Price Discrimination in the Information Age: List Prices, Poaching, and Retention with Personalized Targeted Discounts

Simon Anderson, Alicia Baik, and Nathan Larson*

March 26, 2019

Abstract

We study list price competition when firms can individually target discounts (at a cost) to consumers afterwards. Consumers endogenously separate into two types: captives cannot be profitably poached, while contesteds receive poaching offers and retention counter-offers. In equilibrium, these discount offers are in mixed strategies, but only two firms vie for each contested consumer and the final profits on them are simple and Bertrand-like. More contestable consumers receive more ads and are more likely to buy the wrong product. Poaching exceeds retention when targeting is expensive, but this reverses when targeting is cheap. Firm list pricing resembles monopoly, as marginal consumers are lost to the lowest feasible poaching offer, not to another firm's list price. Cheaper targeting drives list prices higher if captive demand is convex but has no impact on list prices in a spatial context with linear-quadratic transport costs. Targeting improves aggregate consumer surplus if cheap, but may make every consumer worse-off if expensive. Cheaper targeting reduces profits unless demand is log-convex, in which case it can constitute a facilitating practice. In the free-targeting limit, most consumers are better off with targeting than without.

Keywords: targeted advertising, competitive price discrimination, discounting

JEL: D43, L12, L13, M37

*Anderson and Baik: University of Virginia, Department of Economics; Larson: American University, Department of Economics. We thank Régis Renault, Qihong Liu, Jose-Luis Moraga, Kostas Serfes, and Yossi Spiegel for their comments, as well as participants at seminars at MIT Sloan, at the 8th and 9th Economics of Advertising and Marketing Conferences in Tel-Aviv (June 2013) and Nuffield College, Oxford (June 2015) respectively, and the CRETE conference in Naxos (July 2013). The first and second authors gratefully acknowledge support from the NSF and Bankard Fund for Political Economy respectively.

1 Introduction

With advances in information technology, firms are increasingly awash in data about individual consumers' tastes. We study the implications for competition when firms can use this information to send targeted advertisements with personalized price offers to individual consumers. Unusually for this literature, our results do not depend heavily on specific functional form assumptions; indeed, our main results apply to any distribution of consumer tastes (satisfying mild regularity conditions) and any number of (symmetric) firms.

In our model, firms first set public list prices for the differentiated products they sell. But then, at a cost, a firm can identify consumers with specific *taste profiles* and send them individualized discount offers. A consumer's taste profile is the list of her valuations for all the products on sale; thus firms are assumed to be able to target with pinpoint precision. This exaggerates the truth, of course, but by less and less as databases continue to grow and data-mining analytics continue to improve.[1] In practice, a firm can already pay to be connected to consumers who match a list of criteria; our premise is that by specifying a sufficiently detailed list, a firm can pin down the taste profile of the consumers it is matched to.[2]

The exogenous cost of a targeted ad, which is common to all firms, is intended to represent the total costs of identifying a desired consumer, formulating a customized offer, and delivering that offer to her. These details are left in the background, but for motivation one could consider two scenarios. One is that firms are served by competitive data brokers (or ad platforms) who are able to match them to consumers with any particular profile at cost.[3] Another is that firms identify consumers with the desired profiles from their own databases, in which case the targeting cost reflects the internal cost of data processing, formulating an optimal discount offer, and delivering

[1] This growth has been chronicled (often breathlessly) in the popular press. For examples involving Kroger, Target Corporation, and Gilt Group, see Clifford, Stephanie. "Shopper Alert: Price May Drop for You Alone," New York Times, August 9, 2012; Duhigg, Charles. "How Companies Learn your Secrets, "New York Times, February 16, 2012; Vega, Tanzina. "Online Data Helping Campaigns Customize Ads," New York Times, February 20, 2012; Andrews, Lori, "Facebook is Using You," New York Times, February 4, 2012.

[2] For example, TowerData offers to deliver consumers who match up to 34 demographic and taste criteria. See www.towerdata.com

[3] In this case the service might include ad placement (as in the case of paid search advertising), or just contact information (e.g., an email address), leaving the cost of advertising up to the firm. It would be a natural extension to give these intermediaries market power and bring them into the analysis explicitly, but that is not our focus here.

it.[4] Advertising conveys only (discount) price information; consumers are presumed to already know how much they like the products on offer, and list prices are public.

When targeting is cheap enough to be used, equilibrium competition endogenously sorts out which consumers will be *captive* and which will be *contested* with targeted discounts. The former, for lack of better offers, buy their favorite products at list price. Meanwhile, each contested consumer is fought over individually by her top two firms (that is, the firms making her first and second-favorite products): her second-favorite tries to poach her business with undercutting offers, and her favorite advertises to try to retain her. The combination of list prices and discounts amounts to competitive price discrimination, with the novel feature that because of precision targeting, price discrimination over the contested region is first degree.

The expected profits on a consumer who is contested with targeted ads turn out to be Bertrand-like: her favorite firm earns its value advantage over the runner-up, her second-favorite firm earns zero, and no other firm bothers to advertise to her. However, the second-best firm must win the sale with positive probability (since it would not pay to advertise otherwise), so the discounting equilibrium will involve mixed strategies and is allocatively inefficient.

These Bertrand-like outcomes in the discounting stage simplify the firms' reduced-form profits in the first stage, when they set list prices. Our marginal profit expressions are a main novelty of the paper. A firm faces a familiar marginal-inframarginal tradeoff in pricing to its captive consumers, with one catch: the downside of pricing out a marginal captive consumer is not the full profit margin lost on her, but just the cost of the targeted ad that will be needed to win her back (at a small discount). Furthermore, because the buffer zone of contested consumers means that list prices never compete against each other head-to-head, a firm's list price choice simplifies to a (quasi-)monopoly problem.

Because a firm's list price must sometimes compete against rivals' discounts, the analysis hinges on a firm's *captive demand function* $1 - G(y)$: the measure of consumers who prefer its product by at least y dollars over their next best alternatives. This captive demand function may be derived from whatever primitive assumptions

[4] The symmetric advertising cost rules out scenarios in which, for example, a firm has an advantage over its rivals in targeting its own past customers. While this type of informational advantage would be of great interest, we will have more than enough to say as it is. However, it would not be difficult to extend our model so that the cost of targeting varies across firms and across consumer types.

one prefers about the underlying consumer taste distribution.[5] The appeal of our approach is that the fine details of primitive tastes may be left in the background: all of the important features of competition depend only on the captive demand function, and our main qualitative results hold for *any* underlying distribution of tastes satisfying mild conditions on $1 - G(y)$.

Precision targeted advertising is likely to become cheaper and more prevalent, and our main results focus on identifying the winners and losers from that trend. Firms typically suffer, as list prices rise but profits fall (Propositions 2 and 3).[6] Cheaper targeting means that firms cannot resist the temptation to poach deeper into each others' territories – higher list prices are mainly a symptom of the fact that a firm must retreat deeper into its own territory in order to make an uncontested list price sale.

There is an exception: cheaper targeting will soften competition and facilitate higher profits if captive demand is logconvex (Proposition 4). In this case, the option to target can encourage an abrupt shift in strategy to much higher list prices and "discounts" higher than the original list prices. *Inter alia*, this suggests that targeting may tend to look more favorable for firms in models with discrete consumer types, since regions of logconvex demand are an unavoidable byproduct in such models.

We assume throughout that the market is fully covered (see the further discussion in the conclusions). This assumption enables us to focus cleanly on pure competition effects by closing down the well-understood market expansion effect of reaching additional consumers through discriminatory pricing. Hence, equilibrium welfare is first-best when targeted ads are not in use. As targeting becomes cheaper, welfare declines at first but later improves, approaching the first-best level again as targeting costs vanish. This implies that consumers will eventually be better off with targeted advertising than without it (firms' losses become their gains when welfare losses vanish), even though targeting often hurts consumers when it is first introduced (Proposition 8). But there are distributional consequences: as firms shift toward head-to-head competition for individual consumers, consumer surplus shifts away from people with strong favorite products toward those with strong *second*-favorite products.

Welfare losses can be attributed to the direct costs of advertising, and the misal-

[5] We use the primitive tastes associated with Hotelling competition and multinomial choice as running examples to illustrate how to accommodate spatial or non-spatial product differentiation.

[6] These results apply when captive demand is strictly convex, a mild condition that will hold for most taste distributions of interest. For instructive counterexamples, see Section 5.3.

location costs of consumers buying their second-best products. We show that more fickle consumers (those who are closer to indifferent between their top two products) receive more ads and are more likely to buy the "wrong" product than less fickle consumers. In fact, due to a tendency by firms to favor poaching over retention, a subset of consumers will buy the wrong product more often than not.

Our paper is at the juncture of the literatures on informative targeted advertising and competitive price discrimination. In seminal papers (including Butters, 1977, Grossman and Shapiro, 1984, and Stahl, 1994), informative advertising has typically meant that consumers learn about both products and prices from ads; in contrast, we assume away costs of publicizing products and list prices in order to sharpen the focus on discount advertising. Targeting permits firms to address different market segments with different levels of product information, and perhaps different prices. Duopoly examples with homogeneous products include Galeotti and Moraga-González (2008) (with no price discrimination and fixed market segments) and Roy (2000) (with tacit collusion on an endogenous split of the market). Differentiated product models based on Varian's (1980) Model of Sales (with consumers exogenously segmented into captive "loyals" and price-elastic "shoppers") include Iyer et al. (2005) (where targeting saves firms from wasted advertising) and Chen et al. (2001) (where errors in targeting help to soften price competition), and Esteves and Resende (2016) (who break the loyal/shopper dichotomy with consumers who prefer one product but would switch for a sufficiently better price).[7] Several of these papers find that targeting may be profit-enhancing for some model parameters, but the specificity of the models (usually duopolies with restrictive specifications of consumer tastes) makes it difficult to discern general conclusions, and the demand curvature channel that we highlight is novel. In our concluding remarks we offer some thoughts about how to reconcile our conclusions about profits with the varied claims in the literature.

Another branch of the literature examines oligopoly price discrimination when consumers can be informed about prices without costly advertising. One strand, dating to Hoover (1937) and more recently to Lederer and Hurter (1986) and Thisse and Vives (1988) focuses on spatial competition.[8] Thisse and Vives consider duopolists who can charge location-specific prices to consumers. As location is the dimension

[7] See also Brahim et al. (2011). In a monopoly setting with differentiated consumer tastes, Esteban et al. (2001) develops a different notion of targeting precision based on nested subsets of consumers.

[8] See also Anderson and de Palma (1988) and Anderson, de Palma, and Thisse (1989).

along which consumer preferences vary, this permits individualized pricing similar to that in our paper (but without costly advertising), and they reach some similar conclusions (including Bertrand-like competition for contested consumers, with the consequence that competitive price discrimination hurts profits).

In contrast, in Fudenberg and Tirole (2000) and Villas-Boas (1999), it is not a consumer's location that a firm observes but whether she is one of its own past customers. This permits firms to (coarsely) segment consumers and try to poach the rival's past customers with discounts. In broad strokes, this pattern of using discounts to poach is similar to what we describe, but these papers focus more on dynamic effects (namely, how the contracts offered to win customers today are colored by the fact that their purchase histories will be used in pricing tomorrow), where we are focused more on how the costs of targeting affect list prices.[9]

Our two stages of price-setting (list prices followed by discounts) are most similar to prior work on couponing, including Shaffer and Zhang (1995, 2002) and Bester and Petrakis (1995, 1996). In particular, Bester and Petrakis (1996) share our structure of public list prices and costly discount ads and find that the option to send coupons reduces both profits and prices. However in other respects their model is quite different from ours – targeting is coarse (two market segments), there is no retention advertising, and the result on prices is somewhat forced by special assumptions about tastes.[10]

As noted above, we find that discounting and advertising strategies must be mixed, since when several firms advertise, Bertrand competition prevents more than one of them from recovering its ad cost. This idea comes up in other settings where there is a winner-take-all competition in which losers incur participation costs that they cannot recover.[11] In related work (Anderson, Baik, and Larson, 2015), we give a

[9] Extending an older literature on intertemporal price discrimination, Acquisti and Varian (2005) show that a monopolist often will not benefit from price discrimination based on past purchases if consumers are sophisticated and respond strategically. Other work on competitive price discrimination includes Corts (1998) and Armstrong and Vickers (2001).

[10] In particular, consumers have differentiated tastes for the "non-local" firm but homogeneous tastes for the local firm. The latter implies monopoly pricing to the locals when coupons are absent, so coupons have no scope to do anything other than push all prices (list and discounted) down.

[11] Examples include all-pay auctions (Hillman and Riley, 1989) and entry games followed by Bertrand competition (Sharkey and Sibley, 1993). With a suitable interpretation, variations on Varian's Model of Sales also have this structure. (Let the price-sensitive segment be the "prize," and let foregone profits on loyals – due to pricing below their reservation values – be the cost of competing.) See Narasimham (1988) for a duopoly analysis and Koçaş and Kiyak (2006) for oligopoly.

more comprehensive treatment of competition for an individual consumer *via* costly ads. While we are unable to simply cite these results – the presence of list prices introduces some subtle differences that must be addressed – our analysis and results in Section 3 draw heavily on arguments from that paper.

Section 2 describes the model. Section 3 solves the second stage of the game, competition in targeted discounts. The key step toward characterizing an overall equilibrium is the conclusion that profits on a contested consumer will be Bertrand-like. For a reader willing to take this on faith, there is no harm in skipping ahead (and referring back later, as necessary, for the results on advertising and consumer surplus). Using these results, Section 4 analyzes the first-stage competition in list prices and characterizes the symmetric equilibrium. Sections 5 and 6 present our results for prices and profits, welfare, and consumer surplus. Section 8 concludes with suggestions for future work. Proofs omitted from the main text appear in the Appendix.

2 Model

Each of n firms produces a single differentiated product at marginal cost normalized to zero, to be sold to a unit mass of consumers. Each consumer wishes to buy one product; consumer i's reservation value for Firm j's product is r_{ij}. Later we will discuss the primitive distribution of these consumer tastes. For now it will suffice to define a distribution function $G_j(y)$, $y \in [\underline{y}, \bar{y}]$ for each firm, where $1 - G_j(y)$ is the fraction of consumers who prefer product j over their best alternative product (among the $n-1$ other firms) by at least y dollars. (We permit the possibility of $\bar{y} = \infty$, $\underline{y} = -\infty$.) Formally, if $\hat{r}_{i,-j} = \max_{j' \in \{1,\ldots,n\} \setminus j} r_{ij'}$, then

$$G_j(y) = |\{i \mid r_{ij} \leq \hat{r}_{i,-j} + y\}|$$

Later, $1 - G_j(y)$ will be seen to be closely related to Firm j's demand. We will generally impose primitive conditions that ensure the following:

Condition 1 *The functions $1-G_j(y)$ and the densities $g_j(y)$ are strictly logconcave.*

Condition 2 *The functions $G_j(y)$ are symmetric: $G_j(y) = G(y)$ for all $j \in \{1, ..., n\}$.*

There are two stages of competition. In Stage 1, the firms simultaneously set publicly observed list prices p_j^l that apply to all consumers. Then in Stage 2, firms can send targeted discount price offers: for each consumer i, Firm j may choose to send an advertisement at cost A offering her an individualized price $p_{ij}^d \leq p_j^l$. One interpretation is that firms initially know the distribution of tastes, but cannot identify which consumers have which valuations. For example, Firm j understands that consumers with the taste profile $(r_{i1}, r_{i2}, ..., r_{ij}, ...)$ exist, but it does not know who they are or how to reach them. Then A is the cost of acquiring contact information for consumers with this taste profile (through in-house research or by purchase from a data broker), plus the cost of reaching them with a personalized ad.

Finally, each consumer purchases one unit at the firm that offers her the greatest net consumer surplus; consumer i's surplus at Firm j is r_{ij} minus the lowest price offer Firm j has made to her. We assume that if a consumer is indifferent between two list prices, or between two advertised prices, she chooses randomly. However, if she is indifferent between one firm's list price and another's advertised discount price, she chooses the advertised offer. This tie-breaking assumption is motivated the fact that ads are sent after observing list prices, so an advertiser that feared losing an indifferent consumer could always ensure the sale by improving its discount offer slightly. Note that because products are differentiated, an undercutting offer is one that delivers more surplus to a consumer than rival firms' offers.

We assume that consumers' outside options are sufficiently low that they always purchase some product, that is, the market is fully covered. While this assumption is commonly imposed in the literature, it has a bit more bite here because equilibrium list prices may rise as the ad cost A falls. We discuss the implications of allowing outside options to bind in the conclusion. We say that consumer i is *on the turf* of Firm j if it makes her favorite product; that is, if $r_{ij} > r_{ik}$ for all $k \neq j$. She is on a turf boundary if she is indifferent between her two favorite products. Finally, we say that product j is her *default* product if it is the one she would buy at list prices, that is, if $r_{ij} - p_j^l > r_{ik} - p_k^l$ for all $k \neq j$.

To illustrate how the reduced-form distribution $G(y)$ may be derived from underlying consumer tastes, we present two settings that will be used as running examples.

Example 1: Two-firm Hotelling competition (with linear transport costs)

Firms 1 and 2 are at locations $x = 0$ and $x = 1$ on a Hotelling line, with consumers uniformly distributed at locations $x \in [0, 1]$. We refer to a consumer by location x

rather than index i. A consumer's taste for a product at distance d is $R - T(d)$, with $T(d) = td$. Then the set of consumers who prefer Firm 1 by at least y dollars is those to the left of $\bar{x}$, where $\bar{x}$ satisfies $R - t\bar{x} = y + R - t(1 - \bar{x})$. Solving for $\bar{x}$, we have

$$1 - G(y) = \frac{1}{2} - \frac{1}{2t}y$$

The same expression applies for Firm 2, so no subscript on $G(y)$ is needed. In this case, $1 - G(y)$ but not $g(y)$ is strictly logconcave.[12] This setup generalizes easily to the case of n firms located on a circle.

Example 2: n firm multinomial choice (independent taste shocks)

There are n firms, and consumer i's taste r_{ij} for Firm j's product is drawn i.i.d. from the primitive distribution $F(r)$ with support $[\underline{r}, \bar{r}]$.[13] Except where otherwise noted, assume that $F(r)$ and its density $f(r)$ are both strictly logconcave.

Condition on the event that a consumer's best alternative to Firm 1, over products $2, ..., n$, is r. Firm 1 beats this best alternative by at least y (that is, $r_{i1} \geq r + y$) with probability $1 - F(r + y)$. But the consumer's best draw over $n - 1$ alternatives has distribution $F_{(1:n-1)}(r) = F(r)^{n-1}$, so we have:

$$1 - G(y) = \int_{\underline{r}}^{\bar{r}} (1 - F(r + y))\, dF_{(1:n-1)}(r) \tag{1}$$

Conveniently, $1 - G(y)$ inherits the logconcavity of the primitive taste distribution.

Lemma 1 *Strict logconcavity of $f(r)$ implies that the functions $G(y)$, $1 - G(y)$, and $g(y) = G'(y)$ are strictly logconcave as well.*

Without targeted ads, this is a standard multinomial choice model (see e.g. Perloff and Salop, 1985). If the taste shocks are Type 1 extreme value, then we have the multinomial logit model that is widely used in empirical analysis.[14] The novelty in our setting is that a firm does not have to settle for treating these taste shocks as unobserved noise – at a cost, it can target customized offers to consumers with particular taste profiles.

[12] For non-linear transport costs $T(d)$, the analogous condition is that $1 - G(y) = \bar{x}$, where $\bar{x}$ satisfies $r_{\bar{x}1} - r_{\bar{x}2} = T(1 - \bar{x}) - T(\bar{x}) = y$. Thus $G(y)$ is defined implicitly by $T(G(y)) - T(1 - G(y)) = y$. One can confirm that logconcavity of $1 - G(y)$ is satisfied if $x(T'(x) + T'(1 - x))$ is increasing.

[13] We allow for the possibility that $\bar{r} = \infty$ or $\underline{r} = -\infty$.

[14] For theoretical applications see Anderson, de Palma, and Thisse (1992).

The key difference between Examples 1 and 2 is the correlation pattern of consumer tastes across products. In Example 1, consumer tastes for the two products exhibit perfect negative correlation, while in Example 2 tastes are uncorrelated. While our model may be applied to arbitrary distributions of consumer tastes, these two cases encompass many of the settings that are commonly used in the literature.

Next we analyze the targeted advertising subgame in Stage 2.

3 Stage 2: Competition in Targeted Discounts

A firm decides separately for each consumer whether to send a discount ad and, if so, what price to offer. Thus Stage 2 constitutes a collection of independent price competition games for individual consumers. For brevity, we discuss this price competition game for an arbitrary consumer when all, or all but one, of the Stage 1 list prices are symmetric, as these cases govern incentives in the symmetric equilibrium of the full game. A formal analysis of the general case appears in the Appendix.

Consider an arbitrary consumer taste profile $\mathbf{r} = (r_1, r_2, ..., r_n)$. Let $y_j = r_j - \max_{k \neq j} r_k$ be Firm j's value advantage (possibly negative) for this consumer relative to her best alternative product. Firm j chooses a probability a_j of sending an ad to consumers with this taste and a distribution p_j^d over the discount price offered in that ad. For any taste profile, let $r_{(1)} > r_{(2)} > ... > r_{(n)}$ be the relabeling of firms so that Firm (1) makes this consumer's favorite product, Firm (2) makes her second favorite, and so on.[15] An important role is played by this consumer's value advantage for her favorite product, $y_{(1)} = r_{(1)} - r_{(2)}$.

A consumer is said to be *captive* to her default firm if no other firm advertises to her with positive probability. She is *contested* if two or more firms send her ads with positive probability. These will be the only outcomes on the equilibrium path, but off-the-path she could also be *conceded* if exactly one firm advertises to her.

3.1 Targeting with symmetric list prices

Suppose the firms have set list prices $p_1^l = p_2^l = ... = p_n^l = p$ at Stage 1. If $p \leq A$, then no ads will be sent, since any firm advertising a discount below p would not

[15] For smooth taste distributions, consumers who are indifferent between two or more products have zero-measure, and have no impact on profits or list price decisions, so we can ignore them.

recover the cost of sending the targeted ad. Then consumers will remain captive to their default firms which (given symmetric list prices) are also their favorite firms. In this section we drop the subscript and write $y_{(1)} = y$ to reduce clutter.

