

DISCUSSION PAPER SERIES

3412-1560319791

**PRICE DISCRIMINATION IN THE
INFORMATION AGE: PRICES,
POACHING, AND PRIVACY WITH
PERSONALIZED TARGETED
DISCOUNTS**

Simon P Anderson, Alicia Baik and Nathan Larson

INDUSTRIAL ORGANIZATION

PRICE DISCRIMINATION IN THE INFORMATION AGE: PRICES, POACHING, AND PRIVACY WITH PERSONALIZED TARGETED DISCOUNTS

Simon P Anderson, Alicia Baik and Nathan Larson

Discussion Paper 3412-1560319791

Published N/A

Submitted 12 June 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

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

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

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

Copyright: Simon P Anderson, Alicia Baik and Nathan Larson

PRICE DISCRIMINATION IN THE INFORMATION AGE: PRICES, POACHING, AND PRIVACY WITH PERSONALIZED TARGETED DISCOUNTS

Abstract

We study list price competition when firms can individually target discounts (at a cost) to consumers afterwards, and we address recent regulation (such as the GDPR in Europe) that has empowered consumers to protect their privacy by allowing them to choose whether to opt in to data-gathering and targeting. In equilibrium, consumers who can be targeted receive poaching and retention discount offers from their top two firms. These offers are in mixed strategies, but final profits on such a consumer 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. Absent opt-in choice, firm list pricing resembles monopoly, as marginal consumers are lost to the lowest feasible poaching offer, not to another firm's list price. Opt-in choice reintroduces the standard margin too on those who opt out. The winners and losers when targeting is unrestricted (rather than banned) depend on the curvature of demand. For the empirically plausible case (convex but log-concave), targeting pushes up list prices, reduces profits and total welfare, and (if demand is convex enough) hurts consumers on average. Outside of this case, more convex (concave) demand tends to make targeting more advantageous to firms (consumers). We then use our model to study the welfare effects of a policy that forbids targeted advertising to consumers who have not opted in. Consumers opt in or out depending on whether expected discounts outweigh the cost of foregone privacy. For empirically relevant demand structures, allowing opt-in makes all consumers better-off.

JEL Classification: D43, L12, L13, M37

Keywords: targeted advertising, competitive price discrimination, discounting, privacy, GDPR, opt-in

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

Alicia Baik - ab3ja@virginia.edu
University of Virginia

Nathan Larson - nlarson@american.edu
American University

Acknowledgements

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 10th Economics of Advertising and Marketing Conferences in Tel-Aviv (June 2013) and Nuffield College, Oxford (June 2015) respectively, the CRETE conference in Naxos (July 2013), and the 12th Digital Economics conference in Toulouse (Jan 2019). The first and second authors gratefully acknowledge support from the NSF and Bankard Fund for Political Economy respectively.

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

Simon Anderson, Alicia Baik, and Nathan Larson¹

June 11, 2019

¹Anderson 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 10th Economics of Advertising and Marketing Conferences in Tel-Aviv (June 2013) and Nuffield College, Oxford (June 2015) respectively, the CRETE conference in Naxos (July 2013), and the 12th Digital Economics conference in Toulouse (Jan 2019). The first and second authors gratefully acknowledge support from the NSF and Bankard Fund for Political Economy respectively.

Abstract

We study list price competition when firms can individually target discounts (at a cost) to consumers afterwards, and we address recent regulation (such as the GDPR in Europe) that has empowered consumers to protect their privacy by allowing them to choose whether to opt in to data-gathering and targeting.

In equilibrium, consumers who can be targeted receive poaching and retention discount offers from their top two firms. These offers are in mixed strategies, but final profits on such a consumer 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.

Absent opt-in choice, firm list pricing resembles monopoly, as marginal consumers are lost to the lowest feasible poaching offer, not to another firm's list price. Opt-in choice reintroduces the standard margin too on those who opt out.

The winners and losers when targeting is unrestricted (rather than banned) depend on the curvature of demand. For the empirically plausible case (convex but log-concave), targeting pushes up list prices, reduces profits and total welfare, and (if demand is convex enough) hurts consumers on average. Outside of this case, more convex (concave) demand tends to make targeting more advantageous to firms (consumers).

We then use our model to study the welfare effects of a policy that forbids targeted advertising to consumers who have not opted in. Consumers opt in or out depending on whether expected discounts outweigh the cost of foregone privacy. For empirically relevant demand structures, allowing opt-in makes all consumers better-off.

Keywords: targeted advertising, competitive price discrimination, discounting, privacy, GDPR, opt-in

JEL: D43, L12, L13, M37

1 Introduction

Personalized price discounts aimed at poaching and retaining customers are increasingly prevalent.¹ Firms are able to develop a rich picture of individual search and purchase behavior through geo-tracking technology and device stitching, which allows firms to collect information at the personal level by connecting physical and email addresses with phone numbers, cookie ids (on computers), and device ids (on mobile devices). As an example, a firm may observe that a customer who has been browsing its website has entered its competitor's geo-fence. In order to sway the customer to buy the product from its store, the firm may send a personalized discount.² Personalized price discounting places the customer at the center of analysis, so that each consumer may be seen as an individual market (Prat and Valletti, 2019). These markets are linked through list prices paid by those who endogenously do not receive discounts.

Personalized price discounting is a tool within a broader movement toward personalization. Falling costs and the rapid development of advanced technologies including artificial intelligence (AI) and big data are accelerating firms' ability to track customers and personalize advertisements at a large scale.³ In the race to pinpoint consumers' individualized preferences, data collection is becoming more prevalent and precise.

While firms tout that better tracking technology allows them to deliver better

¹Starbucks, Subway, Dunkin' Donuts, Red Box, Kroger, and Giant notably employ personalized discounts to influence customer purchases. See, for example, Marr, Bernard. "Starbucks: Using Big Data, Analytics And Artificial Intelligence To Boost Performance," *Forbes*, May 28, 2018; Kelso, Alicia. "Subway Taps Into Personalized Messaging To Increase Customer Visits," *Forbes*, January 15, 2019; Englund, Shawn. "How Redbox delivers in-store personalized marketing using IoT," *Microsoft Industry Blogs*, September 19, 2017; Clifford, Stephanie. "Shopper Alert: Price May Drop for You Alone," *New York Times*, August 9, 2012.

²Widespread and extensive use (the average customer spends 5 hours on her smartphone) of mobile devices combined with improved tracking technology has spawned a booming geo-tracking industry which is projected to grow to \$1.7 billion by 2024. As retailers scramble to leverage consumer data more effectively, they seek to use discounts to sway customers, see Waggy, Craig. "4 Ways Retailers Can Use Location Intelligence This Holiday Season," *Adweek*, November 19, 2018; St. Louis, Molly. "Marketers Are Getting Up Close and Personal with These New Geofencing Advancements: Geofencing is quickly becoming a must-have for retailers everywhere.," *Inc.*, March 16, 2018; Valentino-deVries et al. "Your Apps Know Where You Were Last Night, and They're Not Keeping It Secret," *New York Times*, December 10, 2018.

³"White Paper Digital Transformation of Industries: Media Industry." World Economic Forum with Accenture. January 2016. <http://reports.weforum.org/digital-transformation/wp-content/blogs.dir/94/mp/files/pages/files/dti-media-industry-white-paper.pdf>.

offers for customers, privacy advocates argue that data collection hurts consumer welfare. A series of significant consumer data breaches have highlighted the vulnerability of the sensitive data firms have collected on consumers.⁴ The number of data breaches more than doubled from 2014 to 2017, with over 1,500 breaches and 178.96 million records exposed in 2017.⁵ Consumers are increasingly aware of the risk they bear when firms collect more information about them. Time magazine just devoted a full quarter of its recent issue (Time, Jan. 28, 2019) to the topic (especially Facebook) and The New York Times initiated a monthlong series beginning in April, 2019 called, “The Privacy Project” to examine better ways to control technology companies.⁶ According to a Pew Research Center study, 91 percent of Americans “agree” or “strongly agree” that people have lost control over how their data is collected even though 74 percent say it is very important to be in control of who can get information about them; of those surveyed, 64 percent say that the government should do more to regulate advertisers.⁷ After the Facebook data breach involving Cambridge Analytica, Facebook consumers have begun to take tangible steps to reduce their risk of a privacy breach with 54 percent adjusting their privacy settings, 42 percent taking a break for several weeks or more, and 26 percent deleting the Facebook app from their phones; in total 74 percent of Facebook users 18 years or older took at least one of these preventative actions.⁸ In the European Union sentiments are generally stronger regarding privacy. According to a European Commission survey in 2016, 92 percent said it is very important that private information on their computers and mobile devices only be accessed with their permission, 60 percent avoided certain websites to avoid being tracked, and 71 percent said it was unacceptable for companies to trade

⁴In November 2018 Marriott announced that sensitive information for 500 million Starwood accounts was exposed to hackers. Cambridge Analytica improperly accessed data from 87 million Facebook users. Facebook also exposed personal data of 50 million, allowing hackers to gain access to user accounts and possibly take control of them. Google acknowledged that from 2015-2018 “hundreds of thousands” of users’ personal data was exposed through using Google+. In 2017, personal information for 143 million Equifax consumers was compromised.

⁵“Annual number of data breaches and exposed records in the United States from 2005 to 2018 (in millions).” Statista. 2018.

⁶As per the cover: “It owns your data. It knows your friends. It has your credit cards. It hears your conversations. It follows you everywhere.” See also “Establishing identity is a vital, risky and changing business” and “How to think about data in 2019” (The Economist, Dec 22, 2018) which discuss how many top tech firms are build on a foundation of tying data to human beings.

⁷“Americans’ complicated feelings about social media in an era of privacy concerns.” Pew Research Center. March 27, 2018.

⁸Perrin, Andrew. “Americans are changing their relationship with Facebook.” Pew Research Center, September 5, 2018.

information about them even if it improved services.⁹

In response to increased privacy concerns, the European Union enacted the General Data Protection Regulation (GDPR) in May 2018 to restore the right to the consumer to privacy and control over individual data. The European Union reshaped the landscape for data collection in Europe, requiring firms to receive opt-in consent from customers for data tracking. While United States regulators have been slower to advocate for privacy regulation, the US Council of Economic Advisers (2015) outlined precursors for privacy regulation and a number of states (California, Vermont, and Colorado) have recently passed privacy legislation mirroring some of the aspects of the GDPR after privacy concerns at Facebook and Google.¹⁰ Research indicates that consumers object to the loss of privacy for psychological reasons (the ‘creepiness’ factor), for fear of having their information used against them in markets, and because of the risks of fraud and identity theft (Tucker, 2015, Turow et al. 2009, White et al. 2008, Acquisti et al. 2016). In light of the shifting regulatory terrain, we analyze the equilibrium impact of privacy regulation that gives customers the right to opt-in to personalized discounts.

We take a forward-looking view of the equilibrium outcome when all firms are perfectly able to employ personalized price discounting. 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. 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. 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.¹¹ We then consider the effects of privacy regulation which allows customers to choose whether they want to opt-in to receiving personalized discounts. We compare targeting with and without regulation

⁹“e-privacy Report” European Commission. December, 2016.

¹⁰See Burt, Andrew. “States are leading the way on data privacy,” The Hill, August 21, 2018; McKinnon, John D. “Lawmakers Push to Rein In Tech Firms After Google+ Disclosure,” The Wall Street Journal, October 10, 2018.

¹¹Data analytics firms are able better understand consumers by integrating online search behavior with online purchase history. This data can be combined with location tracking on smartphones and social media activity. Large-scale data collection combined with artificial intelligence enables marketers to build models with increased predictive capacity.

and assess how advertising and privacy costs shape optimal pricing, firm profits, and consumer surplus.

Two key costs are central to our analysis: the cost of targeted advertisements and a privacy cost associated with receiving targeted ads. 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.¹² 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. 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 it. In both scenarios, technological improvements enable lower costs for targeted offers.

The other key cost, deferred until our evaluation of privacy regulation in Section 7, is an exogenous cost of sacrificed privacy. This cost captures the person-specific burden associated with receiving personalized offers; it can be thought of as an expected cost of resolving identity theft resulting from a security breach or a personal nuisance cost of receiving personalized offers. Under the opt-in policy we consider, customers are empowered to choose whether their expected price-savings from discounting are worth this sacrifice of privacy.

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.¹³ The combination of list prices and discounts amounts to competitive

¹²The 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.

¹³There are many examples of discounts targeted to indecisive or price-sensitive customers. Even before tracking technology, firms assumed price sensitive customers would self-select by spending more time sorting through discounts sent in mass distributions. However, with the ability to observe customer behavior, firms are able to better control to whom discounts are sent. Algorithmic pricing combined with detailed search and purchase history of individual customers enables firms to identify customers who may be searching around between stores. For example Facebook can

price discrimination, with the novel feature being 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. Discount competition tends to shift consumer surplus away from individuals who strongly prefer their favorite product toward those with relatively strong *second*-favorite products. The latter receive the most ads and (as a consequence of mixing) buy their second-best products more often than not.

Our first result characterizes the optimal pricing strategy of the firms. Bertrand-like profits in the discounting stage simplify the firms' reduced form profits in the first stage, when they set list prices. Our marginal profit expressions with targeting are a main novelty of the paper. A firm faces a familiar marginal-inframarginal trade-off 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. When ad costs make targeting prohibitively expensive, firms compete with list prices at the turf boundaries as in classic oligopolistic competition. Interestingly, under privacy regulation, the margin of competition remains at the turf boundary for those who do not opt in, but is at the edge of the buffer zone for the others.

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.

observe if a customer has not completed a purchase. The firm's competitor can then run an ad on the Facebook feed with a cheaper price, see Wallheimer, Brian: "Are you ready for personalized pricing?" Chicago Booth Review, February 26, 2018. A Sears executive outlines how Sears targets discounts to customers who have demonstrated that they are indecisive (either by searching at a competitor's website and then entering their geo-fence or by searching on Sears' website and entering the competitor's geo-fence): see Turow (2017).

This captive demand function may be derived from whatever primitive assumptions one prefers about the underlying consumer taste distribution.¹⁴ 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 $1 - G(y)$.

Our first main policy conclusions concern who gains or loses from targeted advertising, compared to the counterfactual where it is banned. The answers hinge on the curvature of the captive demand function. We argue that the most relevant case is where captive demand is convex but logconcave. Logconcavity is commonly imposed to ensure existence of the standard oligopoly equilibrium, and convexity arises naturally if consumers’ product valuations are independent. In this case, targeting pushes up list prices but leads to lower profits (Propositions 4 and 5). Furthermore, targeting hurts consumers on average if captive demand is sufficiently convex (Proposition 10) and can hurt *every* consumer when the targeting cost is high (Proposition 11).

For firms, the availability of targeting means that rivals cannot resist the temptation to poach deep into each others’ territories – higher list prices are mainly a symptom of the fact that each firm must retreat deeper into its own territory in order to make an uncontested list price sale. Targeting reduces welfare because discount competition is inefficient, due to the cost of ads and ‘misallocation’ when a consumer buys her second-best product.¹⁵ These inefficiency costs are passed through to consumers, who suffer also from higher list prices; under the conditions of Propositions 10 and 11 these harms outweigh the benefits from receiving a discount.

Outside of the ‘convex but logconcave’ demand case, some results change, but in ways that are consistent with the logic above. Consumers tend to benefit on average from targeting when its efficiency costs are small (Result 4) or if captive demand is concave, so that targeting reduces list prices rather than inflating them. Conversely, if demand is logconvex, the option to target can improve profits by facilitating a price hike so steep that even the “discount” prices exceed the list price that would prevail

¹⁴We 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.

¹⁵We assume throughout that the market is fully covered; hence equilibrium welfare is first-best when targeted ads are not in use. 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 – this is discussed further in the conclusions.

without targeting (Proposition 6). *Inter alia*, this suggests that targeting may tend to look more favorable for firms in models with discrete consumer types, since regions of log-convex demand are an unavoidable by-product in such models.

After developing an understanding of outcomes in the laissez-faire regime, we use the model to study whether consumers would be better off with privacy regulation. Under this alternative, consumers decide whether to opt in by rationally weighing expected price discounts against the cost of foregone privacy. We find that under plausible demand conditions similar to those above, *every* consumer benefits from an opt-in policy (compared to unrestricted targeting) regardless of her preferences over products and privacy. In such cases, an opt-in policy reduces both equilibrium list prices and (because they are anchored to list prices) average discounted prices. Then consumers with strong tastes for privacy benefit from the option to opt out, while other consumers benefit from the lower prices. However, we caution that an opt-in policy could push prices up and leave many consumers worse off if captive demand is concave, or if consumers commit to privacy decisions before list prices are set (see Section 7.2). As a practical guide to evaluating such a policy, we suggest that list prices can be an effective barometer: if an opt-in policy induces list prices to fall, then the policy has unambiguously made consumers better off.

