a contest by announcing an *allocation policy* $a_j : [0,1]^n \to [0, q_j n]^n$, i.e. a function from the submitted bids $b_{1j}, ..., b_{nj}$ to an allocation vector $\vec{a}_j = \vec{a}_j(b_{1j}, ..., b_{nj})$, where $\sum_{i=1}^{n} a_{ij} \leq q_j n$.[6]

2. Second, based on the announced allocation policies, buyers simultaneously choose bids to sellers. In his bid b_{ij}, buyer i commits to pay b_{ij} to seller j.

3. Based on the allocation policies and the received bids, sellers allocate items to buyers. The resulting utility of seller j is $\sum_{i=1}^{n} b_{ij}$ and the resulting utility of buyer i is $\sum_{j=1}^{m} a_{ij}(b_{1j}, ..., b_{nj})$.

Regarding terminology, we use *policies* for sellers' strategies and keep *strategies* for buyers. A strategy is either pure or randomized (meaning non-pure). We use the convention that a pure Nash has only pure strategies, and a mixed Nash has at least one randomized strategy. Fixing sellers' allocation policies results in a subgame among the buyers. We refer to this subgame as the (resulting) *buyers' subgame*.[7]

[5] Towards the end of this section we discuss the connection of our specific modeling choices to the motivation of Internet search advertising.

[6] An allocation policy is termed a *contest success function* (CSF) in the standard terminology of Blotto games. We prefer "allocation policy" as it seems more descriptive in our context.

[7] This is a Blotto game where each seller is a battlefield and each buyer is a colonel.

While budget and items are divisible, we assume that there is some finite precision to bids and quantities, so only a finite number of bids and allocation policies is possible. This makes the game finite, and implies the existence of some subgame perfect equilibrium. Note that Nash equilibria strategies in every fixed buyers' subgame can be either pure or mixed. On the other hand, without loss of generality, sellers' policies can be assumed to be pure, namely, involving no randomization, since the allocation of fraction of items combined with the assumption that buyers are expectation maximizers implies that any randomized allocation policy can be replaced by a deterministic one that averages over the random allocations.

Since we have an extensive-form game, buyer i's strategy needs to specify his bid (or randomization over bids) for every tuple of sellers' policies. Denote this as $s_i(a_1, ..., a_m)$. To simplify this complicated object, our analysis will be based on finding (what we term) *an extensive-form safety level policy*:

Definition 1. *A policy a_j^* of seller j provides an extensive-form safety level of x (for some real number x) if the following holds. Fix:*

1. *arbitrary policies a_{-j} of all other sellers, and,*
2. *buyers' strategies $s_1(\cdot), ..., s_n(\cdot)$ such that $s_1(a_j^*, a_{-j}), ..., s_n(a_j^*, a_{-j})$ is a Nash equilibrium of the buyers' subgame when seller policies are (a_j^*, a_{-j}).*

Then, seller j's revenue when sellers play (a_j^*, a_{-j}) and buyers play $s_1(a_j^*, a_{-j}), ..., s_n(a_j^*, a_{-j})$ is at least x.

Finding an extensive-form safety level policy has immediate implications on equilibrium revenue:

Lemma 1. *If there exists a policy a_j^* for seller j that provides an extensive-form safety level of x then j's revenue in any subgame-perfect equilibrium is at least x.*

Proof. Fix an arbitrary subgame-perfect equilibrium (a_j, a_{-j}, s). Therefore s is a Nash equilibrium of every buyers' subgame. In particular, $s_1(a_j^*, a_{-j}), ..., s_n(a_j^*, a_{-j})$ is a Nash equilibrium of the buyers' subgame when seller policies are (a_j^*, a_{-j}). By the definition of an extensive-form safety level it now follows that the revenue of seller j is at least x when players play (a_j^*, a_{-j}, s). The revenue of seller j when players play (a_j, a_{-j}, s) must be at least that as otherwise j will have a benficial deviation, and the claim follows. $\square$

Our analysis in the following sections will proceed by identifying extensive-form safety level policies. As we will see in the sequel, the bounds on equilibrium revenues that this analysis tool will provide will be quite tight, especially when there are two competing sellers. In addition to providing bounds on equilibrium revenues, one can argue that extensive-form safety level policies are conceptually advantageous as they require less rationality assumptions, similarly to arguments in favor of the standard definition of safety levels in normal form games.

Comparing the formal model that we have defined here to our original motivation of Internet search advertising, we wish to discuss three issues. First, recall that the sellers in our model are the search engines, the buyers are the advertisers, and the items are the ad slots. In reality, the allocation policies of the search engines are quite complicated, as allocation policies are made in an online fashion over a period of time. In addition, they are composed of various different levels, from a sequence of auctions at the lowest level to determine the allocation of an ad to a specific user's search result, to the decisions of how to spread an advertiser's budget over a period of time (a month, a week, a day), to various rationing rules, etc. These details boil down to a single allocation policy in our model. Since we allow arbitrary and general allocation policies, in principle our theoretical construct is able to capture any realistic process, as complicated as it may be (though we disregard uncertainties in the number of user searches to be performed over the coming period of time and assume known supply). However, since these processes are so complicated and full of details (to the point that most probably exact specifications of the global process do not currently exist), we will not analyze any specific allocation policy that can be claimed "realistic" in this work. Instead, we look for theoretically close-to-optimal safety level policies, and examine their key principles.

A second issue is our choice to model items as identical across sellers (identical ad slots across search engines). As mentioned in the Introduction, ad slots are characterized by very specific properties, especially in search advertising. Thus, different search words can be treated as being sold on separate markets. A specific abstract market in our model can correspond, for example, to advertisers interested in searches related to "pizza in NYC", or to searches related to "flights SFO JFK". The market does not include ad slots for both of these searches. As a result, all items in a single market are treated as homogeneous. A-priori, as we are not aware of any empirical data that indicates otherwise, it seems reasonable to assume that an ad slot related to a search for "pizza in NYC" on Bing is similar to the same slot on Google.

A third issue is the assumption that buyers do not obtain utility from unspent budgets. This is a standard assumption in models of advertising; it was already posed in the original model of Friedman (1958), and it continues to be a standard assumption regarding advertisers in general and regarding Internet search adevrtisers in particular, see for example a recent relevant paper by Sayedi, Jerath and Srinivasan (2014). More specifically, in the practice of Internet search advertising, advertisers are required to allocate budgets to "campaigns" which is the technical term for the opertaion of assigning a budget (and other parameters) to a specific advertising need. Specific ads are then sold via auctions, where losing bids cannot be confiscated. But campaigns usually terminate when all allocated budget has been exhausted. Our model captures a situation where advertisers assign budgets to campaigns mainly aiming to maximize the number of impressions obtained during the campaign, caring less about any leftover budget.

3 The Proportional Allocation Policy

Since the number of items is equal to the number of buyers (n), and each buyer has the same budget (1), perhaps the first outcome in our abstract market that comes to mind is the unique Walrasian equilibrium outcome: each buyer will be allocated one item, and the price of each item will be 1. In such an outcome the revenue of each seller j will be $q_j n$. We term this the *fair share* revenue of seller j. However, since both buyers and sellers act strategically (and, furthermore, there may be only few buyers and few sellers), it may well be that strategic issues will lead to a different outcome.

In this section we show an allocation policy that obtains a revenue close to her fair share revenue, for every seller j that uses this policy, regardless of the allocation policies chosen by her competitors, and in any pure or mixed NE of the resulting buyers' subgame. (I.e., we show a safety level.) This is the following natural *proportional allocation policy*: Given bids $b_{1j}, ..., b_{nj}$, seller j allocates to buyer i^* a quantity that is proportional to the ratio of his bid to the sum of all bids. More formally,

$$a_{i^*j}(b_{1j}, ..., b_{nj}) = \frac{b_{i^*}}{\sum_{i=1}^{n} b_{ij}} q_j n.$$

This allocation policy is also termed the *lottery contest success function*, and it is in fact one of the most well studied CSFs in the Blotto game literature. In our context, it has several natural interpretations. In particular, given buyers' bids, if a seller is required to fix a price per unit that clears her supply (seeing the buyers' bids), the resulting allocation will be exactly proportional.

