Review of Economic Studies (2022) **0**, 1–31 https://doi.org/10.1093/restud/rdac073

Advance access publication 28 October 2022

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

SIMON ANDERSON, ALICIA BAIK
Department of Economics, University of Virginia

and

NATHAN LARSON
Department of Economics, American University

First version received April 2016*; Editorial decision September* 2022*; Accepted October* 2022 (*Eds.*)

We study list price competition when firms can individually target consumer discounts (at a cost) afterwards, and we address recent privacy regulation (like the GDPR) allowing consumers to choose whether to opt-in to targeting. Targeted consumers receive poaching and retention discount offers. Equilibrium discount offers are in mixed strategies, but only two firms vie for each contested consumer and final profits on them are Bertrand-like. When targeting is unrestricted, firm list pricing resembles monopoly. For plausible demand conditions and if targeting costs are not too low, firms and consumers are worse off with unrestricted targeting than banning it. However, targeting induces higher (lower) list prices if demand is convex (concave), and either side of the market can benefit if list prices shift enough in its favour. Given the choice, consumers opt in only when expected discounts exceed privacy costs. Under empirically plausible conditions, opt-in choice makes all consumers better off.

Key words: Targeted advertising, Competitive price discrimination, Discounting, Privacy, GDPR, Opt-in.

JEL Codes: D43, L12, L13, M37

1. INTRODUCTION

As data analytics and pricing algorithms become common business practice in the digital era, there are growing concerns about the possibility that companies use such tools to engage in personalised pricing.[1] *OECD Secretariat, 2018*

1. Alerts about personalized pricing have also been raised by the OECD Competition Committee, the European Commission, the UK Competition and Markets Authority and Office of Fair Trading, the German Bundeskartellamt, and the US White House—see Anderson, Baik, and Larson (2019).

The editor in charge of this paper was Christian Hellwig.

Motivated by concerns like those voiced above, we develop a theory of targeted, personalized discounting and use it to evaluate the impact of such targeting on consumer surplus, profits, and a policy protecting consumer privacy. Our model engages the strategic linkage between two fronts on which firms compete: at the market level with list prices and at the individual level with customized discount offers. For any demand system with symmetric differentiated products, the key conclusions depend only on whether a reduced form of demand is convex or concave. We study whether privacy-minded consumers benefit from a policy that empowers them to opt into or out of targeted discounting. This policy captures the spirit of protections enacted in the European Union's recent General Data Privacy Regulation; its impact on consumers is not trivial because consumer privacy choices affect equilibrium price levels. We show that for empirically relevant demand systems, list prices fall when more individuals opt out, implying that the policy unambiguously benefits consumers. We next discuss targeting, privacy, and privacy policy, before outlining the workings of the model and its connection to the literature.

Our treatment of discounting amounts to costly, first-degree price discrimination. After publishing its list price, a firm may pay to target consumers who have a specific taste profile (vector of valuations for all firms' products). The per-consumer targeting cost reflects the expense of identifying an individual with the desired preferences and formulating and delivering a personalized offer to her; one could think of it as the cost of in-house research or an (exogenous) fee paid to a data broker. Permitting firms to target an exact preference type is a useful stylization—predicting an individual's willingness to pay will be achieved with increasingly high precision as the volume of consumer data grows and as computing power and forecasting algorithms improve. While it is generally acknowledged that firms have been reluctant to deploy personalized pricing too transparently, for fear of consumer backlash, personalized *discounts* can be both easier to camouflage and more palatable to consumers.[2] In the personalized discounting framework of our model, each targeted consumer becomes an individual market (Prat and Valletti, 2021), but these individual-level markets are linked together by the list prices that consumers pay when a discount is not offered.

We use the model to evaluate a policy intended to protect consumer privacy. Privacy is eroded when firms exploit individualized data, and consumers dislike this 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, fears that have been validated by frequent consumer data breaches (White, Zahay, Thorbjørnsen, and Shavitt, 2008; Turow, King, Hoofnagle, Bleakley, and Hennessy, 2009; Tucker, 2015; Acquisti, Taylor and Wagman *et al.*, 2016).[3] In response to these concerns, the General Data Protection Regulation, enacted by the EU in 2018, has established consumers' rights over their own data, including the requirement that firms obtain opt-in consent for data tracking. While US regulators have been slower to act, five states have passed privacy regulations and the US Council of Economic Advisers (2015) has outlined precursors for a national policy. In the spirit of the GDPR, our evaluation studies an opt-in policy under which consumers rationally trade off the market benefits of permitting their data to be used against the costs of lost privacy. Compared to laissez-faire targeting, the impact of the policy hinges on an externality among consumers: namely, whether

2. Bourreau and de Streel (2018) argue that this caution may reflect the scathing criticism Amazon endured for a personalized pricing incident in 2000, but they note that "there are subtler—and more acceptable, from a consumer viewpoint—ways for a company to achieve the same outcome. First, firms can offer the same uniform prices to all consumers, but with *personalised discounts*." OECD-UK (2018) makes similar arguments.

