

DISCUSSION PAPER SERIES

No. 10464

**PERSONALIZED PRICING AND
ADVERTISING: AN ASYMMETRIC
EQUILIBRIUM ANALYSIS**

Simon Anderson, Alicia Baik and Nathan Larson

INDUSTRIAL ORGANIZATION



Centre for Economic Policy Research

PERSONALIZED PRICING AND ADVERTISING: AN ASYMMETRIC EQUILIBRIUM ANALYSIS

Simon Anderson, Alicia Baik and Nathan Larson

Discussion Paper No. 10464

March 2015

Submitted 22 February 2015

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK

Tel: (44 20) 7183 8801

www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Simon Anderson, Alicia Baik and Nathan Larson

PERSONALIZED PRICING AND ADVERTISING: AN ASYMMETRIC EQUILIBRIUM ANALYSIS[†]

Abstract

We study personalized price competition with costly advertising among n quality-cost differentiated firms. Strategies involve mixing over both prices and whether to advertise. In equilibrium, only the top two firms advertise, earning "Bertrand-like" profits. Welfare losses initially rise then fall with the ad cost, with losses due to excessive advertising and sales by the "wrong" firm. When firms are symmetric, the symmetric equilibrium yields perverse comparative statics and is unstable. Our key results apply when demand is elastic, when ad costs are heterogeneous, and with noise in consumer tastes.

JEL Classification: D43 and L13

Keywords: Bertrand equilibrium, consumer targeting, mixed strategy equilibrium, price advertising and price dispersion

Simon Anderson sa9w@virginia.edu
University of Virginia and CEPR

Alicia Baik ab3ja@virginia.edu
University of Virginia

Nathan Larson nlarson@american.edu
American University

[†] We thank Michael Baye, Dan Kovenock, Qihong Liu, José Luis Moraga-Gonzalez, Regis Renault, Kostas Serfes, Yossi Spiegel, and our referees and advisory editor for their comments, as well as participants at the 6th Economics of Advertising and Marketing in Tel-Aviv (June 2013) and the CRETE conference in Naxos (July 2013). We also thank Ben Leyden for excellent research assistance. The first and second authors gratefully acknowledge support from the NSF and Bankard Fund for Political Economy respectively

1 Introduction

“Recent advances in information technology have...made possible the instantaneous delivery of customized pricing offers to individual consumers.” (Pricing with Precision and Impact, Boston Consulting Group 2002)

Mass marketing made possible through TV, newspapers, and billboards is increasingly evolving into individualized marketing. Firms previously limited to sending messages to heterogeneous groups of consumers (on network TV say) are now able to purchase information on *relatively* homogeneous sets of individuals from intermediary information brokers.¹ With finer levels of categorization, marketing precision is moving to the individual level. Comprehensive purchase history from various retailers can now be merged with demographic and web-site visit data to render very specific individual information on tastes, and firms can deliver individually-tailored price offers based on such information. This means that firms have the potential to compete at the level of the individual consumer. As technological capacity develops and the cost of personalized pricing decreases, the potential for individualized price competition will only increase.

Motivated by these observations, in this paper we develop a model of advertising and price competition in which the individual consumer is the basic unit of analysis. A consumer is characterized by her profile of valuations for the products offered by different firms. While each firm is aware that consumers with this particular taste profile exist, reaching them with personalized-price advertisements is costly. This cost might reflect payment to a data broker to deliver the name and contact information of a consumer with this taste profile. It also reflects the “postage and handling” costs of preparing and delivering an individualized offer. (One could think of these offers as going out by text message, email, or personalized coupons in the mail, as opposed to

¹We abstract away from the strategic role that such brokers may exercise in pricing information.

en masse marketing.) Meanwhile, as in Butters (1977), Grossman and Shapiro (1984), and Stahl (1994), a consumer does not know that a product is available unless she receives an advertised offer from the firm selling it. Among the offers she receives, she chooses the one that yields her the greatest consumer surplus.²

Given a potential consumer with a particular taste profile, firms simultaneously choose whether to advertise to her and if so, what individualized price to offer her. We call the joint price and advertising decisions the Personalized Pricing and Advertising Model, henceforth PPAM. Because sending an offer is costly, equilibria involve mixed strategies over both prices and the decision to advertise.

Though we do not focus on it, the model accommodates the possibility that firms simultaneously compete for a broad range of consumers with different tastes. However, under our assumptions about the precision of targeting, a firm’s strategy with respect to one consumer is completely separable from its strategy with respect to another consumer, so competition for an entire consumer population may be treated as a collection of independent instances of our model.³

Our main focus is on the asymmetric valuation case in which a consumer values some products more highly than others. We find that with n firms, the $n - 2$ “worst” ones sit out and do not advertise at all. The second “best” one advertises with positive probability below one, and earns zero expected profits; while the best one always advertises and earns a rent equaling its social surplus advantage (valuation minus cost superiority) over its closest rival. We also find that social efficiency falls

²The possibility of consumer search is introduced in a later section of Butters (1977) and is an integral part of the model of Robert and Stahl (1993). Shaffer and Zhang (1995, 2002) and Bester and Petrakis (1995, 1996) have considered targeting by location and have included the cost of sending offers to customers. They assume offers are coarse, such as a common discount to a heterogeneous consumer group.

³This separability is reasonable if (as we assume) a firm faces constant marginal costs to produce its product and to reach an additional consumer with a targeted ad. The latter is consistent with firms buying data about blocks of consumers (rather than individual by individual) as long as pricing is on a per-consumer basis and blocks of individuals with the same tastes are on offer.

then rises with advertisement costs, with losses due to wasteful advertisements and non-optimal purchases. These inefficiencies vanish when advertising costs go to zero or when they rise high enough to give the “best” firm a monopoly.

The pattern of our equilibrium results has some precedent in other asymmetric games with discontinuous pay-offs and (non-degenerate) mixed strategy equilibria. One point of resemblance is with the All-Pay-Auction treated in Hillman and Riley (1989) where different bidders have different values from winning. Baye, Kovenock, and de Vries (1996) present a broader set of symmetric and asymmetric combinations to this game by allowing ties in payoffs. A second prominent example is Varian’s (1980) Model of Sales, extended to allow for heterogeneous numbers of “loyal” consumers across firms by Narasimham (1988) for duopoly and by Kocas and Kiyak (2006) for oligopoly.⁴ In both games, there is a winner-take-all prize for the fiercest competitor, but competing incurs costs that “losers” do not recover. In the all-pay auction, the interpretation of the prize and costs is straightforward. In the Model of Sales, the “prize” is sales to the set of informed consumers, while the cost of competing for these consumers by offering a discounted price is the foregone profit on a firm’s loyal consumers.⁵ In both games, only the two players with the highest win value contend the prize, and all other players choose not to (by bidding zero or not discounting, respectively). While the results in these two models and ours share a “family resemblance,” the models themselves have significant differences such that no

⁴Baye, Kovenock, and de Vries (1992) find all the equilibria for the Model of Sales when all firms have the same number of loyal consumers (as in the original). In addition to the symmetric equilibrium analyzed by Varian (1980), there are also asymmetric ones. In these, at least two firms must be active: when there are only two firms in the market the symmetric equilibrium is the unique one, but not otherwise. Of particular interest for what follows in our paper is their result (Example 2, p.500) that with $n > 2$ there are equilibria with $k \geq 2$ firms symmetrically randomizing their prices and the others just charge the consumer reservation price.

⁵As clarified by Janssen and Moraga (2004), the Model of Sales is also at the heart of the literature on firm pricing and consumer search following Stahl (1989). In these search models, “informed” consumers (or “shoppers”) know all prices, while others face a search cost and in equilibrium stop at the first firm sampled, and hence play the role of the “loyal” consumers. Baye and Morgan (2001) successfully expand the basic MoS framework to a two-sided market setting.

pair is formally equivalent (even when reduced to their symmetric versions). Hence, our results cannot be derived from existing ones in the literature.

By taking advertising costs to zero, we can provide a fresh perspective on the long-standing selection problem of multiple equilibria in the classic model of Bertrand competition with asymmetric costs. (That is to say, homogeneous goods, no advertising, and different (constant) marginal costs across firms.) We select the equilibrium where the most efficient firm prices (with probability one) at the cost of its closest rival. Interestingly, the second-best firm makes an offer just often enough to keep the top firm from deviating to its monopoly price.

Analysis of equilibrium price distributions in the literature frequently assumes that firms are symmetric and focuses on a symmetric mixed strategy equilibrium. We argue that the symmetric equilibrium, when it exists in our model, may be seriously misleading. First, we show the striking comparative static prediction that when a consumer views products as homogeneous, the symmetric equilibrium has consumer surplus and welfare decreasing in the number of competing firms. This strong result stems from the indifference condition required to elicit advertising by all n firms. However, this equilibrium is not robust: with any heterogeneity in the consumer's valuations, the set of advertisers collapses down to two firms. Thus the perverse comparative static properties of the symmetric equilibrium may be seen as a symptom of this equilibrium's instability. However, with homogeneous products, the model also has many asymmetric equilibria. When we consider the limit case of heterogeneous firms as they approach homogeneity, then we select the particular asymmetric equilibrium in which only two firms are active (regardless of n) and one always advertises while the other does so with probability strictly less than one. In this limit equilibrium, the number of firms has no impact on welfare, and welfare is weakly higher (strictly for $n > 2$) than under the symmetric equilibrium.

The PPAM model of this paper can also be interpreted (by simply relabeling the

ad cost as an entry cost) as a Bertrand model of pricing and (simultaneous) entry. Previous work in this vein by Sharkey and Sibley (1993) considers symmetric firms and the symmetric equilibrium (we extend their model in section 6.1). Their main result is that an increase in the number of potential firms stochastically raises prices. Stahl (1994) moreover shows that seller entry can decrease social surplus, using a model of price advertising that effectively bridges the Butters analysis to the Sharkey and Sibley one, and for which the limit case when advertising costs are linear in reach corresponds to our model. By contrast, we would select an asymmetric equilibrium, in which case equilibrium price distributions are unchanged when there are more potential firms.

In the following section, we describe the basic set-up of the PPAM and discuss the two key strategic variables: individualized price distributions and advertising. In Section 3 we characterize the equilibrium in terms of the offered surpluses, rolling the decision to advertise into these surplus distributions, and highlight the Bertrand limit. In Section 4 we analyze two sources of competition-induced inefficiency: wasteful advertisements and non-optimal purchases. When firms are symmetric, the model has both symmetric and asymmetric equilibria. We evaluate both in Section 5 and show that the symmetric equilibrium has perverse comparative statics and is not robust to firm heterogeneity, indicating that the asymmetric outcome may be a more reasonable prediction. In Section 6 we argue that a number of our key results still apply if consumers have downward-sloping demand, if targeted advertising costs vary across firms, or if firms' information about consumer tastes is noisy. Section 7 concludes with a discussion of fruitful directions for future work.

2 Model

Each firm's problem will be separable across consumers, so we shall treat competition for an individual consumer as the basic unit of analysis. There are n single-product firms competing for the business of a single consumer who wishes to buy at most one unit from one of them. Each consumer considers the set of price offers she receives and purchases from the firm whose advertised offer gives her the greatest surplus. If she receives no ads or if none of the advertisements offer her weakly positive consumer surplus, she does not make a purchase. We assume that the consumer purchases whenever indifferent and randomizes if she is indifferent among several firms. Aside from this choice, the consumer has no strategic role in the game.

Let r_i be the consumer's individual reservation price for the product offered by firm i . (We assume that r_i is measured relative to some outside option which is normalized to zero.) Let p_i be the price offered to this consumer, and so $\sigma_i = r_i - p_i$ represents the consumer surplus offered by Firm i . This variable will allow us to conflate the advertising and pricing decisions into a single statistic, and we will show that in equilibrium all active firms have the same support for the consumer surplus they deliver.

As in the classic Butters (1977) model, a consumer is unaware of the availability of Firm i 's product unless she receives an advertisement with a price offer from Firm i . Advertising is costly: each firm decides whether to inform the consumer about an individualized price at cost A ; alternatively, a firm can choose not to advertise. In anticipation of mixed strategies, let a firm's cumulative price distribution conditional on advertising be $F_i(p)$. Thus, a strategy for Firm i is a pair $\{a_i, F_i\}$ where a_i is the probability that Firm i advertises. Firms choose these strategies simultaneously.

A firm that does not advertise earns zero profit, while if Firm i advertises price

p_i , its expected profit is given by

$$\pi_i(p_i) = (p_i - c_i) \Pr(i \text{ sells} \mid p_i) - A$$

where c_i is the marginal cost of product i . Firms seek to maximize expected profit. As this is a static model of complete information, the solution concept is simply Nash equilibrium.

The social surplus from a purchase at Firm i is the difference between the consumer's reservation value and the cost of production, $s_i = r_i - c_i$. Throughout most of the paper, we assume that different firms offer different social surpluses, with no ties. A discussion of equilibria when some or all of the products offer the same surplus is reserved for Section 5. Given this assumption, we choose to label firms in decreasing order of social surplus: $s_1 > s_2 > \dots > s_n$. Define the value advantage of Firm i over Firm j to be the difference $\Delta_{ij} = s_i - s_j$, which is strictly positive whenever $i < j$. If $A > s_2$ then there will either be no advertising and no sale in equilibrium (if we also have $A > s_1$), or else Firm 1 will hold a monopoly over the consumer. Thus, the interesting case, which we henceforth consider, is that $A \leq s_2$. Thus, at least two firms would want to advertise if they could earn monopoly profits by doing so.

3 Characterization of Equilibrium

3.1 Participation and Profits

We claim that any equilibrium has the features that the top firm advertises with probability one, the next best firm advertises with positive probability less than one, and no other firm advertises. Furthermore, the top firm earns expected profit equal to Δ_{12} , its surplus advantage over its closest rival, while the second-ranked firm earns 0. These two firms price in mixed strategies; their price supports are such that the consumer faces the same range of possible surplus offers at either firm. The highest

price offered by each firm leaves the consumer with zero surplus, while the lowest price ever offered by each firm leaves the consumer with $s_2 - A$, the full social surplus from a sale (net of the ad cost) at Firm 2. The top firm advertises its monopoly price with positive probability; that is to say, its price distribution has an atom at its upper bound, the consumer's reservation value. The distribution of the second-best firm has no atoms, and (with the exception above) both firms' prices follow Pareto distributions.