Our paper relates to the classical literature on informative targeted advertising and competitive price discrimination. Our main focus is in understanding how competitive firms optimally target price discounts with full and restricted information on customers. 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).¹⁶ 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.¹⁷ Thisse and Vives consider duopolists who can charge location-specific prices to consumers. As location is the dimension 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.¹⁸

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

¹⁶See 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.

¹⁷See also Anderson and de Palma (1988) and Anderson, de Palma, and Thisse (1989).

¹⁸Extending 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).

reduces both profits and prices. However in other respects their model is quite different from ours – targeting is coarse (two market segments), and there is no retention advertising. The results are partly a consequence of this last assumption, as a firm that cannot discount to its ‘home’ segment will try to retain those consumers with a more competitive list price.

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.¹⁹ In related work (Anderson, Baik, and Larson, 2015), we consider competition for an individual consumer via costly ads (one main emphasis there was the equilibrium selection as consumer heterogeneity vanishes). Here the presence of list prices introduces important differences by tying all the individual markets together, but our analysis and results in Section 3 utilize some similar arguments to that paper. We here develop a “bolt-on” (plug-and-play) module that renders analysis of two-stage competition readily accessible for analyzing related problems. We also link this analysis to sub-game equilibrium consumer surplus expressions, enabling a full gamut of results for other applications.

Personalized pricing is an extension of the coarser forms of targeting. Motivated by improvements in tracking technology and precision marketing, a burgeoning body of literature in economics, marketing, law, and computer science evaluates personalized price competition and the implications of restricting personalization.²⁰ Taylor and Wagman (2014) compare firm profits and consumer surplus under uniform pricing and personalized pricing with four standard frameworks. Belleflamme et al. (2017) allow duopolists to compete in price over a homogenous product using profiling technology. They find that the Bertrand paradox disappears if the firms have different precisions while the paradox persists if both have the exact precision or if one firm cannot use any profiling technology. Anderson, Baik, and Larson (2015) study personalized

¹⁹Examples 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.

²⁰See Acquisti et al. (2016) for a comprehensive survey of the personalized targeted and privacy literature.

pricing to an individual consumer, assuming a consumer cannot purchase unless she has received a targeted ad. In contrast, here consumers also have the option to buy at any firm’s published list price, and this creates a strategic linkage for firms between their “macroscopic” competition over list prices and their “microscopic” discount competition over individual consumers.

Belleflamme and Vergote (2016) and Chen et al. (2018) are the papers closest to our opt-in analysis because they evaluate the impact of allowing customers to hide from profiling technology. In Belleflamme and Vergote (2016), the monopolist is able to identify consumer’s willingness to pay through tracking technology and consumers are able to protect their privacy through hiding technology. The authors show that tracking technology lowers consumer surplus because firms are able to price discriminate. However, hiding technology worsens consumer surplus further because firms have incentive to raise regular prices to discourage hiding. In Chen et al. (2018), each firm in a Hotelling model are able to personalize prices for consumers in its target segment and offer a uniform “poaching” price for non-targeted customers. Passive consumers find it too costly to bypass being tracked whereas active consumers can actively manage their identities without cost. Active consumers can evade higher prices reserved for targeted consumers. When all consumers are passive, competition heightens and industry profits are lower. However, active consumers make it harder to poach, softening competition through higher prices for non-targeted consumers. Both papers point to a counterintuitive result of privacy regulation. Our paper similarly highlights how allowing consumers to opt-in to personalized discounts can in fact lower consumer surplus through raising list prices.

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 absent opt-in. Sections 5 and 6 present our results for prices and profits, welfare, and consumer surplus (absent opt-in) and stress the key role of the demand curve shape. Section 7 gives the analysis of consumer opt-in, and 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 $\hat{r}_{i,-j} = \max_{j' \in \{1, \dots, n\} \setminus j} r_{ij'}$, 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 density $g_j(y) = G'_j(y)$ is strictly log-concave.*²¹

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

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 $(r_{i1}, r_{i2}, \dots, r_{ij}, \dots)$ 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

²¹We observe that strict logconcavity of the density $g_j(y)$ implies strict logconcavity of the captive demand function $1 - G_j(y)$ by the Prékopa-Borell theorem. Condition 1 is sufficient for our results, but stronger than necessary in some cases. In particular, our results apply to a running example of Hotelling demand for which $1 - G(y)$ is strictly logconcave but $g(y)$ is only weakly logconcave.

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 log-concave.²² This setup generalizes easily to the case of n firms located on a circle.

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

²²For 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.

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 $[\underline{r}, \bar{r}]$.²³ Except where otherwise noted, assume that $F(r)$ and its density $f(r)$ are both strictly log-concave.

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:

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

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.²⁴ 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. Conveniently, $1 - G(y)$ inherits the log-concavity of the primitive taste distribution. We summarize this with other properties below. Parts (ii) and (iii) will be useful for understanding how targeting affects list prices and how list prices vary with the number of firms.

Lemma 1 *Strict log-concavity of $f(r)$ implies the following:*

- (i) *The functions $G(y)$, $1 - G(y)$, and $g(y) = G'(y)$ are strictly log-concave.*
- (ii) *$1 - G(y)$ is strictly convex for $y > 0$ (for $y \geq 0$ if $n \geq 3$).*
- (iii) *Let $1 - G(y)$ and $1 - \hat{G}(y)$ be captive demand with n and $n + 1$ firms. For $y \geq 0$, $\hat{G}(y) < G(y)$ and $\frac{1 - \hat{G}(y)}{\hat{g}(y)} < \frac{1 - G(y)}{g(y)}$.*

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 sub-game in Stage 2.

²³We allow for the possibility that $\bar{r} = \infty$ or $\underline{r} = -\infty$.

²⁴That is, if the taste distribution is $F(r) = \exp(-e^{-r/\beta})$, then the captive demand function is $1 - G(y) = \frac{1}{1 + (n-1)e^{y/\beta}}$. For theoretical applications see Anderson, de Palma, and Thisse (1992).

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 $\mathbf{r} = (r_1, r_2, \dots, r_n)$. 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 $r_{(1)} > r_{(2)} > \dots > r_{(n)}$ be the relabeling of firms so that Firm (1) makes this consumer's favorite product, Firm (2) makes her second favorite, and so on.²⁵ 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 The targeting sub-game

We give a self-contained analysis of the discounting sub-game in this Sub-Section. The analysis provides a “plug-and-play” synopsis of the key results which can be useful for other applications.²⁶ To this end, we temporarily simplify the notation to consider Stage 2 competition for a consumer with values ranked (without loss of generality) $r_1 > r_2 > \dots > r_n$, with $y_1 \equiv r_1 - r_2$, and suppose that firms have previously set list prices p_j^l , $j = 1, \dots, n$.

First write $s_j^l = r_j - p_j^l$ for the consumer surplus associated with Firm j 's list price offer. Define $P_j = \min(p_j^l, A)$ and let $S_j = r_j - P_j = \max(s_j^l, r_j - A)$, referred to as j 's “last best surplus offer,” be the most generous offer Firm j could conceivably make

²⁵For 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.

²⁶Credit to Davids Myatt and Ronayne for this term.

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 2 (*Captive consumers*) (a) *The consumer is captive to Firm j iff $s_j^l > \bar{S}_{-j}$.*
 (b) *If $p_1^l \leq A$, she is captive to some firm.*
 (c) *If she is not captive to any firm, then $r_1 - A > \bar{S}_{-1}$.*

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

If $p_1^l \leq A$, then $s_1^l \geq r_1 - A > \max_{j > 1} (r_j - A)$. Let $s_d^l = \max_j s_j^l$ be the consumer's default offer. If $d = 1$, then $s_1^l > \bar{S}_{-1}$, and she is captive to Firm 1. Otherwise, $s_d^l > s_1^l \geq \max_j (r_j - A)$ implies she is captive to firm d . Finally, note $\bar{S}_{-1} = \max(r_2 - A, s_{-1}^l)$. Non-captivity implies $r_k - A > s_d^l \geq s_{-1}^l$ 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 Lemma 2c 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 4 below.

Lemma 3 *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. ■

This key Lemma underscores the Bertrand nature of competition carries over despite the fact that (as seen below) more than one firm will be actively advertising and discounting. The next Lemma ties down the Bertrand rent to the “best” firm.

Lemma 4 (*Non-captive profits*) *If the consumer is non-captive, then $\pi_1 = S_1 - \bar{S}_{-1} = \min(y_1, r_1 - A - s_{-1}^l) > 0$, where $s_{-1}^l = \max_{j \geq 2} s_j^l$, and $\pi_j = 0$ for all $j \geq 2$.*

Proof. Non-captivity implies $p_1^l > A$ (by Lemma 2b), 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 $p_1^d = r_1 - \bar{S}_{-1}$, with associated surplus $r_1 - p_1^d = \bar{S}_{-1}$, that undercuts the most competitive offer any rival could conceivably make. Strict positivity of $\bar{\pi}_1$ follows from Lemma 2c. 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 $\dot{s} \geq s_{-1}^l$ (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 3. We conclude that $\pi_1 = \bar{\pi}_1 > 0$; Lemma 3 implies zero profits for all other firms. ■

The best firm thus extracts the difference between the greatest surplus it could offer and the maximal surplus that could be offered by its keenest rival, in traditional Bertrand fashion. Anderson, Baik, and Larson (2015) derive an analogous surplus result in a model in which a customer cannot buy unless she receives some price offer through advertising so that prior list prices are not included. Here this property holds even when the prior list prices are appended. As seen below, the prior list prices entail that the equilibrium discount strategy for the second-best firm involves an atom at a price just undercutting the best firm's list price, which cannot occur in the absence of prior list prices. Notice that the rent earned is independent of the list price of the best firm itself, and is independent too of the list price of the most competitive rival as long as its list price is not below the advertising cost, A (which will be the case in the full equilibrium analysis that follows as long as there is some discounting). The next result indicates which firms advertise.

Lemma 5 (*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 $\bar{S}_{-1} = s_{-1}^l$, 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 $r_2 - A > s_1^l > s_{-1}^l$. We must have $a_1 > 0$, as Firm 2 could strictly

profit by undercutting p_1^l if Firm 1 never advertised, and this is inconsistent with Lemma 4. 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 3. 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 3, 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 $a_2 > 0$. ■

We summarize the results above in the next Proposition.

Proposition 1 Rank firms by consumer values $r_1 > r_2 > \dots > r_n$, and set $y_1 \equiv r_1 - r_2$. Set list prices p_j^l , $j = 1, \dots, n$. In the equilibrium to the discounting sub-game:

- a) the consumer is captive to Firm j iff $s_j^l = r_j - p_j^l > \bar{S}_{-j} = \max_{k \neq j} (s_k^l, r_k - A)$ and j earns p_j^l from the consumer;
- b) if the consumer is not captive (to any firm) then
 - i) if $\bar{S}_{-1} = \max_{j \neq 1} (s_j^l, r_j - A) = r_2 - A$ then $\pi_1 = y_1$ and $\pi_j = 0$ for $j = 2, \dots, n$, with $a_1 \in (0, 1)$, $a_2 \in (0, 1)$ and $a_j = 0$ for $j = 3, \dots, n$;
 - ii) otherwise, if $\bar{S}_{-1} = s_j^l = r_j - p_j^l$, then the consumer is conceded by Firm j to Firm 1, with 1 offering \bar{S}_{-1} via a discount price $p_1^d = r_1 - r_j + p_j^l$,²⁷ so $\pi_1 = p_1^d$ and $\pi_j = 0$ for $j = 2, \dots, n$, with $a_1 = 1$ and $a_j = 0$ for $j = 2, \dots, n$.

We next specialize to the case of symmetric list prices and flesh out the intuition of the discounting sub-game and derive it.

²⁷Here we use the tie-break assumption that an advertised discount price beats a list price with the same surplus.

3.2 Targeting with symmetric list prices

Suppose the firms have set list prices $p_1^l = p_2^l = \dots = p_n^l = p$ at Stage 1. If $p \leq A$, then no ads will be sent, since any firm advertising a discount below p would not 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 $p_{(2)}^d < p - y$. 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 (which 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 (Firm (2)) can ensure a net profit on her of at least its value advantage y by advertising the discount price $p_{(1)}^d = y + A$, 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 $p_{(1)}^d > y + A$, 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 $p_{(1)}^d$ by Firm (1) may be regarded as a consumer surplus offer $s_{(1)} = r_{(1)} - p_{(1)}^d$, 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 $s_j^l = r_j - p$ 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 $s = r_{(2)} - A$.

"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 $\bar{s}_{(1)} = \bar{s}_{(2)} = S_{(2)} = r_{(2)} - A$.²⁸ We also have $\underline{s}_{(1)} = \underline{s}_{(2)} = s_{(1)}^l$: advertising an offer $s_{(1)} \in (s_{(1)}^l, \underline{s}_{(2)})$ 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 $s_{(2)} \in [s_{(1)}^l, \underline{s}_{(1)}]$ 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 $s_{(1)}, s_{(2)} \in [s_{(1)}^l, S_{(2)}]$. 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 $1 - a_{(1)} = B_{(1)}(s_{(1)}^l) = \frac{A}{p-y}$. As this is strictly positive for $y > p - A$, 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 $\pi_{(1)} = \lim_{s_{(2)} \nearrow s_{(1)}^l} B_{(2)}(s) (r_{(1)} - s) = y$, so $\lim_{s_{(2)} \nearrow s_{(1)}^l} B_{(2)}(s) = \frac{y}{p}$. However Firm (1)'s profits on advertised offers for $s_{(1)}$ arbitrarily close to $\underline{s}_{(1)} = s_{(1)}^l$ imply $B_{(2)}(s_{(1)}^l) = \frac{A+y}{p}$, so Firm (2)'s strategy must include a measure $\frac{A}{p}$ atom of

²⁸If, e.g. $\bar{s}_{(1)} < S_{(2)}$, then Firm (2) could strictly improve on its Lemma 4 profit by overcutting $\bar{s}_{(1)}$ and selling with probability one (and similarly for Firm (1) if $\bar{s}_{(2)} < S_{(2)}$).

offers $s_{(2)} = s_{(1)}^l$ just undercutting Firm (1)'s list price.²⁹

Because each firm must earn the same expected profit ($\pi_{(1)}(s) = B_{(2)}(s)(r_{(1)} - s) - A = y$ or $\pi_{(2)}(s) = B_{(1)}(s)(r_{(2)} - s) - A = 0$ respectively) on every surplus offer, we arrive at the following Stage 2 equilibrium strategies for a particular consumer.

Proposition 2 *The equilibrium advertising and surplus discount offers to a consumer with taste advantage $y > 0$ for Firm (1) following a symmetric list price $p > A + y$ are:*

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

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

(c) *No other firm advertises to the consumer.*

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.

3.3 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 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

²⁹Recall the tie-breaking assumption that an advertised discount defeats a list price offering the same surplus.

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, \dots, 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 .³⁰ 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 $1 - G(p_1^l - P_{-1})$ 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:

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

Using $P_1 = \min(p_1^l, A)$ 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:

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

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 $1 - G_j(p_j^l - P_{-j})$ consumers who buy at Firm 1's list price. Also as usual, the trade-off of hiking one's price is that marginal list price sales are lost, in this case at rate $g(p_1^l - P_{-1})$. The second key difference lies in the sacrificed profit P_1 per lost marginal sale which 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

³⁰This 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.

if $p_1^l > A$, the “lost” marginal sale is not truly lost.³¹ 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 .³²

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 log-concavity of $1 - G(y)$ (from Condition 1) there is a standard symmetric equilibrium with price given by the first-order condition

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

More generally, the first-order condition (3) is equivalent to $\frac{1-G(p_1^l - P_{-1})}{g(p_1^l - P_{-1})} - P_1 = 0$. Define a function $\Theta(p)$ equal to the left-hand 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 $\Theta(p) \rightarrow h - A$ as $p - A \rightarrow \bar{y}$.³³ 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

³¹Another way to understand the profit sacrificed per lost marginal sale when $p_1^l > A$ is that the sale is only truly lost if the marginal consumer buys from Firm 2, and that only happens when Firm 2 advertises, which it does with probability $1 - y/p = A/p$. Since a lost sale is worth p , the expected loss on the marginal consumer is $pA/p = A$. We thank a referee for this interpretation.

³²As 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.