3. IdentityForce, a branch of TransUnion, catalogues over 40 incidents compromising more than five billion consumer accounts in 2021 alone (https://www.identityforce.com/blog/2021-data-breaches).

an individual who opts in tends to stimulate fiercer competition and lower list prices for other consumers or the opposite.

We now outline the model. Our first objective is to evaluate who benefits from a laissez-faire regime with unrestricted targeting, compared to a counterfactual with no targeted discounting. In the laissez-faire case, a firm's list price determines the margin between its *captive* and *contested* customers. Captive consumers pay list prices for their favourite products, as they do not like other products well enough to be at risk of being poached by rival firms' discounts. In contrast, a contested consumer delivers a Bertrand-like profit to her favourite firm (equal to her valuation difference between its product and her second-favourite). These Bertrand-like expected profits emerge even though costly targeting implies that discounting strategies must be mixed.

The winners and losers from laissez-faire targeting depend on the interplay between the discounts that contested consumers get and the list prices that captive consumers pay. One novelty of our framework is that, despite the oligopoly setting, a firm's list price decision involves a quasi-monopoly trade-off between marginal and inframarginal customers. The twist is that when a firm hikes its price, it does not lose its marginal captive consumer entirely, but its profit on her drops to the Bertrand level because she becomes contestable.

A firm's captive-contested margin hinges on what we call its *captive demand function* (which may be derived for any primitive distribution of tastes): the measure of consumers who prefer its product by at least y dollars over their next best option. The appeal of our approach is that the main conclusions depend only on the curvature of this captive demand function; in Section 5, we discuss why convex captive demand is the empirically plausible case.[4] This case leads to our headline policy conclusion (Proposition 2): laissez-faire targeted discounting reduces profits, and it also reduces consumer surplus if targeting costs are not too small. Consumers are hurt because the prospect of targeted discounting softens list price competition when captive demand is convex, and higher list prices swamp the benefit from discounts.

In our application to consumer privacy, consumers suffer an explicit foregone-privacy cost from opting into targeting; they weigh this cost against the expected savings from discounts. Under plausible demand conditions similar to those above, *every* consumer benefits from an opt-in policy (compared to unrestricted targeting), regardless of how strong her taste for privacy is. Consumers who choose to opt out benefit from preserved privacy, and by opting out they encourage stronger competition in list prices. This creates a positive spillover for all consumers because average discount prices are anchored to list prices.

Our article relates to the classical literature on informative targeted advertising and competitive price discrimination. In seminal papers (including Butters, 1977; Grossman and Shapiro, 1984; 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 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). Models based on Varian (1980) Model of Sales (with consumers exogenously segmented into captive "loyals" and price-elastic "shoppers") include Iyer, Soberman, and Villas-Boas (2005) (where targeting saves firms from wasted advertising), Chen, Narasimhan, and Zhang (2001) (where errors in targeting help to soften price competition), and Esteves and Resende (2016) (who break the loyal/shopper dichotomy with consumers who

4. Like other oligopoly models, we also impose logconcavity to ensure well-behaved profit functions.

prefer one product but would switch for a sufficiently lower price).[5] Several of these papers find that targeting may be profit-enhancing for some model parameters; our concluding remarks help 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 from Hoover (1937) through to Lederer and Hurter (1986) and Thisse and Vives (1988), focuses on spatial competition.[6] 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 article (but without costly advertising), and they reach some similar conclusions (including that competitive price discrimination hurts profits).

Our two stages of price-setting are most similar to prior work on couponing, including Shaffer and Zhang (1995, 2002) and Bester and Petrakis (1995, 1996). Bester and Petrakis (1996) share our structure of public list prices and costly discount ads but assume coarse targeting (two market segments) and no retention advertising. They find that the option to send coupons reduces list prices and profits; this is driven partly by an assumption that firms cannot discount to their "home" segments, so retaining those consumers requires a more competitive list price.

Personalized pricing (or first-degree price discrimination) is an old concept given new relevance by advances in targeting technology. Taylor and Wagman (2014) tabulate comparisons of profits and consumer surplus under uniform or personalized pricing for several common demand models. Anderson, Baik, and Larson (2015) study competition for an individual consumer when price offers are costly (with an emphasis on equilibrium selection). Using arguments similar to some of those in Section 3, they find that equilibria require mixing, a common theme in other settings with winner-take-all competition and participation costs.[7] However, their scope is limited to a single consumer and a single round of price offers. In contrast, we study a market with many consumers who all have the option to buy at list prices—this option creates a strategic linkage that ties firms' "macroscopic" competition over the entire market to their "microscopic" discount competition over individual consumers. The self-contained presentation of the personalized pricing subgame, with clean reduced-form results for profits and consumer surplus, makes it accessible for "plug-and-play" use in other applications of two-stage competition.