Friedman (1958) was one of the first papers to analyze the proportional allocation policy, focusing on the buyers' side. I.e., he assumed as given that all sellers use proportional allocation, and studied the resulting strategic aspects of the buyers' subgame. In particular, he showed:

Theorem (Friedman (1958))**.** *If all sellers use proportional allocation, there exists a unique Nash equilibrium in the resulting buyers' subgame. In this equilibrium outcome,*

1. *Every buyer i bids $b_{ij} = q_j$ at every seller $j = 1, ..., m$.*[8]
2. *The revenue of every seller j exactly equals her fair share revenue $q_j n$.*

Given the first property, one would indeed expect that if all sellers use proportional allocation, the unique equilibrium outcome will indeed be reached, and the revenue split will be proportional as well. However, looking at the sellers' side, and since sellers are strategic, it is entirely unclear why sellers will indeed choose proportional allocation. The obvious question is whether a strategic seller can extract a higher revenue by using a different strategy. To some extent, this is indeed possible, via nonlinear pricing tools, as the following example demonstrates.

[8] In fact, it can be shown (see Appendix A) that this is a safety level strategy that provides each buyer that uses it in every equilibrium outcome a utility of at least 1, regardless of the strategies of the other buyers.

Suppose two sellers $1, 2$ with $q_1 = q_2 = 0.5$, and an even number of buyers n (and so each seller has $\frac{n}{2}$ items). Seller 1 uses proportional allocation. Seller 2 uses the following policy: if a buyer bids at least $1 - \epsilon$ (ϵ to be determined later) he receives 0.5 item, otherwise he receives nothing. Let us check for which values of ϵ it is a symmetric pure Nash for the buyers to bid ϵ for seller 1 and $1 - \epsilon$ for seller 2. In this case the utility of a buyer is 1 (0.5 from each seller). The only possible deviation for a buyer is to bid his entire budget, 1, for seller 1 (no point to bid more than $1 - \epsilon$ at seller 2 and if he bids less than $1 - \epsilon$ on seller 2 he might as well bid 0 there). If he indeed follows this deviation he receives a fraction of $\frac{1}{1+(n-1)\epsilon}$ of the supply of seller 1 which is $n/2$. We need this to be smaller or equal to 1 so we need $\epsilon \geq \frac{n-2}{2(n-1)}$. Set ϵ to be equal to this expression. This results in a pure NE and seller 1's revenue at this equilibrium is $n\epsilon = n(n-2)/2(n-1) < n/2$.

Note, however, that seller 1's revenue in this example is almost her fair share. In particular, it is larger than $\frac{n}{2} - 1$. As we next show, this is not an accident. When a seller uses proportional allocation, her revenue in any equilibrium outcome will be very close to her fair share:

Theorem 1. *If seller P uses proportional allocation while the other sellers use any arbitrary policies, the revenue of seller P in any Nash equilibrium (either pure or mixed) will be strictly larger $q_P n - 1$.*

Proof. Let $Q = q_P n$. Consider an arbitrary Nash equilibrium (pure or mixed) and assume for the sake of contradiction that the expected revenue of P is $R \leq Q - 1$. Let x be a random variable denoting the sum of bids to P. Then x is non-negative and its expectation satisfies $E[x] = R$. Partition the buyers into three classes, *good* who deterministically place all their budget in P, *bad* who deterministically place no budget in P (though they may randomize how they split their budget outside P), and *flexible*.

Every good buyer has expected payoff $E_x[\frac{Q}{(x-1)+1}] \geq \frac{Q}{E[x]} = \frac{Q}{R} > 1$. The first inequality follows from the fact that for every non-negative constant c, the function $\frac{1}{x+c}$ is convex in the domain $x > 0$. As the total expected payoff of all good buyers is at most Q, there are strictly less than Q good buyers.

Consider now an arbitrary bad buyer b. If b were to become good, his expected payoff would be $E_x[\frac{Q}{x+1}] \geq \frac{Q}{E[x]+1} = \frac{Q}{R+1} \geq 1$. The best response property then implies that in the given Nash every bad player has expected payoff at least 1. As the total expected payoff of all bad buyers is at most $n - Q$, there are at most $n - Q$ bad buyers.

The number of flexible buyers is strictly larger than $n - Q - (n - Q) = 0$. Hence there is at least one flexible buyer. As the expected payoff of every non-flexible buyer is at least 1, there must be at least one flexible buyer whose expected payoff is at most 1. Consider such a flexible buyer f and let $r > 0$ be the expected bid of f in P. If f were to become good his expected revenue would be $E_x[\frac{Q}{(x-r)+1}] \geq \frac{Q}{E_x[x]+1-r} > \frac{Q}{R+1} = 1$, contradicting the best response property of the assumed Nash equilibrium. $\square$

This result can be generalized to any set of sellers that use proportional allocation:

Theorem 2. *Let P be a collection of several sellers $\{P_1, \ldots, P_k\}$ and let $Q = \sum_{i=1}^{k} q_i n$ denote their total supply. If all members of P use proportional allocation then their total expected revenue in every Nash equilibrium is strictly larger than $Q - 1$.*

The proof of this theorem is given in Appendix B. As a direct corollary, we have:

Corollary 1. *There is no allocation policy that provides seller j a total revenue of $Q_j n + 1$ (or more) in all Nash equilibrium outcome.*

Proof. Let the collection of sellers P of Theorem 2 include all sellers besides j. By Theorem 2, the sellers in P can jointly obtain a total revenue strictly larger than $(1 - Q_j)n - 1$ which implies the claim. □

Thus, we have established in this section that the optimal safety level of a seller that has total supply Q lies in the interval $[Q-1, Q+1]$. Furthermore, we have shown that proportional allocation yields a safety level of $Q-1$, and thus provides an almost optimal safety level. We leave for future research the further identification of the optimal policy.

4 Exclusivity

Exclusivity as a pricing tool was already useful in the first example given in the previous section and its importance will become even clearer in the sequel. This section makes some preliminary comments and observations regarding this sales tool. We refer to a bid of value 1 as a *rigid bid*. In such a case, the entire budget of a buyer goes to one of the sellers. Rigidity might be enforced by a seller who attempts to force exclusivity via various technical or marketing actions.[9] It may also come from the buyer side, especially in more traditional markets like advertising in the printed media. Rigidity is often termed "single-homing" (see e.g. Athey, Calvano and Gans (2011) and the references therein). Ashlagi, Edelman and Lee (2011) give an empirical analysis of the number of ad platforms (sellers in our terminology) that online advertisers (buyers in our terminology) use. They show that small advertisers tend to single-home even in online advertising. Rigidity is very natural in many settings in which there is some underlying cost to simultaneously work with several sellers, e.g. because of technological differences (and a seller can deliberately make such additional costs particularly high, by this implementing "rigid" bids). In the more broader context of an abstract

[9] To enforce rigidity, sellers must know buyers' budgets (to infer that they are receiving these in full). Thus, the possibility to exercise rigidity crucially relies on the assumption of complete information. While this assumption is clearly restrictive, we note that it is a widely acceptable assumption in many studies on Internet search advertising. E.g., this assumption was made already in Varian (2007) and in some parts of the analysis of Edelman, Ostrovsky and Schwarz (2007).

exchange economy, rigidity simply captures the case that buyers' endowments are indivisible, e.g., in cases where it is not money.[10]

A natural policy that enforces exclusivity is what we term *proportional allocation among rigid bids*: the seller divides her items equally among the buyers that place rigid bids (and non-rigid bidders are ignored). The next result shows that if seller j uses proportional allocation among rigid bids while the other sellers use any arbitrary policies, the revenue of seller j in any Nash equilibrium (either pure or mixed) is strictly larger than $q_j n - 1$. In particular, her revenue in any pure Nash is at least $\lfloor q_j n \rfloor$, which is slightly better than the revenue of proportional allocation.

Theorem 3. *If seller j uses proportional allocation among rigid bids while the other sellers use any arbitrary policies, the revenue of seller j in any Nash equilibrium (either pure or mixed) will be strictly larger $q_j n - 1$. In particular, her revenue in any pure Nash is at least $\lfloor q_j n \rfloor$.*

The proof of the first part of this theorem is identical to that of Theorem 1 and is therefore omitted. The second part of the theorem is a direct consequence of rigidity, and is not true without it, as the example in the previous section has demonstrated.

We note that Theorem 2 does not hold when a set of sellers P all use proportional allocation among rigid bids (and not proportional allocation). Therefore, Theorem 3 limits by how much the equilibrium revenue might drop below the fair share, but not by how much it may increase beyond the fair share.