³³If $\bar{y} = \infty$, define $h = \lim_{y \rightarrow \infty} \frac{1-G(y)}{g(y)} < \infty$. (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.

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

Proposition 3 *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 $A \in (h, p^{NT})$, the list price solves $\Theta(p^l) = 0$, 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:

$$\frac{1 - G(p^l - A)}{g(p^l - A)} = A. \quad (5)$$

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

$$\Pi^{ad} = p^l (1 - G(p^l - A)) + \int_0^{p^l - A} y dG(y), \quad (6)$$

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 Hotelling Example

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

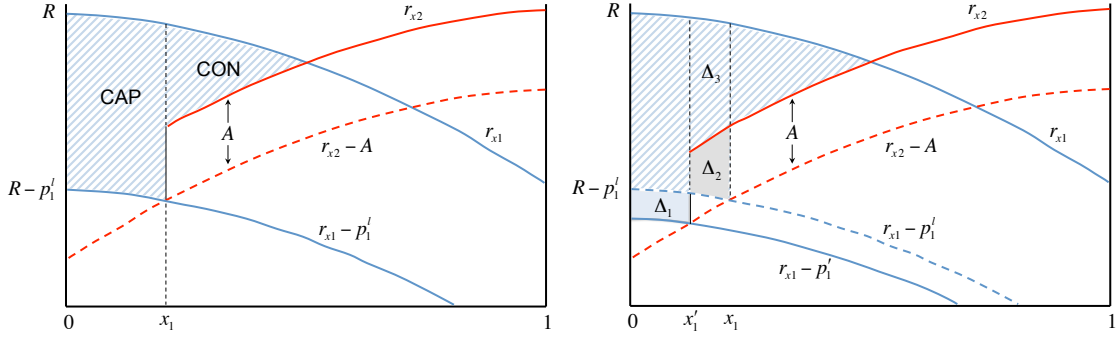


Figure 1: Hotelling competition – profits on captive and contested regions

from Firm 1 (r_{x_1} , blue) or Firm 2 (r_{x_2} , red). Consumer surplus from a list price purchase at Firm 1 is parallel to r_{x_1} 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 $r_{x_2} - A$ 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 $r_{x_1} - r_{x_2}$ 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 $p_1^l = p_1^l + \Delta p$. 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 here concern how the cost of targeted advertising affects firms' list prices, profits, and targeted advertising strategies in the absence of opt-in policies. We emphasize the demand curvature, for this key to understanding the opt-in results.

5.1 List prices

Suppose that $A \in (h, p^{NT})$ 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 4 *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 $\Theta(p^l; A) = \frac{1-G(p^l-A)}{g(p^l-A)} - A = 0$, making the dependence on the parameter A explicit. Differentiate this equilibrium condition implicitly to get

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

But Θ_{p^l} is strictly negative (by Condition 1) so $dp^l(A)/dA$ has the same sign as $g'(p^l-A)$, 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

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

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

$$\frac{\partial^2\Pi_1}{\partial p_1^l\partial(-P_{-1})} = -g(p_1^l - P_{-1}) - P_1g'(p_1^l - P_{-1})$$

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.³⁴ 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 4 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 5 shows that this is generally wrong: cheaper targeting unambiguously makes firms worse off.

Proposition 5 *Suppose there is an interior equilibrium with profit $\Pi(A)$ in a neighborhood of $\hat{A} < p^{NT}$. Then $\Pi'(\hat{A}) > 0$; profits are strictly increasing in the ad*

³⁴This term is unambiguously negative because it equals $\Pi_1''(p_1^l)$.

cost.

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

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 $CON_1 = G(p_1^l - P_{-1}) - G(A - P_{-1})$ is the set of contested 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 $\frac{\partial\Pi_1}{\partial P_{-1}} = Ag(p_1^l - P_{-1}) + CON_1$.³⁵ 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.³⁶

In contrast, Proposition 6 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 6 *Suppose captive demand is strictly log-convex 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 $\Pi^{FT} = \int_0^{\bar{y}} ydG(y)$ (regardless of A), or integrating by parts and using $1 - G(0) = \frac{1}{n}$, we have $\Pi^{FT} = \int_0^{\bar{y}} 1 - G(y) dy = \frac{1}{n}E\left(\frac{1-G(y)}{g(y)} \mid y \geq 0\right)$. As $\frac{1-G(y)}{g(y)}$ is strictly increasing, we have $\Pi^{FT} > \Pi^{NT}$. ■

³⁵Note that P_{-1} appears four times in (2), including the upper and lower limits of the integral over the contested region. The first term in the expression given here for $\partial\Pi_1/\partial P_{-1}$ consolidates the effect of three of those four terms, including the integral limits.

³⁶The overall effect $Ag(p_1^l - P_{-1})$ matches Proposition 5 after applying equilibrium condition (5).

If captive demand is log-convex, 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 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 4 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. In Section 7 we argue that allowing for consumer opt-in choice raises the list price.

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 $x \in [0, \frac{1}{2}]$, 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 7 *If tastes are distributed i.i.d. according to strictly log-concave 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 1 establishes that captive demand is strictly convex for $y > 0$, so Proposition 4 applies. ■

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.

While the role of the number of firms is not a main focus of the paper, we also note that under standard oligopoly competition the equilibrium price p^{NT} falls with n for the independent taste shocks model. This intuitive feature is preserved when there are targeted discounts: holding other parameters constant, the equilibrium list price $p^l(A)$ declines with n , and consumers receiving discounts are better off for the twin reasons that their surplus under discounting is larger with the lower list price and their second best option is stochastically better with more choice.³⁷ These pro-competitive results might help allay misgivings about the mixed strategies in our model since Varian’s (1980) model of sales has been criticized for its property that prices rise with more competition.

³⁷Both claims about list prices follow from Lemma 1.iii, respectively applying $p^{NT} = (1 - G(0))/g(0)$ and $p^l(A) = y^* + A$ with $(1 - G(y^*))/g(y^*) \equiv A$.

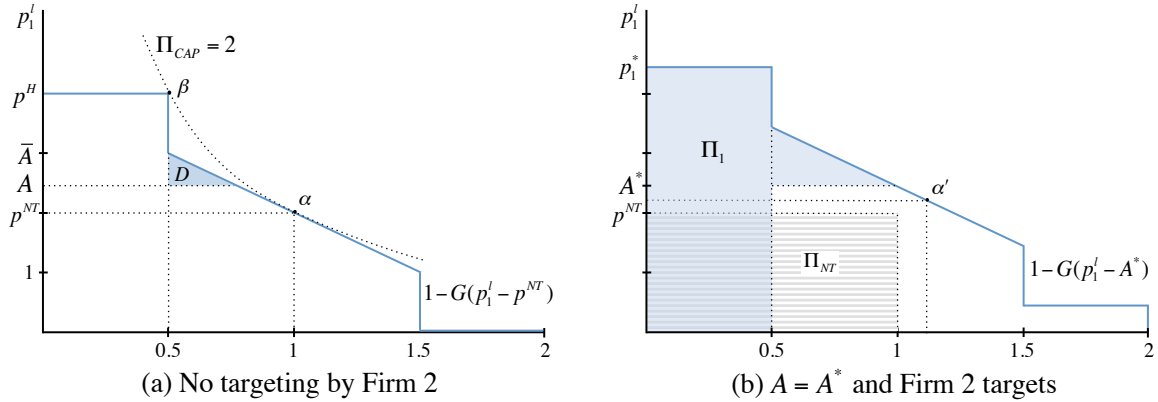


Figure 2: Firm 1's captive demand

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 log-concave. 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 $\bar{y} = 2$.³⁸ Figure 2.a illustrates Firm 1's captive demand $1 - G(p_1^l - p^{NT})$ 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 loyalists. (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 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 6 emerges for $A \leq A^* = \sqrt{6}$. (For $A \in (A^*, \bar{A})$, equilibria involve mixed strategies with partial retrenchment – see the Supplementary Appendix for details.) Figure 2.b shows Firm 1's captive

³⁸Analysis to support the claims made here appears in the Supplementary Appendix.

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 $D = \frac{1}{4}(\bar{A} - A^*)^2 \approx 0.076$ 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

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

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 8 *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 infinitesimal 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.*

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.³⁹ It advertises with probability $1 - \frac{A}{p^l}$ to consumers just on its side of the bound-*

³⁹While the $y = 0$ case has only one symmetric equilibrium, it does have asymmetric equilibria as well. For $y < 0$ there is a unique asymmetric equilibrium; for $y > 0$, there is another unique asymmetric equilibrium. The $y = 0$ case has equilibria at the limits of the two asymmetric cases above, plus the symmetric equilibrium (which “pops up” at $y = 0$ only). The jump in advertising rates on either side of $y = 0$ reflects a jump from the first asymmetric equilibrium to the second one. In the current paper this occurs because the list prices set at Stage 1 induce a discontinuity in the profits each firm would earn if there were no Stage 2. That is, (given equal list prices), Firm 1 would earn strictly positive profits on list price sales to consumers with $y > 0$, but would earn 0 on consumers with $y < 0$ (since it would not sell to them). And conversely for Firm 2. Since these “default payoffs” play an outside-option-like role in modulating how fiercely each firm is willing to compete in Stage 2, it is not unreasonable that a discontinuity in these outside options should “pass through” into a discontinuity in Stage 2 strategies. See Anderson, Baik, and Larson (2015) for analysis on why asymmetric equilibria are more plausible than the symmetric one, based on a stability argument.

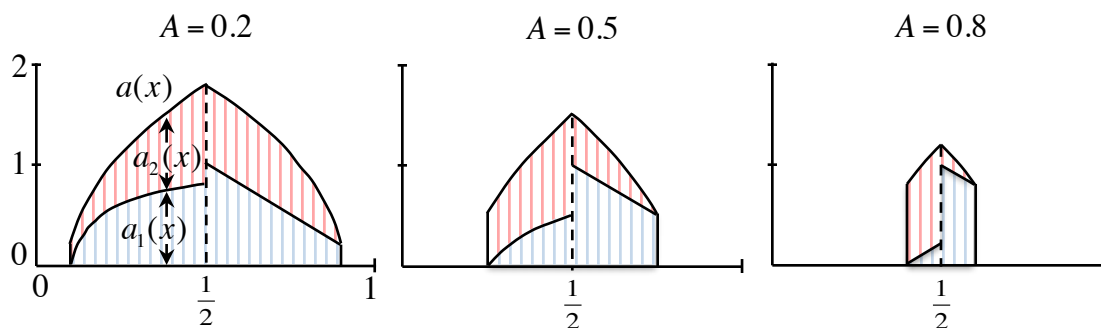


Figure 3: Volume of advertising $a(x)$ by consumer location x . (Linear-Hotelling preferences with $t = 1$.)

ary, 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.⁴⁰ 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.

6 Welfare and Consumer Surplus

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

⁴⁰This 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.

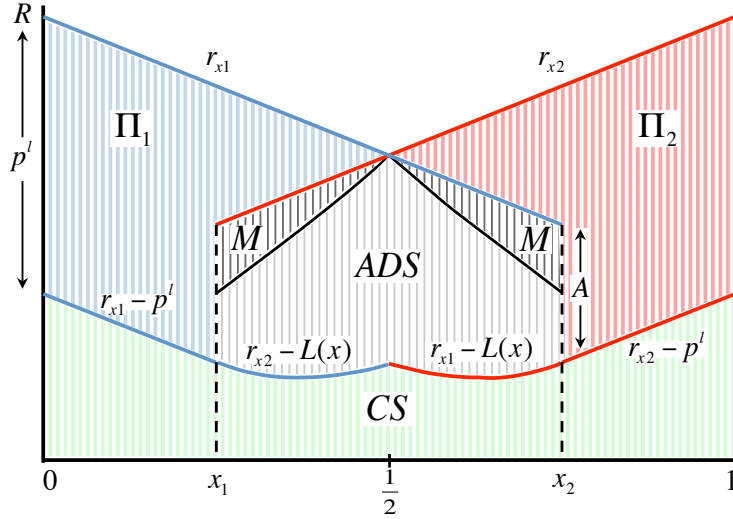


Figure 4: Equilibrium welfare in the linear-Hotelling model. ($R = 1.6$, $t = 1$, $A = 0.5$)

arbitrary taste distributions that we describe below. Note that the analysis here does not account for costs of lost privacy which will be discussed in Section 7.

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 welfare and surplus at the level of an individual consumer. Results about aggregate welfare and consumer surplus are developed in Sections 6.2 and 6.3. Section 6.4 explores misallocation, and Section 6.5 discusses the distributional consequences of targeting for consumers.

6.1 Welfare and Consumer Surplus at an Individual Consumer

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 $r_{(1)} - p^l$; adding in Firm (1)'s profit of p^l on her, we have (first-best) total welfare of $r_{(1)}$. A

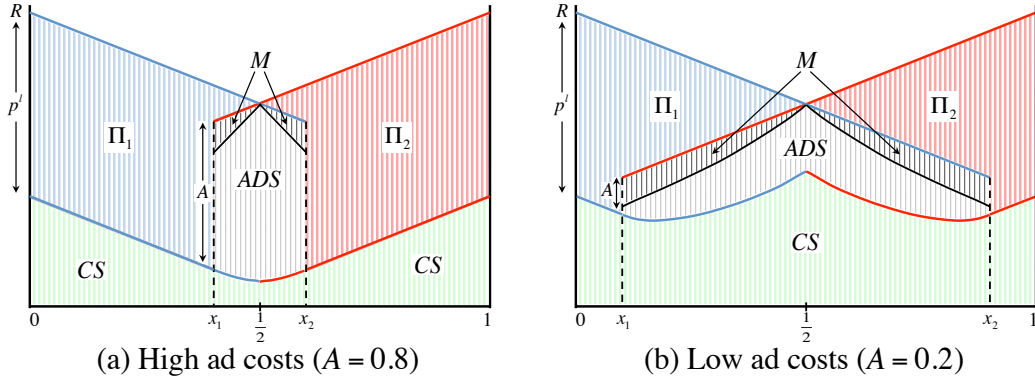


Figure 5: Equilibrium welfare in the linear-Hotelling model. ($R = 1.6$, $t = 1$)

contested consumer takes the better of the final surplus offers from her top two firms. Using the strategies from Section 3.2, this consumer surplus is distributed according to

$$B_{\max}(s) = B_{(1)}(s) B_{(2)}(s) = \frac{A}{r_{(2)} - s} \frac{y + A}{y + r_{(2)} - s} \text{ on } [s_{(1)}^l, S_{(2)}], \quad (8)$$

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

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

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

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

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

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

In Figure 5 the loss function is represented as the sum of the expected ad cost (given by $\mathcal{A}(y) = Aa(y)$ with $a(y)$ in (7)) and the misallocation cost, $M(y)$, discussed

in Section 6.4.

6.2 Impact of Targeting 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 $\bar{L} = n \int_0^{p^l - A} L(y, p^l, A) dG(y)$. In the costless targeting limit, this welfare loss vanishes.

Proposition 9 *The equilibrium welfare loss on each contested consumer vanishes as ad costs vanish: $\lim_{A \rightarrow 0} L(y, p^l(A), A) = 0$ 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 9 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).

6.3 Impact of Targeting on Aggregate Consumer Surplus

It can be useful to frame an individual's surplus as $CS(y) = r_{(1)} - EP(y)$, where $EP(y) = \min(p^l, y + L(y, p^l, A))$ is her expected 'favorite-equivalent' price. The name reflects the idea that buying her second-best product at price $p_{(2)}^d$ may be regarded as paying $p_{(2)}^d + y$ for first-best quality $r_{(1)}$. Then the question of whether she

is better off when targeting is permitted or when it is banned depends on how targeting affects the mix of list and discount prices she is offered. When captive demand is convex, there are countervailing effects: discounts benefit the contested, particularly those with strong second-best choices (low y), but passed-through efficiency costs (ad costs and misallocation) eat into those gains. Meanwhile, higher list prices hurt all consumers – particularly the captives, but also the unlucky contested who happen to receive no ads. Our main results here demonstrate that, due to this list price effect, consumers may be collectively worse off with targeted discounts than without them if captive demand is sufficiently convex. We evaluate “collectively worse off” in two ways, first, by comparing average consumer surplus, and second, with the stronger criterion that every consumer is hurt. Each case requires an appropriate notion of what “sufficiently convex” means.