Belleflamme and Vergote (2016) and Chen, Choe, and Matsushima (2018) are closest to our opt-in analysis because they permit customers to hide from profiling. The former show (for monopoly) that tracking lowers consumer surplus because firms can price discriminate but hiding technology worsens consumer surplus further because the firm raises regular prices to discourage hiding. In a Hotelling setup, Chen *et al.* (2018) have firms personalize offers to targeted consumers and set uniform "poaching" prices to non-targeted ones. Hiding consumers soften competition by making it harder to poach. Both papers suggest, counterintuitively, that privacy regulation empowering consumers may make them worse off. While this is possible in our model, consumers will typically be better off with opt-in choice for empirically plausible demands.

The broader literature on data privacy is well surveyed by Acquisti *et al.* (2016). One recent stream analyses the "social" data externalities of sharing data when consumers' preferences are correlated (see Choi, Jeon, and Kim, 2019; Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2022; Ichihashi, 2021; Bergemann, Bonatti, and Gan, 2022). Another direction studies the

5. See also Brahim, Lahmandi-Ayed, and Laussel (2011). Esteban *et al.* (2001) develop a different notion of targeting precision (under monopoly) based on nested subsets of consumers.

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

7. See e.g. Hillman and Riley (1989), Sharkey and Sibley (1993), Narasimhan (1988), and Koçaş and Kiyak (2006). Myatt and Ronayne (2019) suggest a scheme to purify prices in such settings.

activities of data brokers (see Bergemann and Bonatti, 2019, and the references therein), while more nuanced and complex data revelation and disclosure strategies are addressed by Ishihashi (2020), Montes, Sand-Zantman, and Valletti (2019), Ali, Lewis and Vasserman (2019), and Hidir and Vellodi (2020).

At times, we point the reader toward additional material in the discussion paper version of this article (Anderson *et al.*, 2019), henceforth "the DP". 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 $G_j(y)$, $y \in \left[\underline{y}, \bar{y}\right]$ for each firm, where $1-G_j(y)$ is the fraction of consumers who prefer product j over their best alternative product (among the $n-1$ other firms) by at least y dollars. (We permit the possibility of $\bar{y}=\infty$, $\underline{y}=-\infty$.) Formally, if $\hat{r}_{i,-j}=\max_{j' \in \{1,\ldots,n\} \setminus j} r_{ij'}$, then

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

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.*[8]

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

There are two stages of competition.[9] 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 $\left(r_{i1}, r_{i2}, \ldots, r_{ij}, \ldots\right)$ 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. We assume that paying A enables the firm to target the consumer perfectly; imperfect targeting is discussed in the Conclusion.[10]

8. We 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.

9. An alternative timing assumption would be that all prices (list and discount) are set simultaneously. As we show in the DP, this case has no pure-strategy equilibrium in list prices, making a comparison to our results difficult. One conjecture is that average list prices would be higher if list and discount prices were set simultaneously. In our model, there is an incentive to commit to low list prices (to discourage rivals from stealing customers with undercutting discounts at the second stage); with simultaneous list and discount price-setting, that incentive would be absent.

10. The technological advances making personalization possible and anecdotal evidence of its use are well surveyed in the business and popular press (see e.g. "White Paper Digital Transformation of Industries: Media

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

We assume that consumers' outside options are sufficiently low that they always purchase some product, that is, the market is fully covered. While this assumption is commonly imposed in the literature, it has a bit more bite here because equilibrium list prices may rise as the ad cost A falls. We discuss the implications of allowing outside options to bind in the conclusion. We say that consumer i is *on the turf* of Firm j if it makes her favourite 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 favourite 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.[11] This setup generalizes easily to the case of n firms located on a circle.

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

There are n firms, and consumer i's taste r_{ij} for Firm j's product is drawn i.i.d. from the primitive distribution $F(r)$ with support $[\underline{r}, \bar{r}]$.[12] 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 (i.e. $r_{i1} \geq r+y$) with probability $1-F(r+y)$. But the consumer's best draw over $n-1$ alternatives has distribution $F_{(1:n-1)}(r) = F(r)^{n-1}$, so we have:

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

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; "Shopper Alert: Price May Drop for You Alone," New York Times, 9 August 2012; also, "A Special Price Just for You," Forbes, 17 November 2017.

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

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

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.[13] 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.

Given the symmetric setup, we focus on symmetric equilibria in which all firms set the same list price p^l.[14] We begin with the targeted advertising sub-game.