In fact, we wish to point out that rigidity typically favors larger sellers. If there are many small sellers and very few large sellers, and all sellers use proportional allocation among rigid bids, the increase in revenue of a large seller may be significant. In a pure Nash, if seller 1 has large supply and all other sellers are small, a disproportionately large revenue for seller 1 may result from the gap between $n - \sum_{j=2}^{m} \lfloor q_j n \rfloor$ and $q_1 n$. For example, let $n = 25$, $m = 9$, $q_1 n = 13$ and $q_i n = 3/2$ for $2 \leq i \leq 9$. Then the fair share of seller 1 is 13, but in every pure Nash her revenue is 17. Thus, the use of proportional allocation among rigid bids gives large sellers a revenue significantly larger than their fair share.

Moreover, the expected revenue of the large seller in a mixed Nash equilibrium may be significantly larger than her revenue in a pure Nash. For example, suppose the that the supply of seller 1 is of size about $\sqrt{n}$ while the size of the supply of all other sellers is exactly 1. Then Theorem 3 shows that in any pure Nash the revenue of a small seller is at least 1, implying that the large seller does not obtain more than her fair share. Now consider a symmetric mixed Nash in which each buyer bids at seller 1 with probability q. If the result of the coin toss is that the buyer does not bid at the large seller, he bids at some small seller chosen uniformly at random among all small sellers. Given a set of buyers that bid at the small sellers, we have a "balls in bins" process (see for example Johnson and Kotz (1977)) where balls are uniformly at random being put in x bins. For

[10]One example for such a setting is given in our conclusions below; a second example naturally arises when Blotto games are motivated by actual battlefield scenarios (as actual soldiers are usually indivisible).

this random process, it is known that the expected number of empty bins is $\frac{x}{e}$. Since the supply of the large seller $\sqrt{n}$ is very small compared to the aggregate supply of all small sellers $n - \sqrt{n}$, in equilibrium most buyers must end up at a small seller, and because a fraction of $\frac{1}{e}$ of the small sellers will be empty, the expected utility of every buyer given that he ends up at a small seller will be about $1 - \frac{1}{e}$. Since this is a Nash equilibrium, his expected utility given that he ends up at the large seller must be the same, implying that the expected number of buyers at the large seller will be about $\frac{e}{e-1}$ times her supply, implying a similar increase in her revenue, relative to her fair share.

5 Two Extremely Asymmetric Sellers

When we have two nearly symmetric sellers, Theorem 3 shows that each of them is able to obtain in any equilibrium outcome revenue which is very close to their fair share. However, in extremely asymmetric cases, that is, when $q_j n$ is comparable to 1, the theorem no longer supports such a statement. In particular, when $q_j n \leq 1$ the theorem does not provide any meaningful bound.[11] It turns out that the subtraction of 1 from the bound in the theorem is not a technical limitation of the proof, but, rather, an accurate observation regarding the significant differences between the case of near symmetries in supply and the case of extreme asymmetries in supplies. Consider the following example, where L's supply is $1 - \frac{1}{n}$ and H's supply is $n - (1 - \frac{1}{n})$. In this case, H can offer each buyer a fraction of $\frac{1}{n}$ of the total supply of H, provided that the buyer signs an exclusivity contract with H, and bids his entire budget at H. With such a policy, every buyer will prefer exclusivity with H, regardless of L's policy, since L's total supply $1 - \frac{1}{n}$ is strictly smaller than the number of items that H is offering each buyer. As a result, H will attract the full budget of all buyers and L will be driven out of the market.

How extreme should the differences in supply be, for the large seller to be able to reduce the revenue share of the small seller significantly below her market share? While the example above is fairly simple, the general case is much more complicated. The "first" extremely asymmetric case which is far from trivial (as will soon become evident) is the case when L's supply is exactly 1 and H's supply is $n-1$. On the one hand, Theorem 3 for L is still meaningless in this case. On the other hand, it is no longer clear how the large seller should act, in order to reduce L's revenue share below her market share. In particular, it is easy to verify that the simple exclusivity method suggested above no longer works. In fact, as the following result demonstrates, a large class of selling policies for H (on top of exclusivity) cannot yield a very low revenue for L in all equilibrium outcomes.

Definition 2. *An allocation policy a_j is* **monotone** *if an increase in a buyer's bid (fixing the other bids) does not decrease his allocation.*

[11] As discussed in the introduction, our motivation makes sense also in small markets. For example, $n = 10$ and $q_j = 10\%$ seems like a perfectly plausible scenario.

Definition 3. *The* **anti-competitive** *allocation policy (for L) is as follows. If exactly one buyer places a rigid bid at L, that buyer receives L's item. In any other case, L does not allocate her item.*

Theorem 4. *If L's supply is exactly 1, H's supply is $n-1$, L uses the anti-competitive allocation policy, and H uses any arbitrary monotone allocation policy, there exists a pure Nash equilibrium in the resulting buyers' subgame in which L's revenue is* 1.

Proof. Consider first a situation in which all buyers place their entire budgets at H. In this case, at least one of the buyers must receive strictly less than one item. Let such a buyer be i. Note that i bidding the full budget at L and all others bidding the full budget at H is a pure Nash equilibrium: None of the others can improve their utility as they will not receive any item at L even if bidding there, and any lower bid to H cannot increase their utility by the monotonicity of the policy of H. Buyer i cannot increase her utility as bidding the full budget at H will result in strictly lower utility (by our choice of i), and bidding less than the full budget at H can only result in an even lower utility for i, because of the monotonicity of H. $\square$

This result demonstrates that the case where H has $n-1$ items and L has 1 item is not a-priori clear and simple. Can H limit the revenue of L to be close to 0 in any NE, similarly to the case when L's supply is $1-\frac{1}{n}$? Or maybe L can obtain revenue almost proportional to her fair share revenue, in all equilibrium outcomes or in some equilibrium outcomes? This section studies these questions in detail to obtain a better understanding of this case. In particular, we will identify a strategy the does manage to reduce L's revenue share significantly below her market share in any NE outcome.

5.1 Mixed Nash equilibria with rigid bids

Exclusivity and bid rigidity turn out to be of particular importance to our analysis of this case. In this section we develop important observations regarding their role by analyzing the case where both H and L use proportional allocation among rigid bids. In this case, every pure Nash equilibrium of the resulting buyers' subgame has the following structure: exactly one buyer bids his full budget at L, while all others bid their full budget at H. In this outcome every buyer receives exactly one item, and no buyer can deviate and receive more than one item. The revenue of both sellers in such an equilibrium outcome is exactly their fair share and in particular L's revenue is 1.

However, there are also mixed Nash equilibria in the resulting buyers' subgame, including a symmetric mixed Nash equilibrium in which each buyer submits his full budget to L with some probability p (that depends only on n), and with probability $1-p$ submits his full budget to H. It turns out that L's revenue dramatically drops in every mixed NE outcome compared to her fair share: in every mixed NE L's revenue is about $\frac{1}{n}$ (while in every pure NE L's revenue is 1). We prove this result for the case when H uses proportional allocation among rigid bids and L uses any policy from the following class:

Definition 4. *An allocation policy a_j is:*

- **anonymous** *if for any bids $b_{1j}, ..., b_{nj}$ and any renaming of the buyers (a permutation $\pi : \{1, ..., n\} \to \{1, ..., n\}$), $a_{ji}(b_{1j}, ..., b_{nj}) = a_{j\pi(i)}(\tilde{b}_{1j}, ..., \tilde{b}_{nj})$ where $\tilde{b}_{ij} = b_{\pi(i)j}$.*

- **efficient** *if whenever the seller receives at least one positive bid, all items of the seller are allocated.*

- **cross-monotone** *if an increase in a buyer's bid (fixing the other bids) does not decrease the allocation of another player.*

Proportional allocation satisfies all the above properties (efficiency, anonymity, monotonicity and cross monotonicity). Proportional allocation among rigid bids satisfies anonymity, monotonicity and cross monotonicity but does not satisfy the efficiency property – if a seller receives only non-rigid bids, her items are not allocated. However, it does satisfy a property that we shall call *weak efficiency*, namely, that whenever the seller receives at least one rigid bid, all items of the seller are allocated.