We proceed through a series of lemmas to establish these results. We show first that if any firm is advertising, then all higher ranked firms advertise as well. Next we show that at most one firm makes strictly positive profits in equilibrium. Third, the profits of all active firms are strictly ranked in the natural order. Fourth, using these results, we establish that at most the top two firms are active. Then (fifth and sixth), we show that the second-ranked firm does advertise, with probability less than one, while the top firm advertises with certainty. These results imply (Lemma 7) that equilibrium profits are Δ_{12} for the top firm and zero for all others. These facts make a full characterization of equilibrium strategies relatively straightforward; this characterization is given in Proposition 1. In referring to (candidate) equilibrium profits for a firm i , for brevity we will often simply write π_i (without reference to the particular price offer) rather than $\pi_i(p_i)$ since Firm i will typically be indifferent over a range of optimal prices.

Lemma 1 *In any equilibrium, if $a_i > 0$, and $j < i$, then $a_j > 0$.*

Proof. Suppose toward a contradiction that there is an equilibrium with $a_j = 0$, $a_i > 0$, and $j < i$. Let \hat{p}_i be the lowest price that Firm i ever advertises. (To be careful, we should have \hat{p}_i be the infimum of Firm i 's prices support, which may be degenerate.) Let $\hat{\pi}_i = \hat{q}_i(\hat{p}_i - c_i) - A$ be Firm i 's expected profit when offering \hat{p}_i and $\hat{q}_i > 0$ its probability of making a sale. (Again, for extra care, the limiting profit and

sale probability as $p_i \rightarrow \hat{p}_i$.) Note that $\hat{\pi}_i \geq 0$, otherwise Firm i would not be active. Let $\hat{p}_j = \hat{p}_i + (r_j - r_i)$ be the price from Firm j that would make the consumer equally well off as price \hat{p}_i at Firm i . If Firm j were to advertise price $\hat{p}_j - \varepsilon$, its probability of making a sale would be no less than \hat{q}_i , say $\hat{q}_j \geq \hat{q}_i$, and so it would earn profit

$$\begin{aligned} \hat{\pi}_j^\varepsilon &= (\hat{p}_j - c_j - \varepsilon) \hat{q}_j - A \\ &= (\hat{p}_i - c_i) \hat{q}_j + (\Delta_{ji} - \varepsilon) \hat{q}_j - A \\ &\geq \hat{\pi}_i + (\Delta_{ji} - \varepsilon) \hat{q}_j \end{aligned}$$

But then because $\Delta_{ji} > 0$, for ε small enough, $\hat{\pi}_j^\varepsilon > \hat{\pi}_i \geq 0$, so Firm j could earn strictly positive expected profit by deviating to advertising price $\hat{p}_j - \varepsilon$. ■

Lemma 2 *In equilibrium, at most one firm makes a strictly positive expected profit.*

Proof. Suppose to the contrary that there is an equilibrium with $\pi_i > 0$ and $\pi_j > 0$ for some firms i and j , with $j < i$. Then neither firm is indifferent between advertising and not advertising (as the latter earns zero profit), so both firms must be advertising with probability one. Let \hat{p}_i be the supremum over all prices ever offered by i , with \hat{p}_j the supremum over prices offered by j . We must have $\hat{p}_i = \hat{p}_j + (r_i - r_j)$, so that the consumer is indifferent between prices \hat{p}_i and \hat{p}_j . (Firm i will never advertise any $p_i > \hat{p}_j + (r_i - r_j)$, as this price would lose the sale for sure, earning profit $-A$, and similarly for Firm j .) Furthermore, Firm j 's strategy must place an atom at \hat{p}_j . (If not, then Firm i 's chance of winning the sale would tend to zero for p_i sufficiently close to \hat{p}_i , making it unprofitable to pay to advertise such prices.) Similarly, Firm i 's strategy must place an atom at \hat{p}_i . The firms' profit margins $\hat{p}_i - c_i$ and $\hat{p}_j - c_j$ at these upper bound prices must be strictly positive, since otherwise they could not cover the advertising cost and earn positive profits. But then because Firm i ties Firm j 's atom when offering \hat{p}_i , it could earn a strictly higher profit by deviating

to an undercutting price, contradicting the optimality of including \hat{p}_i in its support. (And similarly for Firm j .) ■

Lemma 3 *If Firm i , ($i > 1$), advertises in equilibrium, then $\pi_i < \pi_j$ for all $j < i$.*

Proof. The argument follows essentially the same lines as Lemma 1. Suppose \hat{p}_i is the lowest price that Firm i ever offers in equilibrium, with profit margin $\hat{p}_i - c_i$. Firm j earns a strictly larger profit margin, $\hat{p}_j - c_j = \hat{p}_i - c_i + \Delta_{ji}$ on the price $\hat{p}_j = \hat{p}_i + (r_j - r_i)$ that would leave the consumer equally well off as buying from i at \hat{p}_i . By offering slightly less than \hat{p}_j , Firm j could sell at least as often as Firm i does at price \hat{p}_i , thereby earning a profit strictly greater than π_i . Firm j 's equilibrium profit must be at least this good; thus $\pi_j > \pi_i$. ■

Lemma 4 *No firm other than the top two advertises in equilibrium: $a_i = 0$ for $i \geq 3$.*

Proof. If Firm $i \geq 3$ were to advertise, then Lemma 1 implies that Firms 1 and 2 would do so as well, and then Lemmas 2 and 3 imply that π_2 must be zero. But then, another application of Lemma 3 would imply that $\pi_i < 0$, so advertising with positive probability cannot actually be a best response for Firm i after all. ■

Lemma 5 *Firm 2 advertises with positive probability less than one: $a_2 \in (0, 1)$.*

Proof. If Firm 2 did not advertise at all, then Firm 1's best response would be to advertise its monopoly price, $p_1 = r_1$, with probability one, leaving the consumer with zero surplus. Firm 2 could offer the consumer the same surplus at price $p_2 = r_2$, with profit margin $r_2 - c_2 = s_2$. Thus, by advertising a price that slightly undercuts Firm 1 by ε , Firm 2 could win the sale with probability one and earn profit $s_2 - A - \varepsilon$. Because $A < s_2$ by assumption, this deviation would be profitable for sufficiently small ε ; thus $a_2 = 0$ is impossible. On the other hand, if $a_2 = 1$, then π_1 is strictly

positive by Lemma 3, and so Firm 1 must also advertise with probability one. But then by arguments similar to Lemma 2, Firm 2's profit margin at the highest price it ever offers must be weakly negative. But then, Firm 2 does not cover its ad cost, and so $\pi_2 < 0$, contradicting the optimality of advertising with probability one. ■

Lemma 6 *Firm 1 advertises with probability one. That is, $a_1 = 1$.*

Proof. Lemmas 3 and 5 imply that $\pi_1 > 0$. But this means that Firm 1 cannot be indifferent to not advertising (and thereby earning zero profit), so $a_1 = 1$. ■

Lemma 7 *Equilibrium profits are $\pi_1 = \Delta_{12}$ for Firm 1 and $\pi_i = 0$ for all $i > 1$.*

Proof. As noted just above, Lemmas 3 and 5 imply that $\pi_1 > 0$. The fact that $\pi_i = 0$ for all $i > 1$ follows from Lemma 2. To pin down π_1 , let \underline{p}_1 and \underline{p}_2 be the lower bounds on the supports of the price distributions used by Firms 1 and 2 respectively. These lower bounds must give the consumer equal surplus – that is, $\underline{p}_1 = \underline{p}_2 + (r_1 - r_2)$ – as if they did not, the firm offering the consumer the better deal could raise its price slightly without affecting its chance of making the sale. Next, we claim that $\underline{p}_2 - c_2 = A$. Clearly we cannot have $\underline{p}_2 - c_2 < A$, as in this case Firm 2 could not recover its ad cost by offering \underline{p}_2 . On the other hand, if $\underline{p}_2 - c_2 > A$, then either (i) Firm 1 has no atom at price \underline{p}_1 , in which case Firm 2 wins for sure by advertising \underline{p}_2 , thereby making strictly positive profit $\underline{p}_2 - c_2 > A$, or (ii) Firm 1 has an atom at \underline{p}_1 , in which case Firm 2 could win for sure and make a strictly positive profit by deviating slightly below \underline{p}_2 . As both cases are incompatible with zero profit for Firm 2 in equilibrium, we have $\underline{p}_2 - c_2 = A$. But this implies that $\underline{p}_1 - c_1 = A + \Delta_{12}$. Furthermore, Firm 2 cannot have an atom at \underline{p}_2 either (or else Firm 1 could do strictly better by deviating below \underline{p}_1), so Firm 1 wins with probability one when it offers \underline{p}_1 , earning profit $\underline{p}_1 - c_1 - A = \Delta_{12}$. Since any other price in the support of Firm 1's price distribution must do equally well, we have $\pi_1 = \Delta_{12}$. ■

3.2 Mixed Strategy Offer Distributions

Notice that when Firm i advertises a price p_i , this is equivalent to offering the consumer a surplus of $\sigma_i = r_i - p_i$, so firms' strategies may be characterized either in terms of the distributions of prices they demand or the distributions of surpluses they offer. It is convenient to roll the decision to advertise into these surplus distributions by regarding a decision not to advertise as an offer of zero surplus. That is, let $G_i(\sigma) = \Pr(\sigma_i \leq \sigma)$ be the probability that the consumer's offer from Firm i is no better than σ , with the event that Firm i does not advertise recorded as $\sigma_i = 0$. Given the probability a_i that Firm i advertises, its price distribution conditional on placing an ad may be recovered from the identity

$$G_i(\sigma) = 1 - a_i + a_i \Pr(p_i \geq r_i - \sigma \mid \text{Firm } i \text{ advertises})$$

That is, an offer weakly worse than σ means that Firm i either did not advertise, or advertised a price weakly higher than $r_i - \sigma$.

Proposition 1 *In equilibrium, the top firm advertises with probability one and makes expected profit equal to Δ_{12} , its surplus advantage over the second-ranked firm. The second-ranked firm advertises with probability $a_2 = \frac{s_2 - A}{s_1} \in (0, 1)$ and earns zero expected profit. No other firm advertises. The surplus distributions offered to the consumer by Firms 1 and 2 are $G_1(\sigma) = \frac{A}{s_2 - \sigma}$ and $G_2(\sigma) = \frac{A + \Delta_{12}}{s_1 - \sigma}$ respectively, with common support $\sigma \in [0, s_2 - A]$.*

Proof. Lemmas 1 through 7 establish that $a_1 = 1$, $a_2 \in (0, 1)$, and $a_3, \dots, a_n = 0$. Let $\bar{\sigma}_i$ and $\underline{\sigma}_i$ be the upper and lower supports on the surplus distribution offered by Firm i , $i \in \{1, 2\}$. Since Firm 2 does not always advertise, we have $\underline{\sigma}_2 = 0$. By standard arguments, these supports are common (with $\bar{\sigma}_1 = \bar{\sigma}_2 = \bar{\sigma}$ and $\underline{\sigma}_1 = \underline{\sigma}_2 = 0$), have no gaps, and have no atoms on $(0, \bar{\sigma}]$. If $\bar{\sigma}_1 > \bar{\sigma}_2$, then Firm 1 could be strictly less

generous than $\bar{\sigma}_1$ and still sell with probability one, and *vice versa*, so $\bar{\sigma}_1 = \bar{\sigma}_2$. If $0 = \underline{\sigma}_2 < \underline{\sigma}_1$, then (i) if Firm 2 makes any offers in the interval $(0, \underline{\sigma}_1)$, they never succeed and thus lose money, or (ii) if Firm 2 makes no offers in $(0, \underline{\sigma}_1)$, then Firm 1 could make a less generous offer than $\underline{\sigma}_1$, sell no less often, and make more money. So $\underline{\sigma}_1 = \underline{\sigma}_2 = 0$. The argument for gaps is completely standard. For atoms, first note that $\bar{\sigma} \leq s_2 - A$ (as Firm 2 would lose money by advertising more generous offers). Thus the gross profit margin (before ad costs) on any offer is at least $s_2 - \bar{\sigma} \geq A > 0$ for Firm 2, and greater for Firm 1. Then standard undercutting arguments apply – by shifting its offer from slightly below to slightly above a rival’s atom, a firm would enjoy a jump in its sales at (essentially) the same, strictly positive gross profit margin. Finally, note that Firm 2 sells with probability one when advertising $\sigma_2 = \bar{\sigma}$, thus earning net profit $(s_2 - \bar{\sigma}) - A$. But $\pi_2 = 0$ by Lemma 7, so $\bar{\sigma} = s_2 - A$.)

Note we have not ruled out atoms at $\sigma = 0$. Firm 2 must have such an atom, because it does not always advertise, while Firm 1 will turn out to have such an atom because it will advertise $p_1 = r_1$ with positive probability. We must be a bit careful in handling these, as advertised offers of $\sigma = 0$ incur ad cost A , while unadvertised offers do not.

When Firm 2 offers surplus $\sigma_2 \in (0, s_2 - A]$, it sells with probability $G_1(\sigma_2)$ and earns profit $(s_2 - \sigma_2)G_1(\sigma_2) - A$. Then, as $\pi_2 = 0$ and Firm 2 must be indifferent over its support, we have $G_1(\sigma) = \frac{A}{s_2 - \sigma}$ for $\sigma \in (0, s_2 - A]$. Similarly, when Firm 1 offers $\sigma_1 \in (0, s_2 - A]$, it sells with probability $G_2(\sigma_1)$ and earns profit $(s_1 - \sigma_1)G_2(\sigma_1) - A = \pi_1 = \Delta_{12}$; thus we have $G_2(\sigma) = \frac{A + \Delta_{12}}{s_1 - \sigma}$ for $\sigma \in (0, s_2 - A]$. Notice that Firm 1 advertises $\sigma_1 = 0$ with positive probability $G_1(0) = \frac{A}{s_2}$. Given this, Firm 2 cannot find it optimal to *advertise* $\sigma_2 = 0$ itself – doing so would tie Firm 1’s atom, while undercutting with a slightly better offer would win twice as often. Thus any probability mass on $\sigma_2 = 0$ reflects Firm 2’s failure to advertise. Since $G_2(0) = \frac{A + \Delta_{12}}{s_1}$, we have $1 - a_2 = \frac{A + \Delta_{12}}{s_1}$ and so $a_2 = \frac{s_2 - A}{s_1}$. ■

The short-cut intuition for some of the key values in the Proposition is as follows. First, because Firm 2 earns zero profit in equilibrium then its lowest price (at which it wins for sure) is A above its unit production cost. This is analogous to the lowest price in Butters' (1977) model: anything lower would not cover the cost of sending the ad. Thus the highest consumer surplus value of $s_2 - A$ is attained when buying at that price. When Firm 1 matches this surplus level, its corresponding price is $r_1 - (s_2 - A)$ and it wins for sure. Subtracting its unit production cost, then 1's gross revenue is $\Delta_{12} + A$. Subtracting from this amount the cost A of sending the ad gives 1's equilibrium profit level as the value of its advantage, Δ_{12} . Firm 1 gets the same profit when it delivers zero surplus to the consumer, pricing at r_1 and earning a gross profit of s_1 when it wins. Firm 1 only wins at this highest price when its rival does not advertise, which happens with probability $(1 - a_2)$, and costs A . This profit indifference property $s_1(1 - a_2) - A = \Delta_{12}$ ties down the *rival's* advertising probability as $a_2 = 1 - \frac{\Delta_{12} + A}{s_1} = \frac{s_2 - A}{s_1}$. Notice here the inherent asymmetry between ad levels, which remain distinctly different even as social surpluses get arbitrarily close. Even for small social surplus differences the dominant firm always advertises while the weaker one rarely contests it if A is a significant fraction of s_2 . We return to this asymmetry below.