3. STAGE 2: COMPETITION IN TARGETED DISCOUNTS

In order to identify the incentive to deviate from a symmetric list price in Stage 1, suppose that all firms besides Firm 1 have set the same list price p. (The extension to arbitrary list prices, as well as proofs for this section, are in the Appendix.) We will focus on Firm 1's profit from targeted discounting to a consumer with tastes satisfying $r_2>r_3>...>r_n$.[15] Then, Firm 1's value advantage is $y_1=r_1-r_2$, the consumer is on Firm 1's turf if $y_1>0$, and Firm 2 makes the most attractive rival product. In Stage 2 competition for this consumer, each firm j chooses a probability a_j of sending her an ad (at cost A) and, if an ad is sent, a distribution over the discount price p_j^d offered. The consumer is said to be *contested* if at least two firms advertise to her with positive probability, or *conceded* if only one firm does; otherwise she is *captive* to her default firm.

As a leading case, suppose that both p and p_1^l exceed the ad cost; this means that paying A to send a targeted discount is potentially profitable for any firm. Proposition 5 in the Appendix shows the following results about equilibrium discount competition for the consumer described above: (i) if $y_1>p_1^l-A$, she is captive to Firm 1, who earns its list price on her; (ii) if $y_1\in\left(0,p_1^l-A\right)$, she is contested by her two most-favoured firms, with expected profit $\pi_1=y_1$ for Firm 1 and $\pi_2=0$

13. That is, if the taste distribution is $F(r)=\exp\left(-e^{-r/\beta}\right)$, 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).

14. Under duopoly there are no asymmetric equilibria. This may be true for $n>2$ as well, but we have not proved it.

15. For smooth taste distributions, consumers who are indifferent between two or more products have zero-measure, and have no impact on profits or list price decisions, so we can ignore them. Relabeling firms so that Firm 2 is the closest rival for the consumer is a matter of convenience.

for Firm 2; (iii) if $y_1 < 0$, Firm 1 earns zero profit on this consumer.[16] While we focus on Firm 1's profits, similar logic applies to any other firm.

Let us trace some of the logic. In case (i), the consumer's preference for Firm 1 is strong enough that the closest competitor would need to advertise an unprofitably low discount $p_2^d < A$ to attract her away from Firm 1's list price (which it will not do). In Case (ii), Bertrand-like profits ensue, even though the equilibrium discounting strategies are mixed. Discount competition will drive other firms' profits on this consumer to zero (Lemmas 2 and 3), implying that Firm 1's lowest advertised price will be $p_1^d = A + y_1$, leaving no room for its closest rival to profitably undercut. As this offer leaves Firm 1 with net profit y_1 on the consumer, any other discount offers it mixes over must do equally well. In Case (iii), some other firm has the value advantage over the consumer, and Firm 1's profit is driven to zero by competition.

On the other hand, if all list prices are equal to or smaller than A, then no firm will pay to send a discount ad, and Firm 1 will earn its usual oligopoly profit: it sells at its list price to only those consumers $y_1 > p_1^l - p$ whose relative preference for Firm 1 exceeds any list price difference. Appendix covers discount competition when the Stage 1 list prices permit Firm 1 to advertise but not other firms ($p_1^l > A \geq p$), or *vice versa*. These cases provide a firm's off-the-path profits, which we note here and will use in Section 4. Firm 1 earns its list price on captive consumers $y_1 > p_1^L - P_{-1}$, where $P_{-1} = \min(p, A)$. If $p_1 \leq A$, these are its only customers; otherwise it earns $\pi_1 = y_1 + P_{-1} - A$ on consumers $y_1 \in \left(A - P_{-1}, p_1^l - P_{-1}\right)$ and zero on everyone else. If $p > A$, then rival firms can potentially send discount offers as low as A, in which case Firm 1's profits collapse to the case discussed above; otherwise rival firms cannot afford to advertise discounts, and Firm 1's profits are constrained by their list price p. For this reason, we refer to P_{-1} as Firm 1's most competitive rival price.

