

Contests for Revenue Share*

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

December 25, 2017

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 result in disproportionately low equilibrium revenues for her smaller opponents.

*This work was conducted at Microsoft Research. A preliminary version of this work, 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 was supported by a Marie-Curie IOF fellowship Advance-AGT 330240 during the preparation of this paper.

[§]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

There could be a variety of reasons for differences in revenue shares. For example, it might be that the “items” are valued differently, and this can be checked empirically.³ Another reason could be that sellers can somehow leverage their market share to obtain a disproportionately higher revenue share. In this paper we focus on this latter hypothesis, and study the connection between market shares and revenue shares in cases where all items are in fact identical.

Our theoretical abstraction captures markets with four 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

¹The “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. (In various literature this is also termed “capacity share”, “unit market share”, and “volume market share”.) Its “revenue share” (also commonly termed “revenue market share”, “value market share”, and “dollar market share”) is the budget spent on ads in this search engine relative to the budgets spent on all search ads.

²Source: 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/>.

³For example, an ad slot for a search “pizza in NYC” on Google might provide a different value than the same ad slot for the same search on Bing. We are not aware of any empirical investigation of this.

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 strategy space of the opponent to defend against. Third, we are interested to study asymmetries, even extreme, among the sellers, while assuming the simpler case of identical buyers. Fourth, we wish to capture realistic utilities of advertisers with budgets who usually do not care about any unspent budget (or might even dislike having a left-over budget as this sometimes implies a reduction in future budget assignments). This is a common assumption in many marketing/advertising models.⁴

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 and on directed search 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 5 provides many references and a detailed discussion.

Our first result identifies a very simple contest that provides a strong safety level to each seller that uses it. 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. This also implies that i 's revenue share in any subgame perfect equilibrium will be at least her market share minus $\frac{1}{n}$.

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.

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.⁵ 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

⁴Anyone with a research grant experiences a similar utility structure, for example.

⁵For 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.

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 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 equilibria outcomes. In the latter case, a large seller may decrease the revenue share of her significantly smaller opponent substantially below her market share.

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 equilibria outcomes.

Despite this, our analysis of this case shows that the revenue share of the small seller will be significantly lower than her market share in all subgame perfect equilibria. To achieve this, the large seller relies on a certain non-monotonicity, alongside exclusivity, to overcome the limitations of monotone contests. Broadly speaking, this non-monotonicity is conceptually related to situations where a seller offers attractive deals to her rival’s customers.

The fact that such non-monotonicity 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. We further discuss this in Section 6.3.

The remainder of the paper is organized as follows. Section 2 lays out our model and terminology. Section 3 discusses the case where market shares are large relative to $\frac{1}{n}$. Section 4 discusses the case of two extremely asymmetric sellers, where the market share of the small seller is exactly $\frac{1}{n}$.

Related literature is surveyed in Section 5. Section 6 provides concluding remarks. Some additional details, examples, and proofs are provided in two appendices.

2 Setup

There are m sellers and n buyers. There is only one type of good, and this good is divisible.⁶ Each seller has a fixed quantity of the good. 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 $B = 1$. 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 quantity of the good that he receives, or in situations in which randomization is involved, to maximize the expected quantity. Buyers do not derive any utility from unspent budgets. The objective of a seller is to maximize her expected revenue – her sum of bids. A seller does not derive any utility from any left-over quantity of the good.

Since the good is divisible, by scaling (changing units) we may assume without loss of generality that the total quantity of the good is n , the number of buyers. We shall refer to this amount as being composed of n identical *items*, with the understanding that each item is divisible as well. For example, if the total quantity of the good (say, sugar) is 10 kilograms, and there are 8 buyers, then the quantity of an item is $\frac{10}{8}$ kilograms. Thus, the overall number of items is 8, which is exactly the number of buyers.⁷ The number of items of seller j is denoted by $q_j n$ (this need not be an integer), where $q_j > 0$ and $\sum_{j=1}^m q_j = 1$.