We begin with average consumer surplus, denoted \overline{CS}^{NT} or \overline{CS}^T in an equilibrium where targeting is forbidden or permitted respectively. Impose Condition 1, so that $\ln(1 - G(y))$ is strictly concave. We say that captive demand is ρ -convex, for $\rho > 0$, if $(1 - G(y))^\rho$ is convex for $y \geq 0$. Note that ρ closer to zero corresponds to a higher degree of convexity; in the limit as ρ goes to zero, captive demand approaches an exponential distribution, which is the boundary case between logconcavity and logconvexity. This condition is only sensible if taste advantages have upper limit $\bar{y} = \infty$, which we assume for Proposition 10 below.

Proposition 10 *Suppose targeting is costly ($A > 0$) and that the no-targeting price is $p^{NT} = (1 - G(0)) / g(0) > A$. Then if captive demand is ρ -convex with $\rho < A/p^{NT}$, average consumer surplus is higher when targeted discounts are banned than when they are permitted.*

The condition on ρ -convexity has the virtue of a simple interpretation – consumers are hurt by targeting when it is sufficiently costly, and there is a broader range of targeting costs for which consumers are hurt when captive demand is more convex. It is a sufficient condition, not a necessary one, and there are demand systems where it is not met but targeted discounting still hurts consumers on average. One notable example is logistic captive demand (generated from independent Type 1 extreme value taste shocks, as discussed in Section 2, and as frequently assumed in empirical work). Suppose demand is logistic and some (but not all) consumers are targeted in equilibrium. Then, for any number of firms and without further qualification,

consumers would be better off on average if targeting were banned.⁴¹

Next we give conditions under which targeting makes *all* consumers worse off. In this case, requiring captive demand to be strictly convex at the turf boundary $y = 0$ is the appropriate notion of ‘sufficiently convex.’ We continue to require Condition 1 but no longer need the assumption that $\bar{y} = \infty$.

Proposition 11 *If $g'(0) < 0$, then if targeting costs are sufficiently high ($A \in (\bar{A}, p^{NT})$ for some \bar{A}), every consumer would be strictly better off if targeting were banned. A sufficient condition for $g'(0) < 0$ is independent tastes drawn from strictly logconcave $f(r)$ with at least three firms.*

The proof involves showing that every consumer’s surplus in the targeting equilibrium is increasing in A for A sufficiently large. Then the result follows because a consumer’s surplus at $A = p^{NT}$ (where targeted are just barely too expensive to use) coincides with her surplus in a no-targeting equilibrium. The argument that consumer surplus rises with A is straightforward for captive consumers, who benefit because they face lower list prices (by Proposition 7). 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 condition on $g'(0)$ ensures that $\frac{\partial p^l}{\partial A}$ remains *strictly* negative even near $A = p^{NT}$. Under certain conditions the conclusion of Proposition 11 also applies with two firms (in either the i.i.d. or Hotelling case), but second-order terms must be consulted

⁴¹Let $1 - G(y) = \frac{1}{1+(n-1)e^{y/\beta}}$. It may be confirmed that the no-targeting equilibrium list price is $p^{NT} = \frac{n}{n-1}\beta$ and that for $A \in (\beta, p^{NT})$ the targeting model has an interior equilibrium with $y^* = \beta \ln\left(\frac{\beta}{A-\beta} \frac{1}{n-1}\right)$ and list price $p^T = y^* + A$. Follow the proof of Proposition 10 to establish $\overline{EP}^{NT} = p^{NT} = \frac{n}{n-1}\beta$ and $\overline{EP}^T \geq \frac{1}{1-G(0)} \int_0^{y^*} 1 - G(y) dy + A = n\beta \ln\left(\frac{p^{NT}}{A}\right) + A$, where the last step follows by direct computation. Then we have $\overline{EP}^T - \overline{EP}^{NT} \geq \phi(p^{NT}) - \phi(A)$, where $\phi(x) := n\beta \ln x - x$. The function $\phi(x)$ is strictly increasing on $(0, n\beta)$, so $\phi(p^{NT}) > \phi(A)$, and therefore $\overline{EP}^T > \overline{EP}^{NT}$.

because both $\frac{\partial CS}{\partial A}$ and $\frac{\partial p^i}{\partial A}$ vanish near $A = p^{NT}$ – Proposition 17 in the Supplementary Appendix gives further details.

The counterpoint to Proposition 11 is that consumers may collectively benefit from targeted discounts if the inflation of list prices is more modest. The sharpest conclusions are obtained when the curvature of captive demand reverses; in this case, consumers see no downside to targeting since it reduces list prices rather than inflating them.

Proposition 12 *If captive demand is strictly concave and targeting is in use ($A < p^{NT}$), then every consumer would be strictly worse off if it were banned.*

Proof. By Proposition 7, $p^T < p^{NT}$. The conclusion follows from the fact that a consumer’s surplus is $r_1 - p^{NT}$ if targeting is banned and at least $r_1 - p^T$ if it is permitted (since purchasing at list price is always an option). ■

The linear-quadratic Hotelling example discussed earlier is an intermediate case: because of list price neutrality, contested consumers are harmed by a ban on targeting, but captive consumers are unaffected. Finally, the counterpart to Proposition 10 is that consumers may benefit from targeting on average if the efficiency losses from discounting are sufficiently small. Notice that average 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, as long as Condition 1 holds, the fact that firms are worse off in the free targeting limit (by Proposition 5), means that consumers must collectively be better off.

Result 4 $\overline{CS}_{A=0} > \overline{CS}_{NT}$.

Continuity implies that consumers are better off on average with targeted ads than without them if the targeted ad cost A is sufficiently small.

6.4 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 $L(y, p^l, A) = M(y) + \mathcal{A}(y)$, 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 $s_1 \in (s_1^l, S_2)$.⁴²

Proposition 13 *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 $m(p^l - A) = \frac{A}{p^l}$ and $m'(y) < 0$, and so $m(y) > \frac{1}{2}$ for all $y \in [0, p^l - A]$ 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 $A \rightarrow 0$, $m(y)$ 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).

6.5 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

⁴²This integral (correctly) excludes any weight on Firm 2's atom of undercutting offers, as these never win when Firm 1 advertises.

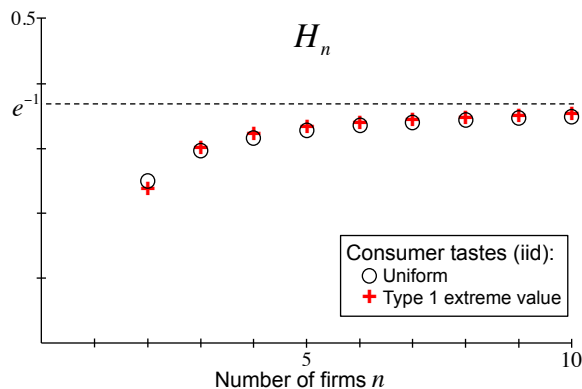


Figure 6: Fraction of consumers harmed by targeted advertising

consumers with a strong favorite product toward those who are more willing to shop 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 $1 - G_n(p_n^{NT})$ 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 $H_n = n(1 - G_n(p_n^{NT}))$. 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 14 *Suppose tastes are i.i.d. with well-behaved, strictly log-concave density $f(r)$ and $\bar{r} = \infty$. Then $\lim_{n \rightarrow \infty} H_n = \frac{1}{e}$.*

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 $1 - G_n(y) = E(1 - F(\tilde{r}_{n-1} + y))$, where \tilde{r}_{n-1} is the (random) largest rival valuation. The expected percentile of that highest rival valuation is $E(F(\tilde{r}_{n-1})) = \frac{n-1}{n}$, so in some sense \tilde{r}_{n-1} is centered around the certain value \hat{r}_{n-1} defined by $F(\hat{r}_{n-1}) = \frac{n-1}{n}$. The crux of the Gabaix et al. result is that for n large we may approximate $1 - G_n(y)$ with $1 - F(\hat{r}_{n-1} + y)$. 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

harmful 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 14 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, the consequent consumer backlash amid data breaches, 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 second 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 chooses whether to pay A per consumer to send a targeted discount offer to a consumer at location y , if she has opted in. 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 c.d.f. $H(c)$, 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.

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 earlier in 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 at a location that receives targeting (this is $y \in [0, p^T - A]$ 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 understood largely through the policy’s effect on list prices. This is not as straightforward as saying that higher prices are bad for consumers, for the list price impacts the discount price distribution, and a lower list price also reduces the set of consumers getting offers. Nonetheless, we can track the different groups of affected consumers and engage our results from Section 6 that CS_g is decreasing in list price while a consumer is targeted, and moreover that CS_g is continuous through

the list price where the consumer at y switches from targeted to captive. We first compare the targeting and opt-in regimes.

Lemma 6 *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.*

Lemma 6 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 7 *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 $\tilde{p} = \min(p_1^l, p_2^l + y)$. 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:

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

The key features of $\Delta(y; \tilde{p})$ are readily proved from (10):

Lemma 8 *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 8 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 underscoring 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 $\Delta(y; p_1^l)$ 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 $y \in (0, p_1^l - p_2^l)$ would buy at Firm 2 in the absence of discounts. An additional price hike by Firm 1 has no effect on their gain from discounting $\Delta(y; p_2^l + y)$, their decision to opt in or out, or any profit that Firm 1 earns on the opt-ins in the discounting sub-game.

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

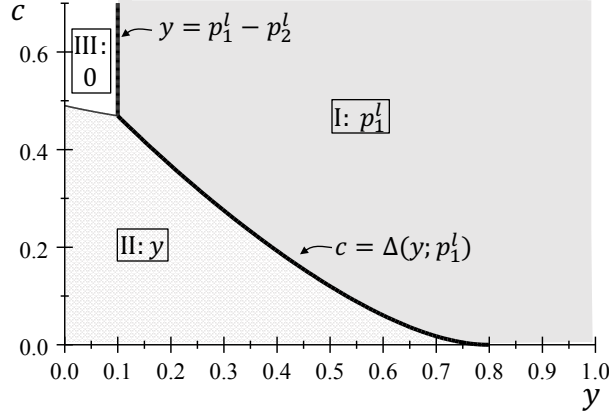


Figure 7: Profit Regions for Firm 1. (Parameters: $A = 0.3$, $p_1^l = 1.1$, $p_2^l = 1$.) The boundary of Firm 1's captive demand is in bold. Consumers in Region I opt out and buy at p_1^l . Consumers in Region II opt in; Firm 1 earns net profit y on each of them. Consumers in Region III opt out and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

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

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

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

Figure 7 illustrates 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 $\int_0^{p_1^l - p_2^l} y H(\Delta(y; p_2^l + y)) dG(y) + \int_{p_1^l - p_2^l}^\infty y H(\Delta(y; p_1^l)) dG(y)$. Then Firm 1's marginal profit may be written as:

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

where $\underline{y} = p_1^l - p_2^l$. The second and third terms represent the two different margins.⁴³ The second term is the conventional oligopoly margin: consumers who are indifferent

⁴³The second term in (13) corresponds to the margin between regions I and II in Figure 7, and the last one is the I-II margin. Note that the II-III margin depends on rival price, and not own price, and so is unchanged.

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. We can rewrite the expression for list price sales as

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

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

$$\frac{d\Pi_1}{dp_1^l} = \int_0^{\infty} \Lambda(c) dH(c), \quad (15)$$

where, suppressing arguments to $y^*(c; p_1^l)$,

$$\Lambda(c) = 1 - G(y^*) - (p_1^l - y^*) g(y^*) \frac{\partial y^*}{\partial p_1^l} \quad \text{for } c < \bar{c} \quad (16)$$

$$= 1 - G(p_1^l - p_2^l) - p_1^l g(p_1^l - p_2^l) \quad \text{for } c > \bar{c}. \quad (17)$$

That is, marginal profit can be expressed as an expectation over the marginal profits $\Lambda(c)$ associated with each privacy cost type.⁴⁴ 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

⁴⁴We can see the idea from Figure 7 by taking $\Lambda(c)$ as the price derivative at any level of c .

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. First we establish some facts about $\partial y^*(c; p_1^l) / \partial p_1^l$, which is the rate at which the opt-in threshold at cost type c rises with Firm 1's list price. For $y \in [0, p - A]$, define the function $D(y, p, A) = -\frac{\Delta_p(y, p, A)}{\Delta_y(y, p, A)}$. Since the opt-in threshold is defined by $\Delta(y^*(c; p_1^l), p_1^l) = c$, the implicit function theorem implies that $\partial y^*(c; p_1^l) / \partial p_1^l = D(y^*(c; p_1^l), p_1^l, A)$.

We first give an instrumental result we will need for bounding (16).

Lemma 9 *For $y \in [0, p - A]$, $D(y, p, A) \in [1 - \frac{p - y - A}{3p + A}, 1)$. Furthermore, $D(p - A, p, A) = 1$ and $\lim_{A \rightarrow 0} D(y, p, A) = 1$ for all $y \in (0, p - A)$.*

Note that the lower bound in the lemma satisfies $1 - \frac{p - y - A}{3p + A} \geq \frac{2}{3}$, so for $c > 0$ (implying $y^*(c; p_1^l) < p_1^l - A$), we have $\partial y^*(c; p_1^l) / \partial p_1^l \in [\frac{2}{3}, 1)$. The main takeaway is that the marginal opt-in location does not rise one-for-one with a price rise. This is in contrast with our benchmark model, where the marginal targettee, $y = p_1^l - A$, does rise one-for-one with the price. The difference here is that the benefits to a consumer from opting in rise more slowly than p_1^l because the inefficiencies of discount competition (the ad costs and misallocation in $L(\cdot)$) are passed through to consumers. This has the effect of making a firm's captive demand a bit less elastic under the opt-in policy than it would be otherwise, creating a slight incentive toward higher list prices that complicates one of the conclusions below. Anticipating this complication, we introduce the following condition.

Condition 3 *The list price with unrestricted targeting represents a markup of at least $\frac{p^T - p^{NT}}{p^{NT}} > \delta(A/p^{NT})$ over the list price when targeting is banned, where $\delta(\alpha) = \frac{1}{2} \left(\sqrt{(\alpha + \frac{1}{2})^2 + 2(1 - \alpha)} - (\alpha + \frac{1}{2}) \right)$.*

The minimum markup $\delta(A/p^{NT})$ in Condition 3 is decreasing in A , with $\delta(0) = \frac{1}{2}$ and $\delta(1) = 0$. Its purpose is to give a condition under which list prices are high enough under targeting that an opt-in requirement unambiguously brings prices down, in spite of the additional inelasticity mentioned above. Situations where Condition 3 is met are discussed below.

Echoing earlier conclusions, our main result shows that the price implications of an opt-in requirement depend on the curvature of captive demand. We adapt our earlier definition to say that captive demand is strictly ρ -convex, for $\rho > 0$, if $(1 - G(y))^\rho$ is strictly convex for $y \geq 0$. Proposition 15 applies to any distribution of consumer privacy costs $H(c)$, excluding the degenerate case where all consumers have $c = 0$.

Proposition 15 *If captive demand is weakly concave ($g' \geq 0$), then any symmetric equilibrium under opt-in satisfies $p^{OI} > p^T$. Alternatively, suppose captive demand is strictly convex ($g' < 0$). If the Mills ratio $\frac{1-G(y)}{g(y)}$ is also convex and Condition 3 is met, then any symmetric equilibrium under opt-in satisfies $p^{OI} \in [p^{NT}, p^T)$. In particular, Condition 3 is satisfied if captive demand is $\frac{2}{3}$ -convex.*

Empirical work often assumes a demand system based on independent taste shocks. In this case, Proposition 7 implies that regardless of the primitive taste distribution (as long as it is log-concave), the captive demand function must be strictly convex. Thus, for empirical applications, the conclusions in Proposition 15 about convex demand may be of greatest interest. The additional condition on the Mills ratio is mild and appears to be satisfied by most, if not all, common demand systems.⁴⁵ Figure 8 illustrates the wedge $1 + \delta(\alpha)$ between the targeting and no-targeting prices that suffices to ensure that imposing an opt-in rule will bring prices down. The figure also shows when Condition 3 is met for logistic demand derived from i.i.d. Type 1 extreme value taste shocks. Under duopoly, Condition 3 is satisfied whenever targeting costs are low enough ($\alpha = A/p^{NT} \lesssim 0.76$), or equivalently, whenever at least 32% of consumers would be contested if targeting were unrestricted. With logistic demand and three or more firms, Condition 3 is always satisfied, regardless of the targeting cost. If taste shocks are i.i.d. uniform, we have the stronger conclusion that Condition 3 is always satisfied (for any A and for any number of firms, including $n = 2$). More generally, any captive demand distribution where the maximal taste advantage \bar{y} is large enough, including the broad class of distributions with $\bar{y} = \infty$, must satisfy Condition 3 when A is small.