3.1. *Competition for a contested consumer: equilibrium discounting strategies*

Consumers' gains from discounting will depend on the equilibrium mixed strategies that underpin the profits discussed above. We discuss those strategies below, restricting attention to when all firms have set the same list price $p > A$. As noted in point (ii) above, a consumer with value advantage $y_1 \in (0, p - A)$ will be contested by her favourite and second-favourite firms only, Firms 1 and 2. Because their competition drives Firm 2's profit to zero, no less-preferred firm could break even if it were to target this consumer (Lemma 4). As a consequence of the positive ad cost in combination with Bertrand undercutting incentives, Firm 1 and 2's targeting strategies must be mixed. We write $B_1(s)$ and $B_2(s)$ for the firms' mixed strategy distributions over discount *surplus* offers, where the surplus offered to the consumer is related to the discount price by $s_1 = r_1 - p_1^d$ and $s_2 = r_2 - p_2^d$. We also make the convention that "not advertising" may be regarded as a surplus offer $s_1^l = r_1 - p$ or $s_2^l = r_2 - p$ at a firm's list price, so the probabilities of sending a targeted ad are $a_1 = 1 - B_1\left(s_1^l\right)$ and $a_2 = 1 - B_2\left(s_2^l\right)$ respectively. This consumer will buy at Firm 1's list price and enjoy surplus s_1^l if she receives no discount, so "advertised" surplus offers will need to offer her an improvement. Thus advertised offers satisfy $s_1 > s_1^l$ for Firm 1 and $s_2 \geq s_1^l$ for Firm 2 (recalling the assumption that ties go to the discount offer). Proposition 6 in the Appendix derives the following equilibrium strategies for Firms 1 and 2 with respect to this contested consumer.

16. The boundary cases $y_1 = 0$ and $y_1 = p_1^l - A$ are omitted for smoother exposition; as they are zero-measure, they do not affect the Stage 1 profits.

Firm 1 Firm 1 sends no ad with probability $1-a_1 = B_1\left(s_1^l\right) = \frac{A}{p-y_1}$. Its advertised offers are distributed $B_1(s) = \frac{A}{r_2-s}$ over support $(s_1^l, r_2-A]$. The corresponding discount prices p_1^d have support on $[A+y, p)$.

Firm 2 Firm 2 sends no ad with probability $1-a_2 = B_2\left(s_2^l\right) = \frac{y_1}{p}$. Otherwise, its advertised offers are distributed $B_2(s) = \frac{A+y_1}{r_2+y_1-s}$ over support $\left[s_1^l, r_2-A\right]$; this includes an atom $\frac{A}{p}$ of advertised offers at Firm 1's list price surplus. The corresponding discount prices p_2^d have support on $[A, p-y]$.

These distributions are dictated by indifference conditions and the firms' equilibrium profits. In particular, the atom of offers undercutting Firm 1's list price is just large enough to provoke a response—if it were smaller, Firm 1 would not find it worthwhile to pay A to advertise small discounts. If ad costs vanish ($A \to 0$), Firm 1's price collapses to the pure strategy that is conventionally assumed for the stronger firm in asymmetric Bertrand competition: it advertises the highest discount price $p_1^d = y_1$ that its rival cannot undercut, corresponding to a surplus offer r_2. Firm 2's strategy remains mixed in this limit: $B_2(s) = \frac{y_1}{r_2+y_1-s}$. While its discount offers never win the consumer, they exert just enough competitive discipline to restrain Firm 1 from pricing higher.

Because a consumer takes the best surplus she is offered, a contested consumer's equilibrium surplus is a draw from the distribution $B_1(s)B_2(s)$. We will calculate her expected consumer surplus and use it to evaluate policies in Section 5.

4. STAGE 1: COMPETITION IN LIST PRICES

4.1. *No-targeting benchmark*

If targeted advertising is impossible or banned, the model collapses to standard differentiated-product price competition, and there is a symmetric equilibrium at common list price p^{NT} characterized by the first-order condition:[17]

$$p^{NT} = \frac{1-G(0)}{g(0)}. \tag{2}$$

This remains the model's unique symmetric equilibrium outcome if targeted ads are available but prohibitively expensive: $A \geq p^{NT}$.

4.2. *Equilibrium list prices when targeting is in use*

From now on, we focus on the case $A < p^{NT}$ where targeting will be used in equilibrium. With an eye toward symmetric equilibrium conditions, we begin with Firm 1's overall profit at list price p_1^l when Firms 2 through n are expected to price at p^l. Using the results of the previous section, that profit may be written:

$$\Pi_1\left(p_1^l, p^l\right) = \begin{cases} p_1^l\left(1-G\left(p_1^l - P_{-1}\right)\right) & \text{if } p_1^l \leq A; \\ p_1^l\left(1-G\left(p_1^l - P_{-1}\right)\right) + \int_{A-P_{-1}}^{p_1^l - P_{-1}} \left(y + P_{-1} - A\right) dG(y) & \text{if } p_1^l > A. \end{cases} \tag{3}$$

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

17. Condition 1 ensures quasiconcave profit functions, so (2) is sufficient as well as necessary.

interior optimum, as:

$$\frac{\partial \Pi_1(p_1^l)}{\partial p_1^l} = 1 - G\left(p_1^l - P_{-1}\right) - P_1 g\left(p_1^l - P_{-1}\right) = 0. \tag{4}$$