The more interesting case is when $p > A$ so that firms can afford to discount. A consumer's favorite product is still her default. Any rival firm would need to discount by at least y to poach her business with a targeted ad, and this cannot be profitable if $p - y < A$, so consumers with large taste advantages, $y > p - A$, will remain captive to their favorite firms. Any consumer with a smaller taste advantage, $y < p - A$, cannot be captive in equilibrium, since her second favorite firm could profitably poach her with an undercutting offer $p^d_{(2)} < p - y$. In this case, the consumer must be contested rather than conceded: the poacher will offer a minimal discount if it does not expect competition, but this would permit her default firm to retain her with a minimal discount of its own. Since winning her at a price of $p - y$ was profitable for the poacher (who is at a value disadvantage), retaining her at a price just below p is surely profitable for her default firm.

When a consumer is contested, the number of firms vying for her and the probabilities that each sends an ad are limited by the firms' need to cover the cost of a targeted ad. First of all, there cannot be two or more firms sending her targeted ads with probability one – standard Bertrand undercutting would rule out an equilibrium at any prices high enough for all of the firms to cover A. This helps to pin down the firms' final expected profits on this consumer. Her default firm can ensure a net profit on her of at least its value advantage y by advertising the discount price $p^d_{(1)} = y + A$, which no rival can profitably undercut. All other firms must earn zero profit on her. (Strictly positive profits would oblige them to advertise with probability one which is not viable.)

Competition from the consumer's second-best firm places limits on the profit of her best firm and rules out discounting by any lower-ranked firm. If any lower-ranked Firm j could make a non-negative profit by advertising with positive probability, then Firm (2) could earn a strictly positive profit by undercutting Firm j's lowest discount offer (winning the sale just as often, but earning a larger profit margin due to its value advantage). Similarly, if a consumer's favorite firm were to earn strictly more than y from her, then its discount offers would all satisfy $p^d_{(1)} > y + A$, but if this were true, Firm (2) could earn a strictly positive profit by undercutting the lowest such offer. Thus only the consumer's top two firms contest her, with (Bertrand-like) profits $\pi_{(1)} = y$ and $\pi_{(2)} = 0$ respectively.

In equilibrium, these top two firms must mix over discount prices (as pure strategies would provoke undercutting that would prevent Firm (2) from covering its ad cost). Note that any discount offer $p^d_{(1)}$ by Firm (1) may be regarded as a consumer surplus offer $s_{(1)} = r_{(1)} - p^d_{(1)}$, and similarly for Firm (2). It is convenient to cast the firms' Stage 2 strategies in terms of these surplus offers, rolling the advertising probability a_j and the distribution of discounts together by regarding "not advertising" as a surplus offer $s^l_j = r_j - p$ at the original list price. Let $B_{(1)}(s)$ and $B_{(2)}(s)$ be the distributions of these surplus offers by firms (1) and (2). A discount just below Firm (1)'s list price corresponds to $s = r_{(1)} - p$, while the most generous surplus that Firm (2) could afford to advertise is $s = r_{(2)} - A$. Standard arguments ensure that the firms' discount offers have common support over this full range without gaps or (with one exception) atoms. Because each firm must earn the same expected profit ($\pi_{(1)}(s) = B_{(2)}(s)\left(r_{(1)} - s\right) - A = y$ or $\pi_{(2)}(s) = B_{(1)}(s)\left(r_{(2)} - s\right) - A = 0$ respectively) on every surplus offer, we arrive at the following Stage 2 equilibrium strategies for a particular consumer.

Firm (1) *sends no ad with probability* $B_{(1)}\left(s^l_{(1)}\right) = 1 - a_{(1)} = \frac{A}{p-y}$. *Its discount offers are distributed* $B_{(1)}(s) = \frac{A}{r_{(2)}-s}$ *on support* $(s^l_{(1)}, r_{(2)} - A]$.

Firm (2) *sends no ad with probability* $B_{(2)}\left(s^l_{(2)}\right) = 1 - a_{(2)} = \frac{y}{p}$. *Its discount offers are distributed* $B_{(2)}(s) = \frac{y+A}{y+r_{(2)}-s}$ *on support* $\left[s^l_{(1)}, r_{(2)} - A\right]$. *These offers include an atom* $\frac{A}{p} = B_{(2)}\left(s^l_{(1)}\right) - B_{(2)}\left(s^l_{(2)}\right)$ *of offers at surplus* $s^l_{(1)}$, *just undercutting Firm* (1)*'s list price.*[16]

One novelty is the atom in Firm (2)'s strategy. Firm (1) must be willing to advertise discounts just below its list price, and to be worth the ad cost, those slight discounts must win substantially more often than its list price does. This is only true if Firm (2) frequently "cherrypicks" just below that list price offer.

3.2 Targeting after a list price deviation

The results above apply on the equilibrium path, but we need also to know how profitable it would be for a firm to deviate to a different list price. Suppose that Firm

[16] Recall the tie-breaking assumption that an advertised discount defeats a list price offering the same surplus.

1's list price is p_1 and all other firms set the same list price p. For our purposes, it will suffice to pin down the final profit of the deviator on an arbitrary consumer.

First consider which consumers will be captive to Firm 1. If $p \leq A$, then rival firms will not advertise. Firm 1's default consumers, and also its captives, will be those for whom its value advantage over the best alternative product exceeds its price differential: $y_1 > p_1 - p$. On the other hand, if $p > A$, then a default consumer is not safe from poaching unless the stronger condition $y_1 > p_1 - A$ holds, since the best alternative firm can offer discount prices as low as A in an attempt to win the consumer. Thus, Firm 1's captive consumers are those for whom $y_1 > p_1 - \min(p, A)$, and it earns profit p_1 on each of these.

Next consider Firm 1's profit on non-captives. If its list price is below A, it earns nothing on them since it cannot afford to advertise a discount. If p_1 and p both exceed A, then arguments very much like those of the last section apply. Any consumer who is not captive to some firm will be contested by the firms making her two favorite products, with Bertrand-like final profits that depend on the difference in their values but not on the original list prices. Thus, on any non-captive consumer who likes Firm 1's product best, $y_1 \in (0, p_1 - A)$, Firm 1 earns its value advantage y_1; otherwise it earns zero profit on this consumer. Finally, if $p_1 > A \geq p$, then Firm 1 will keep any consumers it can afford to poach since rivals cannot afford to retaliate. On a consumer where its value advantage is y_1, Firm 1 does best to set the "undercutting" price $p_1^d = y_1 + p$. This covers the ad cost if $y_1 > A - p$, so consumers with value advantage $y_1 \in (A - p, p_1 - p)$ are ultimately conceded to Firm 1 with net profit $\pi_1 = y_1 + p - A$.

Unifying principles The logic underlying the various cases is as follows. Consider the "almost symmetric" case where list prices are $(p_1, p, ..., p)$. Given list prices, let $P_{-1} = \min(p, A)$ be the *"last best price"* for Firm 1's rivals; this is the most competitive offer (list or discount) a rival could afford to make. Then Firm 1's captives are always those consumers who cannot be tempted away from its list price by their best alternative firm's last best price: $y_1 > p_1 - P_{-1}$. And if Firm 1 can afford to advertise a discount ($p_1 > A$), then its potential profit on a non-captive with value advantage y_1 depends on what it would earn by undercutting the last best price of the consumer's best alternative. That implies a price $y_1 + P_{-1}$, and a net profit $\pi_1 = y_1 + P_{-1} - A$. On non-captives where this potential net profit is positive, namely $y_1 \in (A - P_{-1}, p_1 - P_{-1})$, this is what Firm 1 earns in equilibrium. And on consumers where this potential profit would be negative, Firm 1 earns zero profit. We use this unified characterization in the profit expression (2) developed below.

4 Stage 1: Competition in List Prices

Now we turn to the determination of list prices in Stage 1, focusing on symmetric equilibria with list price p^l.[17] Suppose firms 2 through n are all expected to price at p^l in Stage 1, and examine the incentives of Firm 1 in setting its own list price p_1^l.

Because Firm 1's value advantage y_1 is distributed according to $G(y)$, the summary of the last paragraph of Section 3 implies that it serves $1 - G\left(p_1^l - P_{-1}\right)$ captive consumers at its list price. If $p_1^l \leq A$, these are its only consumers; otherwise it also earns the Stage 2 expected profit $y_1 + P_{-1} - A$ on non-captives for whom $y_1 \in (A - P_{-1}, p_1 - P_{-1})$. In summary, Firm 1's overall expected profit is:

$$\Pi_1\left(p_1^l, p^l\right) = \begin{cases} p_1^l\left(1 - G\left(p_1^l - P_{-1}\right)\right) & \text{if } p_1^l \leq A \\ p_1^l\left(1 - G\left(p_1^l - P_{-1}\right)\right) + \int_{A-P_{-1}}^{p_1^l - P_{-1}} (y + P_{-1} - A)\, dG(y) & \text{if } p_1^l > A \end{cases} \tag{2}$$

Using $P_1 = \min\left(p_1^l, A\right)$ for Firm 1's own last best price, the two piecewise expressions may be consolidated to write Firm 1's marginal profit, and its first-order condition for an interior optimum, as:

$$\frac{\partial \Pi_1\left(p_1^l\right)}{\partial p_1^l} = 1 - G\left(p_1^l - P_{-1}\right) - P_1 g\left(p_1^l - P_{-1}\right) = 0 \tag{3}$$

There is a strong structural resemblance to the marginal profit expressions that are typical of other oligopoly models, but with two key differences. First, the margin at which list price sales are lost is determined by the condition $y = p_1^l - P_{-1}$: this is the consumer who weakly prefers Firm 1's list price to the *last best* offer – list price or lowest advertised discount – of any other firm. As usual, raising one's list price generates a gain on inframarginal consumers, in this case, the $1 - G_j\left(p_j^l - P_{-j}\right)$ consumers who buy at Firm 1's list price. Also as usual, the tradeoff of hiking one's price is that marginal list price sales are lost, in this case at rate $g\left(p_1^l - P_{-1}\right)$. The second key difference lies in the sacrificed profit P_1 per lost marginal sale which

[17] This is natural, given the symmetry of the model. With two firms, it is straightforward to rule out asymmetric equilibria, so the symmetric equilibrium is unique. This seems likely to extend to more than two firms (perhaps under additional regularity conditions), but we do not have a proof.

depends on whether Firm 1 is willing to advertise to win that sale back. It cannot afford to if $p_1^l \leq A$; in this case the sacrifice at the margin is the full list price p_1^l. But if $p_1^l > A$, the "lost" marginal sale is not truly lost. She is lost to a rival who, with its most competitive possible offer, can barely make her happier than she would be at Firm 1's list price. To win her back, Firm 1 need do no more than advertise just below its list price; thus the profit sacrificed by losing her as a captive is simply the ad cost A.[18]

Benchmark with no targeted ads If targeted advertising is impossible or prohibitively expensive, then the last best prices in (3) are simply the list prices p_1^l and p^l, and the model collapses to a standard one-stage game of price competition. Given the strict logconcavity of $1 - G(y)$, there is a standard symmetric equilibrium with price given by the first-order condition

$$p^{NT} = \frac{1 - G(0)}{g(0)} \tag{4}$$

More generally, the first-order condition (3) is equivalent to $\frac{1-G\left(p_1^l - P_{-1}\right)}{g\left(p_1^l - P_{-1}\right)} - P_1 = 0$. Define a function $\Theta(p)$ equal to the lefthand side of this expression, evaluated at the strategy profile in which *all* list prices are equal to p:

$$\Theta(p) = \frac{1 - G(p - \min(p, A))}{g(p - \min(p, A))} - \min(p, A) = \begin{cases} \frac{1-G(0)}{g(0)} - p & \text{if } p \leq A \\ \frac{1-G(p-A)}{g(p-A)} - A & \text{if } p > A \end{cases}$$

$\Theta(p)$ has the same sign as each firm's marginal profit and is strictly decreasing (as monotonicity of $\frac{1-G(y)}{g(y)}$ follows from Condition 1). Letting $h = \frac{1-G(\bar{y})}{g(\bar{y})}$ be the value of the inverse hazard rate at the largest possible value advantage, we have $\Theta(p) \to h - A$ as $p - A \to \bar{y}$.[19] If $h - A$ is negative (as it must be when $h = 0$), then there is always a unique solution to $\Theta(p) = 0$ identifying a symmetric interior equilibrium. Alternatively, if $A < h$ then $\Theta(p)$ is strictly positive whenever firms retain any captive consumers ($p - A < \bar{y}$). In this case, at any common price level at which the firms retain some captive consumers, each firm has an incentive to hike its list price relative

[18] As it happens, firm 1 will choose *not* to advertise to this marginal consumer in order to retain her. But it could – and its indifference about whether or not to try to retain her means that the logic described here remains relevant.

[19] If $\bar{y} = \infty$, define $h = \lim_{y\to\infty} \frac{1-G(y)}{g(y)} < \infty$. (The limit exists by monotone convergence.) Typical demand distributions satisfying Condition 1 will be sufficiently thin-tailed to have $h = 0$. However, if the tails of captive demand look exponential (as in the Type 1 extreme value case of Example 2), then h will be positive but finite.

to its rivals and send targeted ads to a broader set of consumers than they do.

Proposition 1 *Under Condition 1, there is a unique symmetric equilibrium. This is the unique equilibrium of the game if there are two firms. If $A \geq p^{NT}$, the common list price is p^{NT} and targeted discounts are not used. If $A \in \left(h, p^{NT}\right)$, the list price solves $\Theta\left(p^l\right) = 0$, targeting is used, and all non-captive consumers are contested by their top two firms. If $A < h$, then $p^l = \bar{y} + A$, and all but the most captive consumers are contested with targeting by their top two firms.*

We can now write the symmetric equilibrium profits a bit more simply than (2). In a regime where advertising is not used, each firm serves the $1 - G(0)$ fraction of consumers who are on its turf. But since every consumer has some favorite product, symmetry implies that $1 - G(0) = \frac{1}{n}$. Thus when $A \geq p^{NT}$, each firm's profit is $\Pi^{NT} = \frac{1}{n} p^{NT}$. Alternatively, if advertising is used, then the common first-order condition determining the equilibrium list price reduces to:

$$\frac{1 - G\left(p^l - A\right)}{g\left(p^l - A\right)} = A \tag{5}$$

For profits, the lower bound of the integral in (2) collapses to zero and we have

$$\Pi^{ad} = p^l \left(1 - G\left(p^l - A\right)\right) + \int_0^{p^l - A} y \; dG(y) \tag{6}$$

where the list price is given by (5). We refer to this case, where each firm has a positive measure of both captive and contested consumers, as an interior equilibrium. If $A < h$ so that all consumers are contested with targeted discounts, then the first term vanishes and we simply have $\Pi = \int_0^{\bar{y}} y \; dG(y)$ – that is, each firm earns its value advantage on the consumers who like its product best.

4.1 Examples

4.1.1 Hotelling

The left panel of Figure 1 illustrates Firm 1's profit and list price choice in the Hotelling setting for a case ($p_1^l, p_2^l > A$) where both firms can afford to advertise. The upper envelope is the total social surplus when a consumer at location x purchases from Firm 1 (r_{x1}, blue) or Firm 2 (r_{x2}, red). Consumer surplus from a list price

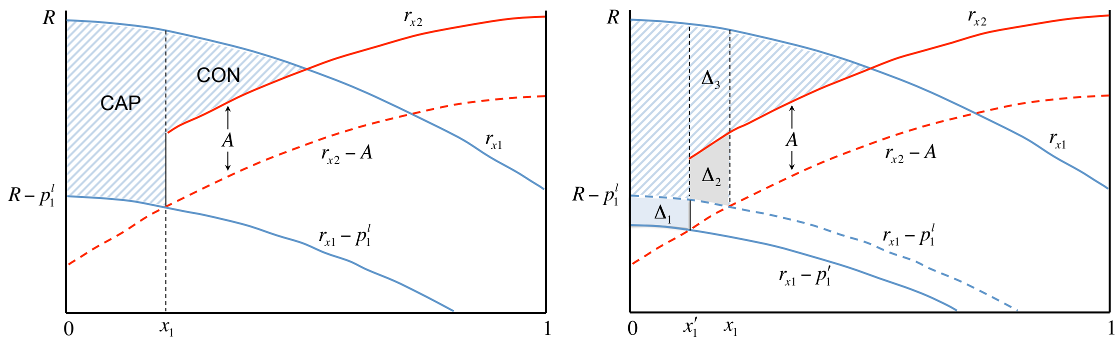

Figure 1: Hotelling competition – profits on captive and contested regions

purchase at Firm 1 is parallel to r_{x1} but shifted downward by p_1^l. Firm 1's marginal captive consumer x_1 is indifferent between this list price surplus and the most competitive surplus offer $r_{x2} - A$ that Firm 2 could make (dashed red). Firm 1 earns total profit $\Pi_1 = CAP + CON$, where $CAP = p_1^l x_1$ is the profit on captives, and CON represents the expected net profit $r_{x1} - r_{x2}$ on those contested consumers who favor Firm 1's product.

The right panel illustrates the change in profits when Firm 1 raises its list price to $p_1' = p_1^l + \Delta p$. The marginal captive consumers located in (x_1', x_1) will now be contested by Firm 2. Firm 1 gives up profit of $\Delta_2 \approx A\Delta x_1$ on these consumers, less than the lost profit $p_1^l \Delta x_1 = \Delta_2 + \Delta_3$ it would suffer if it could not win some of these sales back by discounting, and gains profit of $\Delta_1 = \Delta p \cdot x_1'$ on the inframarginal captives who remain. As drawn, $\Delta_1 \approx \Delta_2$, so Firm 1's initial list price is approximately optimal.

Firm 2's list price is notably absent from the diagram, as it plays no role in Firm 1's profit maximization decision (as long as $A < p_2^l$ so that Firm 2 can afford to discount). In this sense, Firm 1's position is similar to that of a limit-pricing monopolist, as it will use its list price to control how deep into its territory the incursions from rival discounting will be.

5 Equilibrium Prices, Profits, and Targeted Advertising

Our main results concern how the cost of targeted advertising affects firms' list prices, profits, and targeted advertising strategies.

5.1 List prices

Suppose that $A \in \left(h, p^{NT}\right)$ so that targeting is affordable and there is an interior symmetric equilibrium with list price characterized by (5); write this price as $p^l(A)$.

Proposition 2 *The equilibrium list price $p^l(A)$ is strictly decreasing (increasing) in the ad cost A if captive demand $1 - G(y)$ is strictly convex (respectively, strictly concave) for $y > 0$.*

Proof. Given $A < p^{NT}$, the equilibrium condition is $\Theta\left(p^l; A\right) = \frac{1-G\left(p^l-A\right)}{g\left(p^l-A\right)} - A = 0$, making the dependence on the parameter A explicit. Differentiate this equilibrium condition implicitly to get

$$\frac{dp^l(A)}{dA} = -\frac{\Theta_A}{\Theta_{p^l}} = -\frac{g'\left(p^l - A\right)}{\Theta_{p^l}} \frac{1 - G\left(p^l - A\right)}{g\left(p^l - A\right)^2}$$

But Θ_{p^l} is strictly negative (by strict logconcavity of $1 - G(y)$) so $dp^l(A)/dA$ has the same sign as $g'\left(p^l - A\right)$, establishing the claim. ■

Convex captive demand implies $g'(y) < 0$ (for $y > 0$), which means that a firm will tend have more consumers who prefer its product by a little bit than those who prefer it by a lot. This seems more empirically plausible than the alternative (unless tastes are strongly polarized), in which case cheaper targeting will usually tend to push up list prices.

This conclusion would not be very surprising for a monopolist blending list price sales to core customers with price-discriminating offers to a fringe. Cheaper price discrimination should induce it to substitute away from list price sales, thus moving up the demand curve to a higher list price. This substitution effect is present in our model, but with oligopoly there is a competitive effect that makes the final conclusion nontrivial. To illustrate we refer back to the marginal profit expression (3). When ads are in use, the effect of a reduction in A on Firm 1's incentive to hike its list price may be decomposed into a substitution effect operating through the reduction in Firm 1's own last best price P_1 and a competitive effect operating through the corresponding reduction in rivals' last best price P_{-1}. Formally, Firm 1's list price rises if its marginal profit rises, or $\partial^2\Pi_1/\partial p_1^l \partial(-A) = \partial^2\Pi_1/\partial p_1^l \partial(-P_1) + \partial^2\Pi_1/\partial p_1^l \partial(-P_{-1}) > 0$. Using

(3), the substitution effect

$$\frac{\partial^2 \Pi_1}{\partial p_1^l \partial\left(-P_1\right)} = g\left(p_1^l - P_{-1}\right)$$

is positive: the lower ad cost required to win back a marginal consumer lost to a list price increase makes such a price increase more attractive. This is the same incentive that a monopolist would face. However, the competitive effect

$$\frac{\partial^2 \Pi_1}{\partial p_1^l \partial\left(-P_{-1}\right)} = -g\left(p_1^l - P_{-1}\right) - P_1 g'\left(p_1^l - P_{-1}\right)$$

must be negative: a lower ad cost for Firm 1's rival permits it to reach deeper into Firm 1's territory with discount offers, inducing Firm 1 to shore up its flanks by cutting its list price.[20] If captive demand is linear, the second term in this expression drops out, and the substitution and competition effects perfectly offset each other – this is the list price neutrality case mentioned below. By comparison, convex captive demand tends to weaken the competitive effect because rivals find it harder to tempt the less price-sensitive consumers they find deeper in Firm 1's territory. As a result, the substitution effect dominates, and Proposition 2 ensues.

5.2 Profits

Because targeted advertising permits firms to compete on two fronts, one might suspect that it could facilitate higher profits by siphoning off competition for price-sensitive consumers, permitting the firms to maintain high margins on inframarginal consumers. Proposition 3 shows that this is generally wrong: cheaper targeting unambiguously makes firms worse off.

Proposition 3 *Suppose there is an interior equilibrium with profit* $\Pi(A)$ *in a neighborhood of* $\hat{A} < p^{NT}$*. Then* $\Pi'\left(\hat{A}\right) > 0$*; profits are strictly increasing in the ad cost.*

Proof. Write equilibrium profit using $y = p^l - A$ as the firm's strategic variable, with $y(A)$ its optimized level: $\Pi\left(y(A); A\right) = \left(A + y(A)\right)\left(1 - G\left(y(A)\right)\right) + \int_0^{y(A)} y \, dG(y)$. Then by the envelope theorem, $\frac{d\Pi}{dA} = \frac{\partial \Pi}{\partial A} = 1 - G\left(y(A)\right) > 0$. ■

[20] This term is unambiguously negative because it equals $\Pi_1''\left(p_1^l\right)$.