Theorem 5. *Suppose that H uses proportional allocation among rigid bids and L uses any allocation policy that is anonymous and in addition satisfies either one of the following two conditions: (1) L's policy is efficient, or (2) L's policy is monotone, cross monotone, and weakly efficient. Then in every mixed Nash L's revenue is at most $\frac{4}{n-2} + O(1/n^3)$.*

Proof. We shall assume throughout the proof that $n > 4$ (this simplifies computations, and if $n \leq 4$ the theorem does not limit L to below her fair share). We shall present the proof for the case that L's policy is anonymous and efficient, without requiring any monotonicity properties. The other case is easier to prove (the combination of the monotonicity properties of L and rigidness of H allows us to assume that L too only receives rigid bids, simplifying the analysis), and its proof is omitted. Let X be the set of players that have positive probability of going to L. Since we assume a non-pure Nash equilibrium, X is not empty. The proof follows by several Lemmas.

Lemma 2. $|X| \geq 2$.

Proof. Assume towards a contradiction that $X = \{i\}$ for some buyer i, and let $E[i|i \to H]$ denote i's expected utility given that she puts her budget at H (since H is rigid she cannot split her budget). Nash equilibrium implies $E[i] = E[i|i \to H] = E[i|i \to L]$. Since the supply of H is smaller than n, anonymity implies $E[i] = E[i|i \to H] < 1$. However, efficiency (weak efficiency suffices) implies that i can obtain a utility of 1 by deterministically bidding her full budget at L, a contradiction. $\square$

For any $S \subseteq X$, let $Pr[S]$ be the probability that S is exactly the set of players that submit a non-zero bid to L (in this case we say that S is the "colliding set"). Note that if a player submits a non-zero bid to L then she cannot receive anything from H. Let $u[i|S]$ be i's expected utility, given that S is the colliding set. For any $i \in X$, define $\mathcal{S}_i = \{S \subseteq X \text{ s.t. } |S| \geq 2 \text{ and } i \in S\}$, and let $u_i = \frac{\sum_{S \in \mathcal{S}_i} Pr[S]u[i|S]}{\sum_{S \in \mathcal{S}_i} Pr[S]}$. The denominator is the probability that i bids at L and there is a collision at L. Denote this term as $Pr[i]$. Note that u_i is the expected utility of i given that she bids at L and there is a collision at L.

Lemma 3. *If $|X| \geq 2$, there exists a player i such that $u_i \leq 1/2$.*

Proof. For every set of players S, $\sum_{i \in S} u[i|S] \leq 1$, since L has only one item to offer. Thus,

$$\sum_{i \in X} \sum_{S \in \mathcal{S}_i} u[i|S]Pr[S] = \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} Pr[S] \sum_{i \in S} u[i|S] \leq \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} Pr[S]$$

$$\leq \frac{1}{2} \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} |S| Pr[S] = \frac{1}{2} \sum_{i \in X} Pr[i].$$

Therefore there must be a player $i \in X$ such that $\sum_{S \in \mathcal{S}_i} u[i|S]Pr[S] \leq \frac{1}{2} Pr[i]$, and the lemma follows. □

Let $P \in X$ be a player with $u_P \leq 1/2$. Let x be the probability that a player besides P submits a positive bid to L. If P submits all her budget to H she obtains a utility of at least $\frac{n-1}{n}$. Thus, the expected utility of P given that she submits a positive bid to L must be at least that as well. On the other hand, this expected utility is at most $x\frac{1}{2} + (1-x) \cdot 1$. This implies that $x \leq \frac{2}{n}$. Because each player tosses her coin independently, it follows that the expected number of players besides P that go to L is very close to x. This is a standard argument and the next lemma is for completeness.

Lemma 4. *Let L_i for $i = 1...n-1$ be random independent Bernoulli variables where $Pr(L_i = 1) = p_i$ and $Pr(L_i = 0) = 1 - p_i$, and let $L = \sum_i L_i$. Suppose that $Pr(L \geq 1) \leq \frac{2}{n}$. Then $E[L] < \frac{2}{n-2}$.*

Proof. Let $E = E[L] = \sum_i p_i$. We have $\frac{2}{n} \geq Pr(L \geq 1) = 1 - \Pi_i(1 - p_i) \geq 1 - e^{-\sum_i p_i} = 1 - e^{-E}$, where the second inequality follows since $1 - x \leq e^{-x}$. Thus $E \leq \ln(n) - \ln(n-2) < \frac{2}{n-2}$ (recall that $ln(a) = \int_1^a \frac{1}{x} dx$), and the lemma follows. □

Lemma 4 implies that the expected number of players at L (not counting player P) is at most $\frac{2}{n-2}$, and therefore this is at most the expected revenue that L obtains from all players besides P. To complete the proof, it only remains to argue that the probability that P bids at L is small.

Lemma 5. *Let p be the prob. that buyer P (as defined above) submits a positive bid to L. Then $p \leq \frac{2(n-2)}{n(n-4)}$.*

Proof. Recall that we used x to denote the probability that a player besides P submits a positive bid to L. We have shown that $x \leq 2/n$. Consider first the worst possible value for x, namely, $x = 2/n$.

Lemma 4 implies that the probability that there are two or more players besides P that submit a positive bid to L is at most $\frac{2}{n-2} - \frac{2}{n} = \frac{4}{n(n-2)} = \frac{2}{n-2}x$. A similar argument shows (details omitted) that for every $x \leq 2/n$, the probability that there is *exactly* one player beyond P bidding at L is at least $\frac{n-4}{n-2}x$. Consider now the optimal bid for P, conditioned on bidding in L. If no other player bids in L, then every positive bid of P gives him the full item, by the requirement that L's policy is efficient. If exactly one other player bids in L, then from P's point of view, he is in a two player 0-sum game situation, playing against one bid (P does not care which player is giving this other bid), and trying to maximize his share in L's item. By the anonymity of L's policy, this two player game is symmetric, and hence P has a bid that gives him expected value of at least 1/2. Hence in a Mixed Nash equilibrium, P's payoff conditioned on bidding at L is at least $1 - x + \frac{n-4}{2(n-2)}x$. This implies that P's expected payoff conditioned on participating in a collision in L is at least $\frac{n-4}{2(n-2)}$. By averaging over the other players, we conclude that there must be some player P' whose expected payoff when participating in a collision with P at L is at most $1 - \frac{n-4}{2(n-2)} = \frac{n}{2(n-2)}$. Recall that p is the probability that P bids in L. Then the expected payoff of P' when bidding in L is at most $1 - p + p\frac{n}{2(n-2)}$. As the expected payoff of P' at H is at least $1 - 1/n$, it follows that in a Nash equilibrium $p \leq \frac{2(n-2)}{n(n-4)}$. □

To conclude: L's expected revenue is at most the expected number of players besides P that bid at L plus the prob. that P bids at L, which (by the above) is at most $\frac{2}{n-2} + \frac{2(n-2)}{n(n-4)} = \frac{4}{n-2} + \frac{8}{n(n-2)(n-4)}$. This concludes the proof of the theorem. □ □

If L uses a non-anonymous policy, the theorem no longer holds. For example, L can use the following policy: Allocate a fraction of $\frac{n-1}{n}$ of the item to buyer 1, if he bids his full budget at L. In any case do not allocate anything to the other buyers. If H uses proportional allocation among rigid bids, then the following is a mixed Nash: buyer 1 bids his full budget at either H or L with equal probability 0.5, while all other buyers bid their full budgets at H. In this mixed Nash, L's revenue is 0.5, showing that the theorem no longer holds when L uses a non-anonymous policy.

5.2 An Optimal Policy

Two questions now present themselves: Can H achieve a similarly high revenue in all equilibria (and not only in all mixed equilibria), and can H achieve a similarly high revenue when L chooses an arbitrary allocation policy that does not necessarily satisfy the above-mentioned properties. We next show a strategy that eliminates all pure equilibria, and yields a revenue in the order of $n - \frac{1}{n}$ to the large seller (leaving very little revenue to the small seller), in all mixed Nash equilibrium outcomes, regardless of the allocation policy of the small seller.