The main virtue of Proposition 15 is that it permits us to give clear conditions when the imposition of an opt-in policy unambiguously makes consumers better off.

⁴⁵For example, $G(y)$ has a convex Mills ratio for $y \geq 0$ if $G(y)$ is logit, normal, or generated from i.i.d. taste shocks of the form $F(r) = r^a$, for $a > 0$. Furthermore, since we have already assumed $\frac{1-G(y)}{g(y)}$ strictly decreasing, if the support of y is unbounded then the Mills ratio must be convex for y sufficiently large.

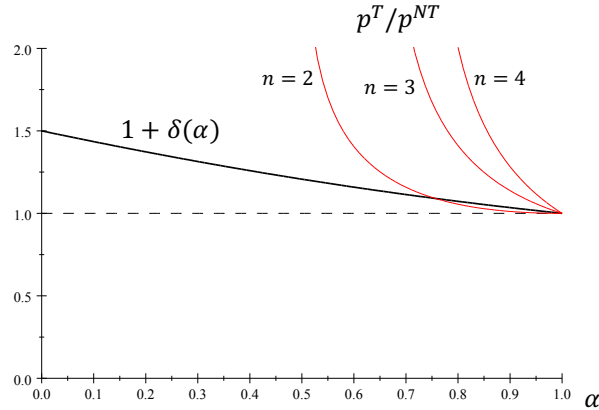


Figure 8: Regions where Condition 3 is satisfied when captive demand is logistic. Demand is $1 - G(y) = \frac{1}{1+(n-1)e^{y/\beta}}$. The red curves give the targeting/no-targeting price ratio p^T/p^{NT} as a function of $\alpha = A/p^{NT}$ for two, three, and four firms.

Corollary 1 *Suppose captive demand is strictly convex and the other conditions of Proposition 15 hold. Then imposing an opt-in policy makes every consumer better off.*

This corollary follows immediately from Proposition 15 and Lemma 6. For policymakers, it may be appealing that the conclusions of Corollary 1 do not depend on welfare trade-offs across consumers, nor do they depend on accurately quantifying the value(s) that consumers place on privacy (which could be difficult in practice). When Corollary 1 applies, list prices across the three regimes are ranked: $p^{NT} \leq p^{OI} < p^T$ by Proposition 15. Corollary 1 serves as a practical guide for evaluating an opt-in policy. Instead of relying on demand concavity estimation to determine the impact on consumers, list prices can be an effective barometer: if an opt-in policy induces list prices to fall, then the policy has unambiguously made consumers better off.

One might speculate that if all consumers are better off with the lower list prices in an opt-in regime, the additional list price reductions from an outright ban on targeting might make them all even better off. However, that is not correct – as Lemma 7 indicates, this will typically hurt a subset of consumers who have the most to gain from discounting and who care the least about privacy.

7.1.4 Equilibrium existence example

While we do not have a general proof of equilibrium existence for the model above (due to the intricacies of the function $\Delta(y; p_1^l)$), we provide a simple and central example

where existence is guaranteed. This is duopoly Hotelling competition (see Example 1 in Section 2) with captive demand $1 - G(y) = \frac{1}{2t}(t - y)$. Assume a two-type privacy cost distribution: fraction λ of consumers have $c = c_L \downarrow 0$, and the remaining consumers have $c = c_H$, with c_H high enough that none of the high-cost consumers opt in. Under either the pure no-targeting or pure targeting models, the equilibrium list price is $p^l = t$ (cf. Proposition 4). In the opt-in model, the condition $\Delta(y^*, p_1^l) = c$ for Firm 1's captive consumer threshold collapses to $y^* = p_1^l - A$ for the low-cost consumer types. Firm 1's marginal profit is $d\Pi_1/dp_1^l = \lambda\Lambda(c_L) + (1 - \lambda)\Lambda(c_H)$ with $\Lambda(c)$ given by (16) and (17). Note that $\Lambda(c_H) = 1 - G(p_1^l - p_2^l) - p_1^l g(p_1^l - p_2^l) = \frac{1}{2t}(t + p_2^l - 2p_1^l)$ is just Firm 1's marginal profit in the standard, no-targeting duopoly. The “spoiler” term in $\Lambda(c_L)$ collapses to $\partial y^*/\partial p_1^l = 1$, so for $p_1^l \geq A$ we have $\Lambda(c_L) = 1 - G(p_1^l - A) - Ag(p_1^l - A) = \frac{1}{2t}(t - p_1^l)$. If $p_1^l < A$, Firm 1's profit on low-cost consumers is just $p_1^l(1 - G(p_1^l - A))$ (as it cannot afford to advertise to those who opt in). Thus for $p_1^l < A$, we have $\Lambda(c_L) = 1 - G(p_1^l - A) - p_1^l g(p_1^l - A) = \frac{1}{2t}(t + A - 2p_1^l)$. Note that $\Lambda(c_L)$ is just Firm 1's marginal profit in the full-targeting model. Both $\Lambda(c_H)$ and $\Lambda(c_L)$ are continuous and strictly decreasing in p_1^l , so the overall marginal profit $d\Pi_1/dp_1^l$ is as well, and thus Firm 1's profit function is concave in p_1^l . The same logic applies to Firm 2, so the symmetric solution $p_1^l = p_2^l = t$ to both firms' profit first-order conditions is necessary and sufficient for an equilibrium.

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

This section studies the opt-in model under the alternative timing assumption that consumers' opt-in decisions are made simultaneously with firms' list price choices.⁴⁶ This timing might be more appropriate if real-world consumers do not 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,

⁴⁶Because an individual consumer has a negligible impact on the firms' decisions, it is equivalent to have consumers make their decisions before list prices are set.

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 8 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

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

where we recall that $I(y) = H(\Delta(y; \tilde{p}^e))$ denotes the fraction of consumers at y who have opted in at symmetric expected list price \tilde{p}^e , and $O(y) = 1 - I(y)$ is the fraction who have not.⁴⁷ Notice that $I(y) = 0$ for $y > \tilde{p}^e - A$, for these consumers do not expect to ever get discounts.⁴⁸

⁴⁷Here $I(y) = H(\Delta(y; \tilde{p}^e))$ for $y > 0$ and an analogous expression for $y < 0$ reflecting the switch in the preferred product.

⁴⁸We can envisage the partition of consumers (and profits) with a diagram analogous to Figure 7. Start with a symmetric list-price situation. The consumers opting in are those below a symmetric tent-shaped function that reflects $\Delta(y; \tilde{p}^e)$ around $y = 0$, and anchored at $\tilde{p}^e - A$ for the highest possible y . The Ins are in the tent, the Outs are out (above it). Firm 1 makes list price sales to the Outs for whom $y = 0$, which is a cut-point (a spike out the top of the tent). Now consider an (unexpected) deviation by Firm 1. If it raises its price, the "spike" moves Right but the tent stays the same, and all the Ins below it with $y > 0$ remain garnering y profit to 1 (when the list price rises, more would like to opt-in because both firms would discount to them, but it is too late). For

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

$$\Pi_1(p_1^l) = p_1^l Q_1 + \int_0^{p_1^l - A} y I(y) dG(y). \quad (18)$$

Then Firm 1's marginal profit is:

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

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.⁴⁹ 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 $\Delta(y; p_1^e)$. This benefit satisfies $\Delta(y; p_1^e) = 0$ 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 $I(y) = H(\Delta(y; p_1^e)) = 0$ 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

$$\left. \frac{d\Pi_1}{dp_1^l} \right|_{p_1^l = p_1^e} = \int_0^\infty O(y) dG(y) - O(0) g(0) p_1^e. \quad (19)$$

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

a price drop, the ‘spike’ moves Left, but then the right-most triangle in the tent (i.e., above new $p_1^l - A$) will not be targeted. These consumers are converted to list price sales.

⁴⁹So $I(p_1^l - A) = 0$ for $p_1^l \geq \tilde{p}^e$.

in response to p_1^l , has dropped out here. In this sense, inattentive consumers do not discipline list prices as effectively as do alert consumers.

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

$$\int_0^\infty \frac{O(y)}{O(0)} dG(y) - g(0) \hat{p}^{OI} = 0. \quad (20)$$

In effect, each firm faces a demand curve for list price sales that has been hollowed out by consumers opting in to discounts. 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 (20) would reduce to the standard oligopoly first-order condition with solution p^{NT} . However, as we noted in the discussion around Lemma 8, $O(y) = 1 - H(\Delta(y; \hat{p}^{OI}))$ 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.

Condition 4 Let $\gamma(p) = \int_0^\infty \frac{1-H(\Delta(y,p))}{1-H(\Delta(0,p))} dG(y) - g(0)p$ be the left-hand side of (20). The equation $\gamma(p) = 0$ has a unique solution \hat{p}^{OI} , and $\gamma'(\hat{p}^{OI}) < 0$.

If Condition 4 holds then there is a symmetric equilibrium at some price level \hat{p}^{OI} and each firm's marginal profit is positive (negative) at any candidate symmetric equilibrium profile of prices lower (higher) than \hat{p}^{OI} .

Proposition 16 Any symmetric equilibrium with inattentive consumers and opt-in has higher list prices than a regime where targeting is banned ($\hat{p}^{OI} > p^{NT}$). Furthermore, suppose Condition 4 holds. Then prices are higher when consumers are inattentive than under any symmetric equilibrium with opt-in and alert consumers ($\hat{p}^{OI} > p^{OI}$).

Proof. The first ranking follows directly from the fact that $O(y)/O(0) > 1$, which (with log-concavity of $g(y)$) implies the left-hand side of (20) is strictly positive for all prices p^{NT} or lower. For the second ranking, a symmetric equilibrium with alert consumers satisfies the first order condition

$$\int_0^\infty (1 - H(\Delta(y, p^{OI}))) dG(y) - (1 - H(\Delta(y, p^{OI}))) g(0) p^{OI} - Z = 0,$$

where

$$Z = \int_0^{p^{OI}-A} (p^{OI} - y) h(\Delta(y, p^{OI})) \frac{\partial \Delta(y, p^{OI})}{\partial p^{OI}} dG(y) > 0$$

reflects the “extra” loss of marginal consumers that is absent from marginal profit when consumers are inattentive. We can write this first order condition as

$$(1 - H(\Delta(0, p^{OI}))) \gamma(p^{OI}) = Z > 0,$$

and so $p^{OI} < \hat{p}^{OI}$ by Condition 4. ■

7.2.1 Equilibrium existence with inattentive consumers: an example

There are no simple conditions that guarantee quasi-concave profit functions when consumers precommit to opting in or out, so the symmetric first-order condition (20) is necessary but not sufficient for a symmetric equilibrium at a common list price \hat{p}^{OI} . This section gives a parametric example where profit quasi-concavity is verified numerically, and so the solution to (20) does identify an equilibrium. The example also illustrates some of the complications that might prevent (20) from identifying an equilibrium in other cases. Details of the analysis are in Supplementary Appendix C.

The demand and privacy cost structure of the example is similar to Section 7.1.4. Consider two firms with linear (Hotelling) demand $1 - G(y) = \frac{1}{2}(1 - y)$ for $y \in [-1, 1]$. (This corresponds to a transport cost parameter $t = 1$.) Fraction $\lambda = \frac{1}{2}$ of consumers have prohibitively high privacy costs, ensuring that they opt out; the remaining $1 - \lambda = \frac{1}{2}$ share of consumers have privacy cost $c = 0.2$. The targeted ad cost is $A = 0.4$. Note that for these parameters the no-targeting and unrestricted targeting prices are both equal to one: $p^{NT} = p^T = 1$.

The analysis in the Supplementary Appendix begins by solving (20) to identify a candidate equilibrium. The solution is at the common list price $\hat{p}^{OI} \approx 1.428$ and involves low privacy cost consumers opting in if they are located on the interval

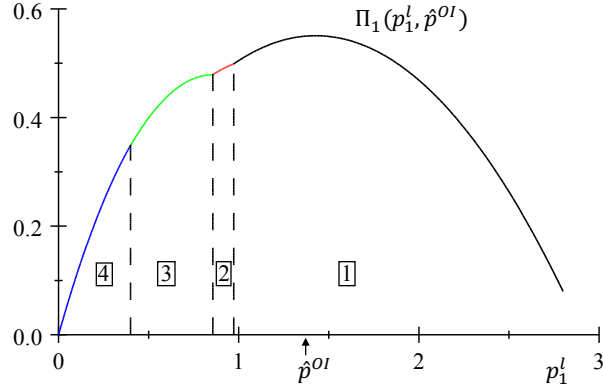


Figure 9: Firm 1's profit function when Firm 2 sets $p_2^l = \hat{p}^{OI} \approx 1.43$. (Labels 1-4 indicate regions where different expressions for Π_1 apply; see the Appendix for details.)

$y \in [-y^*, y^*]$, with $y^* \approx 0.572$. To confirm this solution as an equilibrium, we construct Firm 1's profit $\Pi_1(p_1^l, \hat{p}^{OI})$ from a deviation to list price p_1^l (given $p_2^l = \hat{p}^{OI}$ and that consumers act on the expectation that both firms price at \hat{p}^{OI}). Figure 9 plots this profit and demonstrates that $\Pi_1(p_1^l, \hat{p}^{OI})$ is maximized at $p_1^l = \hat{p}^{OI}$. Given the symmetry of the firms, this confirms the candidate solution as an equilibrium.

Figure 9 labels four regimes for Firm 1's profit from a deviation. Because consumers have committed to their opt in or out decisions and are not able to adjust if Firm 1 sets an unexpected price, the functional form of $\Pi_1(p_1^l, \hat{p}^{OI})$ can change abruptly with p_1^l as Firm 1's marginal list price consumer shifts across the boundaries between opting in or out. Moving from right to left across the figure, in Regime 2, Firm 1 has cut its price enough to capture some consumers who opted in with the expectation of being contested; at the lower than expected p_1^l , Firm 2 cannot profitably send them discount ads. Firm 1 cannot yet make money on the opt-ins who are closer to Firm 2 ($y \in [-y^*, 0]$) – discount competition drives Firm 1's profit on them down to zero. But in Region 3, Firm 1 begins to leapfrog these consumers to capture list price sales from consumer locations $y \leq -y^*$ that opted out (expecting to buy from Firm 2). Finally, in Region 4, Firm 1's list price is low enough that all of its profits are earned at list price.

While in this case the symmetric first-order condition does identify an equilibrium, the plot of Firm 1's marginal profit in Figure 10 illustrates why this might not always be the case. The marginal profit jumps discontinuously at the borders between

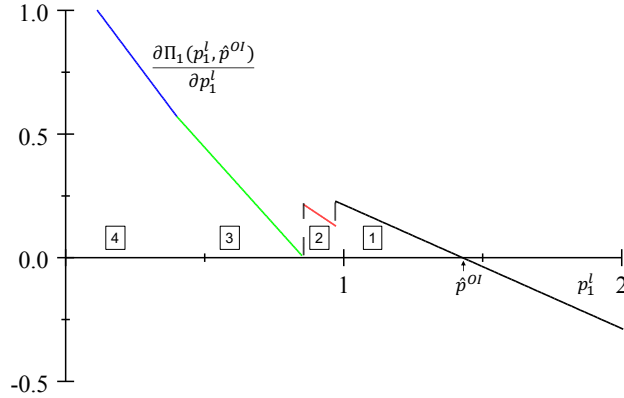


Figure 10: Firm 1's marginal profit function when Firm 2 sets $p_2^l = \hat{p}^{OI} \approx 1.43$.

Regions 3 and 2, and between Regions 2 and 1. In each case, as Firm 1 prices itself out of contention for a group of consumers who are sensitive to its list price, its disincentive to raising its price even further is reduced. Or, to make the same argument in the opposite direction, the marginal benefit from *reducing* p_1^l jumps up at these boundaries as Firm 1 begins to access groups of consumers whose demand is more elastic with respect to its list price. These discontinuities in marginal profits imply that profit functions could be multi-peaked (although that is not the case here), so care must be taken with any candidate symmetric equilibrium to confirm that neither firm would rather undercut the putative equilibrium price.

Consumers do poorly in this equilibrium compared to our usual benchmarks. If targeting is banned or unrestricted, list price neutrality applies, with prices $p^{NT} = p^T = 1$ and per-firm profits $\Pi^{NT} = 0.5$ or $\Pi^T = 0.29$ respectively. With an opt-in policy and inattentive consumers, list prices are 43% higher than either benchmark, and profits are higher as well: $\Pi^{OI} \approx 0.551$. All consumers with high privacy costs and roughly 90% of those with low privacy costs are worse off under this opt-in policy than they would be if targeted discounts were simply banned. Furthermore, 70% of consumers (40% of the high privacy cost types and all of the low-cost ones) are worse off under this opt-in policy than they would be under unrestricted targeting.