Now consider the probability $G_1(0)$ that Firm 1 charges its top price, r_1 , delivering zero consumer surplus. In the mixed strategy equilibrium, this probability must make Firm 2 indifferent between advertising and not. If Firm 2 sets its price just below r_2 , it wins the consumer with probability $G_1(0)$ for a gross profit of s_2 at a cost of A . Thus $G_1(0)s_2 = A$, and so the probability that Firm 1 sets the top price is thus A/s_2 . Notice that this probability goes to 1 as A rises to s_2 , so that Firm 1 sets its monopoly price more frequently as the cost of advertising rises. Indeed, for $A \geq s_2$ (but $A < s_1$) Firm 1 is an uncontested monopolist and always prices at r_1 .

3.3 Bertrand limit as $A \rightarrow 0$

In the usual version of asymmetric Bertrand competition when the consumer is notified about firms' price offers automatically and costlessly, the standard pure strategy equilibrium has the second-ranked firm pricing at cost, $p_2 = c_2$, while the top firm offers $p_1 = c_2 + (r_1 - r_2)$, the highest price at which its product is weakly preferred over Firm 2's, sells with probability one, and earns profit Δ_{12} . However, somewhat *ad hoc* arguments must be made to dispense with technical complications before reaching this intuitive conclusion.⁶ Our model delivers this outcome naturally, as the unique limit of equilibria as the advertising cost vanishes. Corollary 1 follows directly from Proposition 1 by taking $A \rightarrow 0$.

Corollary 1 *In the limit as $A \rightarrow 0$, the top firm advertises with probability one and makes expected profit equal to the social surplus difference Δ_{12} . The second-ranked firm advertises with probability s_2/s_1 and earns zero expected profit. No other firm advertises.*

Furthermore, it is straightforward to verify that consumer surplus tends to s_2 – the full surplus from the second-best option – and total social surplus tends toward its first-best level s_1 (implying the consumer buys from Firm 1 with probability tending to one). These points match the standard Bertrand result as well. While the firms never stop mixing over the full support of prices, as $A \rightarrow 0$ the weight Firm 1 places on its most competitive offer goes to one. In contrast, as $A \rightarrow 0$, Firm 2 advertises

⁶Tirole (1988, p.234) notes two problems: the open-set problem of ϵ -undercutting, and the possibility of an equilibrium price between the two cheaper firms' costs. The former problem is typically solved by invoking an efficient allocation rule to allocate customers to the socially preferable firm when faced with price ties (see, e.g., Lederer and Hurter, 1986). The latter problem can be resolved by eliminating weakly dominant strategies (e.g. Tirole, 1988, p.234, fn 37.), although such recourse would also eliminate the second-best firm pricing at its cost. An alternative solution is to consider a fine grid of prices and again eliminate weakly dominated strategies (see Mas-Collel, Whinston, and Green, 1995, p.430).

just often enough at every surplus level $\sigma \in [0, s_2]$ so that Firm 1 is not tempted to make any offer less generous than its most competitive one.

3.4 Mixed strategy Prices

We can now determine the price distributions for the top two firms, $F_1(p)$ and $F_2(p)$ respectively, conditional on their advertising. These price distributions follow directly from the identity linking prices, advertising, and surplus using $p = r_1 - \sigma$. For Firm 1, we have $G_1(r_1 - p) = \Pr(p_1 \geq p) = 1 - F_1(p) + \Pr(p_1 = p)$. This yields:

$$F_1(p) = \begin{cases} 1 - \frac{A}{(p-c_1)-\Delta_{12}} & \text{if } p \in [c_1 + \Delta_{12} + A, r_1) \\ 1 & \text{if } p \geq r_1 \end{cases}$$

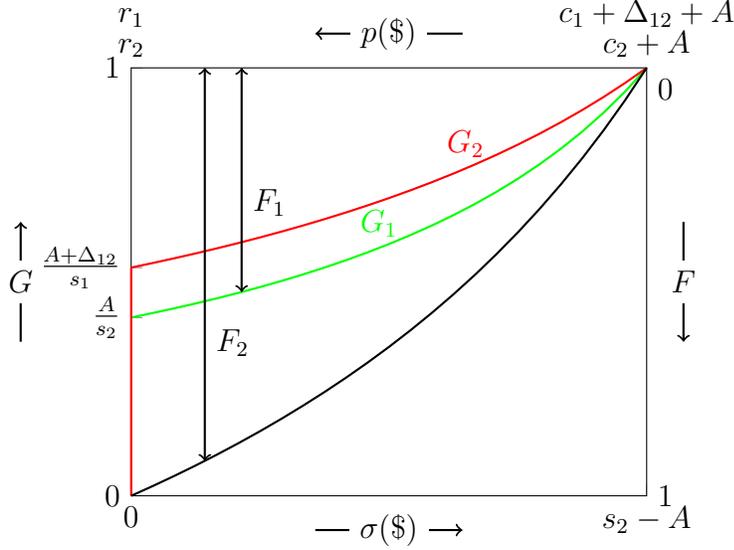
where the atom at zero surplus translates into an atom at the consumer's reservation price because Firm 1 is advertising with probability one. Because Firm 2 does not advertise with positive probability, we have $G_2(r_2 - p) = 1 - a_2 + a_2 \Pr(p_2 \geq p)$, or (using $G_2(0) = 1 - a_2$), $\Pr(p_2 < p) = \frac{1-G_2(r_2-p)}{1-G_2(0)}$. As this distribution is atomless, we may substitute a weak inequality and plug in to get:

$$F_2(p) = \frac{s_1}{s_2 - A} \left(1 - \frac{\Delta_{12} + A}{\Delta_{12} + (p - c_2)} \right) \quad \text{if } p \in [c_2 + A, r_2)$$

As is often the case with price competition in mixed strategies, both firms' price distributions are in the generalized Pareto family with tail exponent 1. Empirical evidence suggests that pricing strategies generally follow a Pareto distribution. A number of well-known papers derive Pareto distributions from their mixed strategy analysis including Butters (1977), Varian (1980), Baye and Morgan (2001), and Stahl (1989).

The construction of the price distributions, $F_1(p)$ and $F_2(p)$ from the surplus distributions $G_1(\sigma)$ and $G_2(\sigma)$ is shown in Figure 1 below. Henceforth we will return to using the surplus distributions in our analysis because of the convenient structure.

Figure 1: Equilibrium Price and Offer Distributions



The lower horizontal axis measures σ for distributions G_1 and G_2 . $G_1(0)$ is the mass point for a zero surplus offer by Firm 1; this equals mass point for Firm 1 setting price r_1 , $1 - F_1(r_1)$. The upper horizontal axis measures price from the right for price distributions, F_1 and F_2 .

4 Consumer Surplus, Social Surplus, Advertising Costs

4.1 Consumer Surplus

A number of facts about equilibrium consumer welfare emerge rather directly from inspection of the surplus distributions G_1 and G_2 . To begin with, we can determine which of the two active firms tends to give the consumer better offers.

Proposition 2 *The consumer's surplus offers from the top firm first order stochastically dominate her surplus offers from the second-ranked firm. (That is, $G_1(\sigma) \succsim_{FOSD} G_2(\sigma)$).*

Proof. Noting that $G_2(\sigma) = \frac{A+\Delta_{12}}{(s_2-\sigma)+\Delta_{12}}$ makes it clear that $G_2(\sigma)$ can be written as a convex combination of $G_1(\sigma) = \frac{A}{s_2-\sigma}$ and $\frac{\Delta_{12}}{\Delta_{12}} = 1$; thus $G_2(\sigma) \geq G_1(\sigma)$ (with

$G_2(\sigma) < G_1(\sigma)$ on the interior of the support: $\sigma < s_2 - A$. ■

This contrasts with the familiar results for asymmetric Bertrand competition when firms' price offers are announced automatically and costlessly. In that case, Firm 2 prices at its cost, and Firm 1 prices at its cost plus Δ_{12} , the markup that makes the consumer indifferent between offers, and the consumer receives surplus s_2 from either firm. Intuitively, because of its higher profit margin, Firm 1 has a greater incentive than does Firm 2 to sweeten its surplus offer to be sure it wins. This logic applies with or without costly advertising; however without advertising, the amount by which Firm 1 needs to sweeten its offer relative to Firm 2 shrinks to zero since it can undercut Firm 2's pure strategy arbitrarily closely.

Realized consumer surplus is just $\sigma_{\max} = \max(\sigma_1, \sigma_2)$, since the consumer picks the best offer she gets. The cumulative distribution function for consumer surplus is then

$$\begin{aligned} G_{\max}(\sigma) &= G_1(\sigma) G_2(\sigma) \\ &= \left(\frac{A}{s_2 - \sigma} \right) \left(\frac{A + \Delta_{12}}{s_1 - \sigma} \right) \\ &= \left(\frac{A}{s_2 - \sigma} \right) \left(\frac{A + s_1 - s_2}{s_1 - \sigma} \right) \end{aligned} \tag{1}$$

Using $G_{\max}(\sigma)$, we determine the impact on consumer surplus from changes in the competitive environment. Several of the highlights are summarized below.

Proposition 3 *The distribution of realized consumer surplus is increasing (in the sense of first order stochastic dominance) in s_2 . It is decreasing in s_1 and in the ad cost A .*

Proof. These properties follow directly from (1). ■

Corollary 2 *Expected consumer surplus increases in s_2 and decreases in s_1 and A .*

It might be tempting to argue that consumer surplus must rise as A declines because a lower barrier to reaching the consumer must surely make the market more competitive. This is not necessarily wrong in the end, but it misses some subtlety. Because the firms' profits do not vary with A , consumer surplus moves in lockstep with total social surplus as A declines. There are two effects on social surplus to consider. First, allocative efficiency – namely, the chance that Firm 1 (the firm with the highest social surplus) gets the sale – generally appears to be U-shaped in A (as Firm 1 wins with probability tending to one for $A \approx s_2$ and $A \approx 0$, but loses with positive probability in between).⁷ Thus when A is large, a reduction in ad costs increases the chance of a sale by the wrong firm, tending to reduce consumer surplus. However, the second effect is that the total cost of advertising, which ends up being borne by the consumer, unambiguously declines with a decline in A . In the end, the second effect dominates.⁸

The fact that a better second-ranked option helps the consumer to carve out more surplus is natural and would hold in the textbook Bertrand setting as well. It is less obvious that an improvement in her best option s_1 should hurt the consumer – after all it would have no effect at all in the textbook Bertrand setting. Here, as one can see from the $G_2(\sigma)$ term within $G_{\max}(\sigma)$ in (1), a stronger best choice s_1 induces the second-ranked firm to back off and compete less vigorously, thereby hurting the consumer.

Expected consumer surplus may be computed directly from $G_{\max}(\sigma)$ in (1):⁹

$$CS = E_{G_{\max}}(\sigma) = s_2 - A \left(1 + \frac{A + \Delta_{12}}{\Delta_{12}} \ln \left(\frac{s_2}{s_1} \frac{A + \Delta_{12}}{A} \right) \right)$$

⁷These points follow from $a_2 \rightarrow 0$ as $A \rightarrow s_2$ and social surplus tending to its first-best level s_1 (implying an efficient allocation with probability one) as $A \rightarrow 0$.

⁸This is not completely trivial, since total advertising volume $a_1 + a_2 = 1 + \frac{s_2 - A}{s_1}$ is decreasing in A . However total ad cost $A(a_1 + a_2)$ is increasing in A over $A \in [0, s_2]$.

⁹An interesting alternative form emerges by recasting the logarithmic expression in terms of profit margins. Let $\bar{\mu}_1 = \bar{p}_1 - c_1$ and $\underline{\mu}_1 = \underline{p}_1 - c_1$ be Firm 1's largest and smallest gross profit margins in

Defining

$$L(A, s_1, s_2) = A \left(1 + \frac{A + \Delta_{12}}{\Delta_{12}} \ln \left(\frac{s_2 A + \Delta_{12}}{s_1 A} \right) \right)$$

we have $CS = s_2 - L(A, s_1, s_2)$ – that is, the consumer earns her asymmetric Bertrand payoff s_2 , minus a loss term that is increasing in the ad cost A . (We know that $L(A, s_1, s_2) \geq A$ because σ_{\max} has upper support $s_2 - A$.) Furthermore, one can show that $\lim_{A \rightarrow 0} L(A, s_1, s_2) = 0$, so as advertising costs vanish, the consumer tends toward her asymmetric Bertrand payoff.

4.2 Advertising and social surplus

Denote expected social surplus as $SS = CS + \pi_1 + \pi_2$. At equilibrium profits,

$$SS = s_1 - L(A, s_1, s_2)$$

First-best social surplus in the absence of advertising costs would just be $SS_{eff} = s_1$, the surplus from allocating the consumer to Firm 1. Thus, $L(A, s_1, s_2)$ also may be interpreted as the shortfall of equilibrium social surplus below its first-best level. If the consumer is unaware of an unadvertised product, the reasonable benchmark is the second-best (constrained-efficient) social surplus that takes the necessity of advertising into account. This is $SS_{2bo} = s_1 - A$, where now the cost of apprising the consumer of her first-ranked option is included. Then we may write

$$\begin{aligned} SS &= SS_{2bo} - (L(A, s_1, s_2) - A) \\ &= SS_{2bo} - A \left(\frac{A + \Delta_{12}}{\Delta_{12}} \right) \ln \left(\frac{s_2 A + \Delta_{12}}{s_1 A} \right) \end{aligned}$$

equilibrium (with $\bar{p}_1 = r_1$ and $\underline{p}_1 = c_1 + \Delta_{12} + A$), and define $\bar{\mu}_2$ and $\underline{\mu}_2$ similarly for Firm 2. Then,

$$CS = \bar{\mu}_2 - \underline{\mu}_2 \left(1 + \frac{\bar{\mu}_1}{\bar{\mu}_1 - \underline{\mu}_2} \ln \left(\frac{\bar{\mu}_2 / \underline{\mu}_2}{\bar{\mu}_1 / \underline{\mu}_1} \right) \right).$$

This expression delivers a simple statistic with which to compute consumer surplus under personalized price competition using only the highest and lowest profit margins for Firms 1 and 2.