Definition 5. *Fix any $\epsilon > 0$. In* **proportional allocation among rigid bids with a gamble bid***, the seller allocates the items equally among all bids whose value is exactly* 1 *as in proportional allocation among rigid bids, but with one exception: If the seller receives $n-1$ bids of value* 1 *and* *one bid of value $1 - \epsilon$ (we term this bid a near rigid bid), give the near rigid buyer $1 + 3/n$ items, and split the remaining items equally among the other $n - 1$ buyers.*

Thus, proportional allocation among rigid bids with a gamble bid is identical to proportional allocation among rigid bids, with one important exception (this is the "gamble bid"): one of the buyers (the one that places a near rigid one) can receive more than one item, but only if all other bids go to this seller as well, and are all rigid. Note that a near rigid bidder receives *zero* items in all other cases (so it is a risky bid).

As mentioned in the Introduction, we interpret this policy as offering an "attractive deal" to a buyer that might otherwise deterministically put his entire budget with L. This is a situation where H receives $n - 1$ bids of value 1 and given this an additional buyer will benefit from bidding his full budget at L. The strategy offers the gamble bid, which we view as an attractive deal, to this buyer. Specifically, if this buyer will move to H he will receive more than all other buyers. Of course, since our setting is of a one-shot game, by "this buyer" we generically mean any buyer that considers deterministically putting his entire budget with L in a situation where all other buyers deterministically put their entire budget with H. The gamble bid makes this option less profitable for any buyer. Because of this, the gamble bid eliminates all pure NE from the game, leaving only mixed NE. This contradicts McAfee's recipe that suggests to foster stability and loyalty in competitive situations. In our specific competitive situation, it turns out that H should promote instability, expressed as a mixed strategy where buyers randomize between H and L. This will significantly decrease L's revenue.

A nice property of this allocation policy is that it is anonymous. Another interesting characteristic of this policy is its non-monotonicity. In a situation where a near rigid bid wins, if this bidder will increase his bid to his full budget, he will decrease the quantity that he receives. Theorem 7 below shows that non-monotonicity is unavoidable if H wants to obtain revenue above $n - 1$ in all equilibrium outcomes. (Our model restricts attention to allocation policies that are functions of bids only, with no extra bits. If we had an extra bit, we could replace the bid $1 - \epsilon$ by a bid "1, and I choose to gamble". In this case, the effect of the extra bit would not be monotone, because depending on the bids of the other buyers it may either cause the payoff of the bidder to increase or to decrease.)

If H uses proportional allocation among rigid bids with a gamble bid and L uses proportional allocation with rigid bids, the outcome in which one buyer deterministically bids at L while the others deterministically bid at H is not a Nash equilibrium, because of the gamble bid: the buyer that submitted his full budget to L can increase his utility by submitting a near rigid bid to H.

More generally, the gamble bid rules out all profitable pure Nash equilibria for the small seller: with less than $n-1$ rigid bids to H, at least one of the remaining bidders can strictly gain by giving a rigid bid to H (the total L can offer to the remaining bidders is too low to prevent defection), and with exactly $n - 1$ rigid bids to H, the remaining buyer must give a near rigid bid to H.

We conclude that if H uses proportional allocation among rigid bids with a gamble bid, then in any pure Nash equilibrium of the resulting buyers' subgame, H receives $n-1$ rigid bids and one near-rigid bid (leaving only ϵ to L). We next show that in every mixed Nash equilibrium the small seller has very little revenue as well:

Theorem 6. *Assume $n \geq 10$. If H uses proportional allocation among rigid bids with a gamble bid, L's revenue in any Nash equilibrium (either pure or mixed) is at most $\frac{8}{n-3} + n\epsilon$, regardless of her policy.*

Proof. We already saw that in every pure Nash the claim holds. We now consider the revenue of L in a mixed Nash, ignoring bids of value ϵ. We start with a useful lemma.

Lemma 6. *If, in an equilibrium strategy, player j declares a near rigid bid to H with positive probability, then the probability that all other players declare a rigid bid to H is at least $1 - \frac{4}{n-1}$. This implies that the expected number of players at L besides j is at most $\frac{4}{n-3}$.*

Proof. Let y be the probability that all players (besides perhaps j) declare a rigid bid to H. The expected utility of j given that she declares a near rigid bid to H is $y(1+\frac{3}{n})+(1-y)0$. On the other hand j can obtain a utility of $\frac{n-1-(1+3/n)}{n-1}$ by placing a rigid bid at H. Thus $y(1+\frac{3}{n}) \geq \frac{n-1-(1+3/n)}{n-1}$ implying $y > 1 - \frac{4}{n-1}$. This proves the first part of the lemma. The second part follows by Lemma 4. □

To prove the theorem, consider the following three cases.

Case 1. Suppose there exist two distinct buyers i, j that declare a near rigid bid to H with positive probability. Then, by Lemma 6, the expected number of players at L besides j is at most $\frac{4}{n-3}$, and the prob. that j bids at L is at most $\frac{4}{n-1}$ (again by Lemma 6, since i also declares a near rigid bid to H with positive prob.). Thus, the expected number of players at L is at most $\frac{8}{n-3}$, which implies the theorem for this case.

Case 2. The second case uses the notation of Theorem 5. Let X be the set of players that have positive probability of going to L, and assume (this is the second case we are considering) that $|X| \geq 2$. In this case, we can use Lemma 3 from Theorem 5, as this lemma does not rely on any of the assumptions made in that theorem, and conclude that there exists a buyer $P \in X$ such that the expected utility of P, given that she bids at L, and given that there is a collision at L, is at most $\frac{1}{2}$. Let x be the probability that at least one buyer besides P bids at L. The expected revenue of P, given that P bids at L, is therefore at most $\frac{1}{2}x + (1-x)1$, and at least $\frac{n-1-(1+3/n)}{n-1}$ (since P can obtain this utility by placing a rigid bid at H). Thus $x \leq \frac{2(n+3)}{n(n-1)}$.

We argue that, in this case, there must exist a buyer $j \neq P$ that bids a near rigid bid at H with positive probability. To see this, assume towards a contradiction that if a player $j \neq P$ bids at H, her bid is rigid. Then, player P can obtain expected utility strictly larger than 1 by deterministically placing a near rigid bid at H, since with this action her utility is $(1-x)(1+3/n) > 1$ (for $n \geq 10$, and using the above inequality upper bounding x). This is a contradiction because P bids at L with a positive probability and her expected utility given that she bids at L is at most 1. By Lemma 6 the expected number of all buyers besides j (but including P) that bid at L is at most $\frac{4}{n-3}$. The probability that j bids at L is at most x. Thus, the expected number of players at L is at most $\frac{4}{n-3} + \frac{2(n+3)}{n(n-1)} < \frac{8}{n-3}$, implying that the theorem holds in this case as well.

Case 3. Finally, suppose that neither of the above two cases hold. That is, at most one player has a positive probability of placing a near rigid bid at H, and $|X| \leq 1$. If X is empty, the theorem immediately follows. Thus assume $X = \{i\}$. If there does not exist any buyer j that declares a near rigid bid to H with positive probability, then since only i may bid at L, all others deterministically place a rigid bid with H. But then i is not best responding by bidding at L, as she can obtain utility strictly larger than 1 by placing a near rigid bid at H. Thus, assume there exists exactly one buyer j that declares a near rigid bid to H with positive probability. If $j = i$, all others are deterministically placing a rigid bid at H and once again i is not best responding when bidding at L. Thus assume that $j \neq i$. By Lemma 6 the probability that i bids at L is at most $\frac{4}{n-1}$ and since i is the only one that bids at L, the theorem follows. □ □

5.3 Restricting H to be Monotone

Non-monotonicity is a crucial ingredient in order to ensure a low revenue for L. To see this, recall the anti-competitive allocation policy from above, which ensures the existence of a pure NE in which L's revenue is 1, as long as H uses any arbitrary monotone allocation policy. The gamble bid manages to escape this problematic equilibrium of the anti-competitive policy, via its non-monotonicity: If i bids $1-\epsilon$ when all others are bidding rigidly at H, he will receive utility strictly larger than 1, and this is sufficient to break the equilibrium.

Given this state of affairs, anti-trust agencies and other regulatory bodies may consider restricting H to use a monotone policy, in order to limit its monopolistic power. In this context, one may wonder whether there exists a policy for L that will obtain a revenue of 1 in all Nash equilibria (and not only in some Nash equilibria), given that H is restricted to be monotone. The next theorem shows that this is unfortunately impossible.