It is illuminating to separate out the own-cost effect from the competitive effect of a change in rivals' ad costs. Using profit expression (2) where own ad costs appear as A and rivals' ad costs appear as P_{-1}, we have $\frac{d\Pi_1}{dA} = \frac{\partial \Pi_1}{\partial A} + \frac{\partial \Pi_1}{\partial P_{-1}}$ (since $\frac{dP_{-1}}{dA} = 1$ when ads are in use). The own-cost effect is $\frac{\partial \Pi_1}{\partial A} = -CON_1$, where $CON_1 = G\left(p_1^l - P_{-1}\right) - G\left(A - P_{-1}\right)$ is the set of constested consumers on whom Firm 1 earns a positive profit. Holding rival prices constant, an increase in A comes out of Firm 1's margin on these consumers. However, the competitive effect is $\frac{\partial \Pi_1}{\partial P_{-1}} = Ag\left(p_1^l - P_{-1}\right) + CON_1$. The second term restores the profit margins on contested consumers, as higher rival ad costs perfectly balance the effect of higher own costs, and the first term softens competition at the captive-contested margin (since rivals cannot penetrate as deeply into Firm 1's territory with their discounts). Given the washout on contested consumers, the competitive effect dominates.[21]

In contrast, Proposition 4 shows that cheaper targeting can benefit firms if profit functions are not single-peaked. We say that a symmetric equilibrium exhibits full-targeting if the common list price is $p^l = \bar{y} + A$ and all interior consumers (those with value advantages $y < \bar{y}$) are contested with targeted discounts.

Proposition 4 *Suppose captive demand is strictly logconvex and (without ads) there exists a no-targeting equilibrium characterized by (4) and profit Π^{NT}. Whenever $A < A^*$ (for some threshold $A^* > p^{NT}$), the unique symmetric equilibrium has full-targeting and profits strictly higher than Π^{NT}.*

Proof. All claims up to the profit ranking are proved in the appendix. Profit in the no-targeting equilibrium is $\Pi^{NT} = \frac{1}{n} p^{NT} = \frac{1}{n} \frac{1-G(0)}{g(0)}$. In the full-targeting equilibrium when $A < A^*$, profit per firm is $\Pi^{FT} = \int_0^{\bar{y}} y dG\left(y\right)$ (regardless of A), or integrating by parts and using $1 - G\left(0\right) = \frac{1}{n}$, we have $\Pi^{FT} = \int_0^{\bar{y}} 1 - G\left(y\right) \; dy = \frac{1}{n} E\left(\frac{1-G(y)}{g(y)} \mid y \geq 0\right)$. As $\frac{1-G(y)}{g(y)}$ is strictly increasing, we have $\Pi^{FT} > \Pi^{NT}$. ■

If captive demand is logconvex, inframarginal captive consumers retain sizeable consumer surplus at uniform prices. As discounting becomes viable, at some point it becomes tempting for a firm to drastically shift strategies, essentially abandoning list price sales in order to capture all of that consumer surplus with targeted offers. By itself, that is not quite enough to explain the rise in profits. But note that $A^* > p^{NT}$, so the transition from no targeting to full targeting occurs while ads would be too expensive to use at the old list prices (but not at the new, higher list prices). This

[21] The overall effect $Ag\left(p_1^l - P_{-1}\right)$ matches Proposition 3 after applying equilibrium condition (5).

means that the transition from no targeting to full targeting as A declines will tend to *soften* rival firms' most competitive prices rather than sharpen them – even the lowest new discount prices will exceed the old list prices. And this permits profits to rise – Section 5.3 gives an example.

5.3 List prices and profits in our leading examples

5.3.1 Hotelling competition

List price neutrality under linear-quadratic transportation costs If transportation costs are $T(d) = \alpha d + \beta d^2$, with $\alpha + \beta = t$, then we have captive demand $1 - G(y) = \frac{1}{2} - \frac{y}{2t}$ and the standard result that $p^{NT} = t$ without advertising (using (4)). But because captive demand is linear, Proposition 2 implies that equilibrium list prices do not change with A when ads come into use: $p^l(A) = t$ regardless of A! *List price neutrality* to the cost of targeted advertising is only possible if captive demand is linear, that is if the density of consumers $g(y)$ who prefer their favorite product by y dollars does not fall with y. This is a rather special property; here it is possible because the primitive taste distributions are uniform and perfectly negatively correlated, so the difference in tastes is also uniformly distributed.

General nonlinear transportation costs Captive demand has the same curvature on $y \geq 0$ as the difference in transportation costs $T(1-x) - T(x)$ does on $x \in \left[0, \frac{1}{2}\right]$, so list prices will fall (rise) with A if $T(1-x) - T(x)$ is strictly convex (concave) on that range. This condition on the difference cannot be reduced (at least, not in a trivial way) to a condition on $T(d)$ itself. For example, consider the family of transportation costs $T(d) = d^\gamma$ for $\gamma > 0$. It is easily confirmed that the curvature of the transportation cost difference switches from convex (if $\gamma \in [0,1]$) to concave ($\gamma \in [1,2]$), then back to convex again ($\gamma \geq 2$). In this case, list prices will be decreasing in the targeting cost when the convexity of transport costs is low or high, but when they are moderately convex the relationship reverses (with list price neutrality at the switch points).

5.3.2 Independent taste shocks

In contrast, whenever tastes are distributed independently across products there is a bright line result.

Proposition 5 *If tastes are distributed i.i.d. according to strictly logconcave density $f(r)$, then captive demand is strictly convex (for $y > 0$), and so the symmetric equilibrium list price $p^l(A)$ is strictly decreasing in the targeted ad cost A whenever advertising is in use.*

Proof. Lemma 10 establishes that $g'(0) \leq 0$ when captive demand is derived from an i.i.d. taste distribution for *any* primitive density $f(r)$ (logconcave or not). Then because strict logconcavity of $f(r)$ is inherited by $g(y)$ (by Lemma 1), we have $g'(y) < 0$ for all $y > 0$. ■

The key difference relative to the Hotelling setting is independence. Given symmetry of tastes across products plus independence, higher densities of consumers will be found where taste differences are smaller – this is due partly to the centralizing effect of taking the difference of independent draws. If tastes are perfectly negatively correlated, as in the Hotelling setting, this centralizing effect is absent, so the density of consumers need not fall as taste differences become more extreme.

5.3.3 Example: rising profits when targeting is adopted

A two-firm example illustrates how firms can benefit from the introduction of targeting when tastes are not logconcave. Suppose there is a unit mass of linear-Hotelling consumers with $t = 1$ (and so value advantages $y \in [0, 1]$) and an additional unit mass of "loyals," split evenly between the firms, who prefer their favorite product by $\bar{y} = 2$.[22] Figure 2.a illustrates Firm 1's captive demand $1 - G\left(p_1^l - p^{NT}\right)$ when Firm 2 prices at the no-targeting equilibrium price $p^{NT} = 2$. Without advertising, Firm 1 is indifferent between also charging p^{NT} (point α) versus "retrenchment" to point β where it charges the higher price $p^H = p^{NT} + \bar{y} = 4$ and serves only its loyals. (This knife-edge is convenient but inessential.) When the ad cost falls below $\bar{A} = p^{NT} + 1 = 3$, the no-targeting equilibrium collapses since Firm 1 can retrench to list price p^H and mop up additional profits (area D) by targeted discounting to some of the regular consumers. By perfectly price discriminating, Firm 1 may extract the entire area beneath its captive demand curve as gross profit (since Firm 2 cannot afford to advertise); area D is its net profit after paying A to reach the consumers from whom it can extract A or more. Notice that both Firm 1's new list price and any

[22] Analysis to support the claims made here appears in the Supplementary Appendix.

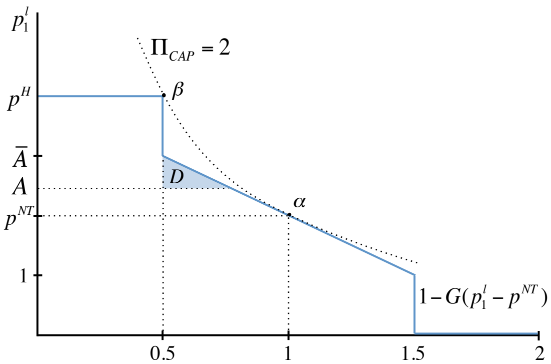

(a) No targeting by Firm 2

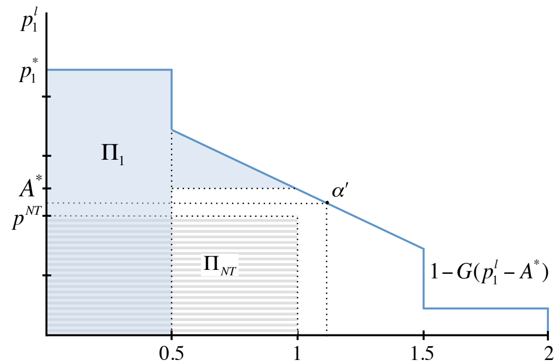

(b) $A = A^*$ and Firm 2 targets

Figure 2: Firm 1's captive demand

discounts $p_1^d \geq A$ represent softer competition for Firm 2 than when Firm 1 charged p^{NT}.

The full-targeting equilibrium discussed in Proposition 4 emerges for $A \leq A^* = \sqrt{6}$. (For $A \in \left(A^*, \bar{A}\right)$, equilibria involve mixed strategies with partial retrenchment – see the Supplementary Appendix for details.) Figure 2.b shows Firm 1's captive demand when $A = A^*$ and Firm 2 has retrenched to a high list price $p_2^l = A^* + \bar{y}$ with discounting. Compared with panel (a), Firm 1's captive demand has shifted upward, as Firm 2's last best price has risen from p^{NT} to A^*. Its most profitable option along the no-targeting portion of this demand curve (list prices below A^*), is at point a', with profit Π_1 strictly exceeding the $\Pi_{NT} = 2$ earned in the no-targeting equilibrium. However, it earns the same profit Π_1 by matching Firm 2's high list price $p_1^* = A^* + \bar{y}$ and discounting (the blue shaded region). This profit includes $\frac{1}{2}p_1^* \approx 2.225$ from captives and $\frac{1}{4}$ from contested consumers, or $\Pi_1 = \frac{1}{2}\sqrt{6} + \frac{5}{4} \approx 2.475$ in total – an improvement of around 24% relative to the no-targeting equilibrium.

In this example, targeting facilitates higher profits mainly because it softens pricing (even discount prices), not because of the benefits of price discrimination *per se*. To illustrate, consider the thought experiment in which the ad cost falls to A^* but Firm 2 continues to price at p^{NT} (so we are in the case of Figure 2.a). Firm 1's optimal profit then involves earning $\frac{1}{2}p^H = 2$ on captives and $D = \frac{1}{4}\left(\bar{A} - A^*\right)^2 \approx 0.076$ with targeted discounts, or roughly 2.076 in total – a gain of only 3.8% relative to Π_{NT}. Thus the lion's share of the profit improvement with targeting – roughly five-sixths – may be attributed to the softer rival pricing it induces.

5.4 Advertising

Recall from Section 3 that a contested consumer receives an ad from her top two firms with probabilities $a_{(1)}(y) = 1 - \frac{A}{p^l - y}$ and $a_{(2)}(y) = 1 - \frac{y}{p^l}$ respectively, if her taste advantage is y. We write

$$a(y) = a_{(1)}(y) + a_2(y) = 2 - \frac{A}{p^l - y} - \frac{y}{p^l} \tag{7}$$

for the total expected number of ads sent to her, at expected cost $\mathcal{A}(y) = Aa(y)$. Let $\bar{a}(A)$ be the total volume of targeted advertising by all firms to all consumers. This total volume may be computed by integrating $a(y)$ over the entire contested region; this is equivalent to n identical copies of the total advertising on Firm 1's turf, so

$$\bar{a}(A) = n \int_0^{y^*} a(y)\, g(y)\, dy \text{ , where } y^* = p^l - A$$

Not surprisingly, we have:

Proposition 6 *Total ad volume $\bar{a}(A)$ is decreasing in the ad cost A.*

What makes this non-trivial is the endogenous response of list prices, which can be a countervailing force on ad volume. Next, holding the ad cost fixed, which consumers are targeted the most? Call a consumer "more contestable" as her taste difference y between her top two options grows smaller, with consumers at a turf boundary being the most contestable.

Result 1 *More contestable consumers receive more ads from both their favorite firms, their second-favorite firms, and in total. That is, $a_{(1)}(y)$, $a_2(y)$, and $a(y)$ are all strictly decreasing in y.*

Proof. These claims follow trivially from (7). ■

In motivating the first-order condition (3), we argued that a firm could retain marginal captive consumers at cost A, simply by advertising an infinitessimal discount. But perhaps surprisingly, this is not what firms actually do.

Result 2 *A firm does not advertise at all to consumers on the boundary of its captive region, even though they are poached with positive probability.*

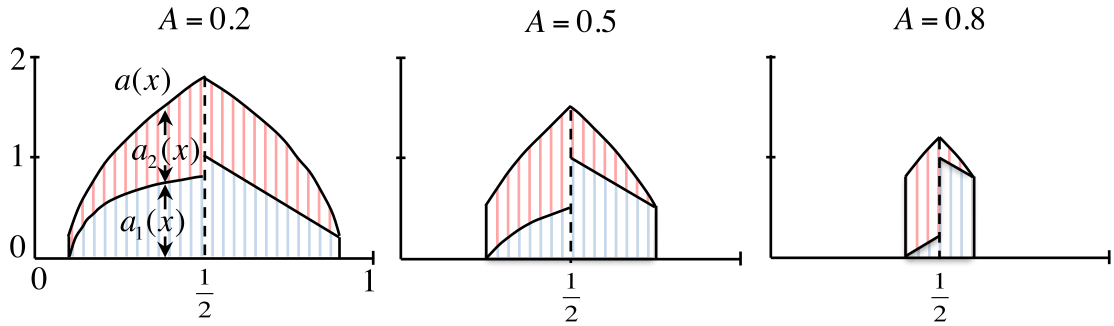

Figure 3: Volume of advertising $a(x)$ by consumer location x. (Linear-Hotelling preferences with $t = 1$.)

Proof. Evaluating at $y = y^* = p^l - A$, we see $a_{(1)}(y^*) = 0$ and $a_{(2)}(y^*) = \frac{A}{p^l}$. ■

The favorite firm does not wish to cannibalize its own list price sales needlessly; in equilibrium, its rivals poach consumers at its captive boundary just often enough that it is on the cusp of responding (but does not). Advertising behavior at the turf boundary between two firms is also a bit curious.

Result 3 *A firm's advertising probability jumps at its turf boundary with another firm. It advertises with probability* $1 - \frac{A}{p^l}$ *to consumers just on its side of the boundary, but with probability* 1 *to consumers just on its rival's side. Consequently, consumers near a turf boundary receive more ads for their second-best products than for their favorites.*

Proof. Evaluate the ad probabilities at $y = 0$ with the firm on its own turf $(a_{(1)}(0))$ and as the second-best option on its rival's turf $(a_{(2)}(0))$. ■

Here too, the intuition relates to cannibalization: because a firm will earn a list price sale from consumers on its side of the boundary in the event that they receive no ads, it has a weaker incentive than its rival to advertise to them.[23] Figure 3 illustrates all of these patterns for Hotelling competition with linear transportation costs. The upper envelope represents total ads $a(x)$ received by consumers at location x. This total is broken into the advertising contributed by Firm 1 (blue) and by Firm 2 (red). When ad costs are high relative to $p^{NT} = 1$, most ads involve poaching by the second-best rival, while as A falls, the contributions of poaching and retention become more balanced.

[23] This may sound incongruous because given $a_{(2)}(0) = 1$, the home turf firm will actually have its list price sale poached every time. One must think of $y = 0$ as the limiting case of competition near the turf boundary, where the incentive to avoid self-cannibalization does apply.

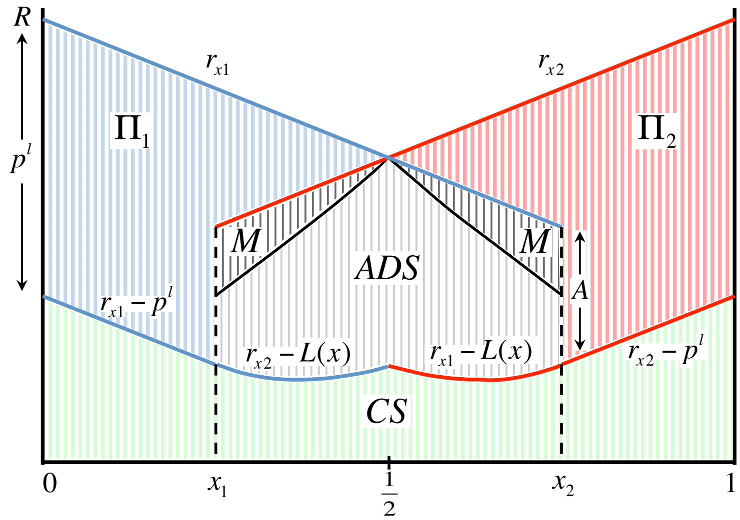

Figure 4: Equilibrium welfare in the linear-Hotelling model. ($R = 1.6$, $t = 1$, $A = 0.5$)

6 Welfare

Figures 4 and 5 summarize the distribution of total welfare and its components across consumers for equilibrium Hotelling competition with linear transportation costs. With a few important exceptions, the patterns depicted reflect the general results for arbitrary taste distributions that we describe below.

In the moderate ad cost case of Figure 4, first-best social surplus is the total area under the upper envelope $\max(r_{x1}, r_{x2})$. Firm profits Π_1 and Π_2 are depicted as in Figure 1. Sales to captive consumers ($x \in [0, x_1]$ and $x \in [x_2, 1]$) are socially efficient, so consumer surplus is the residual after subtracting off profit. However, sales to contested consumers are socially inefficient because of the direct costs of targeted advertising and the misallocation cost when a consumer purchases her second-best product (represented as ADS and M respectively). Figure 5 illustrates how profits, consumer surplus, and the welfare losses change when A is higher or lower. Below we start by characterizing the total welfare loss and its components; conclusions about the distribution of consumer surplus are developed in Section 6.3.

6.1 Consumer surplus and welfare

For the moment, focus again on a single consumer with valuations $r_{(1)}$ and $r_{(2)}$ at her top two firms. As suggested above, a captive consumer enjoys surplus $r_{(1)} - p^l$;

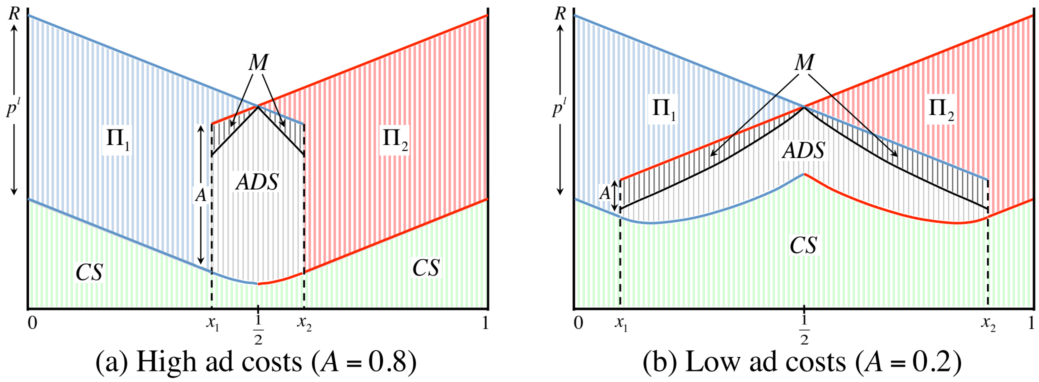

Figure 5: Equilibrium welfare in the linear-Hotelling model. ($R = 1.6$, $t = 1$)

adding in Firm (1) 's profit of p^l on her, we have (first-best) total welfare of $r_{(1)}$. A contested consumer takes the better of the final surplus offers from her top two firms. Using the strategies from Section 3.1, this consumer surplus is distributed according to

$$B_{\max}(s) = B_{(1)}(s)\, B_{(2)}(s) = \frac{A}{r_{(2)} - s}\frac{y + A}{y + r_{(2)} - s} \text{ on } \left[s^l_{(1)}, S_{(2)}\right], \tag{8}$$

including the chance $B_{\max}\left(s^l_{(1)}\right) = \frac{A}{p^l - y}\frac{A+y}{p^l}$ that she receives no offer strictly better than Firm (1)'s list price (where $s^l_{(1)} = r_{(1)} - p^l,\ S_{(2)} = r_{(2)} - A$). Thus her expected consumer surplus is $B_{\max}\left(s^l_{(1)}\right) s^l_{(1)} + \int_{s^l_{(1)}}^{S_2} s\, dB_{\max}(s)$. A straightforward computation shows that this surplus may be written

$$CS(y) = r_{(2)} - L\left(y, p^l, A\right) \tag{9}$$

where her shortfall relative to the full surplus at her second-best firm is given by the loss function

$$L(y, p, A) = A\left(1 + \frac{(A + y)}{y}\ln\left(\frac{A + y}{A}\frac{p - y}{p}\right)\right) \tag{10}$$

Including the firms' expected profits $\pi_{(1)} = r_{(1)} - r_{(2)}$ and $\pi_{(2)} = 0$, total welfare from sales to this consumer is given by

$$SS(y) = r_{(1)} - L\left(y, p^l, A\right)$$

Impact of targeted advertising on aggregate welfare

Aggregate social welfare is maximized if each consumer receives her favorite product and no ad costs are incurred; let SS_{1B} be this first-best aggregate welfare. Consider two bookend cases for the cost of targeted advertising: $A \geq p^{NT}$ (so that targeted ads are not used) and $A \rightarrow 0$ (the costless targeting limit). In the first case, the symmetric equilibrium achieves the first-best welfare. Otherwise, when ads are used, the deviation from first-best is given by integrating the loss term (10) over contested consumers. Given symmetry across firms, aggregate welfare is $SS = SS_{1B} - \bar{L}$, with the total welfare loss (versus the first-best) given by $\bar{L} = n \int_0^{p^l - A} L\left(y, p^l, A\right)\ dG\left(y\right)$. In the costless targeting limit, this welfare loss vanishes.