Shortfalls below these two benchmarks arise from two sources: excessive ad costs and the wrong firm (Firm 2) winning the sale. These can be easily decomposed. Given advertising $a_1 = 1$ and $a_2 = \frac{s_2 - A}{s_1}$, the total social cost of advertising is $A \left(1 + \frac{s_2 - A}{s_1}\right)$. Of this, Firm 1's share A is necessary, in the constrained-efficient sense, while Firm 2's share $A \frac{s_2 - A}{s_1}$ is wasteful. Thus, the "Avoidable inefficiency," or $SS_{2bo} - SS = L(A, s_1, s_2) - A$, may then be broken down as follows. The social cost of misallocation is $L(A, s_1, s_2) - A \left(1 + \frac{s_2 - A}{s_1}\right)$, or

$$\begin{aligned} \text{Cost of wasteful advertising} &= Aa_2 = A \frac{s_2 - A}{s_1} \\ \text{Social cost of misallocation} &= \Delta_{12} \Pr(\text{Firm 2 wins}) \\ &= \frac{A(A + \Delta_{12})}{\Delta_{12}} \ln \left(\frac{s_2}{s_1} \frac{A + \Delta_{12}}{A} \right) - A \frac{s_2 - A}{s_1} \end{aligned}$$

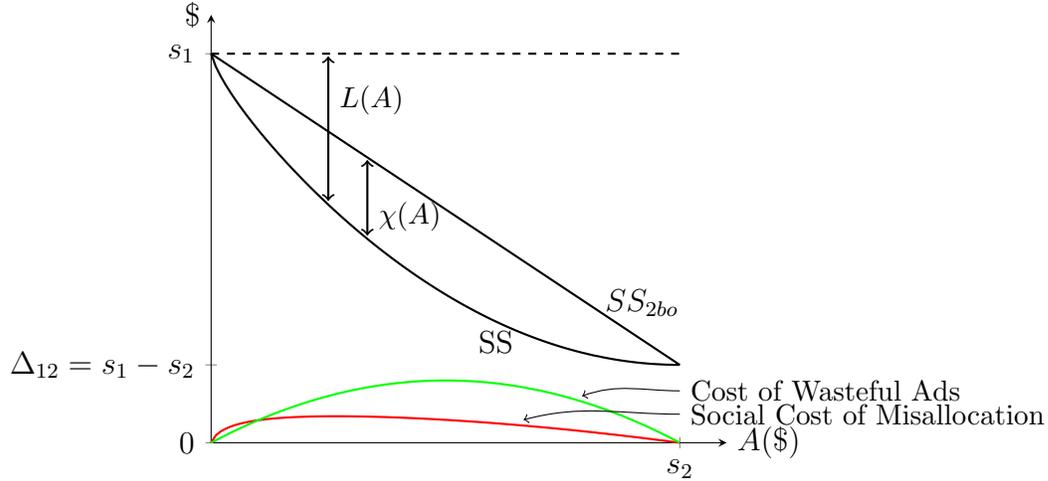
We have already established that $L(A, s_1, s_2)$ is increasing in A , so the gap between equilibrium and first-best social surplus shrinks as ad costs decline. However, this decline is driven in large part by a mechanical effect: the declining cost of Firm 1's certain advertising, $Aa_1 = A$. If we view this cost as unavoidable, as the second-best benchmark does, then the relationship of equilibrium efficiency to ad costs is more nuanced.

Proposition 4 *Avoidable inefficiency $SS_{2bo} - SS$ vanishes at $A = 0$ and $A = s_2$. Furthermore, it is positive, strictly concave, and single-peaked in A over $A \in (0, s_2)$.*

Proof. Let $\chi(A) = SS_{2bo} - SS = \frac{A(A + \Delta_{12})}{\Delta_{12}} \ln \left(\frac{s_2}{s_1} \frac{A + \Delta_{12}}{A} \right)$. It is immediate that $\chi(s_2) = 0$, and $\lim_{A \rightarrow 0} \chi(A) = 0$ follows by taking the limit. Differentiation yields $\chi'(A) = \frac{1}{\Delta_{12}} (2A + \Delta_{12}) \ln \left(\frac{s_2}{s_1} \frac{A + \Delta_{12}}{A} \right) - 1$, so $\chi'(s_2) = -1$ and $\lim_{A \rightarrow 0} \chi'(A) = \infty$. Thus $\chi(A)$ is strictly positive near the endpoints of $(0, s_2)$. Differentiating again yields $\chi''(A) = \frac{2}{\Delta_{12}} \ln \left(\frac{s_2}{s_1} \right) + \frac{1}{\Delta_{12}} \left(2 \ln \frac{A + \Delta_{12}}{A} + \frac{A}{A + \Delta_{12}} - \frac{A + \Delta_{12}}{A} \right)$. The first term is strictly negative because $s_2 < s_1$, so for concavity it will suffice to show the second term is

negative as well. Write the second term as $\frac{1}{\Delta_{12}}\xi\left(\frac{A+\Delta_{12}}{A}\right)$ for $\xi(z) = 2\ln z + \frac{1}{z} - z$. We claim that $\xi(z) < 0$ for all $z > 1$ (and so $\xi\left(\frac{A+\Delta_{12}}{A}\right)$ because $\frac{A+\Delta_{12}}{A} > 1$). To show this claim, observe that $\xi(1) = 0$ and $\xi'(z) = -\left(1 - \frac{1}{z}\right)^2$. ■

Figure 2: Social Surplus and Avoidable Inefficiency



$SS_{2bo} = s_1 - A$ is the second-best social surplus. SS is the equilibrium social surplus and $\chi(A)$ is the avoidable inefficiency, which vanishes at $A = 0$ and $A = s_2$; $\chi(A)$ can be decomposed into the cost of wasteful ads and the social cost of misallocation. $L(A) = \chi(A) + A$ is the difference between equilibrium social surplus and the first-best social surplus.

The fact that the equilibrium is second-best optimal at $A = s_2$ is straightforward, because for $A \geq s_2$ the second-ranked firm cannot afford to enter the market and so the first-ranked firm has a monopoly. Social surplus increases as advertising costs fall below s_2 , permitting the second-ranked firm to enter, though the effect is negligible at first. Social losses due to socially excessive advertising and sales by the wrong firm rise as ad costs decline as shown on Figure 2. In this sense, lower ad costs initially open the door to the second-ranked firm, giving it a chance to win sales (which it should not do, from the standpoint of efficiency), thereby creating an incentive for it to advertise (which it also should not do). Total advertising volume continues to rise as A falls, but eventually the cost of excessive advertising begins to decline as its ad

cost tends to zero. Furthermore, as A falls, the chance of Firm 2 winning rises to a peak before declining to zero.

5 Equilibrium when firms are symmetric

This section considers a symmetric version of our model in which all firms are identical.¹⁰ There is a symmetric equilibrium in which all of the firms advertise with positive probability less than one. However, there is also a natural asymmetric equilibrium obtained by taking limits in Proposition 1 as differences in surplus vanish; we call this the limiting asymmetric equilibrium.¹¹ We characterize both equilibria below and then discuss reasons why the limiting asymmetric equilibrium may be more appealing as a prediction of behavior.

5.1 Symmetric Equilibrium

Suppose that each of the n firms has potential surplus s_1 . In a symmetric equilibrium each firm must earn zero profit. (Strictly positive profits would imply that all firms always advertise, but this is impossible because under the Bertrand competition that would ensue none of the firms would even cover its ad cost.) The best surplus offer made must be $\bar{\sigma} = s_1 - A$ which wins with probability one and earns zero profit. (If $\bar{\sigma}$ were lower, then any firm could earn strictly positive profits by overcutting it.) The worst offer (which wins only if no other firm advertises) must still be the monopoly offer $\underline{\sigma} = 0$; standard arguments rule out atoms or gaps over the support $(0, s_1 - A]$. If $G_i(\sigma)$ is the common mixed strategy distribution over surplus offers, then Firm 1's probability of winning the sale with offer σ is $G_{-1}(\sigma) = G_i(\sigma)^{n-1}$

¹⁰Partially symmetric cases (where only some firms are identical) are treated in the appendix.

¹¹More properly, there is a collection of such equilibria, differing only in the identities of the firms playing the roles of Firm 1 and Firm 2. There are also additional equilibria in which an arbitrary subset of $\tilde{n} < n$ of the firms play a version of the symmetric equilibrium (with $\tilde{n} \geq 2$ replacing n), while the $n - \tilde{n}$ others sit out (never advertise). For details, see the appendix.

(where $G_{-1}(\sigma)$ is the distribution of the best rival offer faced by Firm 1). Since the gross profit on an offer σ is $s_1 - \sigma$, Firm 1's zero-profit indifference condition then becomes $(s_1 - \sigma) G_{-1}(\sigma) - A = 0$ (and similarly for each other firm). This pins down a firm's symmetric equilibrium surplus offer distribution as $G_i(\sigma) = \left(\frac{A}{s_1 - \sigma}\right)^{\frac{1}{n-1}}$, for $i \in \{1, \dots, n\}$. Under these symmetric strategies, the probability mass $G_i(0)$ must reflect a failure to advertise, not an atom on zero-surplus offers (as such an atom would be profitably undercut). Hence we have the following:

Proposition 5 *In the symmetric equilibrium with n firms each delivering potential surplus $s_1 > A$, expected profit for each firm is zero, and the equilibrium offer distribution is $G_i(\sigma) = \left(\frac{A}{s_1 - \sigma}\right)^{\frac{1}{n-1}}$ with support on $[0, s_1 - A]$, $i \in \{1, \dots, n\}$. Each firm refrains from advertising with probability $G_i(0) = \left(\frac{A}{s_1}\right)^{\frac{1}{n-1}}$ and the consumer's best offer has distribution $G_{\max}(\sigma) = \left(\frac{A}{s_1 - \sigma}\right)^{\frac{n}{n-1}}$.*

5.2 Limiting Asymmetric Equilibrium

The limiting asymmetric equilibrium of the symmetric model follows trivially from Proposition 1 by taking limits as $s_i \rightarrow s_1$ for all i . We obtain that Firm 1 always advertises, Firm 2 advertises with probability $a_2 = \frac{s_1 - A}{s_1}$, both earn zero profits, and all other firms sit out. Firm 1's surplus offers follow $G_1(\sigma) = \frac{A}{s_1 - \sigma}$, with $G_1(0) = \frac{A}{s_1}$ representing an atom at the monopoly offer $\sigma = 0$. Firm 2's offers follow $G_2(\sigma) = \frac{A}{s_1 - \sigma} - \frac{A}{s_1}$ – identical to Firm 1, except that $G_2(0) = \frac{A}{s_1}$ represents not advertising. It is straightforward to confirm that these strategies do constitute an equilibrium of the symmetric game.

5.3 Arguments against the symmetric equilibrium

While it is commonplace to focus on symmetric equilibria when a game is symmetric, there is not always a compelling rationale for doing so. In our case, we will argue that

the symmetric equilibrium is less attractive than the limiting asymmetric one because it is unstable and generates perverse comparative statics. Other papers have noted similarly counterintuitive comparative statistics in symmetric equilibria, but we also stress the underlying instability of the symmetric equilibria. The same features hold in related games such as Varian’s Model of Sales.

5.3.1 (In)stability of the symmetric equilibrium

Symmetry can be a useful simplifying assumption, but in reality we would generally expect that firms differ at least a little bit in their costs, qualities, or both. Suppose that we start from a symmetric situation and then perturb Firm i ’s surplus to $s_i = s + \zeta_i$, where ζ_i is a publicly observed idiosyncratic shock. Then generically, the surpluses will all be distinct, and so the unique equilibrium is the one characterized by Proposition 1. Thus, for a small perturbation, outcomes in the perturbed model will not be close to the symmetric equilibrium above, which has no counterpart when surpluses are unequal, but will be close to outcomes in the limiting asymmetric equilibrium.¹² In this sense, the symmetric equilibrium is unstable.¹³

5.3.2 Unappealing welfare implications of the symmetric equilibrium

Because firms earn zero profits in the symmetric equilibrium, expected social surplus and consumer surplus are equal. The following unintuitive and implausible property about the symmetric equilibrium follows from the fact that the consumer’s best offer distribution $G_{\max}(\sigma)$ is increasing in n :

¹²To be more precise, there are $n(n - 1)$ limiting asymmetric equilibria, depending on which firms take the roles of Firm 1 and 2, and the equilibrium of the perturbed game must be close to one of these. Furthermore, equilibrium outcomes, such as profits, total advertising, and the distribution of surplus offers, will be close to the values they take in all of the limiting asymmetric equilibria.

¹³This argument does depends on the assumption that at least some of the differences among firms are public rather than private; we would argue that this is reasonable. Correlated deviations from symmetry do not upend this argument unless they are such that the top two (or more) firms remain identical; it is hard to see why this should be the case.

Proposition 6 *In the symmetric equilibrium, expected social surplus and consumer surplus are decreasing in the number of firms n .*

Thus, more competition among firms reduces welfare and makes consumers worse off. This surprising property is driven by the indifference conditions underlying the equilibrium. As n rises, the distribution of the best offer from any collection of $n - 1$ of the firms must remain constant (so as to keep the remaining firm indifferent to competing). This is only possible if each individual firm competes *less* vigorously as n rises. But then a consumer's overall best offer, the max of the best offer from the first $n - 1$ firms and the final firm's offer, must grow statistically worse with n .¹⁴

In contrast, in the limiting asymmetric equilibrium social surplus and consumer surplus are unaffected by additional firms, since competition by the two active firms is already sufficiently fierce to foreclose the market to everyone else. The result that additional competition *does* improve consumer surplus – but that all of the gains are made in the shift from monopoly to duopoly – is admittedly extreme, but it is typical of Bertrand price-setting games. Just as in other Bertrand-like games, we would expect this conclusion to soften into a more gradual consumer surplus improvement with n under softer price competition (such as discussed in Section 6.3, for example).¹⁵

Thus we argue that the symmetric equilibrium is a less appealing prediction of behavior than the limiting asymmetric one on the grounds of both stability and unintuitive comparative statistics. The two are likely linked. We recall Samuelson

¹⁴This logic has some precedent in mixed strategy equilibria. The result is reminiscent of the Palfrey and Rosenthal (1984) binary public good game whereby acting provides a public value v to all players at cost c to the players that choose to act. Sharkey and Sibley (1993) (for the symmetric case here discussed) already noted the anti-competitive effect of entry on the equilibrium price distribution per firm, and Stahl (1994) shows that social surplus can decrease. Indeed, Stahl (1994) analyzes an advertising cost function that encompasses both the Butters (1977) case and ours (as a limit case) and finds that with a relatively flat marginal cost of advertising, seller entry can decrease social surplus.