This resembles the usual marginal-inframarginal tradeoff one would see in an oligopoly first-order condition. However, the marginal consumer, $y = p_1^l - P_{-1}$ is determined by the most competitive price a rival could offer, which could be as low as an advertised discount price of A. Furthermore, if Firm 1 can discount, it does not lose this marginal consumer entirely when it hikes its list price. It sacrifices only $P_1 = A$, the cost of winning this consumer back with an infinitessimal discount.[18]

A symmetric equilibrium must satisfy (4) at $p_1^l = p^l$. The common list price must exceed A, so $P_1 = P_{-1} = A$, and the necessary condition for equilibrium simplifies to:[19]

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

So long as $A > m = \lim_{y \to \bar{y}} \frac{1-G(y)}{g(y)}$, equation (5) has a unique solution which we denote p^T. In Appendix, we show that (2) and (5) fully characterize the symmetric equilibria of the model.

Existence and uniqueness of symmetric equilibria Under Conditions 1 and 2, the model has a unique symmetric equilibrium. If $A \geq p^{NT}$, the list price is p^{NT} and targeting is not employed. If $A \in \left(m, p^{NT}\right)$, the list price is p^T. If $A < m$, the list price is $p^l = \bar{y} + A$.

5. THE IMPACTS OF TARGETING

The question of who gains or loses from targeted discounting is closely tied to the impact of discounting on list prices. We will show that list prices, in turn, are linked to the curvature of demand. In what we argue is the more compelling case of convex captive demand, targeting pushes list prices up, eroding some of the consumer benefits of discounting.

5.1. *List prices*

We say that captive demand is convex or concave if $1 - G(y)$ is convex or concave over the range of consumer types $y > 0$ who favour a firm's product. Lemma 1 showed that captive demand derived from independent taste shocks will be strictly convex, so this is the relevant case for commonly used empirical specifications like multinomial logit demand. Furthermore, a convex demand function has a decreasing density, $g'(y) < 0$, so 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 reasonable, except where there are persuasive arguments for strongly polarized tastes. For these reasons, we will usually treat convex demand as a leading case in the rest of the article.

18. In equilibrium, Firm 2 advertises to these marginal consumers just often enough ($a_2 = A/p_1^l$) that Firm 1 is indifferent about advertising to retain them. Not advertising means losing its full list price on a fraction a_2 of them, thus a total expected loss of $a_2 p_1^l = A$, matching what it would lose by advertising an infinitesimal discount. We thank a referee for suggesting this clarification.

19. Having assumed $A < p^{NT}$, an equilibrium at $p^l \leq A$ is impossible because the equilibrium condition $\partial \Pi_1 \left(p_1^l\right) / \partial p_1^l \big|_{p_1^l = p^l} = 0$ would reduce to (2).

Equation (5) implicitly identifies a relationship $p^l(A)$ between the equilibrium list price and the targeting cost. Suppose that $A \in (m, p^{NT})$ so that targeting is affordable and there is an interior symmetric equilibrium with list price $p^T = p^l(A)$.

Proposition 1. *If captive demand is strictly convex, then $p^T > p^{NT}$, so list prices are higher when targeting is in use than they would be if it were banned. If captive demand is strictly concave, this reverses: $p^T < p^{NT}$.*

Given Lemma 1, it follows immediately that targeted discounting pushes up list prices (relative to no targeting) for any independent taste shock, multinomial choice demand system of the sort described in Example 2. The proof uses the related result that $p^l(A)$ is decreasing if demand is convex: more costly targeting leads to lower list prices. The ranking follows because the list price tends to p^{NT} as targeting becomes too costly to use ($A \to p^{NT}$). If demand is concave, then $p^l(A)$ is decreasing, and the argument reverses.

To trace out why demand curvature plays this critical role, consider the effect of the ad cost on a firm's marginal profit (4), written $M_1 = \partial \Pi_1 / \partial p_1^l$ here for brevity. When ads are in use, an increase in A affects marginal profit through two channels: $\partial M_1/\partial A = \partial M_1/\partial P_1 + \partial M_1/\partial P_{-1}$. The first term, $\partial M_1/\partial P_1 = -g(p_1^l - P_{-1})$, encourages Firm 1 to cut its list price so as to keep marginal consumers captive (rather than pay the higher cost of advertising to them). But a higher targeting cost also tends to put those marginal consumers out of the range of other firms' discounts; this has a positive effect $\partial M_1/\partial P_{-1} = g(p_1^l - P_{-1}) + P_1 g'(p_1^l - P_{-1})$ on Firm 1's marginal profit and encourages setting a higher list price.[20] Convex demand ($g' < 0$) works counter to this competition-softening effect, allowing the first effect to dominate. Loosely, this is because the decline in $p_1^l - P_{-1}$ pushes the margin into a region of higher consumer density, and hence fiercer price competition.

