

Competition Among Asymmetric Sellers With Fixed Supply*

Uriel Feige[†]

Ron Lavi[‡]

Moshe Tennenholtz[§]

May 14, 2013

Abstract

Motivated by the market for display advertisement over the Internet, we study competition between firms with a fixed supply whose size cannot be changed, and analyze their resulting revenue. We are most interested in studying the asymmetric case in which one large seller dominates the market and competes against a small new entrant seller. We present a model in which sellers announce selling policies, and given these policies buyers distribute their budget in a strategic fashion among sellers so as to maximize the portion of the supply that they receive. As a function of the policies of the sellers, we analyze revenue of sellers in pure and mixed Nash equilibria for the buyers. Our results show an interesting contrast between the near-symmetric case (sellers with similar supply sizes) and the extremely asymmetric case (a very large seller vs. a very small seller). In particular, in the near-symmetric case, simple policies can ensure each seller a revenue almost proportional to her market share. In contrast, in the asymmetric case the large seller has a selling policy that yields disproportionately low revenue for the small seller. Interestingly, in our abstract model, non-monotone selling policies (namely, sometimes giving more of the supply to a buyer who decreases his bid) can offer advantages to the large seller that are (provably) impossible to achieve via monotone selling policies.

1 Introduction

1.1 Background and Motivation

In the last decade the Internet has become a popular platform for display advertising, where thousands of advertisers post hundreds of millions of display ads every day. Revenues of the large Internet companies from selling slots for online ads are huge. Our starting point is the general question of how these companies compete to obtain the largest possible fraction of the overall budget for online advertising. This competition has one important aspect that distinguishes it from the classic competition models of Cournot (Cournot, 1838) and Bertrand (Bertrand, 1883), where firms compete by adjusting the size of their supply. In contrast, the supply of slots for ads can be viewed as fixed, especially in the short and medium run. Increasing the supply of slots requires

*This work was conducted at Microsoft Research.

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

[‡]Faculty of Industrial Engineering and Management, Technion – Israel Institute of Technology. Email: ronlavi@ie.technion.ac.il.

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

increasing the number of users that enter to watch content, which is a difficult challenge and not a simple strategic action available to the pricing managers of these companies. How should Internet publishers (as well as other types of sellers) compete in markets that exhibit such a constraint?

A seminal paper by Friedman (1958) studies competition in advertising markets in a model that captures similar issues. In Friedman’s model, each advertiser strategically partitions her budget among various fixed-size markets, and resulting market shares are then given by the ratio of an advertiser’s budget targeted to that market relative to the overall budget targeted to that market. Over the years, this model has been extended and applied to many settings. For example, Snyder (1989) and others have applied and generalized the model to study how political campaigns distribute advertising budgets among the different legislative districts (the “markets”). Snyder’s model includes a more general formulation of the “market share” (number of votes, in this case) as a function of the budgets allocated to each market. More abstractly, these games are known as Blotto games with lottery contest functions, where markets are battlefields and advertisers are generals that need to assign soldiers (budgets) in order to win as many battlefields as possible. These games have attracted continuous theoretical interest over the years, as they provide a unified theoretical framework that sheds light on a host of important issues in a broad set of environments (see a recent survey by Kovenock and Roberson (2012)).

This framework is a good starting point for our purpose. Internet publishers are battlefields/markets, and an advertiser’s market share is the number of ad slots allocated by the publisher to the advertiser. However, the most important piece is still missing: how do publishers strategically allocate ad slots to advertisers? Both Friedman’s model and the more general framework of Blotto games do not take into account the strategic aspects of publishers’ actions. In Friedman’s model and in its various generalizations, the resulting market share of each advertiser is a fixed *exogenous* function of the budgets targeted to that market (publisher, in our case). In contrast, large Internet publishers like Google and Yahoo! determine endogenously their ad allocation policies, to maximize revenue. They are key players in the game. In the abstract terminology of the Blotto game, we would like to remove the assumption that battlefields have fixed contest functions. Instead, each battlefield strategically sets its contest function to maximize the number of soldiers sent to it.