While budgets and items are divisible, we assume that there is some finite precision δ to bids and allocated quantities of items. This makes the games that we study finite, and ensures the existence of equilibria. Our results hold for any arbitrarily small precision δ , that can depend on the other parameters, e.g., on n and m . In particular, it will be important to choose a small-enough δ so that all the quantities that we mention in the sequel will be multiples of δ .

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 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 buyers and sellers act strategically, it may well be that strategic issues will lead to

⁶We discuss the connection of our specific modeling choices to the motivation of Internet search advertising in Section 6.2 below.

⁷This scaling is without loss of generality in our model, we make it only since we believe that it simplifies notation, our results hold regardless of the relation between the number of items and the number of buyers, as long as the good is divisible.

a different outcome. We therefore consider 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* $a_j : \{0, \delta, 2\delta, \dots, 1\}^n \rightarrow \{0, \delta, 2\delta, \dots, q_j n\}^n$, i.e. a function from the submitted bids b_{1j}, \dots, b_{nj} to an allocation vector $\vec{a}_j = \vec{a}_j(b_{1j}, \dots, b_{nj})$, where $\sum_{i=1}^n a_{ij} \leq q_j n$.⁸
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 $\sum_{j=1}^m a_{ij}(b_{1j}, \dots, b_{nj})$.

Regarding terminology, we use *policies* for sellers' strategies and keep *strategies* for buyers. 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*.⁹ Note that Nash equilibria strategies in every fixed buyers' subgame can be either pure or mixed.

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 $s_i(a_1, \dots, a_m)$. 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. For all:*

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

seller j 's revenue when sellers play (a_j^, a_{-j}) and buyers play $s_1(a_j^*, a_{-j}), \dots, s_n(a_j^*, a_{-j})$ 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 .*

⁸An 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.

⁹This is a Blotto game where each seller is a battlefield and each buyer is a colonel.

Proof. Fix an arbitrary subgame-perfect equilibrium (a_j, a_{-j}, s) . Therefore s is a Nash equilibrium of every buyers' subgame. In particular, $s_1(a_j^*, a_{-j}), \dots, s_n(a_j^*, a_{-j})$ is a Nash equilibrium of the buyers' subgame when seller policies are (a_j^*, a_{-j}) . Therefore 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 beneficial deviation, and the claim follows. \square

Our analysis in the following sections will proceed by identifying extensive-form safety level policies. Section 3 studies the case of sellers with nearly symmetric supply sizes and describes for this case a policy that provides an extensive-form safety-level revenue close to the seller's fair share revenue. Section 4 studies the case of two extremely-asymmetric sellers and describes for this case a policy that provides for the large seller an extensive-form safety-level revenue larger than her fair share revenue. This safety-level strategy implies that in this case of two extremely-asymmetric sellers, the subgame-perfect equilibrium revenue of the small seller will be significantly smaller than her fair share revenue. 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.

3 General Supply Sizes

This section shows a simple policy that provides an extensive-form safety-level revenue which is close to the seller's fair share revenue, for arbitrary supply sizes. This policy, which we term *proportional allocation among rigid bids*, is defined as follows. Let a *rigid bid* be a bid of value 1, i.e., a bid where the full budget of the buyer goes to one seller. In proportional allocation among rigid bids, the seller divides her items equally among the buyers that place rigid bids, and completely ignores non-rigid bids.

Theorem 1. *Proportional allocation among rigid bids provides an extensive-form safety-level strictly larger than $q_j n - 1$, for any seller j .*

Proof. We need to show 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) of the resulting buyers game will be strictly larger $q_j n - 1$.

Let $Q = q_j n$. Consider an arbitrary Nash equilibrium (pure or mixed) of the resulting buyers game and assume for the sake of contradiction that the expected revenue of j in this Nash equilibrium is $R \leq Q - 1$. Let x be a random variable denoting the sum of bids to j . 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 at j , *bad* who deterministically place no budget at j

(though they may randomize outside j), and *flexible* which are all the rest. (Note that some of these classes may be empty.)