Theorem 7. *For any $n \geq 4$ and any allocation policy of L, there exists a monotone and anonymous policy for H for which in at least one Nash equilibrium the revenue of L is at most $3/n$ (showing that L cannot obtain a revenue of 1 in all equilibrium outcomes even if H is restricted to be monotone).*

Proof. Fix an arbitrary policy for L. The policy that we will construct for H will accept only rigid bids, and hence we can assume that a buyer either bids at L (not necessarily a rigid bid) or at H, but not both.

We first examine a different game that involves only L, as follows. Each player chooses which bid to submit to L. For a parameter $p = 3/n^2$, with probability p the bid of a player is indeed submitted to L, and with probability $1 - p$ it is discarded and the player's utility in this case is 0. This is done independently for each player. The utility of each player who's bid is being submitted to L is determined by L's policy, as a function of all bids that were submitted to L. As we assume that buyers have only finitely many choices of bids, there is some Nash for the players on L (specifying for every player a probability distribution over bids, conditioned on bidding at L). Fix an arbitrary such Nash, $s_1, ..., s_n$, and let r_i denote the expected revenue of player P_i conditioned in bidding at L (i.e. conditioned on the result of the coin toss being that the player's bid is submitted to L).

The expected utility of player i in this different game is pr_i and thus the total expected number of items that L allocates is $p\sum_i r_i$. With probability $(1-p)^n$ no item is being allocated, hence even if in all other cases L allocated her full item, still the total expected number that L allocates is at most $1 - (1-p)^n$. We conclude that $\sum_i r_i \leq (1 - (1-p)^n)/p$. Our choice of $p = 3/n^2$ implies that $(1-(1-p)^n)/p \leq n-1$ (to verify this, one can use the inequality $(1-p)^n \geq 1 - np + \binom{n}{2}p^2 - \binom{n}{3}p^3$, which follows from the binomial theorem).

Now consider the following policy for H: every player i that places a rigid bid at H receives r_i items, and nonrigid bids get nothing. This policy is feasible since $\sum_i r_i \leq n - 1$. The following strategy for each buyer i forms a Nash equilibrium in the resulting buyers' subgame: With probability $1-p$ submit a rigid bid to H and with probability p submit a bid to L according to s_i. This is since the expected utility of a buyer given that she bids at H is equal to her expected utility given that she bids at L; she has no better action at H because H receives only rigid bids, and she has no better action at L since our construction implies that s_i maximizes a player's expected utility given that she bids at L and given that other players bid at L according to s_{-i}, each one independently with probability p. In this Nash equilibrium outcome, L's expected revenue is at most $np = 3/n$, and the theorem follows. □

Therefore, no monotone policy of H can obtain in all equilibria of the resulting buyers' subgame a revenue strictly more than $n - 1$, but on the other hand no policy of L can obtain in all equilibria of the resulting buyers' subgame a revenue larger than an order of $\frac{1}{n}$. How much revenue can H obtain in all equilibria of the resulting buyers' subgame, under the restriction of monotonicity? Theorem 4 tells us that such an equilibrium revenue cannot be strictly more than $n - 1$, but can she obtain at least that? The next theorem shows that H can obtain a revenue of almost $n - 1$ in all equilibrium outcomes, but not $n - 1$. In this sense, putting restrictions on H, as the dominant seller, can slightly help the small seller. She will not be able to obtain revenue proportional to her market share in all equilibrium outcomes, and L might even obtain revenue share slightly larger than her fair share in some equilibrium outcomes.

Theorem 8. *If H uses proportional allocation among rigid bids, and regardless of the policy that L uses, the revenue of H in any Nash equilibrium of the resulting buyers' subgame is at least* $n-1-\frac{2}{n-2}$.

On the other hand, for any monotone policy of H there exists a policy of L such that, in at least one Nash equilibrium of the resulting buyers' subgame, the revenue of H is at most $n - 1 - \frac{1}{n^2}$.

We postpone the proof of this theorem to Appendix C.

6 Conclusions

Studying competition among firms with fixed market shares, we have analyzed how extreme supply asymmetries can yield an advantage to large sellers. We have shown that exclusive contracts and bid rigidity tend to work in favor of large sellers, that mixed Nash outcomes (as compared with pure outcomes) tend to work in favor of large sellers, and that non-anonymity can be a counter-measure for the small sellers. We have also shown that non-monotone policies can provide an equilibrium revenue for the large seller that cannot be obtained by monotone policies, and that even if sellers are restricted to monotone policies, the small seller might still land in Nash equilibria that result in low utility for her. In sharp contrast, all these selling tools do not provide significant advantages when sellers have similar supply sizes. In this case, simple selling policies provide near optimal revenues in all equilibrium outcomes.

Our main motivating application is the market for Internet search advertising. This market is divided to many small sub-markets, based on searched keywords. Our results demonstrate that revenue shares in markets with such a structure might be significantly more asymmetric, biased towards the larger seller. The following example demonstrate this point. Consider a grand market with 800 buyers and 800 divisible items (i.e., an overall budget of \$800), where a small firm L offers 10% of the total supply (i.e., 80 items), and a large firm H offers the rest. Theorem 1 shows that the small firm can obtain an equilibrium revenue of \$79 – very close to her fair share revenue which is \$80. If, however, the same grand market is actually composed of 100 small markets with 8 buyers and 8 divisible items each (the same overall number of items, buyers, total budget, and market shares), revenue shares can change drastically. In particular, L's revenue drops to zero if H uses proportional allocation among rigid bids in every small market.[12] Aggregating the revenue over all the small markets, L's total revenue still remains zero, of-course. Thus, with such a structure of the grand market, a large seller that has a 90% market share can obtain a 100% revenue share, leaving no revenues for her smaller opponent. Such examples motivate the second part of our inquiry. Theorem 6 for example demonstrates that H can still cut L's revenue to less than one-fourth of her fair share revenue even if the grand market above is composed of only 40 small markets with 20 buyers and 20 divisible items each.[13]

[12]Specifically, in every small market H offers at least $\frac{7.2}{8} = 0.9$ items to every buyer that puts his entire budget with her. Since L has at most $0.8 < 0.9$ items to offer, all buyers of the small market will choose to go exclusively with H, leaving L with a zero revenue.

[13]Specifically, if H implements the strategy suggested in Theorem 6 in every small market then according to the theorem L's revenue in every such market will be at most $\frac{8}{n-3} = \frac{8}{20-3} < 0.5$. Aggregating over the 40 small markets results in a total revenue of less than \$20 for L.

The important assumption that drives our results for the extremely asymmetric case is that the large seller owns an extremely large fraction of the overall supply in the market. The number of small sellers that compete against such a large seller is less important. If the large seller is large only relative to any other seller in the market, and not relative to the total supply, she may not be able to significantly decrease the aggregate revenue of the small sellers, below their aggregate fair share. In this case, the small sellers can jointly obtain almost their fair share, simply by using proportional allocation (each one in an independent way).

One weakness of our paper is the fact that a significant part of it deals with the clearly non-generic case in which the share of capacity held by the low capacity seller is precisely $\frac{1}{n}$ (where n is the number of buyers). That being said, we do believe that carefully analyzing this case is important for two reasons. First, this case sheds light on interesting (and sometimes surprising and non-intuitive) phenomena that are entailed in our model. Thus, it serves as a motivation for a further more general analysis of the model. Second, the analysis here seems like a necessary step that lays the foundation for the more general case, where the market share of the small seller varies from $\frac{1}{n}$ to $\frac{1}{2}$.

We should also emphasize that there could clearly be a variety of reasons for differences in revenue shares. Our work addresses and investigates only one such possible aspect. For example, another possible reason why Bing's and Google's revenue shares differ from market shares could be that the user characteristics across these search engines are different, which makes for example the value of an ad slot for a search for "pizza in NYC" inherently different on Google and Bing. This question should be investigated empirically. By assuming in our model that items are homogeneous, we have assumed that these values are in fact similar across the two search engines.