To incorporate this issue in Friedman’s model, we study the following setup: There are m sellers (publishers) and n completely symmetric buyers (advertisers), where each buyer has the same budget, B . Without loss of generality we can assume that $B = 1$. There is only one type of good, and this good is divisible. By changing units, we assume without loss of generality that the total quantity of the good is n (we can think of this quantity as being composed of n identical *items*, with the understanding that each item is divisible as well). The amount of supply of seller j is denoted by $q_j n$, where $q_j > 0$ and $\sum_{j=1}^m q_j = 1$. Sellers and buyers engage in a two-stage game. First, each seller announces an *allocation policy*, specifying the number of items that each buyer

will receive, given the fractions of budgets that buyers assign to this seller. Second, given these allocation policies, buyers face a complete information game, and we assume that they play some (either pure or mixed) Nash equilibrium of this game. The utility of a seller equals her resulting revenue (sum of budgets allocated to her), and the utility of a buyer equals the total number of items she obtained.

1.2 Overview of Results

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 that 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 . This is also the unique Walrasian outcome of this market, if one would like to analyze this simple market in such a way. As emphasized above, we rather focus on the strategic aspects of the competition among sellers, especially since we are motivated by a market in which the number of sellers is very small, even just two.

The fundamental question that we study in this paper is whether there exists an allocation policy for seller j that will guarantee revenue close to her fair share revenue (or maybe even more?), regardless of the allocation policies chosen by her competitors. (I.e., we focus on safety level strategies.) First consider the natural, fixed policy originally assumed by Friedman (here, we will refer to it as the *proportional allocation policy*): denote buyer i 's bid to seller j as b_{ij} ; the number of items that seller j allocates to buyer i is exactly $\frac{b_{ij}}{\sum_i b_{ij}} q_j n$, namely, proportional to buyer i 's bid at j . With this allocation policy, the connection between the strategic setting and the Walrasian setting becomes explicit, as the following theorem shows.

Theorem I (Friedman (1958)). *If all sellers use the proportional allocation policy, then the resulting buyers game has a unique Nash equilibrium. This Nash equilibrium is pure, and the revenue of each seller exactly equals her fair share revenue.*

In other words, if all sellers happen to choose the proportional allocation policy, they collectively implement the Walrasian outcome. But can a seller *guarantee* her fair share revenue, even when the other sellers choose different allocation policies? (we use “guarantee” in a strong way: for any policy of the other sellers, and for any Nash equilibrium of the resulting buyers game.) It turns out that proportional allocation can guarantee revenue very close to the fair share revenue:

Theorem II. *If seller j uses proportional allocation 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 $q_j n - 1$.*

This theorem is further discussed and proved in Section 3. It implies that sellers with large supply are able to *guarantee* almost their fair share in revenue (in relative terms). Sellers with small supply (in which $q_j n = O(1)$) may still get only a small fraction of their fair share.

Rigid bids. 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. 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. In the more broader context of an abstract exchange economy, rigidity simply captures the case that the buyers’ endowment is indivisible (in cases where it is not money).

To study our setting when rigidity is enforced, we define the policy of *proportional allocation among rigid bids*: the seller divides her items equally among the buyers that place rigid bids (and non-rigid bidders are ignored). Section 3.1 shows that, similarly to Theorem II, 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 a slightly better guarantee than that of proportional allocation. It also shows examples in which the use of proportional allocation among rigid bids gives large sellers a revenue significantly larger than their fair share. For the case of pure Nash equilibria of the resulting buyers game, this phenomenon is strongly connected to methods that are in common use for allocating parliamentary seats to political parties based on election results. This analogy is explained in Section 5.

Two extremely asymmetric sellers. When we have two nearly symmetric sellers, Theorem II shows that each of them is able to guarantee almost their fair share in revenue (when n is sufficiently large). However, in extremely asymmetric cases this is no longer true. If L’s supply is $1 - \frac{1}{n}$ and H’s supply is $n - (1 - \frac{1}{n})$, H can attract the full budget of all buyers by simply offering each one $1 - \frac{1}{n} + \frac{1}{n^2}$ items (this is a fraction of $\frac{1}{n}$ of the total supply of H, hence this can be implemented e.g. by proportional allocation).