¹⁵In the appendix, we compare welfare in the symmetric and limiting asymmetric equilibria more closely, with a focus on decomposing sources of inefficiency.

(1941): “[T]he problem of stability of equilibrium is intimately tied up with the problem of deriving fruitful theorems in comparative statics.”

6 Extensions

6.1 Downward-sloping demand

While we have assumed the firms face a consumer with unit demand, our key results continue to apply if her demand is downward-sloping. In particular, the top firm always advertises and earns a positive profit, the second-best firm advertises with positive probability and earns zero profit, and all other firms sit out. Below we sketch these results and point out some new wrinkles that emerge.

For simplicity, we assume the consumer has the same demand $D(p)$ for any firm’s product, while the firms have heterogeneous, constant marginal costs $c_1 < c_2 < \dots < c_n$.^{16,17} The demand function $D(p)$ is assumed to be twice continuously differentiable and (-1)-concave.¹⁸ Let p_i^m be Firm i ’s monopoly price against this demand curve. To rule out uninteresting cases, we assume that $(p_1^m - c_2) D(p_1^m) > A$ which ensures that at Firm 1’s monopoly price at least the top two firms could cover their ad costs with a sale. Let $\tilde{\pi}_i(p) = (p - c_i) D(p)$ be the “pseudo-profit” function representing the gross profit Firm i could earn if it faced the consumer as a monopolist. The concavity assumption ensures $\tilde{\pi}_i(p)$ is strictly quasi-concave with $p_i^m \geq p_1^m$, so that Firm i ’s pseudo-profit is unambiguously increasing in price over the interval $p \in (c_i, p_1^m)$.

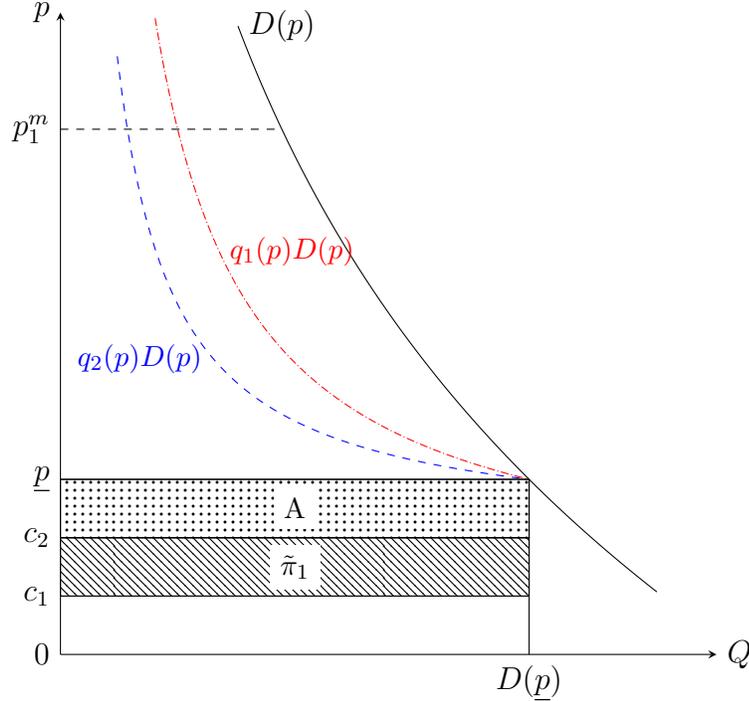
Absent quality differences across firms, the consumer will simply choose her lowest

¹⁶The same analysis would apply, with some relabeling, if the consumer’s demand for different products differed only by a firm-specific quality shift term.

¹⁷Of course, it would be interesting to study more general asymmetries in the demand faced by firms. However, when those asymmetries cannot be summarized by a one-dimensional parameter (implying that firm competitiveness cannot be unambiguously ranked), the analysis becomes quite complex. See the next section, with heterogeneous ad costs, for a sense of the issues that arise.

¹⁸This ensures that $D(p)$ is “more concave” than a rectangular hyperbola, and that marginal revenue slopes down. See Caplin and Nalebuff (1992) for more details on ρ -concavity.

Figure 3: Downward-sloping Demand



Determination of equilibrium sales probabilities $q_i(p)$ for downward-sloping demand.

price offer, so we present the analysis in terms of price offers rather than surplus offers. The main modification to our earlier set-up is that Firm i 's expected profit when advertising p_i is now

$$\pi_i(p_i) = \tilde{\pi}_i(p_i) \Pr(i \text{ sells} \mid p_i) - A = (p_i - c_i) D(p_i) \Pr(i \text{ sells} \mid p_i) - A ;$$

where the only difference is the inclusion of the term $D(p_i)$. Lemma 1 goes through essentially unchanged; namely, if Firm i can break even at its lowest price offer, then a more efficient firm $j < i$ can survive at a slightly lower price. Lemmas 2 through 6 continue to hold as well, as does the argument that Firms 1 and 2 must mix over common support $[\underline{p}, \bar{p}]$ (plus an atom at \bar{p} for Firm 1), establishing the claims we made above.

As usual, the lower bound of the support is pinned down by the requirement that Firm 2 wins for sure and earns zero profit when offering \underline{p} ; thus we have $\tilde{\pi}_2(\underline{p}) = A$ (which uniquely ties down \underline{p}). Since \underline{p} is also in Firm 1's support, this establishes its equilibrium profit as $\pi_1^* = \tilde{\pi}_1(\underline{p}) - A = (c_1 - c_2) D(\underline{p})$ (which reduces to our Lemma 7 result for unit demand). However this is smaller than the rents of $(c_1 - c_2) D(c_2)$ that Firm 1 would receive in the standard Bertrand model without advertising.¹⁹ Furthermore, the upper bound \bar{p} must be Firm 1's monopoly price p_1^m , since this offer wins for Firm 1 only when Firm 2 does not advertise; in contrast with the unit-demand case, this least competitive offer leaves the consumer a strictly positive surplus.

As with our earlier analysis, the firms' profits, the bounds on prices, and the two indifference conditions completely pin down the equilibrium mixed strategies and Firm 2's chance of advertising. These conditions may be succinctly illustrated on a diagram of each firm's residual demand $D_i^r(p) = q_i(p) D(p)$, where $q_i(p)$ is its equilibrium probability of a sale when offering p . As shown in Figure 3, each firm enjoys the full demand $D(\underline{p})$ when offering \underline{p} , with gross profit of $\tilde{\pi}_2(\underline{p}) = A$ for Firm 2, and $\tilde{\pi}_1(\underline{p}) = A + \pi_1^*$ for Firm 1, as depicted. At higher prices, competition from Firm 2 must pull back Firm 1's residual demand just enough so that its gross profits lie on the iso-profit hyperbola $(p - c_1) D_1^r(p) = A + \pi_1^*$. Similarly, Firm 2's residual demand is determined by the iso-profit hyperbola $(p - c_2) D_2^r(p) = A$.²⁰ Then because Firm 2's chance of winning depends on Firm 1's price distribution $F_1(p)$ through the relation $q_2(p) = 1 - F_1(p)$, $F_1(p)$ can be inferred from the ratio of Firm 2's residual

¹⁹Due to the need to cover A , Firm 2 cannot compete down to marginal cost (as would be socially efficient); consequently quantity consumed is inefficiently low. Firm 1's per-unit profit margin is fixed at $c_1 - c_2$.

²⁰As Figure 3 shows, $q_1(p) = D_1^r(p) / D(p) > q_2(p) = D_2^r(p) / D(p)$, so that Firm 1 has a greater chance of winning at any price than Firm 2. This must be the case in equilibrium because $q_1(p)$ is a convex combination of $q_2(p)$ and $D(\underline{p}) / D(p) > 1$. The equilibrium D^r 's can be derived graphically using an analogous device to that used in the Appendix, where we find the equilibrium for the symmetric case.

demand to total demand. Likewise, Firm 1's chance of winning depends on Firm 2's price distribution, conditional on advertising, according to $q_1(p) = (1 - a_2) + a_2(1 - F_2(p))$; thus the ratio $D_1^r(p)/D(p)$ pins down Firm 2's equilibrium strategy.²¹

All of this analysis presumes that each firm must offer a simple linear price. However, since our focus is on personalized offers made to a consumer based on detailed information about his tastes, it is natural to consider the possibility that firms can craft nonlinear price offers. If, for example, firms can offer two-part pricing, with a fixed fee and a per-unit price, it is not hard to see that the analysis is even simpler than with linear prices – in fact, it collapses back to the single-unit analysis. To see why, note that for standard reasons, each firm will always wish to set its per-unit price efficiently, at marginal cost, and take any profits through the fixed fee. But then we may simply apply our standard model, substituting in for the surplus s_i the total consumer surplus generated by purchasing the socially optimal quantity from Firm i , and letting the fixed fee play the role of p_i .

6.2 Competition when firms have different costs of targeted advertising

Thus far we have assumed that all firms face the same cost A to reach the consumer with a targeted ad. However, one can imagine circumstances where this ad cost varies across firms, perhaps because they deliver offers through different channels (such as email versus postal mail) or purchase data from different data brokers.²²

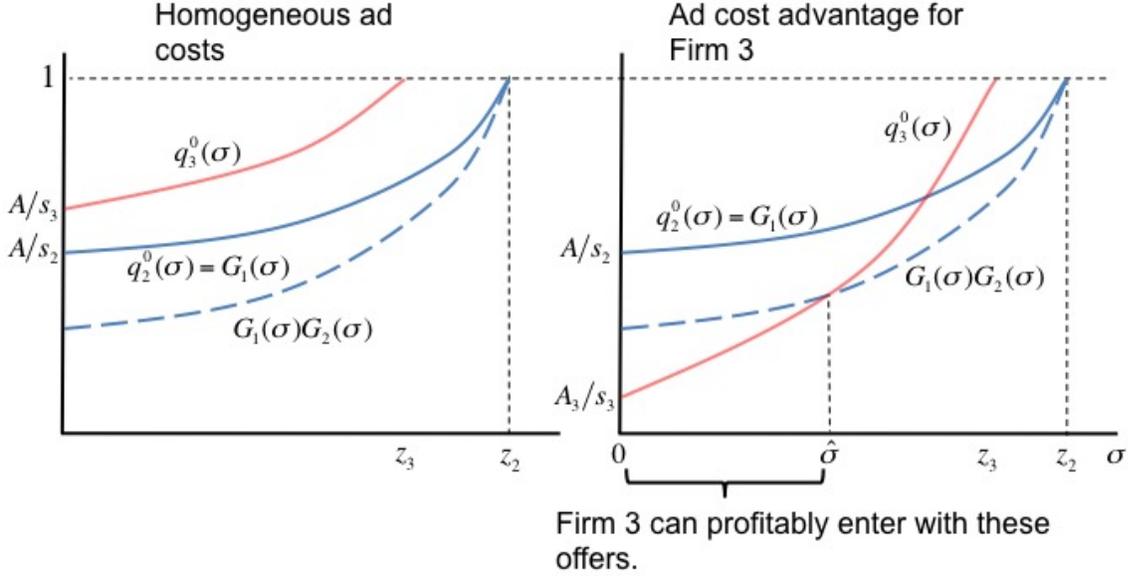
This subsection discusses how competition in our model looks like under asymmetric

²¹In particular, for Firm 1 we have $F_1(p) = 1 - A/((p - c_2)D(p)) = 1 - \tilde{\pi}_2(\underline{p})/\tilde{\pi}_2(p)$ for $p \in [\underline{p}, p_1^m]$, with a mass point of size $A/\tilde{\pi}_2(p_1^m) = \tilde{\pi}_2(\underline{p})/\tilde{\pi}_2(p_1^m)$ at its monopoly price p_1^m . Next, using $q_1(p) = (A + \pi_1^*)/((p - c_1)D(p)) = \tilde{\pi}_1(\underline{p})/\tilde{\pi}_1(p)$, we pin down a_2 (from Firm 1's chance of winning at its monopoly price) as $a_2 = 1 - q_1(p_1^m) = 1 - \tilde{\pi}_1(\underline{p})/\tilde{\pi}_1(p_1^m)$ and therefore Firm 2's conditional-on-advertising price distribution as $F_2(p) = (1 - \tilde{\pi}_1(\underline{p})/\tilde{\pi}_1(p))/(1 - \tilde{\pi}_1(\underline{p})/\tilde{\pi}_1(p_1^m))$ for $p \in [\underline{p}, p_1^m]$. All of these expressions reduce to our earlier results for the case of unit demand.

²²We thank a referee for raising this point.

ad costs. While the main results do not change too much, it does become possible for more than two firms to advertise in equilibrium.

Figure 4: Probability of Third Firm Entry



In a candidate equilibrium where only Firms 1 and 2 are active, $G_{\max}(\sigma) = G_1(\sigma)G_2(\sigma)$ represents Firm 3's chance of winning the sale if it enters with an offer of σ . Firm 3 needs to win with probability $q_3^0(\sigma)$ or greater to break even. If ad costs are homogeneous, as on the left, this is impossible, as the curves will be ranked as shown. But if Firm 3 has a sufficiently large ad cost advantage $A_3 < A_1 = A_2 = A$ (intercept $q_3^0(0) = \frac{A_3}{s_3} < G_1(0)G_2(0) = \frac{A}{s_2} \frac{A+s_1-s_2}{s_1}$ as shown on the right), then it can profitably enter with any surplus offer in the range $[0, \hat{\sigma}]$, disrupting the Firm 1/Firm 2 candidate equilibrium.