5.1.1. List price neutrality under Hotelling competition. The competing effects discussed above will cancel each other out if $g' = 0$. To illustrate the implications, suppose transportation costs are linear-quadratic in the Hotelling model of Example 1: $T(d) = \alpha d + \beta d^2$, with $\alpha + \beta = t$. Then captive demand is linear $1 - G(y) = \frac{1}{2} - \frac{y}{2t}$, and so we have *list price neutrality*: $p^{NT} = p^l(A) = t$ for all ad costs $A \leq t$. It turns out that linear and quadratic costs are both knife-edge cases; for the family of transport costs $T(d) = d^\gamma$, one can confirm that captive demand is strictly convex for $\gamma \in (0,1)$ or $\gamma > 2$, but strictly concave for $\gamma \in (1,2)$.[21]

When list price neutrality obtains in the Hotelling model, it is because large taste differences are exactly as common as smaller ones ($g'(y) = 0$), and this is possible because tastes for the two products are negatively correlated. In contrast, taking the difference of independent draws in the i.i.d. case has a centralizing effect that implies higher densities of consumers at smaller taste differences.

While the role of the number of firms is not a main focus of the article, we also note that under standard oligopoly competition the equilibrium price p^{NT} falls with n for the independent taste shock 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

20. This term is unambiguously positive because it equals $-\Pi_1''(p_1^l)$.

21. More generally, captive demand has the same curvature on $y \geq 0$ as the difference in transportation costs $T(1-x) - T(x)$ does on $x \in [0, \frac{1}{2}]$; this cannot be reduced (at least, not in a trivial way) to a condition on $T(d)$ itself.

lower list price and their second best option is stochastically better with more choice.[22] These pro-competitive results might help allay misgivings about the mixed strategies in our model since Varian (1980) model of sales has been criticized for its property that prices rise with more competition.

5.2. *Impact of targeting on profits and consumer surplus*

We now examine the impact on firms and consumers when improvements in data gathering and analysis make targeted discounting viable. The benchmark is a standard oligopoly equilibrium with no targeting and common list price p^{NT}. We compare this to a scenario where targeting costs have fallen to $A < p^{NT}$, and there is a new equilibrium with targeted discounting at common list price $p^T = p^l(A)$. Unless otherwise stated, we continue to assume that Condition 1 holds (strict logconcavity of captive demand). Let Π^{NT} and Π^T be a firm's profit in the two scenarios, with CS^{NT} and CS^T the respective aggregate consumer surpluses.

Our result for consumer surplus relies on a measure of how convex demand is: 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. The ρ-convexity of many commonly used demand systems can be readily verified. For example, in the independent taste shock formulation of Example 2, captive demand is $\frac{1}{n}$-convex if taste shocks are uniform, or $\frac{1}{n-1}$-convex if they are Type 1 extreme value (the multinomial logit case), where n is the number of firms.[23]

Proposition 2. *Targeting reduces profits: $\Pi^T < \Pi^{NT}$. If captive demand is ρ-convex and $A/p^{NT} > \rho$, then targeting also reduces consumer surplus: $CS^T < CS^{NT}$. However, for A sufficiently close to zero, targeting improves consumer surplus: $CS^T > CS^{NT}$.*

Thus, if demand is sufficiently convex and ad costs are non-negligible, the introduction of targeted discounting hurts both sides of the market. We flesh out the logic behind these results below.

5.2.1. Profits.

In the no-targeting benchmark, each firm serves the $1-G(0)$ fraction of consumers who are on its turf and earns profit $\Pi^{NT} = p^{NT}(1-G(0))$. Meanwhile, from (3) we have the following equilibrium profit for a firm in the scenario with targeting:

$$\Pi^T = p^T\left(1-G\left(p^T - A\right)\right) + \int_0^{p^T - A} y\, dG(y). \tag{6}$$

A firm earns positive profits only on the consumers who like its product best; those with the strongest preference ($y > p^T - A$) pay the list price, and the firm earns its value advantage on the rest. If $A < m$, the equilibrium has all consumers contested with targeted discounts; in this case, the first term vanishes, and we have $\Pi^T = \int_0^{\bar{y}} y\, dG(y)$.

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

23. These examples suggest the plausible but unproven conjecture that captive demand is generally more convex when there are more firms in the market.