An extremely asymmetric case that is much more interesting to analyze is the case when L’s supply is exactly 1 and H’s supply is $n - 1$. On the one hand, the guarantee provided by Theorem II is still meaningless, but on the other hand it is not clear how the sellers should act in such a case. Can H limit the revenue of L to be virtually 0, similarly to the case when L’s supply is $1 - \frac{1}{n}$? Or maybe L can guarantee revenue almost proportional to her fair share revenue? We study this special case in detail to obtain full answers to these questions.

Bid rigidity turns out to be of particular importance to our analysis of this case. When there are two extremely asymmetric sellers H and L, 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 (that depends only on n), and with probability $1 - p$ submits his full budget to H. One may argue that since buyers are completely symmetric, the symmetric mixed Nash is as reasonable as any pure Nash, which cannot be symmetric in our case. It turns out that L’s revenue dramatically drops in every mixed Nash equilibrium outcome compared to her fair share:

Theorem III. *For H and L as above, if H uses proportional allocation among rigid bids whereas L uses proportional allocation (either among rigid bids or not), then in every mixed Nash equilibrium of the resulting buyers game the expected revenue of L is $O(\frac{1}{n})$ and the revenue of H is $n - O(\frac{1}{n})$.*

This theorem is further discussed and proved in Section 4.1. According to this theorem, L’s loss in every mixed Nash equilibrium is significant, and may drive her out of the market. In this respect, the large seller may be motivated to enforce bid rigidity on the buyers. However, the theorem is inconclusive, as it leaves some loopholes for L: the possibility of reaching a pure (rather than a mixed) Nash, and the possibility that allocation policies that are not proportional will offer better revenue for L. We show that H has a different allocation policy that eliminates these loopholes.

The strategy that we suggest for H in this case further modifies proportional allocation among rigid bids, adding a “gamble bid”, as follows. First, H accepts not only a rigid bid (i.e. her full budget), but also a *near rigid* bid (i.e. a bid of $1 - \epsilon$ for some small ϵ) which is the gamble bid. In most cases, H completely ignores near rigid bids, except for one case: if all buyers except one submit rigid bids to H, and the additional buyer submits a near rigid bid to H, then the additional buyer receives $1 + O(\frac{1}{n})$ items, while the other $n - 1$ buyers equally split the remaining items. We emphasize that our model allows buyers to submit any two real numbers (between 0 and 1, that sum to 1) as their two bids to L and H. If H uses proportional allocation among rigid bids with a gamble bid, her allocation policy filters the received bids, effectively causing only rigid and near rigid bids to be “noticed”.

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. In fact no pure Nash equilibrium exists in this case, and the expected revenue of L in every mixed Nash equilibrium is at most $O(\frac{1}{n})$. The following theorem shows that this upper bound on L’s revenue holds regardless of her allocation policy.

Theorem IV. *If H uses proportional allocation among rigid bids with a gamble bid, then regardless*

of L's allocation policy, in every Nash equilibrium (either pure or mixed) of the resulting buyers game, the revenue of L is at most $O(1/n)$.

This theorem is further discussed and proved in Section 4.2.

From all the above, we conclude that in our abstract model there is an interesting contrast between the near-symmetric case and the extremely asymmetric case. In the near-symmetric case, simple policies can ensure each seller a revenue almost proportional to her market share. In contrast, in the extremely asymmetric case (L's supply is at most 1), the large seller has a selling policy that limits the expected revenue of the small seller to almost zero.

The conclusion that a very large seller can have disproportional power over a very small seller is very intuitive, and was already demonstrated many times, e.g., in studies on loyalty discounts and nonlinear pricing (e.g., Armstrong and Vickers, 2001), on exclusive dealing (e.g., Besanko and Perry, 1993), and on switching costs (e.g., Klemperer, 1995). However, this literature mostly studies competition by various price-setting or supply-setting firms, following Bertrand and Cournot. A-priori, it is not clear how a large seller could leverage her supply advantage in our fundamentally different setting. The analysis of our stylized, abstract model confirms that rigidity is a necessary part of the answer. The analysis also reveals that reaching the mixed Nash equilibrium outcome is a second necessary ingredient that enables the large seller to exercise her power.