Every good buyer has expected payoff $E_x[\frac{Q}{x}] \geq \frac{Q}{E[x]} = \frac{Q}{R} > 1$.¹⁰ 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 i and let $r > 0$ be the expected bid of i at j . If i were to become good his expected revenue would be at least $\frac{Q}{E_x[x]+1-r} > \frac{Q}{R+1} \geq 1$, contradicting the best response property of the assumed Nash equilibrium. \square

Theorem 7 in Appendix A shows that there does not exist an allocation policy that provides an extensive-form safety-level of $q_j n + 1$ (or more). In this respect, proportional allocation among rigid bids is close to optimal. We do not know what is the optimal safety-level in this range. Appendix A also discusses a related allocation policy, proportional allocation, that was first defined by Friedman (1958), and is helpful in proving this latter result.

Theorem 1 holds for any arbitrary supply sizes. However, when $q_j n$ is comparable to 1, the theorem is no longer that significant. 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 with two sellers L and H, 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 an allocation 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. This observation motivates the inquiry performed in the next section.

¹⁰The first inequality follows from the fact that for every non-negative constant c , the function $\frac{1}{x+c}$ is convex in the domain $x > 0$.

¹¹As 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.

4 The Case of Two Extremely Asymmetric Sellers

The example at the end of the previous section shows that with extreme supply differences, a large seller H can attract all budgets, leaving the small seller L with no revenue at all. 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? The “first” extremely asymmetric case which is far from trivial (as will soon become evident) is when L’s supply is exactly 1 and H’s supply is $n - 1$. On the one hand, Theorem 1 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. One can verify that simple rigidity no longer works and the following result shows that in fact a large class of selling policies for H (including rigidity) cannot yield a very low revenue for L in all equilibrium outcomes of the buyers subgame.

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.*

Definition 3. *The **anti-predatory** 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 2. *If L’s supply is exactly 1, H’s supply is $n - 1$, L uses the anti-predatory 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 of the buyers subgame, 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 further studies this case and identifies a safety-level strategy that does manage to reduce L’s revenue share significantly below her market share in any subgame-perfect outcome. We discuss how extreme supply asymmetries can be found in Internet search advertising in Section 6.1 below.

Part of the complexity of this case arises since the buyers subgame may have multiple NE, some are pure and some are mixed. For example, when L uses the anti-predatory allocation policy and H

uses proportional allocation among rigid bids, every *pure* Nash equilibrium of the resulting buyers game 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 game, including a symmetric mixed Nash equilibrium in which each buyer submits his full budget to L with some probability p and with probability $1 - p$ submits his full budget to H. Calculating the probability p which makes this strategy a mixed NE, it turns out that L’s revenue dramatically drops to about $\frac{1}{n}$ (while in every pure NE L’s revenue is 1).¹² The safety level strategy that we next construct for H builds upon this observation that H should prefer mixed NE.

Definition 4. *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.*¹³

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). When H uses this policy, the anti-predatory policy no longer provides any benefit for L. In particular, one can easily verify that the pure NE described in the proof of Theorem 2 no longer exists.¹⁴

Theorem 3. *Assume $n \geq 10$, the supply of H is $n - 1$, and the supply of L is 1. Fix any epsilon > 0 . Then, proportional allocation among rigid bids with a gamble bid provides an extensive-form safety-level of at least $n(1 - \epsilon) - \frac{8}{n-5}$. Therefore, L’s revenue in any subgame perfect equilibrium is at most $\frac{8}{n-5}$ (while her fair share revenue is 1).*

Proof. To prove the first part, we fix L’s allocation policy to be any arbitrary policy, H’s policy to be proportional allocation among rigid bids with a gamble bid, and analyze the resulting revenue of L and H in any pure or mixed NE outcome of the resulting buyers game. We will rely on the following lemma.

¹²In a symmetric NE a buyer’s expected utility from (1) submitting his full budget to L is equal to his expected utility from (2) submitting his full budget to H. The utility from (1) is 1 if no other buyer submits to L which happens w.p. $(1 - p)^{n-1}$, and 0 otherwise. The utility from (2) is at least $\frac{n-1}{n}$, regardless of the actions of the other buyers. Thus p satisfies $(1 - p)^{n-1} \geq \frac{n-1}{n}$, implying that $p = O(\frac{1}{n^2})$. L’s expected revenue which is $n \cdot p$ is therefore $O(\frac{1}{n})$.