Proposition 7 *The equilibrium welfare loss on each contested consumer vanishes as ad costs vanish:* $\lim_{A \rightarrow 0} L\left(y, p^l\left(A\right), A\right) = 0$ *for all* $y \geq 0$*. Thus total social surplus tends toward its first-best level* SS_{1B} *as* $A \rightarrow 0$*.*

Of course this means that both components of the welfare loss, total ad spending and misallocation costs, vanish as $A \rightarrow 0$. The former is not too surprising (although it does rely on the equilibrium result that a consumer receives at most two ads). The fact that allocative efficiency is restored in the limit is perhaps less obvious. Although a general characterization is difficult for intermediate values of A, Proposition 7 suggests that welfare is broadly U-shaped in the cost of targeted advertising. While we may also conclude that aggregate welfare is lower when targeted ads are used than when they are not, one should not make too much of this result – it is more or less dictated by the absence in our model of any socially useful function for targeting (such as informing consumers about products, or replacing mass advertising of list prices).

Impact of targeted advertising on aggregate consumer surplus

Aggregate consumer surplus may be written as $\overline{CS} = SS_{1B} - \bar{\Pi} - \bar{L}$, where $\bar{\Pi}$ is total firm profits and $\bar{L}$ is the total welfare loss defined above. In either of the bookend cases ($A \geq p^{NT}$ and no targeting, or the $A \rightarrow 0$ limit), welfare losses vanish, so we have $\overline{CS}_{NT} = SS_{1B} - \bar{\Pi}_{NT}$ and $\overline{CS}_{A=0} = SS_{1B} - \bar{\Pi}_{A=0}$ respectively. But then, because firms are worse off in the free targeting limit (by Proposition 3), consumers must collectively be better off.

Result 4 *If captive demand is strictly logconcave, then* $\overline{CS}_{A=0} > \overline{CS}_{NT}$*.*

Continuity implies that consumers as a whole are better off with targeted ads than without them if the targeted ad cost A is sufficiently small. But this does not necessarily mean that they are also better off at higher levels of A.

Proposition 8 *Suppose tastes are i.i.d., with $f(r)$ strictly logconcave, and $n \geq 3$. Then every consumer's surplus is increasing in A (and hence so is aggregate consumer surplus) in a neighborhood below $A = p^{NT}$.*

Thus when there are at least three firms, the initial introduction of targeted advertising unambiguously makes every consumer worse off. Captive consumers benefit from a rise in A because they face lower list prices (by Proposition 5). The effect on contested consumers may be decomposed into a direct effect and a list price effect: $\frac{dCS}{dA} = \frac{\partial CS}{\partial A} + \frac{\partial CS}{\partial p^l}\frac{\partial p^l}{\partial A}$. The latter is positive, just as it is for captive consumers. However, the direct effect is negative, as an increase in A induces the firms to make less competitive discount offers (in the sense of first-order stochastic dominance – see (8)). The balance of these two effects depends on how often a contested consumer's best offer is equivalent to, versus strictly better than, the list price at her default firm. When ads are just barely affordable (A near p^{NT}) her best offer is unlikely to strictly improve on her default offer, so the list price effect dominates.

The presence of at least three firms ensures that $\frac{\partial p^l}{\partial A}$ remains *strictly* negative even near $A = p^{NT}$. Under certain conditions the conclusion of Proposition 8 also applies with two firms (in either the i.i.d. or Hotelling case), but second-order terms must be consulted because both $\frac{\partial CS}{\partial A}$ and $\frac{\partial p^l}{\partial A}$ vanish near $A = p^{NT}$ – Proposition 12 in the supplementary appendix gives further details.

An exception proves the rule: if list prices do not vary with A, then the only relevant effect is the direct effect on contested consumers and so we have the following.

Result 5 *If list price neutrality obtains, then individual consumer surplus is uniformly (weakly) decreasing in A.*

6.2 Misallocation: consumers buying the wrong product

As seen in Section 5.4, consumers will often be courted more aggressively by their second-favorite firms, and sometimes those efforts will be successful in tempting a consumer to purchase the "wrong" product. Let $m(y)$ be the probability that a consumer with taste difference y purchases her second-favorite product, and let $M(y) = ym(y)$

be the associated welfare cost (as illustrated in Figure 4). While $m(y)$ may be extracted from the accounting identity $L\left(y,p^l,A\right) = M(y) + \mathcal{A}(y)$, computing it directly is more illuminating. Let $m_{not1}(y) = (1 - a_1(y))\, a_2(y)$ be the probability that Firm 2 advertises a discount and Firm 1 does not, and let $m_{both}(y)$ be the probability that both advertise but Firm 2 wins the sale; $m(y)$ is the sum of these two cases. From the ad probabilities, the first term is $m_{not1}(y) = A/p^l$. For the second we have $m_{both}(y) = \int_{s_1^l}^{S_2} (1 - B_2(s_1))\, dB_1(s_1)$, as $1 - B_2(s_1)$ is the chance of a better offer from Firm 2 when Firm 1 advertises discount surplus $s_1 \in \left(s_1^l, S_2\right)$.[24]

Proposition 9 *If $A/p^l > \frac{1}{2}$, all contested consumers buy their second-favorite products more than half the time.*

Proof. Make the change of integration variables $p_2 = r_2 - s_1$ in $m_{both}(y)$ to get $m(y) = m_{not1}(y) + m_{both}(y) = \frac{A}{p^l} + \int_A^{p^l - y} \frac{p_2 - A}{p_2 + y} \frac{A}{p_2^2} dp_2$. From this representation it is immediate that $m\left(p^l - A\right) = \frac{A}{p^l}$ and $m'(y) < 0$, and so $m(y) > \frac{1}{2}$ for all $y \in \left[0, p^l - A\right]$ if $A > \frac{1}{2}p^l$. ■

Thus when targeting is in use but expensive, firms will be relatively successful at poaching consumers outside of their natural markets (although they will not profit by doing so) and relatively unsuccessful at retaining consumers on their own turf. However, as $A \to 0$, $m(y)$ tends to zero (uniformly over y), so when targeting is sufficiently cheap a firm will ultimately retain consumers on its own turf with probability tending to one (while failing to win any others by poaching).

6.3 Consumers: winners and losers from targeted advertising

In general, the introduction of targeted advertising will benefit some consumers and hurt others. To put this contrast in sharpest relief, we compare the bookend cases of no targeting ($A \geq p^{NT}$) and costless targeting ($A = 0$). In the first case, a consumer's surplus is $r_{(1)} - p^{NT}$, while in the latter it is simply $r_{(2)}$, as all consumers are contested and the loss term in (9) vanishes. A consumer "wins" with costless targeting (relative to her surplus without targeted ads) if her value advantage $y = r_{(1)} - r_{(2)}$ satisfies $y < p^{NT}$, and loses otherwise, so the main impact of targeting is to shift surplus from consumers with a strong favorite product toward those who are more willing to shop

[24] This integral (correctly) excludes any weight on Firm 2's atom of undercutting offers, as these never win when Firm 1 advertises.

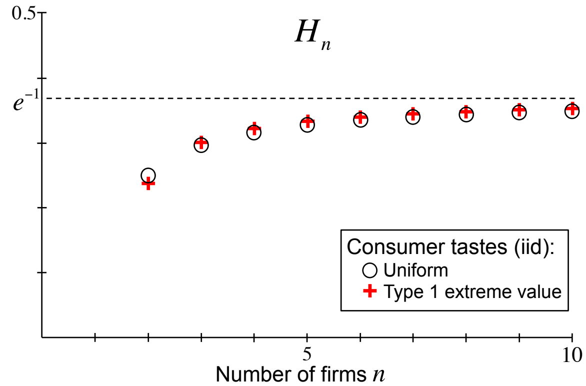

Figure 6: Fraction of consumers harmed by targeted advertising

around. We will pay particular attention to the "competitive limit" as the number of firms grows large, as there are interesting and sharp results in this case.

There are $1 - G_n\left(p_n^{NT}\right)$ consumers on Firm 1's turf who are harmed by costless targeting because their value advantages satisfy $y > p^{NT}$. (The subscript n has been added to emphasize that both G and p^{NT} will depend on the number of firms.) Then given symmetry, the overall fraction consumers harmed by costless targeting (relative to no targeting) is $H_n = n\left(1 - G_n\left(p_n^{NT}\right)\right)$. If tastes are i.i.d. and thin-tailed, we have the striking result that roughly 37% of consumers will be made worse off by the introduction of costless targeted advertising (and roughly 63% will benefit) in the competitive limit.

Proposition 10 *Suppose tastes are i.i.d. with well-behaved, strictly logconcave density $f(r)$ and $\bar{r} = \infty$. Then $\lim_{n\to\infty} H_n = \frac{1}{e}$.*

The proof uses an asymptotic result from Gabaix et al. (2015) characterizing the tails of an oligopolist's demand function as n grows large. We may write $1 - G_n(y) = E\left(1 - F\left(\tilde{r}_{n-1} + y\right)\right)$, where $\tilde{r}_{n-1}$ is the (random) largest rival valuation. The expected percentile of that highest rival valuation is $E\left(F\left(\tilde{r}_{n-1}\right)\right) = \frac{n-1}{n}$, so in some sense $\tilde{r}_{n-1}$ is centered around the certain value $\hat{r}_{n-1}$ defined by $F\left(\hat{r}_{n-1}\right) = \frac{n-1}{n}$. The crux of the Gabaix et al. result is that for n large we may approximate $1 - G_n(y)$ with $1 - F\left(\hat{r}_{n-1} + y\right)$. Consequently, both the no-targeting price and the gap between a consumer's top two draws are governed by the tail behavior of F and its hazard rate, and for these purposes, all thin tails are alike. Figure 6 shows the fraction of consumers harmed for two parametric examples where H_n may be calculated explicitly. The growth of targeted advertising has been accompanied at times by a sense of public unease and questions about whether limits or bans on targeting should be imposed. While a serious consideration of political economy is outside of our scope, Proposition 10 suggests one reason that bans on targeted ads may not be enacted: under broad conditions a majority of consumers would not support them.

7 Opt-in and consumer privacy

The striking rise in the degree of targeting by firms, and the consequent consumer backlash amid data breaches and realization about the extent of incursion into privacy and just how much data is being traded has invoked a major policy innovation in Europe, in the guise of the GDPR introduced in May 2018. We have already in the model the various ingredients needed to evaluate the policy, and it remains to pull them together. In particular, we have the firms' payoffs from targeting a specific consumer from our analysis of the discounting sub-game, we know how the list prices affect the set of contested consumers, and we know what are the individual consumers' benefits from receiving targeted discounts. The latter statistics enable us to determine the individuals' calculus in trading off against costs of privacy whether or not to opt-in to allowing firm access.

We analyze two plausible set-ups. The first has consumers making opt-in decisions after observing firms' list prices ("alert" consumers); the latter has them choosing beforehand ("inattentive" consumers).

7.1 Opt-in with alert consumers

We continue to assume n symmetric firms, but as earlier it suffices to focus on competition between two arbitrary firms – without loss of generality, Firms 1 and 2 – over the consumers for whom these are the top two products. As before, let $y = r_1 - r_2$ be a consumer's value advantage for Firm 1, which we can refer to as her "location." Firms choose list prices p_1^l and p_2^l at Stage 1, and at Stage 3 each firm can choose whether to pay A per consumer to sent a targeted discount offer to a consumer at location y, if she has opted into data collection and processing. Consumers who have opted out cannot be targeted. In between, at Stage 2, each consumer chooses between actions I and O. Action I means that she opts in at all firms, so any firm can send her a targeted discount, while action O means that she opts out at all firms. When choosing action I, the consumer incurs a privacy cost c; privacy costs are distributed according to cdf $H(c)$, independently of a consumer's location y. Except where noted, we assume $H(c)$ has support on the positive real line, with $H(0) = 0$. The all-or-nothing opt-in decision could arise in two plausible scenarios: (i) all firms access the consumer's data through a common data broker, or (ii) consumers can opt in or out at individual firms but take the view that once one firm has their data, the incremental risk from letting additional firms use it is small. This game timing reflects *alert* consumers, meaning that consumers are attentive to prices and are prepared to update their data security choices from time to time.

While we will develop the model with a general privacy cost distribution $H(c)$, to sharpen the conclusions we will often specialize to a two-type distribution where a fraction λ_k of consumers have cost c_k, $k \in \{L, H\}$, with $c_L < c_H$. We assume c_H to be sufficiently large that high-cost consumers never opt in.

7.1.1 Which consumers benefit from an opt-in policy

We analyze below a symmetric equilibrium with common list price p^{OI} (for "opt-in"). We will be interested in how consumers fare under an opt-in policy regime compared with unrestricted targeted discounts (studied in the rest of this paper), and no targeting. We refer to the symmetric equilibrium list prices in these alternative regimes as p^T and p^{NT} respectively. In both cases, we assume that a consumer suffers her privacy cost c if and only if she is located in the range of locations that receive targeting; this is $y \in \left[0, p^T - A\right]$ for the targeting regime, while no consumers incur privacy costs in the no-targeting case.

We will use the term "gross consumer surplus," CS_g, to refer to a consumer's surplus from her purchase – that is, valuation minus price. Her "net consumer surplus" CS_n also accounts for her privacy: $CS_n = CS_g - c$ if the privacy cost is incurred, or $CS_n = CS_g$ otherwise. We will show that the implications of each policy for consumer surplus can be largely understood through the policy's effect on list prices. We first compare the targeting and opt-in regimes.

Lemma 2 *If $p^{OI} < p^T$, every consumer is better-off under the opt-in regime than under unrestricted targeting. If $p^{OI} > p^T$, all consumers are worse-off, except for* *some portion of the consumers who would have been targeted but now choose to opt-out.*

Proof. We divide consumers into four groups.

i) Consumers who are captive under T and opt out under OI, buy their favorite products at list price in both cases. These consumers are clearly better off (in terms of both gross and net consumer surplus) in the regime where list prices are lower.

ii) Consumers who are contested under regime T and opt-in under regime OI receive discounts in either case. Their gain under regime OI is

$$CS_g^{OI} - CS_g^T = CS_n^{OI} - CS_n^T = L\left(y, p^T, A\right) - L\left(y, p^{OI}, A\right),$$

where the last equality uses (9), and recall that the function $L(\cdot)$ (which reflects the efficiency costs of discount competition) is increasing in the list price. Thus even though these consumers typically receive discounts, their expected surplus is anchored to list prices in a way that leads them to prefer the regime where list prices are lower.

iii) Now consider a consumer who is targeted under T but opts out when given the chance under OI. If $p^{OI} < p^T$, she would be targeted under both regimes and (by the argument in (ii)) she would have been strictly better-off with the lower list price in regime OI. Since she chose to opt out instead, revealed preference dictates that she enjoys higher final consumer surplus under regime OI. Less can be said if $p^{OI} > p^T$: arguments similar to case (ii) establish that her gross surplus falls under the opt-in regime (she switches from a lower list price with discounts to a higher list price without discounts.) However, her net surplus may be higher or lower depending on how highly she values preserving her privacy.

When $p^{OI} < p^T$, cases (i)-(iii) exhaust the ways a consumer could be treated under the two regimes, so establishing the first part of the Lemma.[25]

(iv) The final case only arises for $p^{OI} > p^T$. A consumer at location $y \in \left[p^T - A, p^{OI} - A\right]$ would not have been contested under T, but at the higher list price under OI she will receive discount offers if she opts in; if her privacy cost is low enough, she will do so. Conversely, consumers at locations $y \in [0, p^T - A)$ had no choice about being targeted under T but may opt out under OI, the higher list

[25] There will be no consumers who opt in under OI but were not targeted under T – if it was unprofitable to poach a consumer by undercutting the higher list price p^T, it will certainly be unprofitable to poach her when the list price is lower (so she has nothing to gain by opting in).

price notwithstanding, if their privacy costs are high enough. We claim that the first type of consumer is unambiguously worse-off after the regime change – she has $CS_g^{OI} < CS_g^T$, and so *a fortiori* her final consumer surplus is lower in the opt-in regime after accounting for her privacy cost. The claim about her market consumer surplus amounts to a statement that consumers benefit from a lower list price even when it forecloses the possibility of discounting. To verify this claim, fix a location $\hat{y} \in \left[p^T - A, p^{OI} - A\right]$ and consider the fictitious list price $\hat{p} = \hat{y} + A > p^T$ at which this location first becomes poachable. Raising the list price from p^T to $\hat{p}$ reduces the market surplus of a consumer at this location by $\hat{p} - p^T$. If the list price then rises from $\hat{p}$ to p^{OI}, (9) applies and her surplus falls by an additional $L\left(\hat{y}, p^{OI}, A\right) - L\left(\hat{y}, \hat{p}, A\right)$.

In summary, if switching from unrestricted targeting to an opt-in regime leads to higher list prices, we can conclude that all consumers are worse-off, except for some portion of the consumers who would have been targeted but now choose to opt-out.

■

Lemma 2 thus has the strong conclusion that *every* consumer, regardless of location y or privacy cost c, is better-off if list prices are lower with the opt-in regime than under unrestricted targeting; but there remains ambiguity for one consumer group if prices go the other way. Before we determine which way these prices go, we next provide a corresponding result for comparison to a no-targeting regime. The proof is analogous, but more straightforward.

Lemma 3 *If $p^{OI} < p^{NT}$, every consumer is better-off under the opt-in regime than with no targeting. If $p^{OI} > p^{NT}$, all consumers are worse-off, except for some of those consumers who would have been targeted but now choose to opt-out.*

Here, if $p^{OI} < p^{NT}$, consumers who do not switch status like the lower price, while consumers who choose to opt in are better off by revealed preference. Conversely, if $p^{OI} > p^{NT}$, those who opt out (whether targeted or not) are worse off from facing higher prices, while those who opt in are split (those who are worse off face privacy costs that do not offset the discounting gains).

7.1.2 Stage 2 discount competition and the consumer opt-in decision

Consider a consumer at location $y \geq 0$ who has opted in. Firm 2 can afford to send her offers as long as $y \leq p_1^l - A$, and so Stage 2 competition for her is exactly as described in Section 3. We engage the earlier analysis to determine how much the consumer benefits from these discount offers.

While this consumer will prefer Firm 1's list price in a symmetric equilibrium, we must also consider her out-of-equilibrium incentives. With this in mind, refer to $p_2^l + y$ as Firm 2's "normalized" price, and let $\tilde{p} = \min\left(p_1^l, p_2^l + y\right)$. In the absence of discounts, this consumer buys from the firm with the lower (normalized) price, and enjoys surplus $s^l = r_1 - \tilde{p}$. The consumer's total expected gross consumer surplus when she *can* be targeted with discount offers is given by the following (generalized) version of equation (9): $CS_g(y; \tilde{p}) = r_2 - L(y, \tilde{p}, A)$, where $L(y, \tilde{p}, A)$ is given by (10). Then, accounting for the privacy cost, she anticipates final utility s^l if she opts out, or $CS_g(y; \tilde{p}) - c$ if she opts in. Define $\Delta(y; \tilde{p}) = CS_g(y; \tilde{p}) - s^l$ to be her expected surplus gain from receiving discount offers. Then this consumer's optimal decision is simply to opt in if $c < \Delta(y; \tilde{p})$, or stay out if $c > \Delta(y; \tilde{p})$.

The expected benefit from discount offers can be expressed as:

$$\Delta(y; \tilde{p}) = \tilde{p} - y - L(y, \tilde{p}, A). \tag{11}$$

The next Lemma summarizes key features of $\Delta(y; \tilde{p})$, which are readily proved from (10).

Lemma 4 *For $y \geq 0$, the expected surplus improvement from opting in, $\Delta(y; \tilde{p})$, is increasing in $\tilde{p}$ and decreasing in y, with $\Delta(y; \tilde{p})|_{y=p_1^l - A} = 0$.*

As one might expect, consumers with relatively attractive second-best options gain more from opting in, and all consumers find opting in more attractive when list prices rise. We write $I(y) = H(\Delta(y; \tilde{p}))$ and $O(y) = 1 - I(y)$ for the fraction of consumers opting in or out at location y. Lemma 4 shows that $I(y)$ is decreasing in y and that $I(y) = 0$ for $y > p_1^l - A$ (consumers who Firm 2 cannot profitably reach do not opt in.) For the analysis below, it is also worth pointing out that the surplus gain $\Delta(y; \tilde{p})$, and a consumer's decision to opt in or out, is responsive to her *best* list price $\tilde{p}$. This will be important in the analysis of the disequilibrium scenario where Firm 1 charges a higher price than Firm 2. In this case, consumers $y > p_1^l - p_2^l > 0$ face a gain from discounting $\Delta\left(y; p_1^l\right)$ that depends on Firm 1's price; an additional price hike by Firm 1 will induce some of these consumers to opt in rather than pay p_1^l. However, consumers $y \in \left(0, p_1^l - p_2^l\right)$ would buy at Firm 2 in

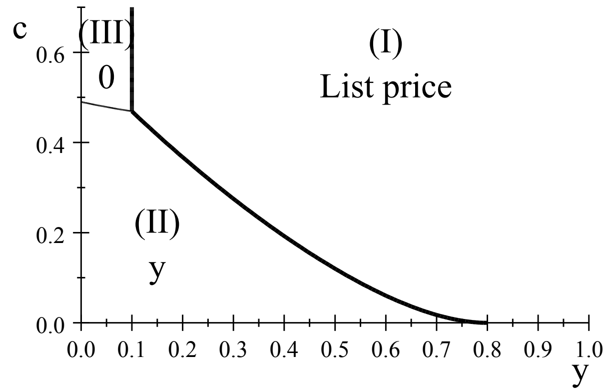

Figure 7: Profit Regions for Firm 1. (Parameters: $A = 0.3$, $p_1^l = 1.1$, $p_2^l = 1$. The boundary of Firm 1's captive demand is in bold. Consumers in Region (I) opt out and buy at p_1^l. Consumers in Region (II) opt in; Firm 1 earns net profit y on each of them. Consumers in Region (III) opt out and buy at Firm 2; Firm 1 earns 0 on them.)

the absence of discounts. An additional price hike by Firm 1 has no effect on their gain from discounting $\Delta\left(y; p_2^l + y\right)$, their decision to opt in or out, or any profit that Firm 1 earns on the opt-ins in the discounting subgame.