Theorem IV requires extreme differences in supply, and it is not clear how common such extreme differences are. While we do not analyze any empirical data in this paper, we wish to remark that in some settings extreme differences may be more common than what initially seems. For example, consider a case where a grand market is composed of many small sub-markets. This is the case in the market for display advertising: it is composed of many sub-markets characterized by the properties (location, age, gender, etc.) of the user to whom the ad will be displayed. In our abstract model, consider for example a 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 II shows that the small firm can guarantee a revenue of \$79 – very close to her fair share which is \$80. If, however, the same 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 of L and H), the situation changes drastically. Since L has 0.8 units in each of the small markets, her revenue might drop to zero in every small market (say, if H uses proportional allocation among rigid bids), and thus her total revenue might be zero. With such a market structure, a large incumbent seller has disproportional power over new entrants.

Thus, even situations that appear to be reasonably symmetric might contain hidden extreme asymmetries. It is therefore important to understand exactly when differences in supply are in fact “extreme”. This is a technically challenging question, even in our simple model. Theorem IV makes a first non-trivial step in this direction. In particular, in the example above, it shows that H might drive L's revenue to nearly zero even if the grand market is composed of only 80 small

markets with 10 buyers and 10 divisible items each.

Back to the technical aspects of Theorem IV, the “gamble bid” introduces non-monotonicity in the allocation policy: fix a buyer, i , and suppose all other buyers submit rigid bids to H. Then, if buyer i submits a non-rigid bid $(1 - \epsilon)$ to H he receives strictly more than one unit, while if he increases his bid to 1 he receives strictly less than one unit. The following theorem clarifies the necessity of non-monotonicity in reducing L’s revenue.

Theorem V.

1. *L has an allocation policy such that for every monotone allocation policy used by H, the resulting buyers game has a pure Nash equilibrium in which L’s revenue is 1.*

2. *However, L cannot guarantee a revenue of 1, even if H is restricted to be monotone: For any allocation policy of L, there exists a monotone allocation policy for H for which the resulting buyers game has at least one Nash equilibrium in which L’s revenue is at most $O(1/n)$.*

3. *H can guarantee a revenue of $n - 1 - O(\frac{1}{n})$ using a monotone allocation policy, but cannot guarantee a revenue of $n - 1 - O(\frac{1}{n^2})$ using a monotone allocation policy. In particular, H cannot guarantee a revenue of $n - 1$ if it is restricted to be monotone.*

This theorem is further discussed and proved in Section 4.3.

One way to generalize the setup of two extremely asymmetric sellers is to have one large seller and *multiple* small sellers. In this case, the important variable is the *aggregate* supply of all small sellers:

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

This theorem is further discussed and proved in Section 3. It demonstrates that 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 guarantee almost their fair share, simply by using proportional allocation (each one in an independent way).

1.3 Paper organization

The remainder of the paper is organized as follows. Section 2 gives a formal description of our model and of the terminology that we use. Section 3 discusses the proportional allocation policy. In particular, it shows that with arbitrary supply sizes, a seller that uses proportional allocation is guaranteed to obtain revenue close to her fair share revenue. Section 4 discusses the case of two extremely asymmetric sellers. It shows that in such a case the large seller can keep the small seller on a revenue level significantly smaller than her fair share. Section 5 summarizes, and describes additional applications of our model.

2 Setup

There are m sellers and n buyers. There is only one type of good, and this good is divisible. Each seller has its own fixed supply of the good. By scaling we may assume without loss of generality that the total amount of the good is n . For convenience we shall sometimes refer to this amount as being composed of n identical *items*, with the understanding that each item is divisible as well. The amount of supply 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 the most general model the budget of buyer i is denoted by B_i . However, in this work we only 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 $B = 1$. 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 action of a seller is to distribute items to buyers that placed a positive bid with her. The amount of items distributed cannot exceed her supply. A seller cannot distribute items to buyers that did not give her a positive bid.

The objective of a buyer 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. Here we assume that the revenue of a seller is the total amount of bids that she received. In particular, the revenue includes also those bids from buyers that did not receive an item from her. (This convention might appear strange, but is the natural limit as ϵ tends to 0 to the situation in which the buyer received an ϵ fraction of an item in return to his bid.) A seller does not derive any utility from left-over items. This setup gives rise to a multiplayer game with $m + n$ players, and we view this game as having three distinct, consecutive steps:

1. First, each seller announces an *allocation policy* $a_j : [0, 1]^n \rightarrow [0, q_j n]^n$, i.e. a function from the submitted bids b_{1j}, \dots, b_{nj} to an allocation vector $\vec{a}_j = a_j(b_{1j}, \dots, b_{nj})$, where $\sum_{i=1}^n a_{ij} \leq q_j n$.