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) underpin patterns that arise quite 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 log-concave captive demand. As we show, that intuition can reverse if demand is log-convex: 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.)

A second, more speculative potential explanation is imperfect targeting. In models like ours, targeting induces 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.⁵⁰ Slightly imperfect targeting would not change our conclusions. In particular, firms will continue to mix over whether to advertise because consumer price sensitivity drives down gross profits in the price-competition game so that both firms cannot cover the ad cost A , though discount prices may be in pure strategies. Importantly, we also note that small amounts

⁵⁰Often (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.

of targeting noise does not cause (say) Firm 1 to advertise to a consumer at the boundary y^* of its captive region since in equilibrium Firm 2 cannot potentially profit from poaching just inside Firm 1’s captive border, even if there are subtle differences in the location of this boundary.⁵¹ In related work, we explain on continuity grounds why competition for contested consumers would continue to be fierce if firms’ information about consumers were a little bit noisy.⁵²

In practice, some firms may have collected proprietary information about consumer tastes. One way to model asymmetrically informed firms in our setting is *via* the targeting cost: suppose a better-informed firm can identify particular types of consumer at lower cost. This is perhaps too reductive to be entirely satisfying, but more sophisticated approaches appear rather challenging. To illustrate, consider the rather natural case where a firm knows an individual consumer’s taste for its own product, but not how that consumer values alternative products. Discount competition for such an individual then resembles an asymmetric independent private-values auction (where the firms “bid” in surplus offers) with costly entry (the targeting cost) and an endogenous outside option (the chance of making a list price sale without advertising if the consumer’s other options end up being sufficiently weak). The latter two features imply that a firm will refrain from targeting consumers with sufficiently low or sufficiently high values for its product (in the first case because a discount is unlikely to succeed, and in the latter because it is unlikely to be necessary). While standard tools from auction theory could be brought to bear on this problem, both the asymmetry and the endogenously top- and bottom-truncated supports of the bidding distributions would pose technical hurdles.

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 marketplace. 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

⁵¹We thank referees for prompting these and subsequent questions.

⁵²See Anderson, Baik, and Larson (2015).

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. Thisse and Vives (1988) find a result of this kind for the dominant firm when the asymmetry between firms is sufficiently large.⁵³

While we have assumed that list prices precede discount offers, one might also consider the case where all prices (list and discount) are set simultaneously. In our setting, with list prices set first, there is a Stackelberg leader effect: by reducing its list price, a firm can discourage its rivals from advertising to some consumers they would have otherwise tried to poach. Since this effect is absent in the simultaneous version of the model, one might expect equilibrium list prices to be higher. Unfortunately this hypothesis is difficult to evaluate because the model with simultaneous price-setting fails to have a pure-strategy equilibrium in list prices.⁵⁴

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). Furthermore, we have not addressed the market in which firms acquire consumer data.⁵⁵

Finally, our results in Section 7 can be read as a strong but conditional defense of consumer opt-in requirements like those mandated by the GDPR. Under assumptions about demand that are common in the empirical literature, mandating opt-in makes

⁵³In 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.

⁵⁴We can sketch the details of why a candidate equilibrium at common list price p^l must fail when list prices and discount offers are chosen simultaneously. Assuming duopoly, for simplicity, if Firm 2 expects $p_1^l = p^l$, it does not advertise to any consumer types $y > y^* = p^l - A > 0$. But then, if Firm 1 were to deviate to $\hat{p}_1^l = p^l + \varepsilon$, with $\varepsilon < y^*$, it would do no worse on the contested consumers $y \in [0, y^*]$, and it would continue to sell at the higher list price to all consumers $y > y^*$ (since their next best option is to buy at $p_2^l = p^l$ from Firm 2). This rules out an equilibrium with captive consumers who pay list prices. Alternatively, consider a candidate equilibrium with list prices high enough that all consumers are contested. Then our Stage 2 analysis indicates that Firm 2 will make an atom of discount offers just undercutting p_1^l at each $y \in [0, \bar{y}]$. But then it must be profitable for Firm 1 to deviate to $p_1^l - \varepsilon$ to win back these atoms of poached consumers without paying the ad cost A on them.

⁵⁵But see Montes, Sand-Zantman, and Valletti (2015).

all consumers better off if targeting is common and consumers are vigilant (adjusting their privacy choices as prices change). Because this conclusion can be overturned if consumers are less agile about updating their privacy choices, it seems important to gather data about how these privacy choices are made in practice. Furthermore, while a case can be made for opting in as an all-or-nothing decision (as we have modeled it), it would be helpful to understand how our conclusions hold up if consumers can choose which personal information to release, and to which firms to release it.

References

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

Acquisti, A., Taylor, C. and Wagman, L., 2016. The economics of privacy. *Journal of Economic Literature*, 54(2), pp.442-92.

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 André 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.

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.

Belleflamme, P., Lam, W.M., and Vergote, W., 2017, Price Discrimination and Dispersion under Asymmetric Profiling of Consumers, mimeo.

Belleflamme, P., and Vergote, W., 2016, Monopoly Price Discrimination and Privacy: The Hidden Cost of Hiding, *Economics Letters*, 149, 141 - 144.

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, Z., Choe, C., and Matsushima, N., 2018, Competitive Personalized Pricing, mimeo.
- 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.
- Prat, Andrea and Valletti, Tommaso, 2019, Attention Oligopoly, mimeo.
- Roy, S., 2000, Strategic Segmentation of a Market, *International Journal of Industrial Organization*, 18, 1279-1290.
- 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.
- Stahl, D., 1994, Oligopolistic Pricing and Advertising, *Journal of Economic Theory*, 64(1), 162-177.
- Taylor, C. and Wagman, L., 2014, Customer Privacy in Oligopolistic Markets: Winners, Losers, and Welfare, *International Journal of Industrial Organization*, 34, 80-84.
- Thisse, J-F. and Vives, X., 1988, On The Strategic Choice of Spatial Price Policy, *American Economic Review*, 78, 122-137.
- Tucker, C.E., 2015. Privacy and the Internet. In *Handbook of Media Economics* (Vol. 1, pp. 541-562). North-Holland.
- Turow, J., 2017, *The Aisles Have Eyes: How Retailers Track Your Shopping, Strip Your Privacy, and Define Your Power*, Yale University Press.

Turow, J., King, J., Hoofnagle, C.J., Bleakley, A. and Hennessy, M., 2009. Americans reject tailored advertising and three activities that enable it. Available at SSRN: <https://ssrn.com/abstract=1478214> or <http://dx.doi.org/10.2139/ssrn.1478214>.

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.

White, T.B., Zahay, D.L., Thorbjørnsen, H. and Shavitt, S., 2008. Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters*, 19(1), pp.39-50.

A Appendix

Proof of Lemma 1

Part (i) We appeal to known properties of log-concave distributions; see the references for further information.⁵⁶ Cumulative distribution functions and their complements are strictly log-concave if their density functions are, so $F(x)$ and $F(x+y)$ are strictly log-concave. Products of strictly log-concave functions are strictly log-concave, so $f_{(1:n-1)}(x)$ is strictly log-concave, 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 log-concave functions are strictly log-concave, so integrating over x , we have $G(y)$ and $1-G(y)$ strictly log-concave. Similar arguments applies to $g(y) = \int f(r+y)f_{(1:n-1)}(r)dr$.

Part (ii) We will prove that $g'(0) \leq 0$, with $g'(0) < 0$ if $n \geq 3$. The claim follows because $g'(y)/g(y)$ is strictly decreasing by part (i).

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 we can write

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

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

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

⁵⁶For example, see Bergstrom and Bagnoli (2005).

where the second term should be understood as $\lim_{r \rightarrow \infty} f(r) f_{(1:n-1)}(r-y) = 0$ if $\bar{r} = \infty$ (since $\lim_{r \rightarrow \infty} f(r) = 0$ 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 $f'(r) f(r) = \frac{1}{2} d(f(r)^2)$, so if $n = 2$ we have $\frac{g'(0)}{n-1} = -\frac{1}{2} (f(\bar{r})^2 + f(\underline{r})^2) \leq 0$. Otherwise, integrate by parts to get

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

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

Part (iii) Note that for $y \geq 0$ we can write $1 - G(y) = \int_{\underline{r}}^{\bar{r}-y} f(r+y) F(r)^{n-1} dr$. The claim $\hat{G}(y) < G(y)$ follows because $\hat{G}(y)$ has a strictly smaller integrand. For the second claim, hold a particular $y \geq 0$ fixed throughout the arguments below. We have $g(y) = f(\bar{r}) F(\bar{r}-y)^{n-1} - \int_{\underline{r}}^{\bar{r}-y} f'(r+y) F(r)^{n-1} dr$, so the hazard rate may be written $g(y) / (1 - G(y)) = B(y) - C(y)$, where

$$B(y) = f(\bar{r}) \left(\int_{\underline{r}}^{\bar{r}-y} f(r+y) \left(\frac{F(r)}{F(\bar{r}-y)} \right)^{n-1} dr \right)^{-1} \text{ and } C(y) = \int_{\underline{r}}^{\bar{r}-y} \frac{f'(r+y)}{f(r+y)} h(r) dr$$

and $h(r)$ is the density function $h(r) = H'(r) = f(r+y) F(r)^{n-1} / \int_{\underline{r}}^{\bar{r}-y} f(r'+y) F(r')^{n-1} dr'$ with cdf

$$H(r) = \frac{\int_{\underline{r}}^r f(r'+y) F(r')^{n-1} dr'}{\int_{\underline{r}}^{\bar{r}-y} f(r'+y) F(r')^{n-1} dr'}$$

Define the corresponding terms $\hat{B}(y)$, $\hat{C}(y)$, and $\hat{H}(y)$ with $n+1$ firms. It will suffice to show $\hat{B}(y) \geq B(y)$ and $\hat{C}(y) < C(y)$. The former is satisfied if $f(\bar{r}) = 0$; otherwise it follows from $F(r)/F(\bar{r}-y) \leq 1$ within the integrand. For the latter, we claim (proved below) that $\hat{H}(r) < H(r)$. Then, as $f'(r+y)/f(r+y)$ is a strictly decreasing function in r (by part (i)), we have $\hat{C}(y) < C(y)$ by properties of first-order stochastic dominance. This completes the proof.

Proof of the claim that $\hat{H}(r) < H(r)$. Let $\eta(r) = f(r+y)F(r)^{n-1}$. Then

$$\frac{H(r)}{1-H(r)} = \frac{\int_r^{\bar{r}} \eta(r') dr'}{\int_r^{\bar{r}-y} \eta(r') dr'} = \frac{\int_r^{\bar{r}} \eta(r') F(r) dr'}{\int_r^{\bar{r}-y} \eta(r') F(r) dr'}$$

while

$$\frac{\hat{H}(r)}{1-\hat{H}(r)} = \frac{\int_r^{\bar{r}} \eta(r') F(r') dr'}{\int_r^{\bar{r}-y} \eta(r') F(r') dr'}$$

Because $F(r') < F(r)$ for $r' \in [\underline{r}, r)$ and $F(r') > F(r)$ for $r' \in (r, \bar{r}-y]$, we can conclude that $\frac{\hat{H}(r)}{1-\hat{H}(r)} < \frac{H(r)}{1-H(r)}$.

Proof of Proposition 3

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. Furthermore, these first-order conditions are also sufficient for profit maximization, as (referring to (3)) strict logconcavity of $1 - G(y)$ and the fact that P_j weakly is increasing in p_j^l imply that each firm's marginal profit $\partial \Pi_j(p_j^l) / \partial p_j^l$ is strictly positive for $p_j^l < p^*$ and strictly negative for $p_j^l > p^*$. This establishes the symmetric equilibrium at p^* . 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 $\Theta(p^l)$ 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 $\Theta(p^l)$ 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.

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 $v(u, v)$ by

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

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

$$v(p_1^l, p_2^l) > v(p_2^l, p_2^l) \geq v(p_2^l, p_1^l) \geq 0.$$

Proof of Proposition 6

Remarks about sufficient conditions for a no-targeting equilibrium. By intention, the proposition *supposes* the existence of a no-ads symmetric equilibrium, rather than deriving existence from primitive conditions. Before beginning the formal proof, we briefly discuss conditions under which such an equilibrium *does* exist with logconvex demand. (A full, formal analysis on this point would take us too far afield.)

When targeted ads are not available, a candidate equilibrium price satisfies the symmetric first-order condition (4); thus $p^{NT} = \frac{1-G(0)}{g(0)}$. Firm j 's marginal profit when all other firms charge p^{NT} is

$$\Pi'_j(p_j) = g(p_j - p^{NT}) \left(\frac{1 - G(p_j - p^{NT})}{g(p_j - p^{NT})} - p_j \right)$$

A sufficient condition for the equilibrium is that $\Pi'_j(p_j) \geq 0$ for all $p_j \leq p^{NT}$. Because we have $\frac{1-G(y)}{g(y)}$ strictly increasing, one simple condition that ensures existence is that $\frac{1-G(y)}{g(y)}$ has slope less than one. One family of demand systems satisfying this condition is $1 - G(y) \propto y^{-a}$ with $a > 1$.

We continue on to the main proof.

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 log-convexity 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) $\Pi_1^{dev}(A)|_{A=p^{NT}} = \Pi^{NT} < \Pi^{FT}$. (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 $p^l = \bar{y} + p^{NT}$ and using targeting yields profit $\tilde{\Pi}_1^{dev}(A) = \int_{A-p^{NT}}^{\bar{y}} (y + p^{NT} - A) dG(y)$ 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 $\tilde{\Pi}_1^{dev}(A) \rightarrow \Pi^{FT} > \Pi^{NT}$ as $A \rightarrow p^{NT}$, 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 10 *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 4, we have

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

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

Proof of Proposition 8

Differentiate to get

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

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 10.

Proof of Proposition 10

Given symmetry, it suffices to aggregate over consumers $y_1 \geq 0$ with favorite product at Firm 1. Define $EP(y_1)$ as in the text, with \overline{EP} its average over $y_1 \geq 0$. It suffices to show $\overline{EP}^T > \overline{EP}^{NT} = p^{NT}$. When targeting is permitted, we have $EP(y_1) = p^T$ if $y_1 > y^*$, or $EP(y_1) = y_1 + L(y_1, p^T, A)$ if $y_1 \in [0, y^*)$, where $p^T = y^* + A$ and y^* satisfies the equilibrium condition $\mu(y^*) = A$ (possibly at $y^* = \infty$ if $\lim_{y \rightarrow \infty} m(y) = m > A$). Because $L(y_1, p^T, A) \geq A$ (see (10)), we have

$$\overline{EP}^T = \int_0^\infty EP(y) \frac{g(y)}{1-G(0)} dy \geq \int_0^{y^*} y \frac{g(y)}{1-G(0)} dy + \int_{y^*}^\infty y^* \frac{g(y)}{1-G(0)} dy + A$$

After integrating by parts this reduces to $\overline{EP}^T \geq \frac{1}{1-G(0)} \int_0^{y^*} 1-G(y) dy + A$.

Using Lemma 11(i) and the fact that $p^{NT} = \mu(0)$, we have $y^* \geq (p^{NT} - A)/\rho$. Then use Lemma 11(ii) to get $\overline{EP}^T \geq \int_0^{\frac{1}{\rho}(p^{NT}-A)} \left(1 - \rho \frac{y}{p^{NT}}\right)^{1/\rho} dy + A$. Integrate to get

$$\overline{EP}^T \geq \frac{p^{NT}}{1+\rho} - \frac{p^{NT}}{1+\rho} \left(\frac{A}{p^{NT}}\right)^{\frac{1+\rho}{\rho}} + A$$

Writing $\alpha = A/p^{NT}$ and using this bound, a sufficient condition for $\overline{EP}^T - \overline{EP}^{NT} > 0$ is $\alpha - \frac{\rho}{1+\rho} - \frac{1}{1+\rho} \alpha^{\frac{1+\rho}{\rho}} > 0$. Rearrange this condition as:

$$\rho < \alpha \left(\frac{1 - \alpha^{1/\rho}}{1 - \alpha} \right) \quad (21)$$

Since $\alpha < 1$, $\rho < \alpha$ suffices to ensure that $\frac{1-\alpha^{1/\rho}}{1-\alpha} > 1$. Thus we conclude that $\rho < \alpha$ is sufficient to ensure (21).