7.1.3 Stage 1 list prices

Firm 1's overall profit may be constructed by evaluating how much it earns on each consumer type (y, c). Firm 1 sells at its list price to consumers who opt out ($c > \Delta(y; \tilde{p})$) and prefer its product sufficiently ($y \geq p_1^l - p_2^l$). Note the second condition implies that $\tilde{p} = p_1^l$ for these consumers. Meanwhile, Firm 1 earns net profit y on any consumer who opts in ($c < \Delta(y; \tilde{p})$) and prefers Firm 1's product ($y \geq 0$). Thus Firm 1's overall profit is

$$\Pi_1\left(p_1^l\right) = p_1^l Q_1 + \int_0^{\infty} y \; H\left(\Delta\left(y; \tilde{p}\right)\right) \; dG\left(y\right) \tag{12}$$

where $Q_1 = \int_{p_1^l - p_2^l}^{\infty} \left(1 - H\left(\Delta\left(y; p_1^l\right)\right)\right) \; dG\left(y\right)$ represents total list price sales.

Figure 7illustrates these sales and profits over the space of consumer types (y, c), in a scenario with $p_1^l \geq p_2^l$. In this case, the benefit from opting in is governed by $\tilde{p} = p_1^l$ for consumers located at $y \geq p_1^l - p_2^l$, or $\tilde{p} = p_2^l + y$ otherwise, and the opt-in portion of profits may be decomposed as $\int_0^{p_1^l - p_2^l} y \; H\left(\Delta\left(y; p_2^l + y\right)\right) \; dG\left(y\right) +$

$\int_{p_1^l - p_2^l}^{\infty} yH\left(\Delta\left(y; p_1^l\right)\right) dG\left(y\right)$. Then Firm 1's marginal profit may be written as:

$$\frac{d\Pi_1}{dp_1^l} = Q_1 - O\left(\underline{y}\right) g\left(\underline{y}\right) p_1^l - \int_{p_1^l - p_2^l}^{\infty} \left(p_1^l - y\right) \frac{\partial I\left(y\right)}{\partial p_1^l}\, dG\left(y\right) \tag{13}$$

where $\underline{y} = p_1^l - p_2^l$. The second and third terms represent the two different margins. The second term is the conventional oligopoly margin: consumers who are indifferent between the two list prices are lost to Firm 2 if they have opted out. The third term represents consumers who are induced to opt in by the price hike; as Firm 1 must now compete for them with discounts, its net profit on such consumers drops from p_1^l to y.

Comparing marginal profit here with marginal profit in our benchmark model is easier with a change of variables. The terms in Firm 1's profit (12) represent double integrals over the space of consumer types (y, c), where we have implicitly integrated over c first, and then over y, but we can reverse this. For list price sales we have

$$Q_1 = \int_0^{\bar{c}} \left(1 - G\left(y^*\left(c; p_1^l\right)\right)\right)\ dH\left(c\right) + \int_{\bar{c}}^{\infty} \left(1 - G\left(p_1^l - p_2^l\right)\right)\ dH\left(c\right) \tag{14}$$

where $y^*\left(c; p_1^l\right)$ is the threshold location at which a consumer with privacy cost c would opt in, defined by $\Delta\left(y^*; p_1^l\right) = c$, and $\bar{c} = \Delta\left(y; p_1^l\right)\big|_{y=p_1^l - p_2^l}$ is the privacy cost above which all consumers opt out. Proceeding similarly for the other terms in the marginal profit, we have:

$$\frac{d\Pi_1}{dp_1^l} = \int_0^{\infty} \Lambda\left(c\right)\ dH\left(c\right), \text{ where} \tag{15}$$

$$\begin{aligned}
\Lambda\left(c\right) &= 1 - G\left(y^*\left(c; p_1^l\right)\right) - \left(p_1^l - y^*\left(c; p_1^l\right)\right) g\left(y^*\left(c; p_1^l\right)\right) \frac{\partial y^*}{\partial p_1^l} \quad \text{for } c < \bar{c} && (16)\\
&= 1 - G\left(p_1^l - p_2^l\right) - p_1^l g\left(p_1^l - p_2^l\right) \quad \text{for } c > \bar{c} && (17)
\end{aligned}$$

That is, marginal profit can be expressed as an expectation over the marginal profits $\Lambda\left(c\right)$ associated with each privacy cost type. For types who opt out regardless of location ($c > \bar{c}$), this is the conventional oligopoly marginal profit. For consumers with $c < \bar{c}$, there is a threshold location y^* at which the gains from discounting exactly compensate for c. Firm 1's marginal profit on these consumers involves an inframarginal sales term $1-G(y^*)$, and losses of $p_1^l - y^*$ at the margin from consumers who switch to opting in.

The necessary first-order condition for a symmetric equilibrium at common list price p^{OI} is then $E_H(\Lambda(c))|_{p_1^l=p_2^l=p^{OI}} = 0$. Earlier we showed that list prices rise or fall when targeting is introduced depending on whether demand is convex or concave. We will leverage the structure of (16) and (17) to show how demand curvature determines whether an opt-in policy will lead to list prices p^{OI} that are higher or lower than under unrestricted targeting. We first give an instrumental result we will need for bounding (16).

Lemma 5 $\frac{\partial y^*}{\partial p_1^l} \in (0,1)$ *for* $y^* \in \left(0, p_1^l - A\right)$. *Furthermore,* $\frac{\partial y^*}{\partial p_1^l} \to 1$ *as* $A \to 0$.

Proof. The threshold $y^* = y^*(c;p)$ is defined implicitly $c = \Delta(y;p)$, so by the implicit function theorem $\frac{\partial y^*}{\partial p} = -\frac{\Delta_p}{\Delta_y}$. While (11) applies, it is convenient to express it in another form. Let $\tilde{s}_1 = s_1 - s_1^l$ and $\tilde{s}_2 = s_2 - s_1^l$ be the amounts by which the firms' discount offers improve a consumer's surplus over Firm 1's list price offer. Then $\Delta\left(y;p_1^l\right) = E(\max(\tilde{s}_1,\tilde{s}_2))$. Overloading notation, it follows from the results in Section 3.1 that these surplus improvements are distributed according to $B_1(\tilde{s}) = \frac{A}{p-y-\tilde{s}}$ and $B_2(\tilde{s}) = \frac{A+y}{p-\tilde{s}}$ respectively, for $\tilde{s} \in [0, p-y-A]$, and so the consumer's best surplus improvement is distributed $B_1(\tilde{s})B_2(\tilde{s})$. Then integration by parts implies $\Delta(y;p) = \int_0^{p-y-A}(1 - B_1(\tilde{s})B_2(\tilde{s}))\,d\tilde{s}$.

Then it is easily confirmed that $\Delta_p = \int_0^{p-y-A} B_1'(\tilde{s})B_2(\tilde{s}) + B_1(\tilde{s})B_2'(\tilde{s})\ d\tilde{s} = B_1(\tilde{s})B_2(\tilde{s})|_{\tilde{s}=0}^{\tilde{s}=p-y-A} = 1 - \frac{A}{p-y}\frac{A+y}{p}$. Meanwhile,

$$\Delta_y = \int_0^{p-y-A} B_1'(\tilde{s})B_2(\tilde{s})\ d\tilde{s} + \int_0^{p-y-A} B_1(\tilde{s})\frac{1}{p-\tilde{s}}\ d\tilde{s}$$

Integrate the first term by parts and consolidate to get

$$\begin{aligned} \Delta_y &= B_1(\tilde{s})B_2(\tilde{s})|_{\tilde{s}=0}^{\tilde{s}=p-y-A} + \int_0^{p-y-A} B_1(\tilde{s})\left(\frac{1}{p-\tilde{s}} - B_2'(\tilde{s})\right)\ d\tilde{s} \\ &= \Delta_p + \int_0^{p-y-A}\frac{1}{p-\tilde{s}}B_1(\tilde{s})(1-B_2(\tilde{s}))\ d\tilde{s} \end{aligned}$$

Then the main result follows because $\Delta_y > \Delta_p$ whenever the second term above does not vanish. The limit follows from $\Delta_y \to \Delta_p \to 1$ as $A \to 0$. ■

In words, at a particular cost type c, the marginal opt-in location does not rise one-for-one with a price rise. This is a contrast with our benchmark model, where the marginal targetee, $y = p_1^l - A$, does rise 1-for-1 with the price. The difference here is that inefficiency (the ad costs and misallocation in $L(\cdot)$) rise with p_1^l. So the benefits to a consumer from opting in rise more slowly than p_1^l.

Proposition 11 *If demand is weakly concave ($g' \geq 0$), then any symmetric equilibrium under opt-in satisfies $p^{OI} > p^T$. Alternatively, if demand is strictly convex ($g' < 0$), then for A sufficiently small, any symmetric equilibrium under opt-in satisfies $p^{OI} < p^T$.*

Proof. Concave demand: Let $\hat{y}(p) = p - A$. The price p^T is the solution to the targeting FOC $1 - G(\hat{y}(p)) - (p - \hat{y}(p))\, g(\hat{y}(p)) = 0$. Equivalently, write $\frac{1-G(\hat{y}(p))}{g(\hat{y}(p))} + \hat{y}(p) - p = 0$. Logconcavity of demand implies the lefthand side is negative for $p > p^T$. Define the auxiliary function $\gamma(y) = \frac{1-G(y)}{g(y)} + y$, so that targeting FOC may be expressed as $\gamma\left(\hat{y}\left(p^T\right)\right) - p^T = 0$, and note that weak concavity of demand implies that $\gamma'(y) \leq 0$.

Next evaluate marginal profit $d\Pi_1/dp_1^l$ in the OI model at $p_1^l = p_2^l = P \leq p^T$. Using (15), this marginal profit is an expecion of $\Lambda(c)$ terms. For $c > \bar{c}$, these terms satisfy $\Lambda(c) = 1 - G(0) - Pg(0) \geq 0$ (because Proposition 2 implies $P \leq p^T \leq p^{NT}$, and p^{NT} satisfies $1 - G(0) - p^{NT} g(0) = 0$). And for $c \in (0, \bar{c})$, we have

$$\begin{aligned}
\Lambda(c)\big|_{p_1^l = P} &= 1 - G(y^*(c;P)) - (P - y^*(c;P))\, g(y^*(c;P)) \frac{\partial y^*}{\partial p_1^l} \\
&> g(y^*(c;P))\left(\gamma(y^*(c;P)) - P\right) \\
&> g(y^*(c;P))\left(\gamma(\hat{y}(P)) - P\right) \\
&> g(y^*(c;P))\left(\gamma\left(\hat{y}\left(p^T\right)\right) - p^T\right) = 0
\end{aligned}$$

where the first inequality uses $\partial y^*/\partial p_1^l < 1$, the second uses $\hat{y}(P) = P - A > y^*(c;P)$ and $\gamma'(y) \leq 0$, and the third inequality uses $P \leq p^T$ and $\gamma'(y) \leq 0$ one more time. The final term in parentheses is equivalent to the targeting FOC. It follows that $\left.\frac{d\Pi_1}{dp_1^l}\right|_{p_1^l, p_2^l = P} > 0$ for any $P \leq p^T$. Thus, if a symmetric equilibrium exists at some price p^{OI}, it must satisfy $p^{OI} > p^T$.

Strictly convex demand – sketch: The idea is to reverse the argument for the concave demand case. However in this case the fact that $\partial y^*/\partial p_1^l < 1$ complicates the argument rather than strengthening it; this is the reason for the "for A sufficiently small" modifier. Define $\gamma(y)$ as above but note that now we have $\gamma'(y) > 0$ and (by Proposition 2) $p^T > p^{NT}$. This time we evaluate marginal profit $d\Pi_1/dp_1^l$ at a symmetric price profile $p_1^l = p_2^l = P \geq p^T > p^{NT}$. As above, $d\Pi_1/dp_1^l\big|_{p_1^l, p_2^l = P}$ is an expectation over $\Lambda(c)$ terms. In this case, we have $\Lambda(c) < 0$ for $c > \bar{c}$ (replicating the argument above, but in this case we have $P > p^{NT}$). Let $-Z = \int_{\bar{c}}^{\infty} \Lambda(c)\ dH(c)$ be the contribution of these terms to the marginal profit.

For the terms with $c \in (0, \bar{c})$, letting $\varepsilon = 1 - \partial y^*/\partial p_1^l$, we have

$$\begin{aligned} \frac{\Lambda(c)|_{p_1^l = P}}{g(y^*(c;P))} &= \gamma(y^*(c;P)) - P + \varepsilon(P - y^*(c;P)) \\ &< \gamma(y^*(c;P)) - P + \varepsilon P \\ &< \gamma(\hat{y}(P)) - P + \varepsilon P \\ &< \varepsilon P \end{aligned}$$

*** INCOMPLETE *** ■

Proposition 11 gives the loose conclusion that consumers tend to benefit from a switch from unrestricted targeting to opt-in iff demand is convex. To be more precise, if A is sufficiently small and demand is strictly convex, then *all* consumers benefit from the policy change. And if demand is weakly concave (regardless of A), then 'most' consumers are harmed by the policy change.

Note that the Hotelling example from earlier is covered in the 'weakly concave' case. Recall that the Hotelling case gives rise to linear demand $1 - G(y)$ and the list-price-neutrality result that $p^T = p^{NT}$. In this case, Proposition 11 implies that list prices under opt-in will be higher than they would be if targeting were either unrestricted or banned entirely: $p^{OI} > p^T = p^{NT}$.proof follows same structure, with appropriate flips in statements possible light discussion with effects of banning targeting.

7.1.4 Equilibrium existence

While we do not have a general proof of equilibrium existence (due to the intricacies of the function $\Delta\left(y; p_1^l\right)$), we provide a simple and central example where existence is guaranteed. This is the Hotelling model (with the two firms located at each end-point of a unit interval and linear transport costs t per unit distance) with two cost levels

$c_L \downarrow 0$ and c_H high enough that no consumer on the high level opts in. In this case the candidate equilibrium for the pure non-targeting model (classic Hotelling) has price $p_1^l = t$ as does the pure targeting model (cf. Proposition 2). The candidate first-order condition for the joint model is a convex combination of the first-order conditions for the two extremes and so has the candidate symmetric equilibrium as $p^l = t$. The profit function for the high cost consumers is the classic Hotelling one and it is quadratic (and hence strictly concave) in $p_1^l \in [0, 2t]$ given $p_2^l = t$ (and profit is zero for higher p_1^l.) Indeed, the profit function is then $\left(1 - \frac{p_1^l}{2t}\right) p_1^l$ with own list-price derivative $\left(t - p_1^l\right)/t$.

For the consumers for whom $c_L \downarrow 0$ (we take the limit with the cost tending to zero so that the market splits at $y^* = p_1^l - A$ when this exceeds 0) then y is uniformly distributed with density $\frac{1}{2t}$ on $[-t, t]$ (to match the uniform density of consumers on the unit interval). For $p_1^l \geq A$, the profit on the low-cost consumers is proportional to $\frac{1}{2t}(t - y^*) p_1^l + \frac{1}{2t}\int_0^{y^*} y dy$ with $y^* = \min\left\{p_1^l - A, t\right\}$, so that this profit is $\frac{1}{2t}\left(t - p_1^l + A\right) p_1^l + \frac{1}{4t}\left(p_1^l - A\right)^2$ for $p_1^l \leq t + A$, which is a quadratic concave function with price derivative $\frac{1}{2t}\left(t - p_1^l\right)$, and profit is $\frac{1}{4t}$ for higher p_1^l. For $p_1^l \leq A$ the firm gets nothing in the discount sub-game and profit is proportional to $\frac{1}{2t}(t - y^*) p_1^l$ with $y^* = \max\left\{p_1^l - A, -t\right\}$. This profit too is quadratic, and note that the derivative is continuously differentiable at the switch-point $p_1^l \geq A$ (where it is $\frac{1}{2t}(t - A)$, which is positive given the condition $t > A$ needed to have an equilibrium with targeting).

Thus the full profit function is concave over the relevant range, and the equilibrium sustains at $p_1^l = t$: profits are rising with p_1^l on both profit parts up to $p_1^l = t$ and falling on both beyond there.

7.2 Alternative timing: inattentive consumers. (Privacy choices before list prices)

This section studies the opt-in model under the alternative timing assumption that consumers' opt in or out decisions are made simultaneously with firms' list price choices.[26] This timing might be more appropriate if real-world consumers do not

[26] Because an individual consumer has a neglibible impact on the firms' decisions, it would also be equivalent to have consumers make their decisions before list prices are set.

revisit their privacy settings very frequently (or at least not after every list price change); consequently we refer to this as the case with "inattentive" consumers.

Under this timing, consumers form expectations p_j^e about the firms' list prices; of course, in equilibrium those expectations must be correct. As earlier, consider a consumer with preference $y \geq 0$ for her favorite product at Firm 1 over her second favorite product at Firm 2. This consumer's decision is just as characterized earlier, except that now her choice is based on the expected list prices p_1^e and p_2^e; she opts in if her privacy cost satisfies $c \leq \Delta(y; \tilde{p}^e)$, where $\tilde{p}^e = \min(p_1^e, p_2^e + y)$. If consumers expect symmetric list prices, then their optimal decisions can be summarized by a privacy cost threshold $\bar{c} = \Delta(0; p_1^e)$ and the threshold location function $y^*(c; p_1^e)$ function defined earlier. Consumers $y \geq 0$ with privacy costs above $\bar{c}$ opt out regardless of location, as they expect insufficient gains from discounting. Consumers with privacy costs below this threshold opt in if $y \in [0, y^*(c; p_1^e)]$ and opt out otherwise. Note that the threshold location for opting in is lower than the threshold at which Firm 2 would target the consumer, given the chance: Lemma XXX implies $y^*(c; p_1^e) \leq p_1^e - A$, with equality only at $y^*(0; p_1^e) = p_1^e - A$. This reflects the fact that a consumer with $c > 0$ requires not just a positive chance at discounts, but sufficiently large gains from them, to find opting in worthwhile.

The key change when consumers are inattentive has to do with the incentives of the firms: consumers who have opted out of discounts before seeing Firm 1's realized p_1^l are vulnerable to being held up. To illustrate the issue, we focus on a candidate equilibrium with symmetric prices $p_1^e = p_2^e = p_2^l$ and consider Firm 1's profit from setting price p_1^l. Firm 1 sells at its list price to all opt-outs for whom $y \geq p_1^l - p_2^l$, and also to any opt-ins $y \geq p_1^l - A$ whom Firm 2 cannot profitably target. (The latter case will not occur in equilibrium, but it needs to be considered in Firm 1's incentives.) This implies list price sales

$$Q_1 = \int_{p_1^l - p_2^l} O(y) \; dG(y) + \int_{p_1^l - A} I(y) \; dG(y)$$

Meanwhile, any opt-ins satisfying $y \leq p_1^l - A$ will be contested by Firm 2, in which case Firm 1 earns a net profit of $\max(y, 0)$ on them. Thus we have overall profit

$$\Pi_1\left(p_1^l\right) = p_1^l Q_1 + \int_0^{p_1^l - A} y \; I(y) \; dG(y)$$

Then Firm 1's marginal profit is:

$$\frac{d\Pi_1}{dp_1^l} = Q_1 - \underbrace{p_1^l O\left(p_1^l - p_2^l\right) g\left(p_1^l - p_2^l\right)}_{\text{List price margin}} - \underbrace{A\ I\left(p_1^l - A\right) g\left(p_1^l - A\right)}_{\text{Poaching margin}}$$

The 'list price margin' term is conventional: a price rise loses consumers located at $y = p_1^l - p_2^l$ to Firm 2's list price. The 'poaching margin' term reflects the prospect of opt-in consumers in the neighborhood of $y = p_1^l - A$ that would begin to be targeted by Firm 2 if p_1^l rose slightly. However, this last term should sound problematic: if these consumers do not receive targeted discounts at current prices, then why did they opt in the first place? And indeed, we will show that this term must vanish in equilibrium.

If consumers expect prices $p_1^e = p_2^e$, then for consumers located at $y \geq 0$ the benefit from discounting is $\Delta\left(y; p_1^e\right)$. This benefit satisfies $\Delta\left(y; p_1^e\right) = 0$ for all $y \geq p_1^e - A$ (because these locations do not expect to receive discount offers). But then since $H(0) = 0$, consumer opt-in decisions must satisfy $I(y) = H\left(\Delta\left(y; p_1^e\right)\right) = 0$ for all $y \geq p_1^e - A$. That is, as there is no atom of consumers with $c = 0$, at locations where there is no expected benefit from discounting, all consumers opt out. This immediately implies that the poaching margin term vanishes from $d\Pi_1/dp_1^l$ for $p_1^l \geq p_1^e$. Likewise, the second term in Q_1 also vanishes for $p_1^l \geq p_1^e$. Then evaluated at a candidate symmetric equilibrium, Firm 1's marginal profit reduces to

$$\left.\frac{d\Pi_1}{dp_1^l}\right|_{p_1^l = p_1^e} = \int_0 O(y)\ dG(y) - O(0)\, g(0)\, p_1^e$$

This is rather close to the marginal profit expression with alert consumers (13) when evaluated at a symmetric equilibrium. Indeed, the only difference is that the second marginal term in (13), representing consumers who switch their opt-in decision in response to p_1^l, has dropped out here. In this sense, inattentive consumers do not discipline list prices as effectively as alert consumers would.

We will write $\hat{p}^{OI}$ for a symmetric equilibrium price level when consumers are inattentive, reserving p^{OI} for the alert consumer timing discussed above. At a symmetric equilibrium, consumers hold correct expectations about list prices and each firm's list price maximizes its profit, given consumer opt-in decisions. This means that (with slight rearranging) any symmetric equilibrium list price $\hat{p}^{OI}$ must satisfy the first-order condition

$$\int_0 \frac{O(y)}{O(0)} dG(y) - g(0)\hat{p}^{OI} = 0 \tag{18}$$

In effect, each firm faces a demand curve for list price sales that has been hollowed out by consumers opting in to discounting. The $O(y)/O(0)$ term in the integrand registers how many consumers remain at each inframarginal location, relative to consumers remaining at the $y = 0$ margin. If this term were absent, then (18) would reduce to the standard oligopoly first-order condition with solution p^{NT}. However, as we noted in the discussion around Lemma XXX, $O(y) = 1 - H\left(\Delta\left(y; \hat{p}^{OI}\right)\right)$ is increasing in y – consumers near $y = 0$ have the most to gain from discounting and opt in at the highest rates, while consumers with a stronger preference for one product expect smaller gains and so are more likely to opt out. The fact that this hollowing out is more severe at the margin than inframarginally makes a firm's list price demand effectively less elastic, encouraging higher prices. This logic leads to some straightforward price comparisons that we summarize below.

CLAIM Any symmetric equilibrium with inattentive consumers and opt-in has higher list prices than the case where consumers are alert ($\hat{p}^{OI} > p^{OI}$) and higher list prices than a regime where targeting is banned ($\hat{p}^{OI} > p^{NT}$).