2. Second, based on the announced allocation policies, buyers simultaneously submit bids to sellers. In his bid b_{ij} , buyer i commits to pay b_{ij} of his total budget to seller j .
3. Third, based on the allocation policies and the received bids, sellers allocate items to buyers.

The allocation policies define a complete information game among the buyers, where the action of each buyer i is a vector of bids b_{i1}, \dots, b_{in} with the restriction that $\sum_{j=1}^n b_{ij} \leq 1$. We term this game the *resulting buyers game*, and analyze its Nash equilibria (either pure or mixed). 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. While the budget of each player is divisible, and thus bids can be fractional, we assume that there is some finite precision to bids, so there is only a finite number of bids possible. This makes the resulting buyers game finite, and guarantees existence of some Nash equilibrium (either pure or mixed). Without loss of generality, sellers' policies may 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.

The process of choosing seller policies can be thought of as a game among sellers. However, the payoffs in this game are not well defined. They depend on the particular Nash equilibrium that is reached in the resulting buyers game, and in general this equilibrium will not be unique. Our main interest is in studying the effects of various seller allocation policies on the revenue of the sellers. Due to the equilibrium selection issue, this study has two aspects, one concerning the revenue in the best equilibrium for the particular seller, and the other concerning the revenue in the worse equilibrium.

3 Proportional Allocation

Recall the proportional allocation 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, $a_{i^*j}(b_{1j}, \dots, b_{nj}) = \frac{b_{i^*j}}{\sum_{i=1}^n b_{ij}} q_j n$. If a seller uses a fixed price per unit that "clears" her supply for the given buyers' bids, the resulting allocation will be exactly proportional. As Friedman argues, this seems to be a natural, fair, and intuitive way to allocate the items. It turns out that this policy has several attractive properties, for the buyers as well as for the sellers. Consider first the buyers.

Theorem 1. *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 is guaranteed to 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' game: 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, the following theorem (which is a restatement of Friedman's theorem, stated above as Theorem I) shows that this is the *unique* Nash equilibrium.

Theorem 2. [Friedman (1958)] *If all sellers use proportional allocation, proportional bidding is the unique Nash equilibrium in the resulting buyers' game. 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' game, 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 1), 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 1 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

Given these properties, one would indeed expect that if all sellers use proportional allocation, the revenue split will be proportional as well. The obvious question is whether a seller can extract a higher revenue by using a different strategy. To some extent, this is indeed possible, as the following example shows.

Suppose two sellers 1, 2 with $q_1 = q_2 = 0.5$, and an even number of buyers n . 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 units, 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 $n\epsilon = n(n-2)/2(n-1) < n/2$.

However, in this example, the revenue is larger than $\frac{n}{2} - 1$. As we next show, this is not by accident. When a seller uses proportional allocation, she *guarantees* almost her fair share:

Theorem 3. *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 nonnegative 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 nonnegative 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 theorem can be generalized to any set of sellers that use proportional allocation (this is Theorem VI from the Introduction). The proof is given in Appendix A.

Theorem 4. *Let P be a collection of several sellers $\{P_1, \dots, P_k\}$ and let $Q = \sum_{i=1}^k q_i$ 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 $Q - 1$.*

3.1 Rigid bids

In **proportional allocation among rigid bids**, the seller allocates the items equally among all bids whose value is exactly 1. Thus, a seller forces each buyer to either give all the budget to her, or not give any budget at all. It 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. A seller can deliberately make such additional costs particularly high, by this implementing “rigid” bids. Proportional allocation among rigid bids provides the same guarantee as proportional allocation, using the exact same proof of Theorem 3.

Theorem 5. *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 second part of the theorem is a direct consequence of rigidity, and is not true without it, as the example above has shown. We note that if a set of sellers P all use proportional allocation among rigid bids, the guarantee of Theorem 4 will not hold, as the examples in the sequel of this section show.