As a final remark, we wish to mention a completely different connection of our abstract model to the issue of seat allocation in a parliament system, based on election results. The question of how to allocate the parliament seats to the various parties is of-course central in political science. A Belgian mathematician Victor D'Hondt suggested in 1878 a popular method that is being widely used even today in many countries (see Balinski and Young (2001) for more details). This method is defined as follows. Suppose that n is the number of seats to be assigned, and q_j is the number of votes received in the election by party j. Considering an arbitrary assignment of seats to parties, a seat that is being assigned to party j is worth the overall number of votes for this party divided by the number of seats assigned to this party in the current assignment. Denote an assignment as *stable* if no seat can increase its worth by being reassigned to a different party. D'Hondt proposes to use a stable assignment. (Ties are extremely unlikely when the number of votes is large, and hence we omit the details of how to break ties if they occur.) Consider the following analogy with our competition model. Parties can be viewed as sellers, and the supply of each seller is the number of votes that the party received. Seats are buyers, each with a budget of 1. In this context, buyers must be rigid, of course. Each party uses proportional allocation among rigid bids. It is not hard to verify that pure Nash equilibria in this setting exactly corresponds to all stable assignments in the D'Hondt method. As discussed in Section 4, Nash equilibria outcomes tend to favor larger sellers. Indeed, the D'Hondt method has a reputation of typically favoring the larger parties.

A A Simple Proof of Friedman's Theorem

We start an auxiliary result:

Theorem 9. *Suppose that all sellers use the proportional allocation policy, and fix a buyer i^*. Then, if i^* bids $b_{i^*j} = q_j$ at every seller $j = 1, ..., m$, his resulting utility will be at least 1. Moreover, if at least one of the other bidders uses a randomized (non-pure) strategy, or if there exists a seller j such that $\sum_{i \neq i*} b_{ij} \neq (n-1)q_j$, then i^*'s utility will be strictly larger than 1.*

Proof. For simplicity of notation we give a proof for two sellers ($m = 2$) and denote $q_1 = q$. Let y_i be a random variable that denotes the bid of player i to seller 1 (so $1 - y_i$ is i's bid to seller 2), and $y = \sum_{i=1}^{n} y_i$. Then,

$$u(y) = \frac{q}{y}qn + \frac{1-q}{n-y}(1-q)n$$

is the resulting utility of i^*. The derivative $u'(y)$ is strictly negative for $y < qn$ and strictly positive for $y > qn$, hence the utility is minimal for $y = qn$ where $u(qn) = 1$. This proves the claim for all pure strategies of the other players. If some of the other players use a randomized strategy and y is a non-degenerate random variable, then since $u(\cdot)$ is convex, $E_y(u(y)) > u(E(y)) \geq 1$ where the first inequality is Jensen's inequality and the second inequality follows from our argument for pure strategies. This completes the proof of the theorem. □

An easy corollary is that when sellers use proportional allocation, proportional bidding $b_{ij} = q_j$ is a Nash equilibrium in the resulting buyers' subgame: every buyer obtains utility of 1 regardless of the bids of the other buyers, and the sum of all utilities is at most n. Thus if a buyer deviates he cannot obtain more than 1. In fact, Friedman (1958) shows that this is the *unique* Nash equilibrium:

Theorem 10. *[Friedman (1958)] If all sellers use proportional allocation, proportional bidding is the unique Nash equilibrium in the resulting buyers' subgame. In this equilibrium, the revenue of every seller j is $q_j n$.*

Proof. For simplicity of notation we again give a proof for two sellers ($m = 2$) and use the same notation as above. First notice that in every Nash equilibrium (mixed or pure) of the resulting buyers' subgame, the expected utility of every buyer is exactly 1: the sum of all resulting utilities is n, hence if some buyer has utility strictly larger than 1, another buyer has utility strictly less than 1, which implies that she can strictly gain by deviating to proportional bidding (by Theorem 9), contradicting the equilibrium property.

Now assume equilibrium bids y_i and fix a buyer i^*. We have $\sum_{i \neq i^*} y_i = (n-1)q$, otherwise by Theorem 9 buyer i^* can obtain a utility strictly higher than 1 by playing proportional bidding. In a similar way we also have that all players play pure strategies. Summing over all n possible choices of i^* we have $(n-1)\sum_{i=1}^n y_i = n(n-1)q$ yielding $\sum_{i=1}^n y_i = nq$. Thus, $y_{i^*} = \sum_{i=1}^n y_i - \sum_{i \neq i^*} y_i = q$, and the theorem follows. □

B Proof of Theorem 2

Let P be a collection of several sellers $\{P_1, \ldots, P_k\}$ and let $Q = \sum_{i=1}^k q_i$ denote their total supply. We prove that if all members of P use proportional allocation then their total expected revenue in every Nash equilibrium is strictly larger than $Q - 1$.

Consider an arbitrary Nash equilibrium (pure or mixed) and assume for the sake of contradiction that the total expected revenue of all members of P is $R \leq Q - 1$. Let x be a random variable denoting the sum of bids to P. Then x is non-negative and its expectation satisfies $E[x] = R$. Partition the buyers into three classes, *good* who deterministically place all their budget in P (though they may randomize how they split their budget within P), *bad* who deterministically place no budget in P (though they may randomize how they split their budget outside P), and *flexible*.

Consider an arbitrary good buyer g. The *balanced strategy* for g is to bid $\frac{q_i}{Q}$ at each seller P_i. We show that the balanced strategy has expected payoff at least $\frac{Q}{R} > 1$. Let x_i be the random variable expressing the total bids of all buyers other than g to seller $P_i \in P$. Then $\sum E[x_i] = R - 1$. The payoff of g at P_i is $E_{x_i}[q_i \frac{q_i/Q}{x_i + q_i/Q}] \geq q_i \frac{q_i/Q}{E[x_i] + q_i/Q}$ (the inequality follows from the fact that for every non-negative constant c, the function $\frac{1}{x+c}$ is convex in the domain $x > 0$). Hence the total payoff for g is $\sum_i q_i \frac{q_i/Q}{E[x_i] + q_i/Q}$. Now a convexity argument similar to that of Theorem 3.1 in our paper [one needs to add the details, assuming that they are correct] implies that the sum is minimized (conditioned on $\sum E[x_i] = R - 1$) when $E[x_i] = (R-1)\frac{q_i}{Q}$ for every i, giving a total of $\sum_i q_i \frac{1}{(R-1)+1} = \frac{Q}{R}$. The best response property then implies that in the given Nash every good player has expected payoff more than 1. As the total expected payoff of all good buyers is at most Q, there are strictly less than Q good buyers.

Consider now an arbitrary bad buyer b. If b were to become good and use the balanced strategy, his expected payoff would be $E_x[\frac{Q}{x+1}] \geq \frac{Q}{E[x]+1} = \frac{Q}{R+1} \geq 1$. The best response property then implies that in the given Nash every bad player has expected payoff at least 1. As the total expected payoff of all bad buyers is at most $n - Q$, there are at most $n - Q$ bad buyers.

The number of flexible buyers is strictly larger than $n - Q - (n - Q) = 0$. Hence there is at least one flexible buyer. As the expected payoff of every non-flexible buyer is at least 1, there must be at least one flexible buyer whose expected payoff is at most 1. Consider such a flexible buyer f and let $r > 0$ be the expected bid of f in P. If f were to become good and use the balanced strategy his expected revenue would be $E_x[\frac{Q}{(x-r)+1}] \geq \frac{Q}{E_x[x]+1-r} > \frac{Q}{R+1} = 1$, contradicting the best response property of the assumed Nash equilibrium.

C Proof of Theorem 8

Proof of First Part. We prove that if H uses proportional allocation among rigid bids, and regardless of the policy that L uses, the revenue of H in any Nash equilibrium of the resulting buyers' subgame is at least $n-1-\frac{2}{n-2}$.

We use the notation of Theorem 5 (who also considers H that uses proportional allocation among rigid bids). If $|X| \leq 1$, the claim immediately follows. Thus assume that $|X| \geq 2$, and let P be the player that Lemma 3 identifies. The proof of Theorem 5 shows that the expected number of players besides P that bid at L is at most $\frac{2}{n-2}$, hence L's expected revenue is at most $1+\frac{2}{n-2}$, and the first part of the theorem follows. □

Proof of Second Part. We prove that for any monotone policy of H there exists a policy of L such that, in at least one Nash equilibrium of the resulting buyers' subgame, the revenue of H is at most $n-1-\frac{1}{n^2}$.

Fix an arbitrary monotone strategy for H. Let p_1 be the utility of buyer 1 when all buyers bid their full budget at H. W.l.o.g., $p_1 \leq 1-1/n$. Let p_2 be the utility of buyer 2 when all buyers except 1 bid their full budget at H. W.l.o.g., $p_2 \leq 1$. Let p_3 be the utility of buyer 1 when all buyers but 2 bid their full budget at H. Clearly, $p_3 \leq n-1$, as H has only $n-1$ items.