Lemma 11 *Let $\mu(y)$ be the Mills ratio $\mu(y) = \frac{1-G(y)}{g(y)}$. If captive demand $1-G(y)$ is ρ -convex on $[0, \infty)$, then for $y \geq 0$, (i) $\mu(y) \geq \mu(0) - \rho y$, and (ii) $1-G(y) \geq (1-G(0)) \left(1 - \rho \frac{y}{p^{NT}}\right)^{1/\rho}$.*

Proof. To establish (i), note that the condition that $\frac{d^2}{dy^2} (1-G(y))^\rho \geq 0$ can be shown equivalent to $\mu'(y) \geq -\rho$ by direct computation. Recall that $p^{NT} = \mu(0)$. Thus the hazard rate $\nu(y) = 1/\mu(y)$ satisfies $\nu(y) \leq (p^{NT} - \rho y)^{-1}$. For (ii), note that $1-G(y) = (1-G(0)) \exp(-\int_0^y \nu(y') dy')$. Using the bound on $\nu(y)$, we have $-\int_0^y \nu(y') dy' \geq \frac{1}{\rho} \ln \frac{p^{NT} - \rho y}{p^{NT}}$, from which (ii) follows directly. ■

Proof of Proposition 11

First note that $g'(0) \leq 0$ and Condition 1 imply $g'(y) < 0$ for all $y > 0$, and thus $dp^l/dA < 0$ for all $A < p^{NT}$ by Proposition 7. As noted in that proof, dp^l/dA has the sign of $g'(y^*)$, where $y^* = p^l(A) - A$. Then because the threshold consumer is $y^* = p^l - A = 0$ at $A = p^{NT}$, the additional condition $g'(0) < 0$ ensures that $dp^l/dA < 0$ holds at $A = p^{NT}$ as well. The claim that $g'(0) < 0$ is satisfied with $n \geq 3$ firms is proved in Lemma 1.

For $A \geq p^{NT}$, targeting is not employed and consumers receive their no-targeting surplus. Thus it suffices to show that there is a neighborhood $A \in (\bar{A}, p^{NT}]$ over which $CS(y)$ is strictly increasing in A for all y . 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, $L(y, p^l(A), A)$ is decreasing in A for all $y \in [0, y^*(A)]$. Because $dL(y, p^l, A)/dA$ 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 $dL(y, p^l(A), A)/dA|_{y=0, A=p^{NT}} < 0$, that is, that $L(y, p^l(A), A)$ is strictly decreasing in A at $A = p^{NT}$ for consumers at the turf boundary. That total derivative is $dL/dA = \partial L/\partial A + \partial L/\partial p^l \cdot dp^l/dA$. At $y = 0$, we have $L(0, p^l, A) = A(a_1 + a_2) = 2A - \frac{A^2}{p^l}$ since there are no social costs of misallocation, so the direct effect is $\partial L/\partial A|_{y=0, A=p^{NT}} = 2 - 2A/p^l|_{A=p^{NT}} = 0$. For the indirect effect, we have $\partial L/\partial p^l|_{y=0, A=p^{NT}} = (A/p^l)^2|_{A=p^{NT}} = 1$. Thus we can conclude that $dL(y, p^l(A), A)/dA|_{y=0, A=p^{NT}} = dp^l/dA|_{A=p^{NT}} < 0$, as claimed.

Proof of Proposition 9

Let $y^*(A)$ be the taste difference at the boundary of the contested region. By later arguments, $\lim_{A \rightarrow 0} y^*(A) = \bar{y} = \bar{r} - \underline{r}$; that is, as $A \rightarrow 0$ 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 $\lim_{A \rightarrow 0} p^l(A) \geq \bar{y}$. 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 $L(y, p^l(A), A)$, with

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

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 $L(y, p^l(A), A)$. At the turf boundary, $\lim_{y \rightarrow 0} L(y, p^l(A), A)$ 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 $L(0, p^l(A), A) = A(a_1 + a_2) = 2A - \frac{A^2}{p^l(A)}$, which tends to 0 with A (given $p^l(A)$ bounded away from zero).

Proof of Proposition 14

Let $h_F(r) = \frac{1-F(r)}{f(r)}$. Strict log-concavity and monotone convergence imply that $\lim_{r \rightarrow \infty} h_F(r)$ exists and is finite; consequently $\lim_{r \rightarrow \infty} h'_F(r) = 0$. 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

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

for some $\hat{y} \in (0, p_n^{NT})$ by the intermediate value theorem. Apply the IVT once more to get $h_n(\hat{y}) = h_n(0) - |h'_n(\hat{y})|\hat{y}$ for some $\hat{y} \in (0, \hat{y})$, using $h'_n(y) < 0$. Since $\hat{y} < \hat{y} < p_n^{NT} = h_n(0)$, we have $h_n(\hat{y})/h_n(0) \in (1 - |h'_n(\hat{y})|\hat{y}, 1)$. With careful application of Theorem 2 of Gabaix et al. (2015), asymptotically as $n \rightarrow \infty$ we have $h'_n(\hat{y}) \sim h'_F(\hat{r}_n + \hat{y})$, where \hat{r}_n is defined by $1 - F(\hat{r}_n) = \frac{1}{n-1}$. Clearly $\hat{r}_n \rightarrow \infty$ with n . But then $\lim_{n \rightarrow \infty} h'_n(\hat{y}) = \lim_{r \rightarrow \infty} h'_F(r) = 0$. Thus we can conclude that $\lim_{n \rightarrow \infty} \ln H_n = -1$, so we are done.

Supporting results for the proof of Proposition 15

Lemmas 9 and 12 are used in the proof of Proposition 15. Lemma 9 establishes some facts about the term $\partial y^*(c; p_1^l) / \partial p_1^l$ in equation (16). This is the rate at which the opt-in threshold at cost type c rises with Firm 1's list price; the fact that this rate of increase is less than one-for-one works against our conclusions about convex demand in Proposition 15. Thus the main purpose of Lemma 9 is to bound $\partial y^*(c; p_1^l) / \partial p_1^l$ sufficiently close to one. This leads to a condition on p^{NT} , p^T , and A (equation (22)) under which Proposition 15 goes through; Lemma 12 shows this condition is implied by ρ -convexity of demand. Toward these ends, for $y \in [0, p - A]$, we define the function $D(y, p, A) = -\frac{\Delta_p(y, p, A)}{\Delta_y(y, p, A)}$. Since the opt-in threshold is defined by $\Delta(y^*(c; p_1^l), p_1^l) = c$, the implicit function theorem implies that $\partial y^*(c; p_1^l) / \partial p_1^l = D(y^*(c; p_1^l), p_1^l, A)$.

Proof of Lemma 9

While $\Delta(y, p)$ can be obtained from (11), 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 $\Delta(y; p_1^l) = E(\max(\tilde{s}_1, \tilde{s}_2))$. Overloading notation, it follows from the results in Section 3.2 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 $B_1(\tilde{s})B_2(\tilde{s})$. 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

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

Meanwhile,

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

using the top line expression for $\Delta_p(y, p, A)$, where

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

Then $D(y, p, A) = \frac{\Delta_p}{\Delta_p + K}$. Then $K > 0$ for $y < p - A$ implies the fact that $D(y, p, A) < 1$ for $y < p - A$. A straightforward application of l'Hôpital's Rule establishes that $D(p - A, p, A) = 1$, and the $A \rightarrow 0$ limit follows from $\lim_{A \rightarrow 0} \Delta_p = 1$ and $\lim_{A \rightarrow 0} K = 0$.

It remains to establish the lower bound on $D(y, p, A)$. Direct integration yields

$$K = \frac{A}{y} \left(\frac{p-y-A}{p} - \frac{A}{y} \ln \left(\frac{A+y}{p} \frac{p-y}{A} \right) \right)$$

Note that $D(y, p, A) = \frac{1}{1+K/\Delta_p(y, p, A)}$, so a lower bound on $D(y, p, A)$ will correspond to an upper bound on $K/\Delta_p(y, p, A)$. Working toward an upper bound on K , let

$M = \frac{A}{y} \ln \left(\frac{A+y}{p} \frac{p-y}{A} \right)$ be the second term in parentheses above. Rearrange this as

$$\begin{aligned} M &= \frac{A}{y} \ln \left(\left(1 - \frac{p-y-A}{p} \right) \left(1 + \frac{p-y-A}{A} \right) \right) \\ &= \frac{A}{y} \ln \left(1 + \frac{y}{A} \frac{p-y-A}{p} \right) \end{aligned}$$

Over the relevant range $y \in [0, p-A]$, the argument to the natural log is weakly greater than one. Then, because a second-order Taylor series expansion establishes that $\ln(1+x) \geq x - \frac{1}{2}x^2$ for $x \geq 0$, we have

$$M \geq \frac{p-y-A}{p} - \frac{1}{2} \frac{y}{A} \left(\frac{p-y-A}{p} \right)^2$$

It follows that

$$K \leq \frac{1}{2} \left(\frac{p-y-A}{p} \right)^2$$

and so

$$\frac{K}{\Delta_p(y, p, A)} \leq \frac{1}{2} \frac{p-y}{p(p+A)} (p-y-A)$$

Then

$$\begin{aligned} D(y, p, A) &= 1 - \frac{K/\Delta_p(y, p, A)}{1 + K/\Delta_p(y, p, A)} \\ &\geq 1 - \frac{p-y}{2p(p+A) + (p-y)(p-y-A)} (p-y-A) \\ &\geq 1 - \frac{1}{2\frac{p}{p-y}(p+A) + (p-y-A)} (p-y-A) \\ &\geq 1 - \frac{1}{2\frac{p+y}{p}(p+A) + (p-y-A)} (p-y-A) \\ &\geq 1 - \frac{1}{3p+A+y(1+2A/p)} (p-y-A) \\ &\geq 1 - \frac{p-y-A}{3p+A} \end{aligned}$$

This proves the claim. The second line uses the Taylor series bound derived above, and the fourth line uses $\frac{p}{p-y} \geq \frac{p+y}{p}$.

Lemma 12 *If captive demand is strictly $\frac{2}{3}$ -convex for $y \geq 0$, then the list price with unrestricted targeting represents a markup of at least $\frac{p^T - p^{NT}}{p^{NT}} > \delta (A/p^{NT})$ over the list price when targeting is banned, where $\delta(\alpha) = \frac{1}{2} \left(\sqrt{(\alpha + \frac{1}{2})^2 + 2(1 - \alpha)} - (\alpha + \frac{1}{2}) \right)$. This is equivalent to the condition:*

$$p^T \left(2 \frac{p^T + A}{3p^T + A} \right) > p^{NT} . \quad (22)$$

Proof. We will demonstrate that (22) holds; the markup version follows by solving (22) for p^T . Write LHS for the lefthand side of (22). If captive demand is strictly ρ -convex, then (following Lemma 11 and the proof of Proposition 10) the threshold captive under unrestricted targeting satisfies $y^* > (p^{NT} - A) / \rho$, with $p^T = y^* + A$. Thus here we have $p^T > (3p^{NT} - A) / 2$. Using this bound, we have $LHS > \widehat{LHS} = (3p^{NT} - A) \left(\frac{3p^{NT} + A}{9p^{NT} - A} \right)$. But $\widehat{LHS} = p^{NT} + \frac{A(p^{NT} - A)}{9p^{NT} - A} \geq p^{NT}$, so we are done. ■

Proof of Proposition 15

We will treat the cases of concave and convex captive demand separately. As earlier, let $\mu(y) = (1 - G(y)) / g(y)$ be the Mills ratio.

Concave demand:

Evaluate marginal profit $d\Pi_1 / dp_1^l$ in the OI model at a candidate symmetric equilibrium with list prices $p_1^l = p_2^l = P \leq p^T$. Using (15), this marginal profit is an expectation of $\Lambda(c)$ terms; to rule out the candidate equilibrium, it suffices to show that $\Lambda(c, P) > 0$ for all $c \geq 0$. Let $\bar{c} = \Delta(0, P, A)$ be the privacy cost above which a consumer opts out regardless of location y . We will treat higher privacy cost and lower privacy cost consumers separately.

High privacy costs: $c > \bar{c}$

For $c > \bar{c}$, we have $\Lambda(c, P) = g(0) (\mu(0) - P) = g(0) (p^{NT} - P)$. But concave demand implies $p^{NT} \geq p^T$, with strict inequality if demand is strictly concave. Thus we have $\Lambda(c, P) \geq 0$ (with strict inequality if demand is strictly concave).

Low privacy costs: $c \in (0, \bar{c}]$

In this case the marginal profit at cost type c is $\Lambda(c, P) = g(y^*(c, P)) \lambda(y^*(c, P), P)$, where $\lambda(y, p)$ is defined by

$$\lambda(y, p) = \mu(y) - (p - y) D(y, p, A)$$

$$\Lambda(c, P) = g(y^*) (\mu(y^*) - (P - y^*) D(y^*, P, A))$$

For $c \in (0, \bar{c}]$, we have $y^*(c, P) \in [0, P - A)$ and hence $D(y^*(c, P), P, A) < 1$ by Lemma 9. Thus $\Lambda(c, P) > g(y^*(c, P)) \hat{\lambda}(y^*(c, P), P)$, where

$$\hat{\lambda}(y, p) = \mu(y) + y - p$$

Note that $\hat{\lambda}(p^T - A, p^T) = 0$. (This is equivalent to the first-order condition that identifies p^T .) Then $\mu(y)$ decreasing (by strict log-concavity of $g(y)$) implies $\hat{\lambda}(P - A, P) \geq \hat{\lambda}(p^T - A, p^T)$. Then the facts that $y^*(c, P) < P - A$ and $\frac{\partial \hat{\lambda}}{\partial y} \leq 0$ (by concavity of $1 - G(y)$) imply that $\hat{\lambda}(y^*(c, P)) \geq 0$. It follows that $\Lambda(c, P) < 0$ as claimed.

Convex demand: Note that strict convexity of $1 - G(y)$ and condition (22) follow from strict $\frac{2}{3}$ -convexity. Thus it suffices to show the conclusions hold when (22) is satisfied.

Convex demand, upper bound ($p^{OI} < p^T$):

Consider a candidate symmetric pure strategy equilibrium at list prices $p_1^l = p_2^l = P \geq p^T$. Firm 1's marginal profit, evaluated at the candidate equilibrium price, is given by $\left. \frac{d\Pi_1}{dp_1^l} \right|_{p_1^l=P} = \int_0^\infty \Lambda(c, P) dH(c)$. To rule out the candidate equilibrium, it suffices to show that $\Lambda(c, P) < 0$ for all $c > 0$. There are two cases to consider.

High privacy costs: $c > \bar{c} = \Delta(y, P, A)|_{y=0}$

As consumers with $c > \bar{c}$ opt out regardless of location, we have $\Lambda(c, P) = 1 - G(0) - Pg(0)$. Noting that $p^{NT} = \mu(0)$, we can write $\Lambda(c, P) = g(0)(p^{NT} - P) < 0$ (because $P \geq p^T$ and for strictly convex demand $p^T > p^{NT}$).

Low privacy costs: $c \leq \bar{c}$

In this case, we once again have $\Lambda(c, P) = g(y^*(c)) \lambda(y^*(c, P), P)$, with $\lambda(y, p)$ as defined above and $y^*(c, P) \in [0, P - A]$ implied by $c \in [0, \bar{c}]$. Using Lemma 9, we have $\lambda(y, p) \leq \bar{\lambda}(y, p)$ for $y \in [0, p - A]$, where

$$\bar{\lambda}(y, p) = \mu(y) - \left((p - y) \left(1 - \frac{p - y - A}{3p + A} \right) \right)$$

The term in parentheses is concave in y , and we have assumed $\mu(y)$ convex; thus

$\bar{\lambda}(y, p)$ is convex. Furthermore, we have the following:

$$(i) \quad \bar{\lambda}(y, p)|_{y=P-A, p=P} = \mu(P-A) - A \leq 0$$

$$(ii) \quad \bar{\lambda}(0, P) = p^{NT} - 2P \left(\frac{P+A}{3P+A} \right) < 0$$

Condition (i) follows because p^T is defined by $\mu(p^T - A) - A = 0$, plus $P \geq p^T$ and the fact that $\mu(y)$ is strictly decreasing (by strict log-concavity of $g(y)$). Condition (ii) follows by $\mu(0) = p^{NT}$, condition (22), and the fact that $2p \left(\frac{p+A}{3p+A} \right)$ is increasing in p . Given $\bar{\lambda}(y, p)$ convex, (i) and (ii) imply $\bar{\lambda}(y^*(c), P) < 0$ for all $c \in (0, \bar{c}]$. A fortiori, we have $\lambda(y^*(c), P) < 0$ and so $\Lambda(c, P) < 0$ for all $c \in (0, \bar{c}]$.