This theorem limits by how much the revenue might drop below the fair share, but not by how much it may increase beyond the fair share. We wish to point out that rigidity typically favors the 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. Moreover, the revenue of the large seller in a mixed Nash equilibrium may be significantly larger than her revenue in a pure Nash. For example, suppose 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 II shows that in any pure Nash the revenue of a small seller is at least 1, implying that the large seller does not get 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 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.

4 Two Extremely Asymmetric Sellers

In this section we focus on the case of two sellers, H and L , where seller H has $n - 1$ items and seller L has only one item. We may think of H as an established seller in the market, and of L as a newcomer. We ask whether the small seller can guarantee a revenue of 1, matching her proportional share of the supply. We observe that in this extreme case, the guarantees of Theorem 5 are not very decisive. If the small seller uses proportional allocation with rigid bids, and regardless of the policy of the large seller, in any pure Nash equilibrium of the resulting buyers' game (if there exists any) the revenue of the small seller will be at least 1. However the bound that the theorem provides regarding revenue in a mixed Nash equilibrium is completely meaningless (as the theorem guarantees the small seller a revenue of at least zero). A mixed Nash equilibrium in our case makes sense, for example, because it is the only symmetric equilibrium (pure Nash cannot be symmetric). This makes the special case an interesting case to study. Given Theorem 5, it is still possible that the large seller can keep the revenue of the small seller close to zero (or at least significantly lower than 1), creating a significant barrier for small sellers that attempt to enter the market.

4.1 Mixed Nash equilibria with rigid bids

We first show a separation between pure and mixed Nash equilibria: if both H and L use proportional pricing among rigid bids, then in every pure Nash L 's revenue is 1, while in every mixed Nash it is $O(1/n)$. Theorem III from the introduction extends to additional allocation policies of L as defined below.

Definition 1. *An allocation policy a_j is:*

- **anonymous** if for any bids b_{1j}, \dots, b_{nj} and any renaming of the buyers (a permutation $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$), $a_{ji}(b_{1j}, \dots, b_{nj}) = a_{j\pi(i)}(\tilde{b}_{1j}, \dots, \tilde{b}_{nj})$ 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.
- **monotone** if an increase in a buyer's bid (fixing the other bids) does not decrease his allocation.

- **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 6. *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 rigidity 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 1. $|X| \geq 2$.

Proof. Assume towards a contradiction that $X = \{i\}$ for some buyer i , and let $E[i|i \rightarrow H]$ 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 $E[i] = E[i|i \rightarrow H] = E[i|i \rightarrow L]$. Since the supply of H is smaller than n , anonymity implies $E[i] = E[i|i \rightarrow H] < 1$. However, efficiency (weak efficiency suffices) implies that i can guarantee 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 2. *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,

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

Therefore there must be a 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. \square

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 3. *Let L_i for $i = 1 \dots n-1$ 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 $\frac{2}{n} \geq Pr(L \geq 1) = 1 - \prod_i (1 - p_i) \geq 1 - e^{-\sum_i p_i} = 1 - e^{-E}$, where the second inequality follows since $1 - x \leq e^{-x}$. Thus $E \leq \ln(n) - \ln(n-2) < \frac{2}{n-2}$ (recall that $\ln(a) = \int_1^a \frac{1}{x} dx$), and the lemma follows. \square

Lemma 3 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 4. *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 3 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 $p \leq \frac{2(n-2)}{n(n-4)}$. \square

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. \square

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.

4.2 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 guarantees a revenue in the order of $n - \frac{1}{n}$ to the large seller (leaving very little revenue to the small seller), *regardless of the allocation policy of the small seller*.

Definition 2. *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 B as in proportional allocation among rigid bids, but with one exception: If the seller receives $n - 1$ bids of value B and one bid of value $B - \epsilon$ (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).

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. In Theorem 9 below we show that non-monotonicity is unavoidable if H wants to guarantee revenue above $n - 1$. (In our work we restrict attention to allocation policies that are functions the bids only, with no extra bits. If we had an extra bit, we could replace the bid $B - \epsilon$ by a bid “ B , 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.)