Suppose Firm i has targeted ad cost $A_i > 0$, and let $z_i = s_i - A_i$ denote Firm i 's surplus from a sale net of ad costs. These net surpluses will play an important role when ad costs are heterogeneous; without loss of generality, relabel the firms so that Firms 1 and 2 have the largest and second-largest *net* surpluses: $z_1 > z_2 > z_3 > \dots > z_n$. We claim that under this reordering, the following features of our earlier analysis are preserved: 1) Firm 1 always advertises and Firm 2 does with positive probability less than one; 2) Firm 1 earns a positive profit equal to its *net* surplus advantage over Firm 2, $\pi_1 = z_1 - z_2$, while all other firms earn zero profit; and 3) the support

of surpluses offered to consumers is $[0, z_2]$. The logic mirrors our earlier analysis.^{23,24}

The non-identical ad cost case begins to look different when Firm 3 (or any other firm) can profitably enter the market above (i.e., advertise with positive probability). To explore this question, consider a firm’s “break-even win probability” $q_i^0(\sigma)$: the chance of winning the sale with surplus offer σ at which Firm i earns zero expected profit. This break-even probability is defined implicitly by $q_i^0(\sigma)(s_i - \sigma) - A_i = 0$. Given our results above, competition from Firm 1 implies that Firm 2’s win probabilities in a candidate equilibrium will lie on its break-even line: $G_1(\sigma) = q_2^0(\sigma) = \frac{A_2}{s_2 - \sigma}$. If ad costs are identical, then Firm 3’s break-even line lies strictly above Firm 2’s as in the left panel of Figure 4 – at any surplus offer it would need a strictly better chance of winning than $q_2^0(\sigma)$ to be profitable. But if Firm 3 were to enter with an offer of σ , it would actually win with probability $G_1(\sigma)G_2(\sigma)$ (the dashed line in the figure). This is strictly worse than $q_2^0(\sigma) = G_1(\sigma)$ (since Firm 3 would compete against not just Firm 1 but also Firm 2) and so Firm 3 would lose money. However, if Firm 3 can make up for a smaller surplus by advertising sufficiently more cheaply – that is, if $s_2 > s_3$ and $z_2 > z_3$ but $A_3 < A_2$ – then its break-even line $q_3(\sigma) = \frac{A_3}{s_3 - \sigma}$ may cross that of Firm 2 from below (as in the right panel of Figure

²³Firm 1 can earn at least $z_1 - z_2$ by overcutting the best surplus offer among its rivals and winning for sure. But it cannot earn strictly more, or else Firm 2 could use the same trick to earn a positive profit, and positive profits for any firm other than Firm 1 are impossible for the reasons laid out in Lemma 2. Thus Firm 1 earns $\pi_1 = z_1 - z_2$, always advertises, and its best surplus offer (which wins with probability one) must be $\bar{\sigma} = s_1 - A_1 - \pi_1 = z_2$. Firm 1 would not make such a generous offer unless it faced competition all the way up to $\bar{\sigma} = z_2$; only Firm 2 can offer that much surplus without losing money, so Firm 2 must advertise with positive probability. For the same reasons as earlier, the one firm that does not always face competition (Firm 1) must have an atom at the worst surplus offer $\underline{\sigma}$ consumers receive (so that other advertising firms – who do always face competition – win often enough with their worst offers to cover their ad costs). Since this offer wins only if there are no competing ads, it will be set at the monopoly level, $\underline{\sigma} = 0$.

²⁴If only Firms 1 and 2 are active in equilibrium (as will be the case if no lower-ranked firm has a sufficiently large ad cost advantage), it is straightforward to confirm that a consumer’s best-surplus-offer distribution generalizes to $G_{\max}(\sigma) = G_1(\sigma)G_2(\sigma) = \frac{A_2}{s_2 - \sigma} \frac{A_2 + s_1 - s_2}{s_1 - \sigma}$. Thus the consumer surplus conclusions of Proposition 4 generalize with one nuance: while consumers benefit from a reduction in A_2 , they are unaffected (at the margin) by changes in A_1 .

4). In this case, Firm 3 cannot compete with Firm 2 on generous surplus offers that are likely wins, but it has a comparative advantage in making miserly offers that earn large profit margins but rarely win – it does not need to win as often as Firm 2 would to break even. Firm 3’s ad cost advantage must be great enough to make such offers profitably despite the fact that it can win only with probability $G_1(\sigma)G_2(\sigma)$, due to the combined competition of Firms 1 and 2. In this case, as illustrated in the right panel of Figure 4, Firm 3 can enter the market (i.e. advertise) over a range of high-price, low-surplus offers. In the appendix, we provide an example of an equilibrium like this: Firm 3 mixes over low-surplus offers (displacing Firm 2), Firm 2 mixes over high-surplus offers, and collectively they provide competitive discipline for Firm 1, which mixes over the entire range. Under the right conditions many additional firms could enter the market in this way, each carving out its own niche, although they must all earn zero profits.

6.3 Competition when information about the consumer’s tastes is noisy

We have assumed throughout that firms know the consumer’s tastes perfectly, insofar as her reservation prices at each firm are common knowledge in the price competition game – this simplifying assumption has been very helpful in obtaining the clean, sharp results presented thus far. As data mining improves, we may indeed be approaching such a brave new world in which every consumer’s heart is laid bare to the market, but we are probably not there quite yet. A more realistic assumption might be that the consumer’s reservation prices have an observable component, which firms can infer from data mining, and an idiosyncratic component that the consumer knows but the firms do not.²⁵ We will argue informally that our model provides a reasonable approximation of this more complex setting when the firms’ information is good but

²⁵We thank a referee for suggesting this line of thought.

not perfect.

Suppose the consumer’s reservation value for Firm i ’s product is actually $\tilde{r}_i = r_i + \mu\varepsilon_i$, where r_i is common knowledge, $\mu \geq 0$ is a scale parameter, and ε_i is a taste shock observed only by the consumer. In the spirit of discrete choice modeling, suppose these taste shocks are drawn i.i.d. from a distribution known to the firms. If there were no targeted advertising costs to consider, we would then have a straightforward discrete choice model of price competition as studied by Anderson, de Palma, and Thisse (1992). A general intuition is that the “noise” in the consumer’s preferences tends to soften price competition, permitting firms with lower observable quality to nevertheless carve out some market share and profits. In effect, ad costs would append a simultaneous entry decision to the price competition game: Anderson and de Palma (2001) study a two-stage game with entry decisions preceding pricing. In broad terms, for μ large enough we would expect to see similar results to other entry games: with soft price competition, a larger number of actively advertising firms can be supported in equilibrium as ad costs go down, and there may be multiple equilibria in which different combinations of firms enter.

When the noise in consumer tastes is small, an argument from upper-hemicontinuity of the equilibrium correspondence suggests outcomes will approximate those in our standard model.²⁶ As $\mu \rightarrow 0$ any sequence of equilibria of the μ -game must converge to some equilibrium of the $\mu = 0$ game, but the latter is our standard model which has a unique equilibrium.²⁷

²⁶In their classic paper on equilibrium in games with discontinuous payoffs, Dasgupta and Maskin (1986, p38) make essentially the same point, arguing that equilibria of a limit game can be constructed as a limit of equilibria of a perturbed game in which expected payoffs have been made continuous by adding exogenous uncertainty.

²⁷To flesh out the intuition a bit, for μ small enough, price competition is sharp enough that there is only room for one firm to make a positive profit. In particular, a version of Lemma 2 still applies: if two firms were profitable, both would always advertise. But each firm’s residual demand is sufficiently elastic that, at their least competitive price offers, it is impossible for each of them to win often enough to cover the ad cost and yet neither has an incentive to undercut the other. Thus we have Firm 1 always advertising and earning a profit close to $s_1 - s_2$, and Firm 2 advertising with

7 Conclusion

Our results indicate that much of our intuition about standard asymmetric Bertrand competition (e.g. Lederer and Hurter, 1986) still applies when firms must pay to advertise: only the most and second-most competitive firms are relevant, and only the former earns a profit (equal to its competitive advantage over the latter). However, some of the underlying details are quite different: pricing is in mixed strategies, and equilibrium is inefficient due to both wasteful advertising and sales to the second-best firm. Our equilibrium gracefully limits to the conventional Bertrand outcome as ad costs vanish. When firms are homogeneous and ad costs are positive, our equilibrium selection offers an (arguably) more appealing alternative to the counter-intuitive comparative statics of the usual symmetric equilibrium.

In assuming that firms must pay to advertise and in the result that equilibrium prices follow Pareto distributions, our framework resembles Butters' (1977) celebrated model.²⁸ However, where Butters assumes that ads are matched to individual consumers randomly, we assume a firm knows exactly whom it is targeting with an offer. A secondary distinction is that Butters focuses on symmetric firms, while we stress the importance of asymmetries.

In closing, we note a few directions for future research. In some settings, it may be reasonable to think that consumers are already aware of firms' list prices, but firms can pay to send them targeted ads with personalized discounts. In this case, firms must anticipate the outcome of this targeting game (including which consumers will be fought over twice) when setting their initial list prices. In a companion paper (2014), we study this setting, incorporating the model in this paper as a second stage.

positive probability and earning zero profit. Any strictly lower-surplus firm would make strictly less than Firm 2's gross profit were it to advertise, and so could not cover its ad cost.

²⁸Similar assumptions and results appear in models based on Varian's (1980) Model of Sales, such as Baye and Morgan (2001) and in various search models in the vein of Stahl (1989).

Our assumptions make sense when firms have access – at a price – to the same extensive data about individual consumer tastes. However, some of the most interesting applications of targeting arise when one firm has better information about a consumer than its rivals do. For example a grocery store may be able to use loyalty card data (that its rivals do not see) to link a consumer to his record of past purchases, giving the store an advantage in crafting personalized offers for him. In our model, consumer taste is summarized by willingness to pay for a single product, and so targeting reduces to a personalized price. In reality, consumer tastes are more complex than this, and effective targeting might involve understanding which product to offer, when to offer a discount (based on forecasting when the consumer will need to restock), whether to bundle products together, and so forth. The supply side of targeted ad provision (which we have treated as a reduced form cost) deserves additional attention. Web search, social media, and advertising platforms are all important spaces where a consumer’s tastes are partially observed (*via* cookies on web-sites or by linking social network data to consumption choices, for example) and matched to firms wishing to reach him (often through real-time auctions for advertising placement). While the literature on matching firms to consumers is growing rapidly (see e.g. Athey and Ellison, 2011), much of it suppresses firms’ competition in product prices in order to focus on how they compete in markets for an advertising platform’s targeting service. Because the targeted ad prices generated by such markets can vary across firms (as in the case of position auctions), our preliminary skirmish with product-price competition in this case (Section 6.2) deserves a deeper look.²⁹ Of course, these suggestions are just the tip of the iceberg – data-rich processing is quickly changing the landscape of advertising and pricing, and there is much room for future work.

²⁹In addition to the differences in ad prices that firms may face, they can sometimes opt into different *types* of ad pricing, such as cost-per-click, cost-per-impression, or cost-per-action.

References

- [1] Anderson, Simon, Alicia Baik, and Nathan Larson (2015). "How Targeted Advertising Affects Price Competition, Profits, and Consumers."
- [2] Anderson, Simon and André de Palma. 1988. "Spatial Price Discrimination with Heterogeneous Products," *Review of Economic Studies*, 55, 573-592.
- [3] Anderson, Simon and André de Palma. 2001. "Product diversity in asymmetric oligopoly: Is the quality of consumer goods too low?," *The Journal of Industrial Economics*, 49, 113-135.
- [4] Anderson, Simon, de Palma, André, and Jacques-François Thisse. 1992. *Discrete Choice Theory of Product Differentiation*, MIT Press.
- [5] Athey, Susan and Glenn Ellison. 2011. "Position Auctions with Consumer Search," *Quarterly Journal of Economics*, 126, 1213-1270.
- [6] Baye, Michael R., Dan Kovenock, and Casper G de Vries. 1992. "It Takes Two to Tango: Equilibria in a Model of Sales," *Games and Economic Behavior*, 4, 493-510.
- [7] Baye, Michael R., Dan Kovenock, and Casper G de Vries. 1996. "The All-Pay Auction with Complete Information," *Economic Theory*, 8, 291-305.
- [8] Baye, Michael R. and John Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91, 454-474.
- [9] Bester, Helmut and Emmanuel Petrakis. 1995. "Price Competition and Advertising in Oligopoly," *European Economic Review*, 39, 1075-1088.

- [10] Bester, Helmut and Emmanuel Petrakis. 1996. "Coupons and Oligopolistic Price Discrimination," *International Journal of Industrial Organization*, 14, 227-242.
- [11] Butters, Gerard R. 1977. "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 44, 465-491.
- [12] Dasgupta, Partha and Eric Maskin. 1986. "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications," *Review of Economic Studies*, 53, 27-41.
- [13] Grossman, Gene and Carl Shapiro. 1984. "Informative Advertising with Differentiated Products," *Review of Economic Studies*, 51, 63-81.
- [14] Hillman, Arye L. and John G. Riley. 1989. "Politically Contestable Rents and Transfers," *Economics and Politics*, 1, 17-39.
- [15] Hotelling, Harold. 1929. "Stability in Competition," *Economic Journal*, 39, 41-57.
- [16] Janssen, Maarten C.W., José Luis Moraga-Gonzalez and Matthijs R. Wildenbeest. 2005. "Truly Costly Sequential Search and Oligopolistic Pricing," *International Journal of Industrial Organization*, 23, 451-466.
- [17] Kocas, Cenk and Tunga Kiyak. 2006. "Theory and Evidence on Pricing by Asymmetric Oligopolies," *International Journal of Industrial Organization*, 24, 83-105.
- [18] Lederer, Phillip J. and Arthur P. Hurter, Jr. 1986. "Competition of Firms: Discriminatory Pricing and Location," *Econometrica*, 54, 623-640.
- [19] Mas-Colell, Andreu, Michael D. Whinston and Jerry R. Green. 1995. *Microeconomic Theory*, Oxford: Oxford University Press.

- [20] Narasimhan, C. 1988. "Competitive Promotional Strategies," *Journal of Business*, 61, 427-449.
- [21] Palfrey, Thomas and Howard Rosenthal. 1984. "Participation and the Provision of Discrete Public Goods: a Strategic Analysis," *Journal of Public Economics*, 24, 171-193.
- [22] Robert, Jacques and Dale Stahl. 1993. "Informative Price Advertising in a Sequential Search Model," *Econometrica*, 61, 657-686.
- [23] Samuelson, Paul. 1941. "The Stability of Equilibrium: Comparative Statics and Dynamics," *Econometrica*, 9, 97-120.
- [24] Shaffer, Greg and Z. John Zhang. 1995. "Competitive Coupon Targeting," *Marketing Science*, 14, 395-416.
- [25] Shaffer, Greg and Z. John Zhang. 2002. "Competitive One-to-One Promotions," *Management Science*, 48, 1143-1160.
- [26] Sharkey, William W. and David Sibley. 1992. "A Bertrand Model of Pricing and Entry," *Economics Letters*, 41, 199-206.
- [27] Stahl, Dale. 1989. "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 79, 700-712.
- [28] Stahl, Dale. 1994. "Oligopolistic Pricing and Advertising," *Journal of Economic Theory*, 64, 162-177.
- [29] Tirole, Jean. 1988. *The Theory of Industrial Organization*, The MIT Press.
- [30] Varian, Hal R. 1980. "A Model of Sales," *American Economic Review*, 70, 651-659.