¹³The value $1 - \epsilon$ needs to be a multiple of the precision δ , so only ϵ ’s that satisfy this can be chosen. In addition δ should be determined such that the quantity $1 + 3/n$ will be a multiple of δ .

¹⁴Note that proportional allocation among rigid bids with a gamble bid is not a monotone policy, see further discussion in Section 4.1 below.

Lemma 2. *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 buyers besides j who place a positive bid at L is at most $\frac{4}{n-5}$.*

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.

For the second part, let L_i for $i = 1 \dots n-1$ be random independent Bernoulli variables that represent the fact that buyer i submitted a positive bid to L , where $Pr(L_i = 1) = p_i$ and $Pr(L_i = 0) = 1 - p_i$, and let $L = \sum_i L_i$. From the first part of this claim we know that $Pr(L \geq 1) \leq \frac{4}{n-1}$. Let $E = E[L] = \sum_i p_i$. We have $\frac{4}{n-1} \geq Pr(L \geq 1) = 1 - \prod_i (1 - p_i) \geq 1 - e^{-\sum_i p_i} = 1 - e^{-E}$, where the second inequality follows since $1 - x \leq e^{-x}$. Thus $E \leq \ln(n-1) - \ln(n-5) < \frac{4}{n-5}$ (using $\ln(a) = \int_1^a \frac{1}{x} dx$), and the lemma follows. \square

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 2, the expected number of players at L besides j is at most $\frac{4}{n-5}$, and the probability that j bids at L is at most $\frac{4}{n-1}$ (again by Lemma 2, since i also declares a near rigid bid to H with positive prob.). Thus, the expected number of players at L is at most $\frac{8}{n-5}$, which implies the theorem for this case.

Case 2. Let X be the set of buyers that have positive probability of submitting a positive bid to L , and assume (this is the second case we are considering) that $|X| \geq 2$. For any $S \subseteq X$, let $Pr[S]$ be the probability that S is exactly the set of buyers that submit a bid larger than ϵ to L (in this case we say that S is the ‘‘colliding set’’). Note that if a buyer submits such a bid to L then he 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 he bids at L and there is a collision at L .

Lemma 3. *If $|X| \geq 2$, there exists a buyer 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,

$$\begin{aligned} \sum_{i \in X} \sum_{S \in \mathcal{S}_i} u[i|S] Pr[S] &= \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} Pr[S] \sum_{i \in S} u[i|S] \leq \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} Pr[S] \\ &\leq \frac{1}{2} \sum_{S \subseteq X \text{ s.t. } |S| \geq 2} |S| Pr[S] = \frac{1}{2} \sum_{i \in X} Pr[i]. \end{aligned}$$

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

Let $P \in X$ be such that $u_P \leq \frac{1}{2}$. 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 buyer $j \neq P$ bids at H, his bid is rigid. Then, buyer P can obtain expected utility strictly larger than 1 by deterministically placing a near rigid bid at H, since with this action his 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 his expected utility given that he bids at L is at most 1. By Lemma 2 the expected number of all buyers besides j (but including P) that bid at L is at most $\frac{4}{n-5}$. The probability that j bids at L is at most x . Thus, the expected number of players at L is at most $\frac{4}{n-5} + \frac{2(n+3)}{n(n-1)} < \frac{8}{n-5}$, 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 2 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. \square

4.1 Restricting H to be Monotone

Non-monotonicity is a crucial ingredient in the previous analysis. To see this, recall the anti-predatory 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-predatory 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. One might wonder if restricting H to use only a monotone policy will help limiting her monopolistic power. For example, does there exist a policy for L that provides a safety-level of 1, if H is restricted to be monotone? The next theorem shows that this is unfortunately not true.

Theorem 4. *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 $3/n$. Therefore, L cannot obtain an extensive-form safety-level of 1 even if H is restricted to be monotone.*

Remark. If H chooses the policy described below in the proof of this theorem and L uses the anti-predatory allocation policy described above, the resulting buyers subgame will have at least two Nash equilibria: a pure NE (as described in Theorem 2) in which L 's revenue is 1, and a mixed NE (as described in the proof below) in which L 's revenue will be at most $3/n$.