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 game, 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 7. *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 5. *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 guarantee 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 3. \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 5, 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 5, 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 6. 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 2 from Theorem 6, 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 guarantee 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 5 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 5 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 \square

4.3 Restricting H to be Monotone

Non-monotonicity is a crucial ingredient in order to ensure a low revenue for L. To see this, consider a completely different allocation policy for L:

Definition 3. *The anti-competitive allocation policy 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 8. *If L uses the anti-competitive allocation policy and H uses any arbitrary monotone allocation policy, there exists a pure Nash equilibrium in the resulting buyer's game in which L 's revenue is 1. (Showing necessity of non-monotonicity in Theorem 7).*

Proof. Consider first the situation in which all buyers place all their 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 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

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, one may wonder whether there exists a policy for L that will guarantee 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 impossible.

Theorem 9. *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 guarantee a revenue of 1 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 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' game: 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

Therefore, no monotone policy of H can guarantee (in all equilibria of the resulting buyers' game) a revenue strictly more than $n - 1$, but on the other hand no policy of L can guarantee revenue larger than $O(1/n)$. One may wonder how much revenue can H guarantee, under the restriction of monotonicity. Theorem 8 tells us that she cannot guarantee a revenue of strictly more than $n - 1$, but can she guarantee at least that? The next theorem shows that H can guarantee a revenue of almost $n - 1$, but not $n - 1$. In this sense, putting restrictions on H , as the "almost monopolist" seller, can slightly help the small seller. She will not be able to *guarantee* revenue proportional to her supply, but the best revenue that H can guarantee to herself will still leave an option to L to receive even slightly more than L 's fair share.

Theorem 10. *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' game 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' game, the revenue of H is at most $n - 1 - \frac{1}{n^2}$.*

The proof of this theorem is given in Appendix B.

5 Summary and Discussion

Using a stylized competition model, we have shown that extreme supply asymmetries yield an advantage to large sellers, that rigidity tends to work in favor of large sellers, that mixed Nash tends 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 for the large seller guarantees 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.

Our main motivating application is the market for display advertising, where the short run supply is indeed fixed, since in order to increase their supply, Internet companies need to solve challenging technical and marketing difficulties. This market is divided to many small sub-markets, based on various user characteristics such as physical location and surf history. For example, a display ad to be placed on a page that is being viewed by a user that showed recent interest in planning her vacation is a qualitatively different product than a display ad to be placed on a page that is being viewed by user that showed recent interest in buying a new car. These small sub-markets are many times extremely asymmetric, and give an interesting practical motivation to the second part of our inquiry.

As a final remark, we wish to mention an interesting connection of our 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. An assignment is 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. 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 explained in Section 3.1, Nash equilibria outcomes tend to favor larger sellers. Indeed, the D'Hondt method has a reputation of typically favoring the larger parties. (For example, the Israeli parliament, the Knesset, has 120 seats, and the large Likud party got 2 extra seats beyond its fair share in the 2003 elections.)

A Proof of Theorem 4

Let P be a collection of several sellers $\{P_1, \dots, P_k\}$ 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 nonnegative 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 nonnegative 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.

B Proof of Theorem 10

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' game is at least $n - 1 - \frac{2}{n-2}$.

We use the notation of Theorem 6 (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 2 identifies. The proof of Theorem 6 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. \square

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' game, 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. \square

References

- Armstrong, M. and J. Vickers (2001). "Competitive price discrimination." *RAND Journal of Economics*, pp. 579–605.
- 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.
- Bertrand, J. (1883). "Theorie mathematique de la richesse sociale." *Journal des Savants (1883)*, pp. 499-508.
- Besanko, D. and M. K. Perry (1993). "Equilibrium incentives for exclusive dealing in a differentiated products oligopoly." *The RAND Journal of Economics*, pp. 646–667.

- Cournot, A. (1838). “Recherches sur les principes mathematiques de la theorie des richesses.” *English translation, N. Bacon, trans., Researches into the Mathematical Principles of the Theory of Wealth, New York: Macmillan and Company, 1897.*
- Friedman, L. (1958). “Game-theory models in the allocation of advertising expenditures.” *Operations Research*, **6(5)**, 699 – 709.
- Johnson, N. and S. Kotz (1977). *Urn models and their application: an approach to modern discrete probability theory*. Wiley New York.
- Klemperer, P. (1995). “Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade.” *The Review of Economic Studies*, **62(4)**, 515–539.
- 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.
- Snyder, J. M. (1989). “Election goals and the allocation of campaign resources.” *Econometrica*, **57(3)**, 637 – 660.