A Appendix

A.1 Supporting analysis for the symmetric model

A.1.1 Welfare in the symmetric and limiting asymmetric equilibria

We start by computing expected social surplus for the symmetric equilibrium when all firms have surplus s_1 , using the equilibrium strategies in Proposition 5. Because it does not matter which product the consumer buys, expected social surplus can be decomposed into the expected social value of receiving at least one offer minus total expected advertisement costs:

$$\begin{aligned} SS &= s_1 \Pr(\text{consumer gets an offer}) - \text{Ad costs} \\ &= s_1 (1 - G_i(0)^n) - An (1 - G_i(0)) \\ &= s_1 \left(1 - \left(\frac{A}{s_1} \right)^{\frac{n}{n-1}} \right) - An \left(1 - \left(\frac{A}{s_1} \right)^{\frac{1}{n-1}} \right) \end{aligned}$$

Next, we claim the following.

Proposition 7 *With $n = 2$ firms, the symmetric and limiting asymmetric equilibria have the same expected social surplus and the same expected consumer surplus. With more than two firms, expected social and consumer surplus are both strictly higher in the asymmetric equilibrium.*

Recall that in both cases, social and consumer surplus are equal because firms earn zero profits. The first part follows from the observation that when $n = 2$, the distribution of a consumer's best offer is the same under either equilibrium: $G_{\max}(\sigma) = \left(\frac{A}{s_1 - \sigma} \right)^2$. The rest follows from Proposition 6 (and the fact that the asymmetric equilibrium does not change with n). Regardless of n , second-best social surplus would be $s_1 - A$: the consumer must get some product (it doesn't matter which) and this requires sending at least one ad. For both equilibria (symmetric

and asymmetric), inefficiency relative to this second-best benchmark may be decomposed into two components, excessive advertising costs and misallocation. For the symmetric equilibrium these costs are

$$\begin{aligned} \text{Social cost of wasteful advertising (Symm.)} &= An(1 - G(0)) - A \\ \text{Social cost of misallocation (Symm.)} &= s_1 G(0)^n. \end{aligned}$$

The former is just total expected advertising minus the necessary amount A ; the latter reflects failures to make a sale (the only type of misallocation, since products are identical).³⁰ Under the asymmetric equilibrium, there is no misallocation since the consumer always receives at least one ad and makes a purchase. Thus all inefficiency is due to (socially unnecessary) advertising by Firm 2; this has cost $a_2 A$; hence

$$\text{Social cost of wasteful advertising (Asymm.)} = A\left(1 - \frac{A}{s_1}\right)$$

With two firms, the symmetric and asymmetric equilibria both have total avoidable inefficiency equal to $A\left(1 - \frac{A}{s_1}\right)$, but for slightly different reasons: the latter has more wasteful advertising, but under the former, a sale may be lost due to miscoordinated advertising. It is not a coincidence that these two effects happen to balance out. At a technical level, competitive forces ensure that the consumer's best offer distribution must be $G_1(\sigma)G_2(\sigma) = G(\sigma)^2$ under either equilibrium. The difference between equilibria amounts to an interpretation of $G_1(0)$: if this is the probability of an advertised zero-surplus offer, we have the asymmetric equilibrium; if it reflects a failure to advertise at all, we have the symmetric one. In this situation, Firm 1 must be indifferent between advertising $\sigma = 0$ and not advertising (since it must earn zero profit either way). Furthermore, its private incentives on a $\sigma = 0$ offer are aligned with social welfare (since it would capture the full surplus from a sale), so, its private indifference implies that the shift would be welfare-neutral.

³⁰Thus using $G(0) = A/s_1$, total avoidable inefficiency may be written $\chi(A) = s_1 (A/s_1)^{\frac{n}{n-1}} + An\left(1 - (A/s_1)^{\frac{1}{n-1}}\right) - A$.

A.1.2 Asymmetric Equilibria under Symmetry

There are additional equilibria in which an arbitrary subset of $\tilde{n} < n$ of the firms play a version of this equilibrium (with $\tilde{n} \geq 2$ replacing n), while the $n - \tilde{n}$ others sit out (never advertise). As the argument that leads up to (5) makes clear, the equilibrium offer distribution for the remaining (potentially active) firms is symmetric. There remains the possibility that at most one of them advertises a zero-surplus offer with positive probability. Indeed, if two or more firms were to advertise a zero-surplus offer with positive probability then one could profitably undercut and gain a positive sales increase probability from an infinitesimal price cut. To see that one firm could use a zero-surplus advertisement, recall that the probability mass $G_i(0) = \left(\frac{A}{s_1}\right)^{\frac{1}{\tilde{n}-1}}$ in (5) may include a zero-surplus advertisement for some i . This leaves an indeterminacy. For arbitrary $a_i \in \left[1 - \left(\frac{A}{s_1}\right)^{\frac{1}{\tilde{n}-1}}, 1\right]$, any strategy profile in which Firm i refrains from advertising with probability $1 - a_i$, advertises a zero surplus offer with probability $a_i - \left(1 - \left(\frac{A}{s_1}\right)^{\frac{1}{\tilde{n}-1}}\right)$, and the remaining firms refrain from advertising with probability $\left(\frac{A}{s_1}\right)^{\frac{1}{\tilde{n}-1}}$, is an equilibrium. Thus it remains true and consistent with our earlier analysis that at most one firm can have an atom of ads at $\sigma = 0$, but it is no longer necessary that any firm does so, since they all earn zero profit and so are indifferent between advertising and not. Notice though that this indeterminacy has no bearing on equilibrium payoffs.

Pulling this together, there is thus an equilibrium under symmetry at which only two firms are active: one advertises with probability 1, the other with probability $1 - \left(\frac{A}{s_1}\right)$, and for both the offer distribution is $G(\sigma) = \frac{A}{s_1 - \sigma}$. But this is identical to the limiting equilibrium, under asymmetric costs and valuations, as those asymmetries vanish. That is to say, a perturbation approach of beginning with strictly differentiated firms and taking limits as the gaps among the top n firms vanish will select this asymmetric two-firm equilibrium in the symmetric limit, not the symmetric

n -firm equilibrium.

A.2 Other symmetric cases

We have previously considered the case in which there are no ties in the surpluses s_i that firms offer and the symmetric case in which all surpluses are equal. This section evaluates the remaining cases in which some subsets of the firms are identical; thus we set $s_1 \geq s_2 \geq \dots \geq s_n$. There are three main cases to consider, depending on the highest rank at which any firms tie.

A.2.1 Low ties

The easiest to dispense with is the case in which any ties are among firms at the level of Firm 3 or worse; that is, $s_1 > s_2 > s_3 \geq \dots \geq s_n$. It should be clear that this will not affect the equilibrium outcome – a few of the supporting lemmas must be amended slightly, but Proposition 1 still applies.

A.2.2 Dominant Firm and Fringe Firms

Next suppose that m firms tie for the second-ranked spot (whether or not there are ties below the second-ranked position will be irrelevant) : $s_1 > s_2 = s_3 = \dots = s_{m+1} > s_{m+2} \geq \dots$. It is straightforward to prove that *any equilibrium must have strictly positive profits for Firm 1, zero profits for the other firms, including the m runners-up, and only Firm 1 and some subset of the runners-up advertising with positive probability*. As earlier, let $\Delta_{12} = s_1 - s_2$ be the advantage of Firm 1 over the runners-up, and let $G_i(\sigma)$ be the distribution of the surplus offered by Firm i , with a failure to advertise included as an offer of $\sigma_i = 0$. Likewise, define $G_{-i}(\sigma) =_{\{j \leq m+1 : j \neq i\}} G_j(\sigma)$, the distribution of the best opponent surplus offer faced by Firm i . Arguments similar to those earlier can be used to establish that each of these “best opponent” distributions has support on $[0, s_2 - A]$. Similar arguments establish that Firm 1’s

equilibrium profit is $\pi_1 = \Delta_{12}$: any firm can win with probability one by advertising the upper bound surplus and when Firm 1 does so it charges a price that is Δ_{12} higher than the other firms, thereby earning $\pi_2 + \Delta_{12} = \Delta_{12}$. Before examining other possibilities, first consider the candidate equilibrium in which the m runners-up behave symmetrically. Firm 1's indifference over its mixed strategy support implies that its probability of winning with an offer of σ_1 is no different now that it has m rivals than it was when it faced one (under the assumptions of Proposition 1); that is,

$$G_{-1}(\sigma) = (G_2(\sigma))^m = \frac{A + \Delta_{12}}{s_1 - \sigma}$$

and so $G_2(\sigma) = \left(\frac{A + \Delta_{12}}{s_1 - \sigma}\right)^{1/m}$. Similarly, indifference for each runner-up implies that it must face the same best-opponent distribution that Firm 2 did in Proposition 2; this implies $G_{-i}(\sigma) = G_1(\sigma) (G_2(\sigma))^{m-1} = \frac{A}{s_2 - \sigma}$ for each $i \in \{2, \dots, m+1\}$, and so

$$G_1(\sigma) = \left(\frac{A}{s_2 - \sigma}\right) \left(\frac{A + \Delta_{12}}{s_1 - \sigma}\right)^{-\frac{m-1}{m}}$$

The consumer's best offer is then distributed according to $G_{\max}(\sigma) = \left(\frac{A}{s_2 - \sigma}\right) \left(\frac{A + \Delta_{12}}{s_1 - \sigma}\right)^{1/m}$.

Notice that at both the top and second-ranked firms the consumer has a positive chance of not being offered a strictly positive offer. To complete the description of equilibrium, we must establish whether the probability $G_1(0) > 0$ reflects Firm 1 advertising a zero surplus offer or not advertising, and similarly for $G_2(0)$. Because Firm 1 earns positive profits, it must advertise with probability one, and so the probability mass $G_1(0)$ must represent an atom of advertised zero surplus offers. There cannot be more than one firm advertising an atom of zero surplus offers, as each would have a strict incentive to undercut, and so we must have $G_2(0) = 1 - a_2$ for each of the second-ranked firms.

Having established this template for a symmetric equilibrium, and noting that each of the tied firms is indifferent to not advertising, it is straightforward (cf. Section

A.1.2) to show that there is a family of additional equilibria in which a subset $\tilde{m} < m$ of the tied firms advertise using the strategies above (with $\tilde{m} \geq 2$ substituted for m), and the remainder “sit out.” *A priori*, it is not clear which of these equilibria should be preferred over the others; absent a reason to distinguish between the tied firms, one might argue for the “equal treatment” – and hence symmetric – equilibrium in which they all advertise. However, once again such an equilibrium is unstable. Our preferred approach is to begin with the generic case of unequal $\{s_2, \dots, s_{m+1}\}$ and select the limiting equilibrium as differences between these firms vanish. As per our earlier analysis, this approach selects a limit equilibrium in which one firm (Firm 2) advertises and the other $m - 1$ runners-up sit out. These two alternative equilibrium selections agree on firm profits, but disagree on price distributions, probabilities of advertising for the runners-up, and consumer surplus. In particular, the consumer is better off in the equilibrium where only Firms 1 and 2 are active.

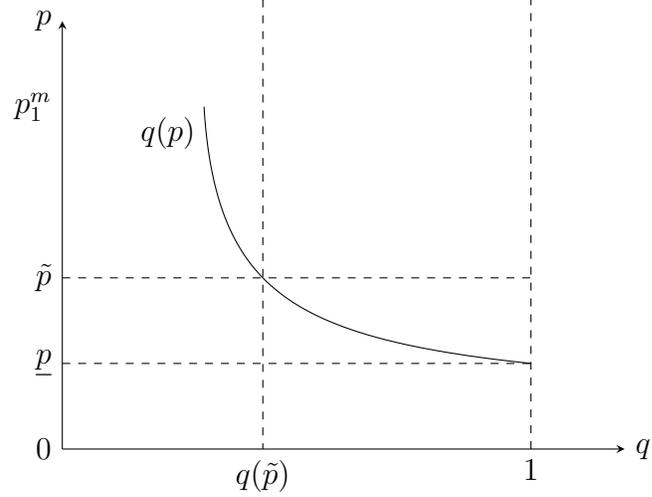
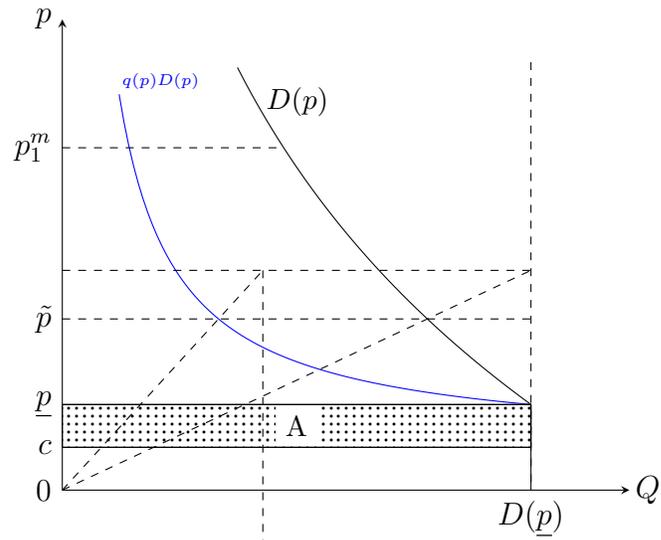
A.2.3 Top tie

Finally, suppose that m firms tie for the top spot: $s_1 = s_2 = \dots = s_m > s_{m+1} \geq \dots$. In this case, only firms at the top will ever advertise, and they all must earn profit zero. Indeed, if one of the lower-ranked firms j were to advertise in equilibrium, then it would have to be the case that all m top firms earn strictly positive profits. (If not, a top firm earning zero could profitably deviate to undercutting j 's best offer.) But then all m top firms would have to be advertising with probability one, and this is impossible for the reasons laid out in Lemma 2.³¹ Consequently, ties below the top level will be irrelevant. Thus the analysis of Section 5 covers this case.

³¹The arguments in Lemma 2 rule out strictly positive profits for more than one of the top firms. Furthermore, if any single firm, say Firm 1, were to earn strictly positive profits in equilibrium, then any of the other top firms could undercut Firm 1's lowest advertised price and earn strictly positive profits as well.