Consider the following strategy for L. It considers only rigid bids, and only from buyers 1 and 2. If 1 bids alone at L, she receives the full item. If buyer 2 bids at L (either with or without buyer 1), she receives p_2 items and buyer 1 (if at L) receives the leftover $1-p_2$.

We will show that the following is a mixed Nash. Buyer 1 bids her full budget at L, buyers $3,...,n$ bid their full budget at H, and buyer 2 bids her full budget at L with probability $q=\frac{1}{n^2}$ and at H with probability $1-q$. Buyers $3,...,n$ will not deviate to a lower bid at H since H is monotone, and they will not deviate to L since L offers them nothing. Buyer 2 is indifferent between H and L since she obtains utility p_2 in each one. Therefore any mix between the two sellers is a best response of buyer 2.

It remains to verify that buyer 1 is best responding. Buyer 1 is bidding deterministically at L, and obtains expected utility $(1-q)+q(1-p_2) \geq (1-q)$. The only deviation that may be profitable for her is to bid the full budget at H (since L accepts only rigid bids and H is monotone). If she bids her full budget at H she obtains expected utility $(1-q)p_1 + qp_3 \leq (1-q)(1-\frac{1}{n}) + q(n-1)$. Since $q=\frac{1}{n^2}$ we have $(1-q) > (1-q)(1-\frac{1}{n}) + q(n-1)$, which implies that buyer 2 is best responding as well.

In the outcome of this equilibrium play, L's expected revenue is $1+q = 1+\frac{1}{n^2}$, and the second part of the theorem follows. □

References

Acemoglu, D., K. Bimpikis and A. Ozdaglar (2009). “Price and capacity competition.” *Games and Economic Behavior*, **66**(1), 1–26.

Martínez-de Albéniz, V. and K. Talluri (2011). “Dynamic price competition with fixed capacities.” *Management Science*, **57**(6), 1078–1093.

Armstrong, M. (2013). “A more general theory of commodity bundling.” *Journal of Economic Theory*, **148**(2), 448–472.

Armstrong, M. and J. Vickers (2010). “Competitive non-linear pricing and bundling.” *The Review of Economic Studies*, **77**(1), 30–60.

Ashlagi, I., B. Edelman and H. Lee (2011). “Competing ad auctions.” Tech. Rep. 10-055, Harvard Business School.

Athey, S., E. Calvano and J. Gans (2011). “The impact of the internet on advertising markets for news media.” Tech. Rep. 2180851, Rotman School of Management.

Balinski, M. and H. Young (2001). *Fair representation: meeting the ideal of one man, one vote.* Brookings Inst Pr.

Barut, Y. and D. Kovenock (1998). “The symmetric multiple prize all-pay auction with complete information.” *European Journal of Political Economy*, **14**(4), 627–644.

Bernheim, B. D. and M. D. Whinston (1998). “Exclusive dealing.” *Journal of Political Economy*, **106**(1), 64–103.

Borel, E. (1921). “La theorie du jeu les equations integrales a noyau symetrique.” *Comptes Rendus del Academie*, **173**, 1304 – 1308. English translation by Savage, L. (1953). The theory of play and integral equations with skew symmetric kernels. *Econometrica*, 21, 97 – 100.

Calzolari, G. and V. Denicolò (2013). “Competition with exclusive contracts and market-share discounts.” *American Economic Review*, **103**(6), 2384–2411.

Che, Y.-K. and I. Gale (2003). “Optimal design of research contests.” *The American Economic Review*, **93**(3), 646–671.

Che, Y.-K. and I. L. Gale (1998). “Caps on political lobbying.” *American Economic Review*, pp. 643–651.

Chone, P. and L. Linnemer (2015). “Nonlinear pricing and exclusion: I. buyer opportunism.” *The RAND Journal of Economics*, **46**(2), 217–240.

Clark, D. J. and C. Riis (1998). “Competition over more than one prize.” *American Economic Review*, pp. 276–289.

Davidson, C. and R. Deneckere (1986). “Long-run competition in capacity, short-run competition in price, and the cournot model.” *The Rand Journal of Economics*, pp. 404–415.

Dechenaux, E., D. Kovenock and R. M. Sheremeta (2014). "A survey of experimental research on contests, all-pay auctions and tournaments." *Experimental Economics*, pp. 1–61.

Dudey, M. (1992). "Dynamic Edgeworth-Bertrand competition." *The Quarterly Journal of Economics*, pp. 1461–1477.

Edelman, B., M. Ostrovsky and M. Schwarz (2007). "Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords." *American Economic Review*, **97(1)**, 242 – 259.

Friedman, L. (1958). "Game-theory models in the allocation of advertising expenditures." *Operations Research*, **6(5)**, 699 – 709.

Ghemawat, P. and A. M. McGahan (1998). "Order backlogs and strategic pricing: the case of the us large turbine generator industry." *Strategic Management Journal*, **19**(3), 255–268.

Hart, S. (2008). "Discrete Colonel Blotto and General Lotto games." *International Journal of Game Theory*, **36**(3-4), 441–460.

Johari, R., G. Y. Weintraub and B. Van Roy (2010). "Investment and market structure in industries with congestion." *Operations Research*, **58**(5), 1303–1317.

Johnson, N. and S. Kotz (1977). *Urn models and their application: an approach to modern discrete probability theory*. Wiley New York.

Konrad, K. A. (2009). *Strategy and Dynamics in Contests*. OUP Catalogue. Oxford University Press.

Kovenock, D. and B. Roberson (2012). "Conflicts with multiple battlefields." In M. R. Garfinkel and S. Skaperdas, eds., *The Oxford Handbook of the Economics of Peace and Conflict*. Oxford University Press, New York.

Kvasov, D. (2007). "Contests with limited resources." *Journal of Economic Theory*, **136**(1), 738–748.

Laslier, J.-F. and N. Picard (2002). "Distributive politics and electoral competition." *Journal of Economic Theory*, **103**(1), 106–130.

Levitan, R. and M. Shubik (1972). "Price duopoly and capacity constraints." *International Economic Review*, **13**(1), 111–122.

McAfee, R. P. (2005). "The price is right mysterious." *Engineering and Science*, **68**(3), 32–42.

Michaels, R. (1988). "The design of rent-seeking competitions." *Public Choice*, **56**(1), 17–29.

Moldovanu, B. and A. Sela (2006). "Contest architecture." *Journal of Economic Theory*, **126**(1), 70–96.

Myerson, R. B. (1993). "Incentives to cultivate favored minorities under alternative electoral systems." *American Political Science Review*, **87**(04), 856–869.

O'Brien, D. P. and G. Shaffer (1997). "Nonlinear supply contracts, exclusive dealing, and equilibrium market foreclosure." *Journal of Economics & Management Strategy*, **6**(4), 755–785.

Roberson, B. (2006). "The Colonel Blotto game." *Economic Theory*, **29**(1), 1–24.

Roberson, B. and D. Kvasov (2012). "The non-constant-sum Colonel Blotto game." *Economic Theory*, **51**(2), 397–433.

Robson, A. R. (2005). *Multi-item contests.* Australian National University, Faculty of Economics and Commerce, School of Economics. Working Paper No. 446.

Sayedi, A., K. Jerath and K. Srinivasan (2014). "Competitive poaching in sponsored search advertising and its strategic impact on traditional advertising." *Marketing Science*, **33**(4), 586–608.

Siegel, R. (2009). "All-pay contests." *Econometrica*, **77**(1), 71–92.

Snyder, J. M. (1989). "Election goals and the allocation of campaign resources." *Econometrica,*, **57(3)**, 637 – 660.

Tirole, J. (1988). *The theory of industrial organization.* MIT press.

Tullock, G. (2001). "Efficient rent seeking." In *Efficient Rent-Seeking*, pp. 3–16. Springer.

Varian, H. R. (2007). "Position auctions." *international Journal of industrial Organization*, **25**(6), 1163–1178.

Weinstein, J. (2012). "Two notes on the Blotto game." *The BE Journal of Theoretical Economics*, **12**(1).

Whinston, M. (2008). *Lectures on Antitrust Economics.* MIT Press.