The second comparison follows directly from the fact that $O(y)/O(0) > 1$, which implies the lefthand side of (18) is strictly positive for all prices p^{NT} or lower.

8 Concluding Remarks

Some of our results – for example, the redistribution of consumer surplus from individuals with high values for their favorite product toward those with high values for their second-best product – give an underpinning for patterns that arise rather consistently throughout the literature on targeting. The impact of targeting on profits is a less settled question. The prevailing view is probably that competitive price discrimination stiffens competition and leaves firms worse-off, and this matches our main finding with logconcave captive demand. As we show, that intuition can reverse if demand is logconvex: the introduction of targeting can soften competition and raise profits. Below we suggest several alternative reasons for targeting to be profitable that may help to explain other results in the literature.

One is simply accounting. If advertising to a consumer is a prerequisite for selling to her, and if targeted and mass advertising are equally costly per consumer reached, then firms may reap cost savings from targeting by consolidating their ad spend on the consumers who are most likely to purchase. This effect is absent in our model because we do not include any cost of publicizing list prices. If we did, it seems fairly clear that this would probably temper our conclusions about profits as long as publicizing list prices involved a per-consumer cost that could be scaled back. (If publishing prices involved a fixed cost instead, it is less clear that anything would change.)

Another is monopoly power. If a firm has consumers whom its rivals cannot effectively compete for, either because of a substantial cost or value asymmetry or because the consumers have binding outside options (see below), then greater flexibility in pricing to these consumers should improve the firm's profits. Thisse and Vives (1988) find a result of this kind for the dominant firm when the asymmetry between firms is sufficiently large.[27]

A third, more speculative potential explanation is imperfect targeting. In models like ours, targeting permits head-to-head Bertrand competition for a contested consumer – it is generally hard for this to be good for firms. In those papers where firms benefit from targeting, the technology usually has some imperfection or limitation that softens price competition over those targeted.[28] Slightly imperfect targeting would not change our conclusions – in related work, we explain on continuity grounds why competition for contested consumers would be continue to be fierce if firms' information about consumers were a little bit noisy.[29] However, the effect of more substantial targeting imperfections on our model is less certain.

As firms collect ever more detailed information about consumers' tastes, individualized price offers are likely to become increasingly common. Our paper provides a theoretical framework for understanding the repercussions of this shift in the market-

[27] In a rare empirical study on this subject, Besanko, Dubé, and Gupta (2003) use a multinomial choice model calibrated from data to simulate a duopoly equilibrium under price discrimination. They find an improvement in profits for one of the firms (over uniform pricing), which they suggest may be connected to a quality advantage for its product.

[28] Often (Galeotti and Moraga-González (2008), Iyer et al. (2005), and Esteves and Resende (2016)) this is because firms cannot be sure which consumers within a targeted group will receive their ads, so they price with a glimmer of hope at *ex post* monopoly power. Or, as in Chen et al. (2001) ads sometimes reach the "wrong" consumers rather than those who were targeted. Alternatively, convex advertising costs (Esteves and Resende, 2016) may prevent all-out competition.

[29] See Anderson, Baik, and Larson (2015).

place. While our approach is quite general in many respects, it is worth discussing our simplifying assumptions and directions for extension.

Because we assume the market to be fully covered, a consumer's next-best option is always some rival firm rather than the outside option of not purchasing. This permits us to treat next-best options symmetrically, which is particularly helpful in keeping the n-firm case tractable. However it also implies that a discounting firm always faces competition. If outside options were to bind, then targeting would also have a market-expanding effect: each firm would be able to make monopoly price-discriminating offers to some consumers who otherwise would not have purchased. In this case, cheaper targeting would likely have a more positive impact on profits than our results suggest, perhaps at the expense of consumer surplus; the implications for list prices seem likely to be the same.

While symmetry is convenient, our framework can be readily adapted to accommodate differences in advertising cost, production cost, or the consumer taste distribution across firms (although broad, tractable conclusions might be harder to obtain). Asymmetries in the information that firms have about consumers (such as superior information about one's own past customers or about preferences for one's own product) deserve study but would probably require a different modeling framework. Finally, we have not addressed the market in which firms acquire consumer data.[30] A treatment of these issues is left to future work.

References

Acquisti, A. and Varian, H., 2005, Conditioning Prices on Purchase History, *Marketing Science*, 24(3), 367-381.

Anderson, Simon, Alicia Baik, and Nathan Larson, 2015, Personalized Pricing and Advertising: An Asymmetric Equilibrium Analysis, *Games and Economic Behavior*, 92, 53-73.

Anderson, Simon and Andre de Palma. 1988, Spatial Price Discrimination with Heterogeneous Products, *Review of Economic Studies*, 55, 573-592.

Anderson, Simon, de Palma, André, and Jacques-François Thisse, 1989, Spatial Price Policies Reconsidered, *The Journal of Industrial Economics*, 38, 1-18.

[30]But see Montes, Sand-Zantman, and Valletti (2015).

Anderson, Simon, de Palma, André, and Jacques-François Thisse, 1992, *Discrete Choice Theory of Product Differentiation*, MIT Press.

Armstrong, M. and Vickers, J., 2001, Competitive Price Discrimination, *Rand Journal of Economics*, 32, 1-27.

Bagnoli, M. and Bergstrom, T., 2005, Log-concave probability and its applications, *Economic Theory*, 26, 445-469.

Besanko, D., Dubé, J-P., and Gupta, S., 2003, Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data, *Management Science*, 49(9), 1121-1138.

Bester, Helmut and Emmanuel Petrakis, 1995, Price Competition and Advertising in Oligopoly, *European Economic Review*, 39, 1075-1088.

Bester, Helmut and Emmanuel Petrakis, 1996, Coupons and Oligopolistic Price Discrimination, *International Journal of Industrial Organization*, 14, 227-242.

Brahim, N., Lahmandi-Ayed, R., and Laussel, D., 2011, Is targeted advertising always beneficial?, *International Journal of Industrial Organization*, 29, 678-689.

Butters, Gerard R. 1977. Equilibrium Distributions of Sales and Advertising Prices, *Review of Economic Studies*, 44, 465-491.

Chen, Y., Narasimhan, C., and Zhang, Z., 2001, Individual Marketing with Imperfect Targetability, *Marketing Science*, 20, 23-41.

Corts, K., 1998, Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment, *Rand Journal of Economics*, 29(2), 306-323.

Esteban, L., Gil, A., and Hernández, J., 2001, Informative Advertising and Optimal Targeting in a Monopoly, *The Journal of Industrial Economics*, 49(2), 161-180.

Esteves, R-B., and Resende, J., 2016, Competitive Targeted Advertising with Price Discrimination, *Marketing Science*, in press.

Fudenberg, D. and Tirole, J., 2000, Customer Poaching and Brand Switching, *Rand Journal of Economics*, 31, 634-657.

Gabaix, X., Laibson, D., Li, D., Li, H., Resnick, S., and de Vries, C., 2015, The Impact of Competition on Prices with Numerous Firms, mimeo.

Galeotti, A., and Moraga-González, J., 2008, Segmentation, Advertising, and Prices, *International Journal of Industrial Organization*, 26, 1106-1119.

Grossman, Gene and Carl Shapiro, 1984, Informative Advertising with Differentiated Products, *Review of Economic Studies*, 51, 63-81.

Hillman, Arye L. and John G. Riley, 1989, Politically Contestable Rents and Transfers, *Economics and Politics*, 1, 17-39.

Hoover, E., 1937, Spatial Price Discrimination, *Review of Economic Studies*, 4, 182-191.

Hotelling, Harold, 1929, Stability in Competition, *Economic Journal*, 39, 41-57.

Iyer, G., Soberman, D., and Villas-Boas, J., 2005, The Targeting of Advertising, *Marketing Science*, 24(3), 461-476.

Koçaş, Cenk and Tunga Kiyak, 2006, Theory and Evidence on Pricing by Asymmetric Oligopolies, *International Journal of Industrial Organization*, 24, 83-105.

Lederer, P. and Hurter, A., 1986, Competition of Firms: Discriminatory Pricing and Location, *Econometrica*, 54, 623-640.

Montes, R., Sand-Zantman, W., and Valletti, T., 2015, The value of personal information in markets with endogenous privacy, mimeo.

Narasimhan, C., 1988, Competitive Promotional Strategies, *Journal of Business*, 61, 427-449.

Perloff, J. and Salop, S., 1985, Equilibrium with Product Differentiation, *Review of Economic Studies*, 52, 107-120.

Shaffer, Greg and Z. John Zhang, 1995, Competitive Coupon Targeting, *Marketing Science*, 14, 395-416.

Shaffer, Greg and Z. John Zhang, 2002, Competitive One-to-One Promotions, *Management Science*, 48, 1143-1160.

Sharkey, William W. and David Sibley, 1992, A Bertrand Model of Pricing and Entry, *Economics Letters*, 41, 199-206.

Thisse, J-F. and Vives, X., 1988, On The Strategic Choice of Spatial Price Policy, *American Economic Review*, 78, 122-137.

Varian, Hal R. 1980, A Model of Sales, *American Economic Review*, 70, 651-659.

Villas-Boas, M., 1999, Dynamic Competition with Customer Recognition, *Rand Journal of Economics*, 30, 604-631.

A Appendix

A.1 Supporting analysis for Stage 2

This section gives rigorous general results to support the analysis in Section 3. Consider Stage 2 competition for a consumer with values ranked (without loss of generality) $r_1 > r_2 > ... > r_n$, with $y_1 = r_1 - r_2$. As earlier, write $P_j = \min\left(p^l_j, A\right)$ and $s^l_j = r_j - p^l_j$ for the consumer surplus associated with Firm j's list price offer. Also, let $S_j = r_j - P_j = \max\left(s^l_j, r_j - A\right)$, referred to as j's "last best surplus offer" be the most generous offer Firm j could conceivably make to the consumer. Let $\bar{S}_{-j} = \max_{k\neq j} S_k$ be the most competitive last best surplus offer by any rival to Firm j.

Lemma 6 *(Captive consumers) The consumer is captive to Firm j iff $s^l_j > \bar{S}_{-j}$. If $p^l_1 \leq A$, she is captive to some firm. If she is not captive to any firm, then $r_1 - A > \bar{S}_{-1}$.*

Proof. Suppose $s^l_j > \bar{S}_{-j}$. Then $s^l_j > \max_{k\neq j} s^l_k$, so Firm j is the consumer's default, and $s^l_j > \max_{k\neq j} r_j - A$, so no rival firm can poach the consumer at a discount price $p^d_k \geq A$. Conversely, suppose $s^l_j < S_k$ for some $k \neq j$. If $s^l_j < s^l_k$, the consumer prefers k's list price and so is not captive to j. Alternatively, if $s^l_j < r_k - A$, then she cannot be captive to Firm j since Firm k could profitably poach with a discount offer $p^d_k = r_k - s^l_j > A$.

If $p^l_1 \leq A$, then $s^l_1 \geq r_1 - A > \max_{j\geq 1}(r_j - A)$. Let $s^l_d = \max_j s^l_j$ be the consumer's default offer. If $d = 1$, then $s^l_1 > \bar{S}_{-1}$, and she is captive to Firm 1. Otherwise, $s^l_d > s^l_1 \geq \max_j (r_j - A)$ implies she is captive to firm d. Finally, note $\bar{S}_{-1} = \max\left(r_2 - A, s^l_{-1}\right)$. Non-captivity implies $r_k - A > s^l_d \geq s^l_{-1}$ for some $k \neq d$. But then $r_1 > r_j\ \forall j \geq 2$ implies $r_1 - A > \bar{S}_{-1}$. ■

We say the consumer is non-captive if she is not captive to any firm. Let π_j be Firm j's equilibrium expected profit from this consumer. The implication of the last part of Lemma 6 is that if the consumer is non-captive, her favorite firm will have a competitive advantage in discounting to her, and so π_1 will be positive – see Lemma 8.

Lemma 7 *If the consumer is non-captive, then at most one firm earns a strictly positive expected profit from her.*

Proof. Suppose toward a contradiction that $\pi_j > 0$ and $\pi_k > 0$. One firm, say k, is not the consumer's default choice; then we must have $a_k = 1$ (since advertising to the consumer is strictly more profitable than not doing so). But then Firm j would earn zero profit by not advertising (as all of k's offers will beat its list price), so $\pi_j > 0$ implies we must also have $a_j = 1$. But if both firms were to advertise with probability one, the lower-ranked firm would fail to cover its ad cost (as Bertrand competition drives its discount price to zero), contradicting the positivity of profits. ■

Lemma 8 *(Non-captive profits) If the consumer is non-captive, then* $\pi_1 = S_1 - \bar{S}_{-1} = \min\left(y_1, r_1 - A - s^l_{-1}\right) > 0$*, where* $s^l_{-1} = \max_{j\geq 2} s^l_j$*, and* $\pi_j = 0$ *for all* $j \geq 2$.

Proof. Non-captivity implies $p^l_1 > A$ (by Lemma 6), so Firm 1 can afford to discount and $S_1 = r_1 - A$. Firm 1 can guarantee $\pi_1 \geq \bar{\pi}_1 = r_1 - A - \bar{S}_{-1} > 0$ by advertising a discount $p^d_1 = r_1 - \bar{S}_{-1}$, with associated surplus $r_1 - p^d_1 = \bar{S}_{-1}$, that undercuts the most competitive offer any rival could conceivably make. Strict positivity of $\bar{\pi}_1$ follows from Lemma 6. Suppose (toward a contradiction) that $\pi_1 > \bar{\pi}_1$, in which case the supremum $\dot{s}$ over Firm 1's surplus offers must satisfy $\dot{s} < \bar{S}_{-1}$. We have $\dot{s} \geq s^l_{-1}$ (or else this offer would not beat the best rival list price), so $\dot{s} < r_2 - A$. But this permits Firm 2 to make a strictly positive profit. Specifically, no Firm $j > 2$ can possibly offer a discount surplus greater than $s' = r_3 - A$, so Firm 2 may offer overcutting surplus $s_2 = \max(\dot{s}, s')$, win with probability one, and earn $\pi_2 = \min(r_2 - A - \dot{s}, r_2 - r_3) > 0$, contradicting Lemma 7. We conclude that $\pi_1 = \bar{\pi}_1 > 0$; zero profits for all other firms follow from Lemma 7. ■

Lemma 9 *(Who advertises?) If the consumer is non-captive and Firm 1's closest competition is discounting by Firm 2 (* $\bar{S}_{-1} = r_2 - A$ *), then the consumer is contested by firms 1 and 2 only (*$a_1 > 0$*,* $a_2 > 0$*,* $a_j = 0\ \forall j \geq 3$*). Otherwise, if* $\bar{S}_{-1} = s^l_{-1}$*, the consumer is conceded to Firm 1 by her default firm.*

Proof. First suppose Firm 1 is the consumer's default choice. Because she is not captive, we have $r_2 - A > s^l_1 > s^l_{-1}$. We must have $a_1 > 0$, as Firm 2 could strictly profit by undercutting p^l_1 if Firm 1 never advertised, and this is inconsistent with Lemma 8. Furthermore, we must have $a_j > 0$ for some $j \geq 2$, as otherwise Firm 1 would have no incentive to advertise itself. We cannot have $a_k > 0$ for any $k \geq 3$. (Suppose otherwise for some firm $\hat{k} \geq 3$, and let $\hat{p}^d$ be Firm $\hat{k}$'s lowest discount price, earning $\pi_{\hat{k}} = 0$ by Lemma 7. But then Firm 2 could earn $\pi_2 > 0$ by undercutting

$\hat{p}^d$ with $p_2^d = \hat{p}^d + (r_2 - r_{\hat{k}})$, winning with at least the same probability as firm $\hat{k}$ but earning a larger profit margin per sale. As this contradicts Lemma 7, we have $a_k = 0$ for all $k \geq 3$.) Thus we have $a_2 > 0$ and $a_k = 0$ $\forall k \geq 3$. If Firm 1 is not the consumer's default, then we have $a_1 = 1$ (since Firm 1 earns zero profit if it does not advertise). Absent competing ads, it will simply undercut the best rival list price offer s_{-1}^l. If $r_2 - A < s_{-1}^l$, then Firm 2 (and *a fortiori*, lower-ranked firms) cannot profitably improve on this surplus offer and the consumer is conceded. If $r_2 - A > s_{-1}^l$, then Firm 2 has room to strictly profit by undercutting Firm 1. In this case, the argument proceeds just as above: some firm $j \geq 2$ must advertise with positive probability, but $a_k = 0$ $\forall k \geq 3$ by the same argument, so $a_2 > 0$. ■

Symmetric list prices: equilibrium strategies on a contested consumer

Suppose $p_j^l = p > A$ for all $j \geq 1$. The results above imply (i) the consumer's default choice is Firm 1 (her favorite product), (ii) she is captive to Firm 1 if $s_1^l > r_2 - A$ ($\Leftrightarrow y_1 > p - A$), (iii) otherwise she is contested by her top two firms (1 and 2) only, with $a_1 > 0$, $a_2 > 0$, and $a_j = 0$ $\forall k \geq 3$, and (iv) if contested, final expected profits on the consumer are $\pi_1 = y_1$ and $\pi_2 = \pi_3 = ... = 0$. Write the combined strategies of firms 1 and 2 (over advertising and discount offers) as surplus offer distributions $B_1(s)$ and $B_2(s)$ as in the text. We will be content to fill in gaps in the argument from the text. "Advertised" surplus offers are those with $s_1 > s_1^l$ for Firm 1 or $s_2 \geq s_1^l$ for Firm 2; let $\bar{s}_j$ ($\underline{s}_j$) be j's supremum (infimum) over advertised offers. We have $\bar{s}_1 = \bar{s}_2 = S_2 = r_2 - A$. (If, e.g. $\bar{s}_1 < S_2$, then Firm 2 could strictly improve on its Lemma 8 profit by overcutting $\bar{s}_1$ and selling with probability one (and similarly for Firm 1 if $\bar{s}_2 < S_2$).) We also have $\underline{s}_1 = \underline{s}_2 = s_1^l$. (Advertising an offer $s_1 \in \left(s_1^l, \underline{s}_2\right)$ wins only if Firm 2 does not advertise, in which case Firm 1's list price would have won anyway (at a higher price and without spending A). So $\underline{s}_1 \geq \underline{s}_2$. Then any offer $s_2 \in \left[s_1^l, \underline{s}_1\right]$ wins iff. Firm 1 does not advertise. Since there is no reason not to make the most profitable such offer, we have $\underline{s}_2 = s_1^l$. And $\underline{s}_1 > \underline{s}_2$ is impossible, as Firm 2 would have no incentive to make offers in the gap $(\underline{s}_2, \underline{s}_1)$, but then Firm 1 could reduce its lowest offer without winning less often.) So advertised surplus offers satisfy $s_1, s_2 \in \left[s_1^l, S_2\right]$. The arguments against gaps and atoms on the interior of this interval are standard, as are the indifference conditions pinning down $B_1(s)$ and $B_2(s)$ on this interval. We have $1 - a_1 = B_1\left(s_1^l\right) = \frac{A}{p - y_1}$. As this is strictly positive for $y_1 > p - A$, Firm 1 must be indifferent between its advertised offers and not advertising. Its list price wins if $s_2 < s_1^l$, but not against an undercutting offer

$s_2 = s_1^l$, so the profit to not advertising is $\pi_1 = \lim_{s_2 \nearrow s_1^l} B_2(s)(r_1 - s) = y_1$, so $\lim_{s_2 \nearrow s_1^l} B_2(s) = \frac{y_1}{p}$. However Firm 1's profits on advertised offers for s_1 arbitrarily close to $\underline{s}_1 = s_1^l$ imply $B_2\left(s_1^l\right) = \frac{A+y_1}{p}$, so Firm 2's strategy must include a measure $\frac{A}{p}$ atom of offers $s_2 = s_1^l$ just undercutting Firm 1's list price. The remaining results are straightforward.

A.2 Proofs

Proof of Lemma 1

We appeal to known properties of logconcave distributions; see the references for further information.[31] Cumulative distribution functions and their complements are strictly logconcave if their density functions are, so $F(x)$ and $F(x+y)$ are strictly logconcave. Products of strictly logconcave functions are strictly logconcave, so $f_{(1:n-1)}(x)$ is strictly logconcave, as are the integrands $F(x+y) f_{(1:n-1)}(x)$ and $(1 - F(x+y)) f_{(1:n-1)}(x)$. Marginals of strictly logconcave functions are strictly logconcave, so integrating over x, we have $G(y)$ and $1 - G(y)$ strictly logconcave. Similar arguments applies to $g(y) = \int f(r+y) f_{(1:n-1)}(r)\, dr$.

Proof of Proposition 1

If $A > h$, let p^* solve $\Theta(p) = 0$. By the arguments in the text, $p_j^l = p^*$ is the unique symmetric solution to the firms' first-order necessary conditions for profit maximization. Referring to (3), each firm's marginal profits $\partial \Pi_j \left(p_j^l\right) / \partial p_j^l$ are decreasing in p_j^l (by strict logconcavity of $1 - G(y)$ and P_j weakly increasing in p_j^l), so these first-order conditions are also sufficient. The features of equilibrium follow from arguments in the text.

If $A < h$, there is no symmetric equilibrium at any list price satisfying $p^l - A < \bar{y}$, since $\Theta\left(p^l\right)$ strictly positive implies any firm would gain by deviating to a higher list price. At $p^l = \bar{y} + A$, all consumers with value advantage $y < \bar{y}$ are contested, and consumers with the largest possible taste advantage $\bar{y}$ are on the captive contested border. As the latter are zero-measure, each firm's profit is $\Pi = \int_0^{\bar{y}} y dG(y)$. Deviating to a lower list price $p_j^l < p^l$ is ruled out by $\Theta\left(p^l\right)$ strictly positive. Deviating to a higher list price ensures that consumers at the upper bound $\bar{y}$ will be contested for sure, and does not change profits on other consumers; as the former are zero-measure, this cannot be a strict improvement.

[31] For example, see Bergstrom and Bagnoli (2005).