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 whose 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, \dots, 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. \square

We have shown that L cannot obtain a safety-level of 1, even if H is restricted to be monotone. On the other hand, Theorem 2 shows that the safety-level that H can obtain under the restriction of monotonicity cannot be strictly larger than $n - 1$. Can H obtain at least that? In an additional result that we do not include in this paper, we show that proportional allocation among rigid bids (which is monotone) provides H an extensive-form safety-level of at least $n - 1 - \frac{2}{n-2}$, and that any monotone allocation policy can provide H a safety-level of at most $n - 1 - \frac{1}{n^2}$. Thus, when H must be monotone she 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.

5 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). Models on directed search and the Bertrand paradox, following Peters (1984), are related to our results in similar ways. Our work significantly differs from all this large body of literature in two main aspects. First, on the sellers' side, this literature mostly focuses on pure price competition while our approach captures a wide array of non-linear pricing aspects of the competition. Some recent directed search models have recognized the importance of generalized pricing methods. For example, Geromichalos (2014) allows the sellers to post item prices that depend on the ex-post realized demand, i.e., the item price can be a function of the number of buyers that decide to buy from this seller. Our model covers this particular pricing method and is clearly more general. 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.¹⁵

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 (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 usefulness 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.

Competition has also been studied in the literature on auctions and mechanism design, for example, McAfee (1993); Peters and Severinov (1997); Burguet and Sakovics (1999); Ellison, Fudenberg and Mobius (2004); Pai (2010); Ashlagi, Monderer and Tennenholtz (2011b). While these papers handle the complications and challenges of incomplete-information, they usually analyze standard auction formats.

¹⁵Earlier 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.

6 Concluding Remarks

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 bid rigidity tend to work in favor of large sellers, that mixed Nash outcomes (compared to 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. Keeping in mind our motivating example, we next give few concluding remarks.

6.1 Remarks on situations with extreme supply asymmetries

The market for search advertising is divided to many small and only loosely related 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 demonstrates 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 at least \$79 – very close to her fair share revenue which is \$80. If, however, the same grand market is actually composed of 100 small unrelated 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.¹⁶ 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 3 demonstrates that H can still cut L’s revenue to about 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.¹⁷

¹⁶Specifically, 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.

¹⁷Specifically, if H implements the strategy suggested in Theorem 3 in every small market then according to the theorem L’s revenue in every such market will be at most $\frac{8}{n-5} = \frac{8}{20-5} < 0.54$. Aggregating over the 40 small markets results in a total revenue of less than \$22 for L.

6.2 Remarks on our modelling choices

Comparing the formal model that we have studied to our original motivation of Internet search advertising raises three issues.

First, recall that sellers in our model are search engines, buyers are advertisers, and items are 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 advertisers 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 operation 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.

6.3 Some possible interpretations of our allocation policies

As already remarked, bid rigidity can be viewed as an exclusivity contract. This might be enforced by a seller, via various technical or marketing actions.¹⁸ 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 (2011a) 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 exchange economy, rigidity simply captures the case that buyers’ endowments are indivisible, e.g., in cases where it is not money. One concrete example for such a setting is given in Section 6.4 below, other examples naturally arise when Blotto games are motivated by scenarios with indivisibilities, like voting scenarios where a voter must fully assign his vote to one of the parties.

The use of a “gamble bid” is conceptually related to a situation where a seller offers a “new customer discount” to a buyer that might otherwise deterministically put his entire budget elsewhere. Specifically, the gamble bid becomes relevant when H receives $n - 1$ bids of value 1 and given this the remaining buyer will prefer to receive one item from L rather than bidding at H. In this situation, our allocation policy offers the gamble bid, which views the last buyer as a new customer and aims to attract him by offering a special deal (a slightly larger quantity) to this buyer. Of course, our setting is one-shot and not sequential, hence formally no buyer is labeled as the “remaining buyer”. This term generically refers to 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 such a buyer. Thus, conceptually, pure NE in our model can be viewed as situations where sellers reward and encourage loyalty and stability, while the introduction of the gamble bid promotes instability in the game, pushing towards a mixed Nash equilibrium outcome. This seems to contradict McAfee’s recipe (described in the Introduction) that suggests to foster stability and loyalty in competitive situations. In our specific competitive situation, it turns out that H should promote instability.