Convex demand, lower bound ($p^{OI} \geq p^{NT}$):

Fix a candidate symmetric pure strategy equilibrium with $p_1^l = p_2^l = P < p^{NT}$. The argument parallels those above; we will show that Firm 1's marginal profit on each privacy cost type satisfies $\Lambda(c, P) > 0$, and this suffices to show that setting $p_1^l = P$ cannot be a best response for Firm 1. For cost types $c > \bar{c} = \Delta(0, P, A)$, the argument is just as above: we have $\Lambda(c, P) = g(0)(p^{NT} - P) > 0$. For cost types $c \in (0, \bar{c}]$, the same argument used for concave demand implies $\Lambda(c, P) > g(y^*(c, P)) \hat{\lambda}(y^*(c, P), P)$. But in this case, strict convexity of $1 - G(y)$ implies $\frac{\partial \hat{\lambda}}{\partial y} > 0$, and therefore $\Lambda(c, P) > g(y^*(c, P)) \hat{\lambda}(0, P)$. But $\hat{\lambda}(0, P) = \mu(0) - P = p^{NT} - P \geq 0$. This implies $\Lambda(c, P) > 0$, as claimed. Thus $d\Pi_1/dp_1^l|_{p_1^l=P} > 0$, and so there cannot be a symmetric equilibrium at common list price P .

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

Section C Equilibrium existence example for the opt-in model with inattentive consumers

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 6 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 $p_1 \in (A, p^{FT})$ 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 $A \in (\frac{3}{2}, A^*)$, firms are better off than under no-targeting – in particular, profits are 24% higher at $A = A^*$.

The transition If $A \in (A^*, \bar{A})$, 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

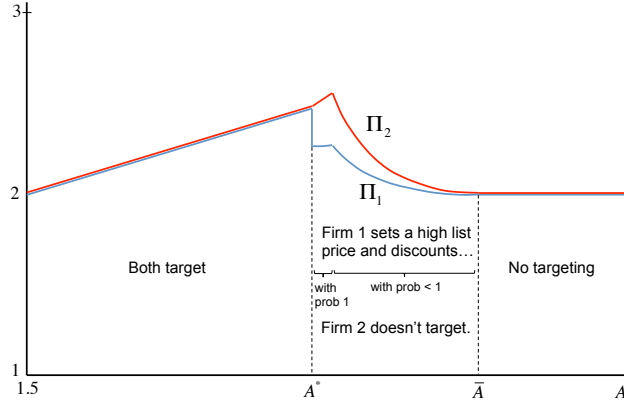


Figure 11: Equilibrium profits. Π_1 and Π_2 coincide in the symmetric no-targeting ($A \geq \bar{A}$) and full-targeting ($A \leq A^*$) regimes. For $A \in (A^*, \bar{A})$, Π_1 (blue) and Π_2 (red) are the asymmetric equilibrium profits described in the text.

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 \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 $A \in (2.517, \bar{A})$ are then $\Pi_2 = \frac{1}{2}(p^*)^2$ and $\Pi_1 = \Pi_{1L} = \frac{1}{2}(1 + \frac{1}{2}p^*)^2$. We have $q \rightarrow 1$ as $A \rightarrow 2.517$, 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 11 plots the firms' equilibrium profits for $A \in [1.5, 3.5]$ with this equilibrium selection on $A \in (A^*, \bar{A})$. 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 $A \in (A^*, \bar{A})$, 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 11 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 11 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 11 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 11, 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 17 *Proposition 11 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 $T'''(\frac{1}{2}) > 8T'(\frac{1}{2}) = 8p^{NT}$ 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).*

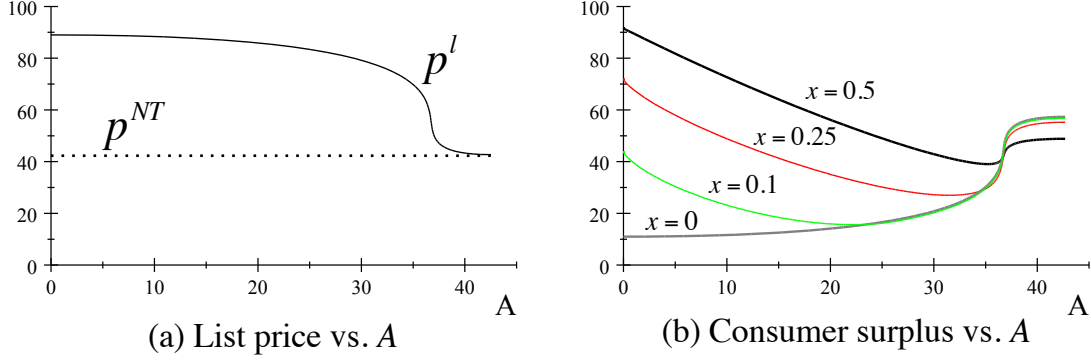


Figure 12: Targeting initially harms all consumers: two-firm example

Figure 12 gives a Hotelling example with nonlinear transportation costs.⁵⁷ As A declines below $p^{NT} \approx 42.7$ the list price rises, with a particularly sharp increase 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 $L(y, p^l, A)$ over all contested consumers $y \in [0, p^l - A]$; 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 18 *The welfare loss $L(y, p^l, A)$ on contested consumers is strictly concave in y (for $y \in [0, p^l - A]$). 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 $L(y, p^l, A)$ is strictly decreasing in y . If the targeted ad cost is low ($\frac{A}{p^l} < \sqrt{2} - 1$), then $\hat{y} \in (0, p^l - A)$: welfare losses are largest for contested consumers strictly between the turf boundary and the captive-contested boundary, and $L(y, p^l, A)$ is inverse-U-shaped in y .*

⁵⁷Note that the transportation cost $T(d) = e^{4.5d} - 1$ is increasing and convex, with $T(0) = 0$.

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.⁵⁸

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

For strictly log-concave 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 19 *Suppose tastes are distributed i.i.d. according to strictly log-concave, well-behaved density $f(r)$, and $h_F = \lim_{r \rightarrow \bar{r}} \frac{1-F(r)}{f(r)}$. Then as the number of firms in the market increases, $\lim_{n \rightarrow \infty} p^{NT} = h_F / \Gamma(2 + \gamma)$. Then the fraction of contested consumers satisfies*

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

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 ($p^{NT} \rightarrow 0$), then

⁵⁸Gabaix et al. (2015) also explicitly require the existence of $\lim_{r \rightarrow \bar{r}} \frac{1-F(r)}{f(r)}$. But strict logconcavity suffices for this (by monotone convergence), so to simplify the exposition we simply restrict attention to strictly logconcave densities.

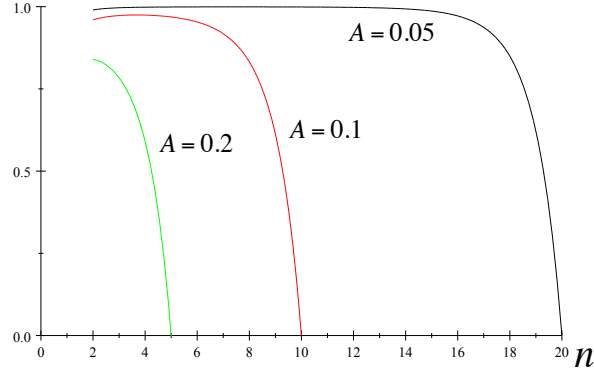


Figure 13: Fraction of consumers who are contested vs. n . (Tastes i.i.d. $r \sim U[0, 1]$)

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 13 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 13 is used in the proof of Proposition 17.

Lemma 13 *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 $CS = r_{(2)} - L(y, p^l, A)$. 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

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

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,

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

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

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

■

Proof of Proposition 17

If the density of $f(r)$ is strictly positive at \bar{r} or \underline{r} then (consult the proof of Lemma 1) $g'(0) < 0$ and Proposition 11 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 $CS^\varepsilon - CS^{NT} = r_{(2)} - L(0, p^\varepsilon, A_\varepsilon) - (r_{(1)} - p^{NT}) = p^{NT} - L(0, p^\varepsilon, A_\varepsilon)$. Recall that $L(0, p^l, A) = A(a_1 + a_2) = 2A - \frac{A^2}{p^l}$ (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 left-hand 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 (5) 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

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

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): $1 - G(p^l - A) = Ag(p^l - A)$. Differentiate totally with respect to A to get

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

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

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

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

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

Putting the pieces together, we have

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

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

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

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

Lemmas 14, 15, 16, and 17 are used in the proof of Proposition 18.

Lemma 14 *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 15 *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 16 *At the captive-contested boundary, $L_y(y, p, A)|_{y=p-A} = -1$.*

Proof. Evaluate L_y from the previous lemma. ■

Lemma 17 *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 $L_y(0, p, A) \geq 0$ iff $A/p \leq \sqrt{2}-1$. ■

Proof of Proposition 18

We appeal to Lemmas 15, 16, and 17. Concavity of $L(y, p, A)$ in y is given in Lemma 15. If $A/p > \sqrt{2}-1$, then $L(y, p, A)$ is decreasing in y at $y=0$ by Lemma 17; 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 17 and decreasing at $y=p-A$ by Lemma 16; along with strict concavity of L , this suffices for the second result.

Proof of Proposition 19

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 $A < \lim_{n \rightarrow \infty} \frac{1-G(0)}{g(0)} \leq \lim_{n \rightarrow \infty} \frac{1-G(y)}{g(y)}$ for all $y \geq 0$, so for n sufficiently large the symmetric equilibrium has $p^l = \infty$ and all consumers contested.

C Example of equilibrium with opt-in and inattentive consumers

This section gives supporting analysis for the example in the main text of equilibrium existence with opt-in and inattentive consumers. We identify a candidate symmetric equilibrium at common list price \hat{p}^{OI} using the sufficient condition (20). Then we verify that each firm cannot do better than to set its list price at \hat{p}^{OI} when the other firm does so. This involves constructing Firm 1's total profit $\Pi_1(p_1^l, \hat{p}^{OI})$ when consumers expect $(p_1^l, p_2^l) = (\hat{p}^{OI}, \hat{p}^{OI})$ and $p_2^l = \hat{p}^{OI}$. We verify numerically that $\Pi_1(p_1^l, \hat{p}^{OI})$ is single-peaked, with a maximum at $p_1^l = \hat{p}^{OI}$. This suffices to show that the candidate solution is an equilibrium. The analysis also illustrates some of the complications that might prevent (20) from identifying an equilibrium in other cases.

Consider two firms with linear (Hotelling) demand $1 - G(y) = \frac{1}{2}(1 - y)$ for $y \in [-1, 1]$. Fraction $\lambda = \frac{1}{2}$ of consumers have prohibitively high privacy costs, ensuring that they opt out; the remaining $1 - \lambda = \frac{1}{2}$ share of consumers have privacy cost $c = 0.2$. The targeted ad cost is $A = 0.4$. Note that for these parameters the no-targeting and unrestricted targeting prices are both equal to one: $p^{NT} = p^T = 1$.

We start by deriving the candidate symmetric equilibrium price \hat{p}^{OI} that solves (20). Then we verify this as an equilibrium by constructing Firm 1's profit function $\Pi_1(p_1^l, \hat{p}^{OI})$ and showing that it is quasi-concave and maximized at $p_1^l = \hat{p}^{OI}$. Condition (20) can be expressed as

$$\lambda(1 - G(0)) + (1 - \lambda)(1 - G(y^*)) - \lambda g(0)\hat{p}^{OI} = 0$$

where the threshold location at which low privacy cost consumers switch between opting in or out satisfies $\Delta(y^*, \hat{p}^{OI}) = c$. This can be simplified to the condition $\hat{p}^{OI} = \frac{1}{\lambda}(1 - (1 - \lambda)y^*)$, which has the solution

$$\hat{p}^{OI} \approx 1.428 \quad \text{and} \quad y^* \approx 0.572$$

This characterizes the candidate equilibrium; suppose that consumers opt in or out as specified by y^* in anticipation of common list price \hat{p}^{OI} .

Now fix $p_2^l = \hat{p}^{OI}$ and consider deviations to alternative list prices by Firm 1.

As long as $p_1^l \geq y^* + A \approx 0.972$ – call this Regime 1 – Firm 2 can afford to target all of the opt-in consumers at locations $y \in [0, y^*]$, so Firm 1 earns net profit y on each such consumer. Given the assumptions about demand, this implies total profit $(1 - \lambda)(G(y^*) - G(0))E(y | y \in [0, y^*]) = \frac{1-\lambda}{4}(y^*)^2$ on these consumers. Meanwhile, Firm 1 sells at its list price to a quantity $Q_1 = \lambda(1 - G(p_1^l - \hat{p}^{OI})) + (1 - \lambda)(1 - G(y^*))$. Consequently we have

$$\Pi_1(p_1^l, \hat{p}^{OI}) = p_1^l Q_1 + \frac{1-\lambda}{4}(y^*)^2 \quad \text{for } p_1^l \geq y^* + A \approx 0.972$$

Next consider Regime 2: $p_1^l \in [\hat{p}^{OI} - y^*, y^* + A)$. If Firm 1's price satisfies $p_1^l - A < y^*$, then there are locations $y \in [p_1^l - A, y^*]$ where low-cost consumers opted in expecting to get discounts, but Firm 1's unexpectedly low price prevents Firm 2 from targeting them. We revise Firm 1's list price demand to $Q_1 = \lambda(1 - G(p_1^l - \hat{p}^{OI})) + (1 - \lambda)(1 - G(p_1^l - A))$; revising its contested profits accordingly, we have

$$\Pi_1(p_1^l, \hat{p}^{OI}) = p_1^l Q_1 + \frac{1-\lambda}{4}(p_1^l - A)^2 \quad \text{for } p_1^l \in [\hat{p}^{OI} - y^*, y^* + A)$$

Next we come to Regime 3: $p_1^l \in [A, \hat{p}^{OI} - y^*)$, where the upper limit is $\hat{p}^{OI} - y^* \approx 0.856$. Remember that among opt-outs, Firm 1 sells to locations $y \geq p_1^l - \hat{p}^{OI}$. Among low-cost consumers, locations $y \in [-y^*, y^*]$ opted in, while locations $y < -y^*$ opted out, expecting to buy from Firm 2. When Firm 1 prices competitively enough – specifically $p_1^l - \hat{p}^{OI} < -y^*$ – it begins to win some of the latter consumers at its list price. In this case, Firm 1 sells at its list price to low-cost consumers located at $y \in [p_1^l - A, 1]$ and $y \in [p_1^l - \hat{p}^{OI}, -y^*]$, and to high-cost consumers located $y \in [p_1^l - \hat{p}^{OI}, 1]$. It earns net profit y on low-cost consumers $y \in [0, p_1^l - A]$, just as in Regime 2. Its list price demand and total profit are then

$$\begin{aligned} Q_1 &= \lambda(1 - G(p_1^l - \hat{p}^{OI})) + (1 - \lambda)(1 - G(p_1^l - A) + G(-y^*) - G(p_1^l - \hat{p}^{OI})) \\ \Pi_1(p_1^l, \hat{p}^{OI}) &= p_1^l Q_1 + \frac{1-\lambda}{4}(p_1^l - A)^2 \quad \text{for } p_1^l \in [A - y^*, \hat{p}^{OI} - y^*) \end{aligned}$$

Finally we have Regime 4: $p_1^l \in [0, A)$. As previously, Firm 1 captures opt-in consumers at locations $y \geq p_1^l - A$ at its list price. But now, because $p_1^l - A < 0$, all of the remaining contested consumers are on Firm 2's turf, and so the contested consumer portion of the profit expression above vanishes. The Regime 3 expression

for Firm 1's list price demand remains valid, and we have

$$\Pi_1(p_1^l, \hat{p}^{OI}) = p_1^l Q_1 \quad \text{for } p_1^l \in [0, A)$$

Figure 9 plots $\Pi_1(p_1^l, \hat{p}^{OI})$ over these four regions. It is clear that $\Pi_1(p_1^l, \hat{p}^{OI})$ has a unique maximum at $p_1^l = \hat{p}^{OI} \approx 1.428$. Given the symmetry of the game, this confirms that there is a symmetric equilibrium at the common list price \hat{p}^{OI} , with common profit $\Pi^{OI} \approx 0.551$. Notice that both the list price and profits are higher than under either a targeting ban ($\Pi^{NT} = 0.5$) or unrestricted targeting ($\Pi^T = 0.29$). All consumers with high privacy costs and roughly 90% of those with low privacy costs are worse off under this opt-in policy than they would be if targeted discounts were simply banned.