For uniqueness with two firms, suppose toward a contradiction that there exists an equilibrium with list prices $p_1^l < p_2^l \leq \bar{y} + A$, so Firm 1's first-order condition must be satisfied, and Firm 2's marginal profit must be weakly positive. Define a function $\upsilon(u, v)$ by

$$\upsilon(x, y) = \frac{1 - G(u - \min(v, A))}{g(u - \min(v, A))} - \min(u, A)$$

so the first-order conditions imply $\upsilon\left(p_1^l, p_2^l\right) = 0$ and $\upsilon\left(p_2^l, p_1^l\right) \geq 0$. But $\upsilon(u, v)$ is strictly decreasing in u and weakly increasing in v (by strict logconcavity of $1 - G(y)$). So if $p_1^l < p_2^l$, we have

$$\upsilon\left(p_1^l, p_2^l\right) > \upsilon\left(p_2^l, p_2^l\right) \geq \upsilon\left(p_2^l, p_1^l\right) \geq 0.$$

Lemma 10 *Let $G(y) = \int_{\underline{r}}^{\bar{r}} F(r + y) f_{(1:n-1)}(r)\, dr$ as in the text, with $g(y) = G'(y)$ and $f_{(1:n-1)}(r) = (n-1) f(r) F(r)^{n-2}$. Then $g'(0) \leq 0$. Furthermore, $g'(0) < 0$ if $n \geq 3$.*

Proof. We allow for the possibility that the upper limit of the support $\bar{r}$ is either finite or infinite. If the former, then for $y \geq 0$, we have $F(r + y) = 1$ and (by convention), $\frac{dF(r+y)}{dy} = f(r + y) = 0$ wherever $r + y \geq \bar{r}$. Then for we can write

$$g(y) = \int_{\underline{r}}^{\bar{r}-y} f(r + y) f_{(1:n-1)}(r)\, dr \quad \text{for } y \geq 0$$

where the upper limit collapses to ∞ if $\bar{r} = \infty$. Differentiating once more,

$$g'(y) = \int_{\underline{r}}^{\bar{r}-y} f'(r + y) f_{(1:n-1)}(r)\, dr - f(\bar{r}) f_{(1:n-1)}(\bar{r} - y)$$

where the second term should be understood as $\lim_{r\to\infty} f(r) f_{(1:n-1)}(r - y) = 0$ if $\bar{r} = \infty$ (since $\lim_{r\to\infty} f(r) = 0$ if the distribution is unbounded). Our aim is to sign $g'(0)$; using the definition of $f_{(1:n-1)}(r)$, we have

$$\frac{g'(0)}{n-1} = \int_{\underline{r}}^{\bar{r}} f'(r) f(r) F(r)^{n-2}\, dr - f(\bar{r})^2 F(\bar{r})^{n-2}$$

But $f'(r) f(r) = \frac{1}{2} d\left(f(r)^2\right)$, so if $n = 2$ we have

$$\frac{g'(0)}{n-1} = -\frac{1}{2}\left(f(\bar{r})^2 + f(\underline{r})^2\right) \leq 0$$

Otherwise, integrate by parts to get

$$\frac{g'(0)}{n-1} = -\frac{1}{2}\left(\left(f(\bar{r})^2 F(\bar{r})^{n-2} + f(\underline{r})^2 F(\underline{r})^{n-2}\right) + (n-2)\int_{\underline{r}}^{\bar{r}} f(r)^3 F(r)^{n-3}\, dr\right)$$

The first term inside the parentheses is weakly positive, and the second is strictly positive, so $g'(0) < 0$ as claimed. ■

Proof of Proposition 4

Existence of the full-targeting equilibrium. At the candidate equilibrium, Firm 1's marginal profit at list price p_1 is $\Pi_1'(p_1) = g(p_1 - A)\gamma(p_1)$ where $\gamma(p_1) = \frac{1-G(p_1-A)}{g(p_1-A)} - \min(p_1, A)$. Thus $\gamma(p_1)$ has the same sign as marginal profits. Consider deviations to $p_1 < \bar{y} + A$. We may restrict attention to list prices $p_1 \leq A$, as any local optimum $\Pi'(\hat{p}_1) = 0$ with $\hat{p}_1 \in (A, \bar{y} + A)$ has $\Pi''(\hat{p}_1) > 0$ and must be a local minimum. ($\Pi'(\hat{p}_1) = 0$ implies $\gamma(\hat{p}_1) = 0$. Then $\Pi''(\hat{p}_1) = g'(\hat{p}_1 - A)\gamma(\hat{p}_1) + g(\hat{p}_1 - A)\gamma'(\hat{p}_1) = g(\hat{p}_1 - A)\gamma'(\hat{p}_1)$. And $\gamma'(p_1) > 0$ for $p_1 > A$ by logconvexity of $1 - G(y)$.) As such prices do not permit advertising, the best deviation profit is

$$\Pi_1^{dev}(A) = \max_{p_1 \leq A} p_1 (1 - G(p_1 - A))$$

$\Pi_1^{dev}(A)$ is strictly increasing in A, with (by definition) $\Pi_1^{dev}(A)\big|_{A=p^{NT}} = \Pi^{NT} < \Pi^{FT}$. (The last ranking is proved in the text.) By continuity, there exists $A' > p^{NT}$ such that $\Pi_1^{dev}(A) \leq \Pi^{FT}$ for all $A \leq A'$.

Uniqueness of the full-targeting equilibrium. Candidate symmetric equilibria with $p^l \in (A, \bar{y} + A)$ may be ruled out by the same failure of second-order conditions as above. The only candidate equilibrium at list prices below A is $p^l = p^{NT}$, with profit Π^{NT} per firm. A deviation to $p_1^l = \bar{y} + p^{NT}$ and using targeting yields profit $\tilde{\Pi}_1^{dev}(A) = \int_{A-p^{NT}}^{\bar{y}} \left(y + p^{NT} - A\right) dG(y)$ for Firm 1. Integrate by parts to get $\tilde{\Pi}_1^{dev}(A) = \int_{A-p^{NT}}^{\bar{y}} 1 - G(y)\ \, dy$. Observe that $\tilde{\Pi}_1^{dev}(A) \to \Pi^{FT} > \Pi^{NT}$ as $A \to p^{NT}$, so there exists $A'' > p^{NT}$ such that the no-targeting equilibrium fails for $A < A''$.

Then for existence and uniqueness, set $A^* = \min(A', A'') > p^{NT}$.

Lemma 11 *At an interior symmetric equilibrium,* $\frac{dp^l(A)}{dA} < 1$*. That is, the equilibrium list price rises no faster than the ad cost.*

Proof. If captive demand is convex this is trivial; the interesting case is when $\frac{dp^l(A)}{dA}$ is positive. Extending the analysis of Proposition 2, we have

$$\frac{dp^l(A)}{dA} = \left.\frac{g'(y)(1-G(y))}{g'(y)(1-G(y))+g(y)^2}\right|_{y=p^l-A} = 1 - \left.\frac{g(y)^2}{g'(y)(1-G(y))+g(y)^2}\right|_{y=p^l-A}$$

The denominator is positive by Condition 1, so the result follows. ■

Proof of Proposition 6

Differentiate to get

$$\frac{d\bar{a}(A)}{dA} = n\left(\frac{\partial p^l}{\partial A} - 1\right)\frac{A}{p} + n\int_0^{y^*}\frac{1}{p^l-y}\left(\frac{A}{p^l-y}\frac{\partial p^l}{\partial A} - 1\right)g(y)\,dy$$

Within the integral, we have $\frac{A}{p^l-y} = \frac{A}{A+y^*-y} \leq 1$, so $\frac{A}{p^l-y}\frac{\partial p^l}{\partial A} - 1 \leq \frac{\partial p^l}{\partial A} - 1$. But then both terms in $\frac{d\bar{a}(A)}{dA}$ are strictly negative by Lemma 11.

Proof of Proposition 7

Let $y^*(A)$ be the taste difference at the boundary of the contested region. By later arguments, $\lim_{A\to 0} y^*(A) = \bar{y} = \bar{r} - \underline{r}$; that is, as $A \to 0$ all consumers are contested, and so the boundary is at the largest possible taste difference (possibly infinite). At any interior equilibrium, $p^l(A) = y^*(A) + A > y^*(A)$, so $\lim_{A\to 0} p^l(A) \geq \bar{y}$. Defer the special case of consumers at a turf boundary. For $y \in (0, \bar{y})$, once A is small enough to contest this consumer, the welfare loss is $L\left(y, p^l(A), A\right)$, with

$$0 \leq L\left(y, p^l(A), A\right) \leq A\left(1 + \frac{A+y}{y}\ln\frac{A+y}{A}\right)$$

where the upper bound follows from $\frac{p^l(A)-y}{p^l(A)} \leq 1$. Because this upper bound tends to 0 with A, so does $L\left(y, p^l(A), A\right)$. At the turf boundary, $\lim_{y\to 0} L\left(y, p^l(A), A\right)$ is well-behaved, but the limits involved are tedious. A simpler path is to observe that for consumers at $y = 0$ there is no social cost of misallocation, so the only costs to account for are total advertising expenditures. Thus $L\left(0, p^l(A), A\right) = A(a_1 + a_2) = 2A - \frac{A^2}{p^l(A)}$, which tends to 0 with A (given $p^l(A)$ bounded away from zero).

Proof of Proposition 8

An increase in A unambiguously improves consumer surplus of captive consumers since it reduces list prices, so we need only show the result for contested consumers. As the consumer surplus of contested consumers moves inversely to the welfare loss function, it suffices to show that, for $p^{NT} - A$ sufficiently small, $L\left(y, p^l(A), A\right)$ is decreasing in A for all $y \in [0, y^*(A)]$. Because $dL\left(y, p^l, A\right)/dA$ is continuous in y and A, and because $y^*(A)$ can be made arbitrarily close to 0 by choosing A sufficiently close to p^{NT}, it suffices to show that $\left.\frac{dL\left(y,p^l(A),A\right)}{dA}\right|_{y=0,A=p^{NT}} < 0$, that is, that $L\left(y, p^l(A), A\right)$ is strictly decreasing in A at $A = p^{NT}$ for consumers at the turf boundary. At $y = 0$, we have $L\left(0, p^l, A\right) = \mathcal{A}(y) = 2A - \frac{A^2}{p^l}$ since there are no social costs of misallocation. We have

$$\begin{aligned}
\frac{dL\left(0, p^l(A), A\right)}{dA} &= \frac{\partial L\left(0, p^l(A), A\right)}{\partial A} + \frac{\partial L\left(0, p^l(A), A\right)}{\partial p^l}\frac{dp^l}{dA} \\
&= \left(2 - 2\frac{A}{p^l}\right) + \left(\frac{A}{p^l}\right)^2 \frac{dp^l}{dA}
\end{aligned}$$

where $\frac{dp^l}{dA} = \frac{Ag'(y^*)}{Ag'(y^*)+g(y^*)}$ (using (5)). Evaluating at $A = p^{NT}$, we have $\left.\frac{dL\left(0,p^l(A),A\right)}{dA}\right|_{A=p^{NT}} = \left.\frac{dp^l}{dA}\right|_{A=p^{NT}} = \frac{p^{NT}g'(0)}{p^{NT}g'(0)+g(0)} \leq 0$ by Proposition 5. Then because $g'(0) < 0$ for $n \geq 3$ (by Lemma 10), we have the stronger result that $\left.\frac{dL\left(0,p^l(A),A\right)}{dA}\right|_{A=p^{NT}} < 0$, and so we are done.

Proof of Proposition 10

Let $h_F(r) = \frac{1-F(r)}{f(r)}$. Strict logconcavity and monotone convergence imply that $\lim_{r\to\infty} h_F(r)$ exists and is finite; consequently $\lim_{r\to\infty} h_F'(r) = 0$. Next let $h_n(y) = \frac{1-G_n(y)}{g_n(y)}$ and note that $p_n^{NT} = h_n(0)$. Noting $1 - G_n(0) = \frac{1}{n}$, we can write

$$\ln H_n = \ln\left(1 - G_n\left(p_n^{NT}\right)\right) - \ln\left(1 - G_n(0)\right) = -\frac{p_n^{NT}}{h_n(\hat{y})} = -\frac{h_n(0)}{h_n(\hat{y})}$$

for some $\hat{y} \in \left(0, p_n^{NT}\right)$ by the intermediate value theorem. Apply the IVT once more to get $h_n(\hat{y}) = h_n(0) - |h_n'(\dot{y})|\,\hat{y}$ for some $\dot{y} \in (0, \hat{y})$, using $h_n'(y) < 0$. Since $\dot{y} < \hat{y} < p_n^{NT} = h_n(0)$, we have $h_n(\hat{y})/h_n(0) \in (1 - |h_n'(\dot{y})|, 1)$. With careful application of Theorem 2 of Gabaix et al. (2015), asymptotically as $n \to \infty$ we have $h_n'(\dot{y}) \sim h_F'(\hat{r}_n + \dot{y})$, where $\hat{r}_n$ is defined by $1 - F(\hat{r}_n) = \frac{1}{n-1}$. Clearly $\hat{r}_n \to \infty$

with n. But then $\lim_{n\to\infty} h_n'(\dot{y}) = \lim_{r\to\infty} h_F'(r) = 0$. Thus we can conclude that $\lim_{n\to\infty} \ln H_n = -1$, so we are done.

B Supplementary Appendix

Section B.1 Example: profits rise with cheaper targeting when demand is not single-peaked

Section B.2 Supplementary results about welfare and consumer surplus

Section B.2.4 Proofs for Section B.2

B.1 Example with profits rising as targeting is adopted

This section provides supporting analysis for the example with rising profits presented in Section 5.3.

As the ad cost A declines, Proposition 4 establishes that the transition from a no-targeting regime to a full-targeting one can be profitable for firms. To illustrate this transition in more detail we work through an example. Start with the two firm linear-Hotelling setup with $t = 1$, so $1 - G(y) = \frac{1}{2}(1 - y)$ and consumers at the two endpoints have value advantage $y = 1$ for their favored products. Augment the model by giving each firm an additional mass of size $L = \frac{1}{2}$ of "loyals" who prefer its product by $\bar{y} = 2$. Thus the total mass of consumers is now 2.

When does the no-targeting equilibrium collapse? If targeting is impossible, there is a no-targeting equilibrium with $p^{NT} = 2$ and $\Pi^{NT} = 2$. Figure 2.a in the main text plots Firm 1's captive demand when Firm 2's last best price is $P_2 = p^{NT} = 2$. If Firm 1 prices $p_1 \leq P_2 + 1$ so as to retain both its loyals and some "interior" consumers, its captive demand will be $L + (1 - G(p_1 - P_2)) = \frac{1}{2}(2 + P_2 - p_1)$, with profit $\pi_1^{CAP}(p_1, P_2) = \frac{1}{2} p_1 (2 + P_2 - p_1)$, while if it prices out interior consumers with $p_1 > P_2 + 1$, it can retain the loyals up to $p_1 = P_2 + \bar{y} = P_2 + 2$. While $p_1 = p^{NT} = 2$ is a weak best response (so the no-targeting equilibrium is valid), it would be equally profitable to "retrench" – that is, deviate to $p_1 = 4$ and serve only the loyals. The knife-edge construction will be convenient in a moment, when ads come in, but it is not essential. The deviation to retrenchment would be a strict improvement if Firm 1 could supplement its loyal profits with any profits at all from discounting to interior consumers. Winning back those consumers will require $p_1^d \leq p^{NT} + 1 = 3$, and this becomes affordable as soon as $A \leq \bar{A} = 3$. Thus for $A < \bar{A}$, the no-targeting equilibrium collapses.

When does the full-targeting equilibrium become viable? In a symmetric full-targeting equilibrium, the firms sell at list price $p^{FT} = \bar{y} + A = 2 + A$ to their loyals and all other consumers are contested. Captive profits are $\pi^{CAP} = p^{FT}L = 1 + \frac{1}{2}A$, while contested profits are $\pi^{CON} = \int_0^1 y dG(y) = \frac{1}{4}$, so total profits are $\Pi^{FT} = \frac{5}{4} + \frac{1}{2}A$. Modest deviations $p_1 \in \left(A, p^{FT}\right)$ to a lower list price with targeting can be easily ruled out as unprofitable, as it takes a price cut of at least 1 to begin winning any additional captives. (Details omitted.) However, a large deviation $p_1 < A$ back to list-price-only sales yields profit $\pi_1^{CAP}(p_1, A) = \frac{1}{2}p_1(2 + A - p_1)$, using last best price $P_2 = A$ now for the rival. The best such deviation is $p_L = 1 + \frac{1}{2}A$ with profit $\Pi_L = \frac{1}{2}p_L^2 = \frac{1}{8}(2 + A)^2$. Thus there is a symmetric full-targeting equilibrium whenever $\Pi^{FT} \geq \Pi_L$, which is true for $A \leq A^* = \sqrt{6} \approx 2.45$. Profits in this regime are declining in A since the limit price p^{FT} required to keep loyals captive must drop, and they eventually (for $A < \frac{3}{2}$) fall below the no-targeting profits. However, for $A \in \left(\frac{3}{2}, A^*\right)$, firms are better off than under no-targeting – in particular, profits are 24% higher at $A = A^*$.

The transition If $A \in \left(A^*, \bar{A}\right)$, neither of these symmetric equilibria (no-targeting or full-targeting) exists. Instead, there is a pair of asymmetric equilibria and a symmetric mixed strategy equilibrium; for simplicity we focus on the former. Let Firm 2 use $p_2 = \bar{p}_2 < A$ and not discount. Let Firm 1 mix between a high price $p_{1H} = \bar{y} + \bar{p}_2 > A$ with discounting and a low price $p_{1L} < A$ without discounting, with $q = \Pr(p_{1H})$. (We shall see why this mixing is necessary.) Let $\bar{p}_1 = E(P_1) = qA + (1 - q)p_{1L}$ be Firm 1's expected last best price. Given the linearity of captive demand, Firm 2's expected profit is simply $\pi_2^{CAP}(p_2, \bar{p}_1)$, and so to be a best reply its list price must satisfy $\bar{p}_2 = 1 + \frac{1}{2}\bar{p}_1$. (One must also rule out deviations to $p_2 > A$, but these are not problematic.) Firm 2's equilibrium profit is $\Pi_2 = \frac{1}{2}\bar{p}_2^2$. By the same token, Firm 1's best no-discounting reply to Firm 2 is $p_{1L} = 1 + \frac{1}{2}\bar{p}_2$, with profit $\Pi_{1L} = \frac{1}{2}p_{1L}^2 = \frac{1}{8}(2 + \bar{p}_2)^2$. Alternatively, it could price to its loyals at p_{1H}, earning captive profit $\Pi_{1H}^{CAP} = \frac{1}{2}p_{1H}$, and poach back consumers $y \in (A - \bar{p}_2, 1)$, earning conceded profit $\Pi_{1H}^{CON} = \int_{A-\bar{p}_2}^1 (y + \bar{p}_2 - A)\, dG(y) = \frac{1}{4}(1 + \bar{p}_2 - A)^2$. Total profit from this strategy is $\Pi_{1H} = 1 + \frac{1}{2}\bar{p}_2 + \frac{1}{4}(1 + \bar{p}_2 - A)^2$. Since there is no equilibrium with both firms pricing below A – (the only candidate is the symmetric equilibrium at p^{NT} which fails for $A < \bar{A}$) – it must be that $\Pi_{1H} \geq \Pi_{1L}$ so that Firm 1 is willing to play p_{1H}. Now we come to the reason mixing is necessary. If $\bar{p}_2$ is too soft, $\Pi_{1H} \geq \Pi_{1L}$ will fail (as Firm 1 can do better by undercutting $\bar{p}_2$), and

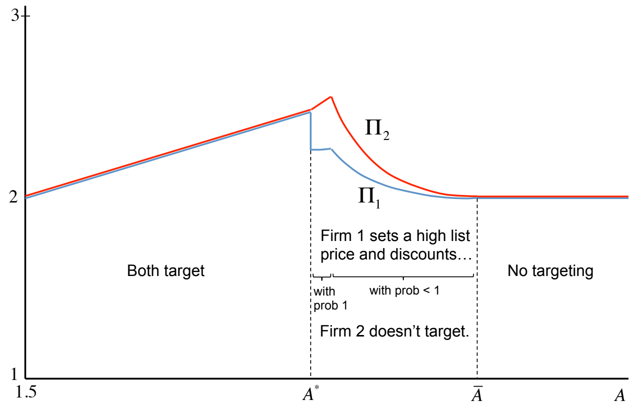

Figure 8: Equilibrium profits. Π_1 and Π_2 coincide in the symmetric no-targeting $(A \geq \bar{A})$ and full-targeting $(A \leq A^*)$ regimes. For $A \in (A^*, \bar{A})$, Π_1 (blue) and Π_2 (red) are the asymmetric equilibrium profits described in the text.

$\bar{p}_2$ generally will be too soft if Firm 1 always prices high. Specifically, $\Pi_{1H} \geq \Pi_{1L}$ requires $\bar{p}_2 = 1 + \frac{1}{2}\bar{p}_1 \leq p^* = 2(A-1) - \sqrt{2A^2 - 4A - 2}$. Firm 1 playing p_{1H} with probability one implies $\bar{p}_1 = A$, and this satisfies $1 + \frac{1}{2}\bar{p}_1 \leq p^*$ only if $A \lesssim 2.517$. Otherwise Firm 1's expected last best price must be depressed below A to keep $\bar{p}_2$ sufficiently competitive, and this means that Firm 1 must play the low price p_{1L} with some probability. Firm 1's indifference pins down $\bar{p}_2 = 1 + \frac{1}{2}\bar{p}_1 = p^*$ in terms of A alone, which in turn pins down p_{1L}, p_{1H}, and the mixing probability q. Equilibrium profits for $A \in \left(2.517, \bar{A}\right)$ are then $\Pi_2 = \frac{1}{2}(p^*)^2$ and $\Pi_1 = \Pi_{1L} = \frac{1}{2}\left(1 + \frac{1}{2}p^*\right)^2$. We have $q \to 1$ as $A \to 2.517$, and for $A \in (A^*, 2.517)$ the equilibrium has Firm 1 setting $p_{1H} = 2 + \bar{p}_2$ with probability one and earning Π_{1H}, and Firm 2 setting $\bar{p}_2 = 1 + \frac{1}{2}A$ and earning $\Pi_2 = \frac{1}{2}\bar{p}_2^2 = \frac{1}{8}(2+A)^2$.