¹⁸To 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).

6.4 A connection to apportionment methods

Our analysis has implications to the issue of seat allocation in a parliament system, based on election results. This is termed apportionment: how to allocate the parliament seats to the various parties. 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. The D'Hondt method has a reputation of typically favoring the larger parties, in line with some of our analysis.

A The Proportional Allocation Policy

The *proportional allocation policy* is the following natural policy: Given bids b_{1j}, \dots, 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,

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

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 game. In particular, he showed:

Theorem 5 (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 $j = 1, \dots, m$.*¹⁹

¹⁹In fact, it can be shown (see Appendix B) that this is a safety level strategy that provides each buyer that uses

2. The revenue of every seller j exactly equals her fair share revenue $q_j n$.

For completeness, we provide a simple proof of this theorem in Appendix B.

Given the first property in the theorem, 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.

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, as proportional allocation provides a safety level of $q_j n - 1$.

Theorem 6. *Let P be a collection of (one or more) sellers $\{P_1, \dots, P_k\}$ and let $Q = \sum_{i=1}^k q_i n$ denote their total supply. If all members of P use proportional allocation, their total expected revenue in every Nash equilibrium of the resulting buyers game is strictly larger than $Q - 1$, regardless of the allocation policies of sellers not in P .*

Proof. 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

it in every equilibrium outcome a utility of at least 1, regardless of the strategies of the other buyers.

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 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. \square

As a direct corollary, we have:

Theorem 7. *There is no extensive-form safety-level policy that provides seller j a total revenue of $q_j n + 1$ (or more).*

Proof. Let the collection of sellers P of Theorem 6 include all sellers besides j . By Theorem 6, the sellers in P can jointly obtain a total revenue strictly larger than $(1 - q_j)n - 1$ in all NE of the resulting buyers game, which implies the claim. \square

Thus, we have established that the optimal safety level of a seller that has total supply Q lies in the interval $[Q - 1, Q + 1]$.

B A Simple Proof of Friedman's Theorem

We start with an auxiliary result:

Theorem 8. *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 $j = 1, \dots, m$, 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. \square

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. [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 8), 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 8 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. \square

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 (2011a). “Competing ad auctions.” Tech. Rep. 10-055, Harvard Business School.
- Ashlagi, I., D. Monderer and M. Tennenholtz (2011b). “Simultaneous ad auctions.” *Math. Oper. Res.*, **36**(1), 1–13.
- 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.
- 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.
- Burguet, R. and J. Sakovics (1999). “Imperfect competition in auction designs.” *International Economic Review*, **40**(1), 231–247.
- 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.
- Chone, P. and L. Linnemer (2015). “Nonlinear pricing and exclusion: I. buyer opportunism.” *The RAND Journal of Economics*, **46**(2), 217–240.
- 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.
- 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.
- Ellison, G., D. Fudenberg and M. Mobius (2004). “Competing auctions.” *Journal of the European Economic Association*, **2**(1), 1–12.
- Friedman, L. (1958). “Game-theory models in the allocation of advertising expenditures.” *Operations Research*, **6**(5), 699 – 709.
- Geromichalos, A. (2014). “Directed search and the bertrand paradox.” *International Economic Review*, **55**(4), 1043–1065.

- 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.
- 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. (1993). “Mechanism design by competing sellers.” *Econometrica*, **61**, 1281–1312.
- McAfee, R. P. (2005). “The price is right mysterious.” *Engineering and Science*, **68**(3), 32–42.
- 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.
- Pai, M. (2010). “Competition in mechanisms.” *SigEcom Exchanges*, **9**(1).
- Peters, M. (1984). “Bertrand equilibrium with capacity constraints and restricted mobility.” *Econometrica*, pp. 1117–1127.
- Peters, M. and S. Severinov (1997). “Competition among sellers who offer auctions instead of prices.” *Journal of Economic Theory*, **75**, 141–179.
- 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.
- Sayed, 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.

- 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.