A.3 Downward-sloping demand: additional analysis

Figure 5: Symmetric Downward-sloping Demand



We next derive the symmetric equilibrium under symmetric costs, c , and we show graphically how to tie down the equilibrium sales probability that underpins the equilibrium price distribution. Because all firms are indifferent between advertising and not, all earn zero profit. This ties down the lowest price in the support as the (unique) solution to $(\underline{p} - c) D(\underline{p}) = A$. Let $q(p)$ be the common equilibrium probability of a sale, which is determined by

$$(p - c) D(p) q(p) = A \text{ for all } p \in [\underline{p}, p^m],$$

where p^m is the (common) monopoly price. The solution is illustrated in Figure 5. The top panel shows the price support as determined from the demand curve, along with the iso-profit line passing through $(\underline{p}, D(\underline{p}))$ that determines the equilibrium $q(p)$. The value of q as a function of p is shown in the lower panel by the device of associating $D(\underline{p})$ to the maximum possible value $q = 1$. Then the other q values can be found by drawing a ray from the origin through the demand $D(p)$ associated to any specific price in the support, finding where the ray reaches the value of $D(\underline{p})$ and then finding q from where the ray through the rectangular hyperbola reaches the same height. That is, we simply use similar triangles to find the ratio q which is the ratio of the horizontal distance to the rectangular hyperbola over the distance to the demand curve. (Notice that a similar device can be used in the asymmetric case of the Figure to tie down the equilibrium q 's there.)

Hence, we tie down the equilibrium conditional price distribution by $q(p) = (1 - aF(p))^{n-1}$. Note that the common equilibrium advertising probability, a , is tied down from the condition that the monopoly price returns zero profit, i.e.,

$$(1 - a)^{n-1} \pi^m = A$$

or $a = 1 - \left(\frac{A}{\pi^m}\right)^{\frac{1}{n-1}}$. Then the probability that there is no ad at all (and so the market

is not served) is

$$(1 - a)^n = \left(\frac{A}{\pi^m} \right)^{\frac{n}{n-1}}$$

which increases in n , reflecting the earlier result (see Proposition 6) that social surplus worsens with more competition. As we argued for the rectangular demand case, the symmetric equilibrium is unstable with respect to cost heterogeneity.

A.4 Advertising by three firms when ad costs differ

While it is beyond our scope to give a general analysis of strategies when more than two firms are active, a worked example may help to suggest what the general case looks like. There are three firms with surpluses $s_1 > s_2 > s_3$. Firms 1 and 2 have common ad cost $A_1 = A_2 = A$, while the low surplus Firm 3 has a cost advantage in advertising: $A_3 < A$. We assume this cost advantage is not too great, so that the net surplus ranking $z_1 > z_2 > z_3$ mirrors the gross surpluses. Thus Firms 1 and 2 will be active in equilibrium, and it remains to be seen whether Firm 3 will be.

As a stepping stone toward finding the equilibrium, it is helpful to set up the equilibrium that would prevail if only Firms 1 and 2 were present and then ask whether Firm 3 can disrupt it. Consulting earlier results, the distributions of surplus offers by Firms 1 and 2 respectively would then be $G_1(\sigma) = \frac{A}{s_2 - \sigma}$ and $G_2(\sigma) = \frac{A + \Delta s}{s_1 - \sigma}$, where $\Delta s = s_1 - s_2$, with the distribution of the best surplus offer between them given by $G_{\max}(\sigma) = G_1(\sigma)G_2(\sigma) = \frac{A}{s_2 - \sigma} \frac{A + \Delta s}{s_1 - \sigma}$. This tells us how often Firm 3 would win if it were to enter with a surplus offer σ ; if $G_{\max}(\sigma)$ ever lies above Firm 3's zero-profit line $q_3(\sigma) = \frac{A_3}{s_3 - \sigma}$, then Firm 3 can enter profitably and disrupt the two-firm equilibrium. A sufficient condition for this is $q_3(0) < G_{\max}(0)$, or $\frac{A_3}{s_3} < \frac{A}{s_2} \frac{A + \Delta s}{s_1}$, as illustrated in Figure 4. For the purpose of the example, suppose this condition does hold. As indicated in the figure, let $\hat{\sigma}$ be the largest surplus offer Firm 3 could make without losing money, given the strategies Firms 1 and 2 would use if Firm 3 were

not present.

In the full game with three firms, write $H_i(\sigma)$ for the distribution of surpluses that Firm i offers, with the understanding as before that $H_i(0) > 0$ represents an atom at $\sigma = 0$ for Firm 1 and declining to advertise for Firms 2 and 3. We claim there is an equilibrium with the following features:

1. Firm 2 specializes in low-price, high-surplus offers with support $\sigma \in [\hat{\sigma}, z_2]$, while Firm 3 specializes in high-price, low-surplus offers with support $\sigma \in [0, \hat{\sigma}]$. Firm 1 mixes over the entire range $\sigma \in [0, z_2]$.
2. At surplus offers above $\hat{\sigma}$, where Firm 3 does not compete, Firms 1 and 2 behave just as they would have done if Firm 3 were absent. That is, $H_1(\sigma) = G_1(\sigma)$ and $H_2(\sigma) = G_2(\sigma)$ for $\sigma \in [\hat{\sigma}, z_2]$.
3. Any probability that Firm 2 would have assigned to low-surplus offers $\sigma \in [0, \hat{\sigma}]$ is simply reassigned to not advertising. Consequently, Firm 2 advertises less often than it would have done if Firm 3 were absent.
4. Firm 1 is forced to compete more aggressively over the low-surplus offers than it would have done otherwise; that is, $H_1(\sigma) < G_1(\sigma)$ for $\sigma \in [0, \hat{\sigma}]$. In particular, it competes just hard enough to drive Firm 3's profit on such offers down to zero. This requires it to more than compensate for the loss of competition from Firm 2 over this range (since Firm 3 could have made positive profits against the competition of both Firms 1 and 2, under their original strategies).
5. Collectively, Firms 2 and 3 provide Firm 1 exactly as much competition at every surplus level as Firm 1 would have faced in the absence of Firm 3. That is, $H_2(\sigma)H_3(\sigma) = G_2(\sigma)$ for $\sigma \in [0, z_2]$. In particular, the probability that no competitor to Firm 1 advertises (given by $H_2(0)H_3(0)$) does not change when Firm 3 is present.

6. While the presence of Firm 3 does not affect any firm's profit, it does make consumers better-off. This is driven entirely by the first-order stochastic improvement in Firm 1's least competitive ($\sigma \in [0, \hat{\sigma})$) offers. Consequently, the presence of Firm 3 improves welfare as well.

In order to verify these claims, first recall that in the candidate equilibrium with only Firms 1 and 2, the equilibrium strategies can be summarized by the surplus offer distributions $G_1(\sigma) = \frac{A}{s_2 - \sigma}$ and $G_2(\sigma) = \frac{A + \Delta s}{s_1 - \sigma}$, with $\Delta s = s_1 - s_2$ (see Proposition 1). Defining $G_{\max}(\sigma) = G_1(\sigma)G_2(\sigma)$ and $q_3^0(\sigma) = \frac{A_3}{s_3 - \sigma}$ as in the text, we have $q_3^0(0) < G_{\max}(0)$ and $q_3^0(z_3) = 1 > G_{\max}(z_3)$. Then let $\hat{\sigma}$, defined by $q_3^0(\hat{\sigma}) = G_{\max}(\hat{\sigma})$ be the point on $[0, z_2]$ at which the two functions cross; one can confirm that this crossing is unique. We claim the following surplus offer distributions represent equilibrium strategies for Firms 1, 2, and 3.

$$\begin{aligned} H_1(\sigma) &= \begin{cases} \frac{q_3^0(\sigma)}{G_2(\hat{\sigma})} = \frac{A}{s_2 - \hat{\sigma}} \frac{s_3 - \hat{\sigma}}{s_3 - \sigma} & \text{if } \sigma \in [0, \hat{\sigma}) \\ G_1(\sigma) & \text{if } \sigma \in [\hat{\sigma}, z_2] \end{cases} \\ H_2(\sigma) &= \begin{cases} G_2(\hat{\sigma}) & \text{if } \sigma \in [0, \hat{\sigma}) \\ G_2(\sigma) & \text{if } \sigma \in [\hat{\sigma}, z_2] \end{cases} \\ H_3(\sigma) &= \begin{cases} \frac{G_2(\sigma)}{G_2(\hat{\sigma})} = \frac{s_1 - \hat{\sigma}}{s_1 - \sigma} & \text{if } \sigma \in [0, \hat{\sigma}) \\ 1 & \text{if } \sigma \in [\hat{\sigma}, z_2] \end{cases} \end{aligned}$$

To confirm this, write $H_{-1}(\sigma) = H_2(\sigma)H_3(\sigma)$ for the distribution of the most generous rival surplus offer faced by Firm 1, and similarly for $H_{-2}(\sigma)$ and $H_{-3}(\sigma)$. By construction, $H_{-1}(\sigma) = G_2(\sigma)$ for $\sigma \in [0, z_2]$. But we know from the Firm 1/Firm 2 equilibrium that this makes Firm 1 indifferent over all surplus offers in $[0, z_2]$ (and thus willing to mix over this range). Firm 3 faces best opponent offer distribution

$$H_{-3}(\sigma) = \begin{cases} q_3^0(\sigma) & \text{if } \sigma \in [0, \hat{\sigma}) \\ G_1(\sigma)G_2(\sigma) & \text{if } \sigma \in [\hat{\sigma}, z_2] \end{cases}$$

But as we argued earlier, any offer σ by Firm 3 that wins with probability $q_3^0(\sigma)$ earns it zero profit, while for $\sigma > \hat{\sigma}$, an offer that wins with probability $G_1(\sigma)G_2(\sigma)$

loses money. Thus mixing over $[0, \hat{\sigma})$ is a best reply for Firm 3. Finally, Firm 2 faces best opponent offer distribution

$$H_{-2}(\sigma) = \begin{cases} \frac{q_3^0(\sigma) G_2(\sigma)}{G_2(\hat{\sigma}) G_2(\hat{\sigma})} & \text{if } \sigma \in [0, \hat{\sigma}) \\ G_1(\sigma) & \text{if } \sigma \in [\hat{\sigma}, z_2] \end{cases}$$

From the Firm 1/Firm 2 equilibrium, we know that any offer $\sigma \geq \hat{\sigma}$ will earn Firm 2 zero profit. Furthermore, the threshold $\hat{\sigma}$ is defined by the fact that $q_3^0(\sigma) < G_1(\sigma) G_2(\sigma)$ for $\sigma < \hat{\sigma}$. Thus for $\sigma < \hat{\sigma}$,

$$H_{-2}(\sigma) = \frac{q_3^0(\sigma) G_2(\sigma)}{G_2(\hat{\sigma}) G_2(\hat{\sigma})} < G_1(\sigma) \left(\frac{G_2(\sigma)}{G_2(\hat{\sigma})} \right)^2 < G_1(\sigma)$$

Because Firm 2 needs to win with probability $G_1(\sigma)$ to break even, any offer $\sigma < \hat{\sigma}$ will lose money. So mixing over $[\hat{\sigma}, z_2]$ is a best reply for Firm 2. Thus the stipulated strategies are a Nash equilibrium, as claimed.

The probabilities that Firm 2 and 3 advertise can be recovered as $a_2 = 1 - H_2(0) = 1 - G_2(\hat{\sigma})$ and $a_3 = 1 - H_3(0) = \frac{\hat{\sigma}}{s_1}$ respectively. This represents a cutback for Firm 2 since it would have advertised with probability $1 - G_2(0)$ if Firm 3 were absent. However, the probability that neither competitor to Firm 1 advertises, $(1 - a_2)(1 - a_3) = G_2(\hat{\sigma}) q_3^0(0) = \frac{A + \Delta s}{s_1} = G_2(0)$, is exactly what it would have been if Firm 3 were absent. (Of course this should not be too surprising since the aggregate competition for Firm 1 is pinned down by its indifference condition, regardless of how many rivals it has.) For Firm 1, because $q_3^0(\sigma) < G_1(\sigma) G_2(\sigma)$ for $\sigma < \hat{\sigma}$, we have

$$H_1(\sigma) = \frac{q_3^0(\sigma)}{G_2(\hat{\sigma})} < G_1(\sigma) \frac{G_2(\sigma)}{G_2(\hat{\sigma})} < G_1(\sigma) \quad \text{for } \sigma < \hat{\sigma}$$

So over the range of high-price, low-surplus offers $\sigma \in [0, \hat{\sigma})$, the presence of Firm 3 induces Firm 1 to shift weight toward more competitive offers. In particular, the probability that Firm 1 advertises its monopoly price declines from $G_1(0)$ to $H_1(0)$.

Because $H_2(\sigma) H_3(\sigma) = G_2(\sigma)$, the best offer from Firms 2 and 3 is statistically equivalent to the best offer from Firm 2 in a game where Firm 3 is absent. Consumer

surplus can be derived from the distribution of the consumer's best offer, $H_{\max}(\sigma) = H_1(\sigma)H_2(\sigma)H_3(\sigma) = H_1(\sigma)G_2(\sigma)$. But then since $H_1(\sigma) \leq G_1(\sigma)$ (strictly on $[0, \hat{\sigma})$), $H_{\max}(\sigma) \leq G_1(\sigma)G_2(\sigma) = G_{\max}(\sigma)$. Thus consumer surplus improves, and this improvement can be attributed to more competitive offers by Firm 1 in the range $\sigma \in [0, \hat{\sigma})$ where it must compete with Firm 3. Since total firm profits do not change, welfare must rise as well. This improvement involves several effects. Allocative efficiency tends to rise, as Firm 1 competes harder and wins more often, but there is a countervailing effect because some of Firm 2's wins shift to the lower-surplus Firm 3. It is not too hard to show that total advertising is greater when Firm 3 is present. (Simple algebra establishes that if $(1 - a_2)(1 - a_3) = 1 - a_2^{old}$, then $a_2 + a_3 > a_2^{old}$.) However, advertising shifts from the higher cost Firm 2 to the lower cost Firm 3.

The equilibrium has a certain intuitive appeal: in effect the runner-up firms keep the top firm honest and discipline its profits by providing a sort of competitive upper envelope, each one turning up the heat on Firm 1 along the range where it has a comparative advantage. Furthermore, the implication that adding additional firms to the mix can only have a neutral to positive effect on welfare seems likely to be general: we know that total firm profits cannot change, nor can the most competitive non-Firm 1 offers (by Firm 1's indifference condition), but the additional competition may induce better offers out of Firm 1, and hence greater consumer surplus. This would not be obvious *a priori*, since additional active firms also could be associated with misallocation of the sale to a lower-surplus firm and socially excessive advertising.