Figure 8 plots the firms' equilibrium profits for $A \in [1.5, 3.5]$ with this equilibrium selection on $A \in \left(A^*, \bar{A}\right)$. Both firms benefit when one of them begins to use targeted discounts, but the non-targeter (Firm 2) gains more. In this sense, competition has the flavor of a game of Chicken where targeting – retrenchment to a high list price and "discount" prices $p_1^d \geq A$ higher than the original list price $p^{NT} = 2$ – is the concession strategy. This makes it clear that the main profit gains are coming not from price discrimination *per se* but from the softening of one's rival's prices. This softening is fueled by the rise in q as A declines – Firm 1 shifts increasing weight onto its softer strategy. At $A \approx 2.517$, we reach $q = 1$ and so the opportunities for further softening have been exhausted. From this point forward, reductions in A make Firm 1's pricing more competitive, not less, and profits begin to fall. There is one last fillip for Firm 1 : at $A = A^* \approx 2.45$, Firm 2 finally concedes and switches from $\bar{p}_2 \approx 2.22$ to a high list price with discounting. This softens its most competitive price from 2.22 up to A^*, permitting a corresponding jump in Π_1.

In the symmetric mixed strategy equilibrium on $A \in \left(A^*, \bar{A}\right)$, both firms mix between a high price with discounting and a low price without. The analysis is similar, and the equilibrium profit rises similarly to the asymmetric profits in Figure 8 as A declines below $\bar{A}$.

B.2 Welfare and consumer surplus: supplementary results

Proofs for the following results are in Section B.2.4.

B.2.1 When does targeting initially make all consumers worse off? Two firm results

Proposition 8 gave conditions under which the initial introduction of targeted ads (that is, a decrease in A when $A \approx p^{NT}$) makes all consumers uniformly worse off. Proposition 8 applies with i.i.d. tastes and at least three firms; here we show that the same result may apply with two firms under either i.i.d. or Hotelling tastes. As in Proposition 8, the result hinges on whether rising list prices initially swamp any benefit from smaller A. However with two firms both effects vanish to first-order near $A = p^{NT}$, so the extra conditions below arise from the need to compare second-order terms.

Proposition 12 *Proposition 8 also applies to i.i.d. tastes with two firms if $f(\bar{r}) > 0$ or $f(\underline{r}) > 0$. Otherwise, let $n = 2$ in either the Hotelling or i.i.d. setting. Suppose $A = p^{NT}$, so that targeted ads are on the cusp of being used. A marginal reduction in A, leading to the introduction of targeting, will harm consumers sufficiently close to the turf boundary iff $T'''\left(\frac{1}{2}\right) > 8T'\left(\frac{1}{2}\right) = 8p^{NT}$ in the Hotelling model, or iff $8g(0)^3 + g''(0) < 0$ in the i.i.d. model. A fortiori, all consumers will be harmed by this reduction in A (as all other consumers are captives who suffer a list price increase).*

Figure 9 gives a Hotelling example with nonlinear transportation costs.[32] As A declines below $p^{NT} \approx 42.7$ the list price rises, with a particularly sharp increase

[32] Note that the transportation cost $T(d) = e^{4.5d} - 1$ is increasing and convex, with $T(0) = 0$.

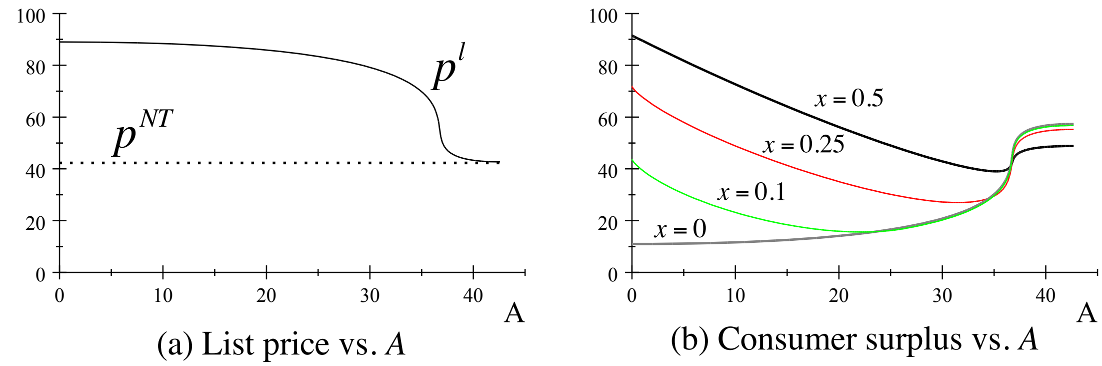

Figure 9: Targeting initially harms all consumers: two-firm example

around $A = 36$. Compared to the no-targeting equilibrium, consumer surplus initially declines as A falls for consumers at the labeled locations, including the most contested consumers at the turf boundary $x = 0.5$. As A continues to decline, consumer surplus eventually begins to recover at most locations, but consumers further from the middle never recover to their no-targeting utility (and those furthest from the middle never recover at all).

B.2.2 Distribution of welfare losses

In equilibrium, targeted advertising is always socially wasteful, but the waste is greater for some consumers than others. Let $\hat{y}$ be the taste advantage that maximizes $L\left(y, p^l, A\right)$ over all contested consumers $y \in \left[0, p^l - A\right]$; we call this the least efficiently served consumer. The pattern of welfare losses across consumers turns out to depend on whether the targeted ad cost is high or low.

Proposition 13 *The welfare loss $L\left(y, p^l, A\right)$ on contested consumers is strictly concave in y (for $y \in \left[0, p^l - A\right]$). If the targeted ad cost is high ($\frac{A}{p^l} > \sqrt{2} - 1$), then $\hat{y} = 0$: welfare losses are largest for consumers at the turf boundary, and $L\left(y, p^l, A\right)$ is strictly decreasing in y. If the targeted ad cost is low ($\frac{A}{p^l} < \sqrt{2} - 1$), then $\hat{y} \in \left(0, p^l - A\right)$: welfare losses are largest for contested consumers strictly between the turf boundary and the captive-contested boundary, and $L\left(y, p^l, A\right)$ is inverse-U-shaped in y.*

B.2.3 Does greater competition (more firms) imply more consumers are contested?

Here we examine whether targeted discounts become more or less prevalent as the market becomes more competitive. We specialize to the i.i.d. taste shock case, and rather than study total ad volume $\bar{a}(A)$ we focus on the the fraction of all consumers who are captive or contested (as these quantities are more tractable). Write $CAP(n) = n(1 - G(y^*))$ for the total fraction of consumers who are captive in an equilibrium with n firms (with $y^* = p^l - A$ determined by (5)). The fraction of consumers who receive ads with positive probability is then $CON(n) = 1 - CAP(n)$.

Gabaix et al. (2015) develop powerful asymptotic results for oligopoly markups that we can apply here. Following them, we impose a mild regularity condition on the primitive taste distribution; it will be satisfied by any commonly used distribution.[33]

Definition 1 *(Gabaix et al.) Suppose $F(r)$ has strictly logconcave density $f(r)$. We say that $f(r)$ is well-behaved iff $f(r)$ is differentiable in a neighborhood of $\bar{r}$ and $\gamma = \lim_{r\to\bar{r}} \frac{d}{dr}\left(\frac{1-F(r)}{f(r)}\right)$ exists and is finite.*

For strictly logconcave distributions with unbounded upper support ($\bar{r} = \infty$), the tail exponent γ will be zero. For the uniform distribution, we have $\gamma = -1$. For these common cases, the denominator below simplifies to $\Gamma(2+\gamma) = 1$.

Proposition 14 *Suppose tastes are distributed i.i.d. according to strictly logconcave, well-behaved density $f(r)$, and $h_F = \lim_{r\to\bar{r}} \frac{1-F(r)}{f(r)}$. Then as the number of firms in the market increases, $\lim_{n\to\infty} p^{NT} = h_F/\Gamma(2+\gamma)$. Then the fraction of contested consumers satisfies*

$$\lim_{n\to\infty} CON(n) = \begin{cases} 1 & \text{if } A < \frac{h_F}{\Gamma(2+\gamma)} \\ 0 & \text{if } A \geq \frac{h_F}{\Gamma(2+\gamma)} \end{cases}$$

In particular, if $h_F = 0$, then p^{NT} tends to zero and regardless of the ad cost, all consumers are captive for n sufficiently large.

So the impact of competition on targeted advertising depends on how much market power firms retain as n grows. If they retain no market power ($p^{NT} \to 0$), then

[33] Gabaix et al. (2015) also explicitly require the existence of $\lim_{r\to\bar{r}} \frac{1-F(r)}{f(r)}$. But strict logconcavity suffices for this (by monotone convergence), so to simplify the exposition we simply restrict attention to strictly logconcave densities.

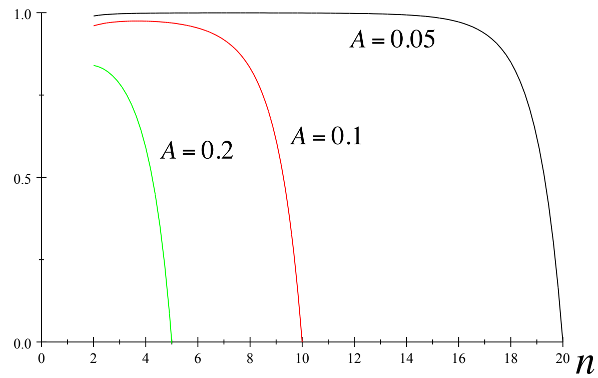

Figure 10: Fraction of consumers who are contested vs. n. (Tastes i.i.d. $r \sim U[0,1]$)

competition eventually drives out targeted advertising. Conversely, if they retain some market power, and targeting is cheap enough, then all sales are through targeted ads, even as n grows large. This asymptotic market power is determined by the tails of the taste distribution – if they are thinner than exponential (for example, uniform or normal), the first case applies; otherwise the second case does. Figure 10 illustrates the fraction of contested consumers when tastes are uniform; notice that for small A, this fraction initially increases with n before eventually declining.

B.2.4 Welfare and consumer surplus: Proofs

Lemma 12 is used in the proof of Proposition 12.

Lemma 12 *In the Hotelling model, equilibrium consumer surplus is continuously differentiable in location, list prices, and the ad cost everywhere except the turf boundary.*

Proof. This is immediate on the interior of a captive region, where $CS = r_{(1)} - p^l$, and on the interior of the contested region (barring $y = 0$), where $CS = r_{(2)} - L\left(y, p^l, A\right)$. It suffices to show that $\frac{\partial CS}{\partial y}$, $\frac{\partial CS}{\partial A}$, and $\frac{\partial CS}{\partial p^l}$ are continuous in y at $y^* = p^l - A$, the boundary between captive and contested consumers. Without loss of generality, consider the captive-contested boundary on Firm 1's turf. On the captive region, let $\zeta_{y+} = \lim_{y \searrow y^*} \frac{\partial CS}{\partial y} = \left.\frac{dr_1}{dy}\right|_{y=y^*}$. On the contested region, let $\zeta_{y-} = \lim_{y \nearrow y^*} \frac{\partial CS}{\partial y}$. We have

$$
\begin{aligned}
\zeta_{y-} &= \left.\frac{\partial r_2}{\partial y} - \frac{\partial L\left(y,p^l,A\right)}{\partial y}\right|_{y=y^*} \\
&= \left.\frac{\partial r_2}{\partial y} + \frac{A^2}{y^2}\ln\left(\frac{A+y}{A}\frac{p^l-y}{p^l}\right) - \frac{A\left(A+y\right)}{y}\left(\frac{1}{A+y} - \frac{1}{p^l-y}\right)\right|_{y=y^*=p^l-A} \\
&= \left.\frac{\partial r_2}{\partial y} + 1\right|_{y=y^*}
\end{aligned}
$$

So $\zeta_{y+} - \zeta_{y-} = \frac{\partial(r_1 - r_2)}{\partial y} - 1 = 0$ as claimed.

Similarly, let $\zeta_{A+} = \lim_{y\searrow y^*}\frac{\partial CS}{\partial A} = 0$, and $\zeta_{A-} = \lim_{y\nearrow y^*}\frac{\partial CS}{\partial A}$. For the latter,

$$
\begin{aligned}
\zeta_{A-} &= \left.-\frac{\partial L\left(y,p^l,A\right)}{\partial A}\right|_{y=y^*} \\
&= \left.-\left(1 + \frac{2A+y}{y}\ln\left(\frac{A+y}{A}\frac{p^l-y}{p^l}\right) + \frac{A\left(A+y\right)}{y}\left(\frac{1}{A+y} - \frac{1}{A}\right)\right)\right|_{y=y^*=p^l-A} \\
&= 0
\end{aligned}
$$

Next, for list prices, let $\zeta_{p^l+} = \lim_{y\searrow y^*}\frac{\partial CS}{\partial p^l} = -1$ and $\zeta_{p^l-} = \lim_{y\nearrow y^*}\frac{\partial CS}{\partial p^l}$. The limit from the contested region is

$$
\begin{aligned}
\zeta_{p^l-} &= \left.-\frac{\partial L\left(y,p^l,A\right)}{\partial p^l}\right|_{y=y^*} \\
&= \left.-\frac{A\left(A+y\right)}{y}\left(\frac{1}{p^l-y} - \frac{1}{p^l}\right)\right|_{y=y^*=p^l-A} \\
&= -1
\end{aligned}
$$

■

Proof of Proposition 12

If the density of $f(r)$ is strictly positive at $\bar{r}$ or $\underline{r}$ then (consult the proof of Lemma 10) $g'(0) < 0$ and Proposition 8 goes through unchanged. Otherwise we have $g'(0) = 0$ for the i.i.d. case. Furthermore, $g'(0) = 0$ holds for the Hotelling case as well; this can be seen by differentiating the identity $T(G(y)) - T(1 - G(y)) = y$ twice and using $G(0) = \frac{1}{2}$.

Let $A_\varepsilon = p^{NT} - \varepsilon$, with p^ε the equilibrium price under ad cost A_ε. For a consumer $y = 0$ at the turf boundary, the surplus difference between the no-targeting equilibrium and being contested in the A_ε equilibrium is $CS^\varepsilon - CS^{NT} = r_{(2)} - L(0, p^\varepsilon, A_\varepsilon) - \left(r_{(1)} - p^{NT}\right) = p^{NT} - L(0, p^\varepsilon, A_\varepsilon)$. Recall that $L\left(0, p^l, A\right) = A(a_1 + a_2) = 2A - \frac{A^2}{p^l}$ (since there are no misallocation costs at $y = 0$), so $CS^{\varepsilon=0} - CS^{NT} = 0$. We aim to provide conditions under which $CS^\varepsilon - CS^{NT}$ is strictly positive or negative for ε small. Continuity of consumer surplus in y then ensures that the same ranking holds for consumers in a neighborhood of the turf boundary.

A Taylor series expansion of consumer surplus yields $CS^\varepsilon - CS^{NT} = -\frac{dCS}{dA}\Big|_{\varepsilon=0} \varepsilon + \frac{1}{2}\frac{d^2CS}{dA^2}\Big|_{\varepsilon=0} \varepsilon^2 + O(\varepsilon^3)$. We claim – to be shown shortly – that the first derivative vanishes, so for ε sufficiently small, $CS^\varepsilon - CS^{NT}$ has the same sign as $\frac{d^2CS}{dA^2}\Big|_{\varepsilon=0}$.

Claim 1 $\frac{dCS}{dA}\Big|_{\varepsilon=0} = 0$

Proof: The total derivative is $\frac{dCS}{dA} = \frac{\partial CS}{\partial A} + \frac{\partial CS}{\partial p^l}\frac{dp^l}{dA}$. (Note that this is the relevant lefthand derivative; the effect of increases in A above $A = p^{NT}$ are identically zero.) The first term is $\frac{\partial CS}{\partial A} = -\left(2 - 2\frac{A}{p^l}\right)$ which vanishes at $A = p^l = p^{NT}$. For the second term, we have $\frac{\partial CS}{\partial p^l}\Big|_{\varepsilon=0} = -\left(\frac{A}{p^l}\right)^2\Big|_{\varepsilon=0} = -1$. For the third term, using (**??**) we have $\frac{dp^l}{dA} = \frac{Ag'(y^*)}{Ag'(y^*)+g(y^*)}$. But evaluated at $A = p^{NT}$, the boundary of the contested region is $y^* = 0$, so $\frac{dp^l}{dA}\Big|_{\varepsilon=0} = \frac{p^{NT}g'(0)}{p^{NT}g'(0)+g(0)} = 0$ since $g'(0) = 0$.

Next we establish the sign of $\frac{d^2CS}{dA^2}\Big|_{\varepsilon=0}$. We have

$$\frac{d^2CS}{dA^2} = \frac{\partial}{\partial A}\left(\frac{dCS}{dA}\right) + \frac{\partial}{\partial p^l}\left(\frac{dCS}{dA}\right)\frac{dp^l}{dA}$$

The second term vanishes at $A = p^{NT}$ because $\frac{dp^l}{dA}\Big|_{\varepsilon=0} = 0$, so

$$\begin{aligned}
\frac{d^2CS}{dA^2}\Big|_{\varepsilon=0} &= \frac{\partial}{\partial A}\left(\frac{dCS}{dA}\right)\Big|_{\varepsilon=0} \\
&= \frac{\partial^2 CS}{\partial A^2} + \frac{\partial^2 CS}{\partial A \partial p^l}\frac{dp^l}{dA} + \frac{\partial CS}{\partial p^l}\frac{d^2p^l}{dA^2}\Big|_{\varepsilon=0} \\
&= \frac{\partial^2 CS}{\partial A^2} + \frac{\partial CS}{\partial p^l}\frac{d^2p^l}{dA^2}\Big|_{\varepsilon=0}
\end{aligned}$$

From above, we have $\left.\frac{\partial CS}{\partial p^l}\right|_{\varepsilon=0} = -1$, and $\left.\frac{\partial^2 CS}{\partial A^2}\right|_{\varepsilon=0} = \left.\frac{2}{p^l}\right|_{\varepsilon=0} = \frac{2}{p^{NT}}$. For the price effect, we go back to (5): $1 - G\left(p^l - A\right) = Ag\left(p^l - A\right)$. Differentiate totally with respect to A to get

$$Z\frac{dp^l}{dA} = Ag'\left(p^l - A\right)$$

where $Z = \left(Ag'\left(p^l - A\right) + g\left(p^l - A\right)\right)$, and then a second time to get

$$\frac{dZ}{dA}\frac{dp^l}{dA} + Z\frac{d^2p^l}{dA^2} = g'\left(p^l - A\right) + A\left(\frac{dp^l}{dA} - 1\right)g''\left(p^l - A\right)$$

Then evaluate at $\varepsilon = 0$, $A = p^{NT}$, using $g'(0) = 0$ and $\left.\frac{dp^l}{dA}\right|_{\varepsilon=0} = 0$, to get

$$\left.\frac{d^2p^l}{dA^2}\right|_{\varepsilon=0} = -p^{NT}\frac{g''(0)}{g(0)}$$

Putting the pieces together, we have

$$\left.\frac{d^2CS}{dA^2}\right|_{\varepsilon=0} = \frac{2}{p^{NT}} + p^{NT}\frac{g''(0)}{g(0)}$$

Now note that the no-targeting list price is $p^{NT} = \frac{1-G(0)}{g(0)} = \frac{1}{2g(0)} = \Delta'(0)$. For the i.i.d. case this gives

$$\left.\frac{d^2CS}{dA^2}\right|_{\varepsilon=0} = \frac{1}{2g(0)^2}\left(8g(0)^3 + g''(0)\right)$$

For Hotelling, repeated differentiation of the identity $T(G(y)) - T(1 - G(y)) = y$ (again using $G(0) = \frac{1}{2}$) yields $g(0) = \frac{1}{2T'\left(\frac{1}{2}\right)}$ and $g''(0) = -\frac{T'''\left(\frac{1}{2}\right)}{T'\left(\frac{1}{2}\right)}g(0)^3$; substitute to get the representation in terms of transport costs.

Lemmas 13, 14, 15, and 16 are used in the proof of Proposition 13.

Lemma 13 *The welfare loss function $L(y, p, A)$ may also be written $L(y, p, A) = A + \int_A^{p-y} \frac{A}{z}\frac{y+A}{y+z}dz$.*

Proof. This is a straightforward computation. ∎

Lemma 14 *The welfare loss function $L(y, p, A)$ is strictly concave in y for $y \in [0, p - A]$.*

Proof. Using the version of $L(y,p,A)$ from the previous lemma, we have $L_y = -\frac{A}{p-y}\frac{y+A}{p} + \int_A^{p-y}\frac{A}{z}\frac{z-A}{(y+z)^2}dz$ and

$$L_{yy} = -\left(\frac{A(A+p)}{p(p-y)^2} + \frac{A}{p-y}\frac{p-y-A}{p^2} + \int_A^{p-y} 2\frac{A}{x}\frac{x-A}{(x+y)^3}dx\right)$$

The first two terms are strictly positive, and the third weakly positive, on $y \in [0, p-A]$, so $L_{yy} < 0$. ■

Lemma 15 *At the captive-contested boundary,* $L_y(y,p,A)|_{y=p-A} = -1$.

Proof. Evaluate L_y from the previous lemma. ■

Lemma 16 *At the turf boundary,* $L_y(y,p,A)|_{y=0}$ *is strictly positive if* $A/p < \sqrt{2}-1$, *or strictly negative if* $A/p > \sqrt{2}-1$.

Proof. From the expression for L_y we have

$$\begin{aligned} L_y(y,p,A)|_{y=0} &= -\left(\frac{A}{p}\right)^2 + A\int_A^p \frac{1}{z^2} - \frac{A}{z^3}dz \\ &= -\left(\frac{A}{p}\right)^2 + \frac{1}{2}\left(1-\frac{A}{p}\right)^2 \end{aligned}$$

so $L_y(0,p,A) \gtrless 0$ iff $A/p \lessgtr \sqrt{2}-1$. ■

Proof of Proposition 13

We appeal to Lemmas 14, 15, and 16. Concavity of $L(y,p,A)$ in y is given in Lemma 14. If $A/p > \sqrt{2}-1$, then $L(y,p,A)$ is decreasing in y at $y=0$ by Lemma 16; with strict concavity of L, this suffices for the first result. If $A/p < \sqrt{2}-1$, then we have L increasing at $y=0$ by Lemma 16 and decreasing at $y=p-A$ by Lemma 15; along with strict concavity of L, this suffices for the second result.

Proof of Proposition 14

For the limiting price, apply Theorem 1 of Gabaix et al. Then if $A > h_F/\Gamma(2+\gamma)$, we will have $A \geq p^{NT}$ (and so no use of targeted ads) for n sufficiently large. Conversely, if $A < h_F/\Gamma(2+\gamma)$, then we have $A < \lim_{n\to\infty}\frac{1-G(0)}{g(0)} \leq \lim_{n\to\infty}\frac{1-G(y)}{g(y)}$ for all $y \geq 0$, so for n sufficiently large the symmetric equilibrium has $p^l = \infty$ and all consumers contested.