For concave demand, the profit ranking is not surprising, since targeting implies lower list prices plus additional discounting. However, if demand is convex, then firms enjoy higher margins on their list price sales when they can target. Proposition 2 implies that these gains must be overshadowed by the loss in profit when consumers who would have otherwise paid p^{NT} become contested. To demonstrate the profit ranking, write the equilibrium profit under targeting as a function of the ad cost: $\Pi^T\left(p^l(A),A\right)$. As targeting becomes uneconomic, $A \to p^{NT}$, these equilibrium targeting profits tend toward Π^{NT}; that is, $\Pi^T\left(p^l(A),A\right)\big|_{A=p^{NT}} = \Pi^{NT}$. (This is clear from inspection of (5) and (2).) Then the claim that $\Pi^T < \Pi^{NT}$ follows because $\Pi^T\left(p^l(A),A\right)$ is strictly increasing in A: $d\Pi^T\left(p^l(A),A\right)/dA = Ag\left(p^T - A\right)$.[24]

For intuition about why profits are increasing in the ad cost, we turn back to the profit expression (3) where the effects of Firm 1's own ad cost, and its rivals' ad cost (written as P_{-1}), can be distinguished. When all firms face higher targeting costs, the net profit $y+P_{-1}-A$ on a contested consumer does not change. The only remaining effect boosts Firm 1's profits: consumers at the $p_1^l - P_{-1}$ margin shift from contested to captive, since Firm 1's rivals can no longer afford to target them.

If Condition 1 is violated, it is possible for firms to benefit from targeted discounting; the DP provides examples and a general result. Without logconcave demand, the average consumer preference for her favourite firm, $E(y \mid y \geq 0)$, may exceed the usual oligopoly price p^{NT}, assuming the latter exists.[25] In this case, firms may be better off setting very high (and irrelevant) list prices, and selling to all consumers through personalized price offers.

5.2.2. Consumer surplus.

It is straightforward to see that consumers benefit from targeted discounting if captive demand is concave, as they face both lower list prices and the possibility of a discount. In the convex demand case that is our main focus, targeting benefits the consumers getting the steepest discounts but hurts those who pay list prices. To take the balance of these two effects, we must investigate how large discounts are on average. We shall see that the effect of higher list prices dominates when demand is relatively convex and ad costs are non-negligible, but this reverses as $A \to 0$.

Consider an equilibrium at common list price p^T and a contested consumer type $y_1 \geq 0$ whose most-preferred products are at Firms 1 and 2. As shown in Section 3, her best offer at each firm can be represented as a surplus draw $s_1 \sim B_1(s)$ or $s_2 \sim B_2(s)$, and she takes the better of these two offers. To focus on how much she gains relative to the surplus $s_1^l = r_1 - p^T$ from purchasing at Firm 1's list price, we introduce the "surplus improvement" variables $\tilde{s}_1 = s_1 - s_1^l$ and $\tilde{s}_2 = s_2 - s_1^l$. Making the change of variables, these surplus improvement offers are distributed according to $\tilde{B}_1(\tilde{s}) = \frac{A}{p^T - y_1 - \tilde{s}}$ and $\tilde{B}_2(\tilde{s}) = \frac{A+y_1}{p^T - \tilde{s}}$, with common support on $[0, p - y_1 - A]$. We define this consumer's expected discount $\Delta\left(y_1, p^T\right)$ to be her expected surplus improvement relative to a list price purchase at her most-preferred firm; thus $\Delta\left(y_1, p^T\right) \equiv E_{\tilde{B}_{\max}}(\tilde{s})$, where $\tilde{B}_{\max}(\tilde{s}) = \tilde{B}_1(\tilde{s})\tilde{B}_2(\tilde{s})$ is the distribution of her best improvement offer. After computing this expectation, we have:[26]

$$\Delta\left(y_1, p^T\right) = p^T - y_1 - L\left(y_1, p^T, A\right), \text{ where} \tag{7}$$

24. Because the equilibrium price $p^l(A)$ is defined by a first-order condition, a version of the Envelope Theorem applies, and we have $d\Pi^T/dA = \partial\Pi^T/\partial A$.

25. Without logconcave demand, a well-behaved oligopoly equilibrium is not guaranteed.

26. After integrating by parts to get $\Delta(y_1, p) = \int_0^{p-y_1-A}\left(1 - \tilde{B}_{\max}(\tilde{s})\right)d\tilde{s}$, this is a straightforward computation.

$$L(y,p,A) = A\left(1 + \frac{(A+y)}{y}\ln\left(\frac{A+y}{A}\frac{p-y}{p}\right)\right). \tag{8}$$

Recall that Firm 1 earns $p-y_1$ less on this consumer than it would if she were captive, matching the first term in the discount. The consumer does not capture this full profit reduction because she bears the cost $L(y_1, p^T, A)$ of the inefficiencies that targeting introduces; these include expected ad costs and the fact that she sometimes ends up with her second-best product. It can be shown that $\Delta_y < 0$, so as one would expect, the consumers who are most flexible about which product to buy get the largest discounts. Furthermore, $\Delta_p \in [0,1)$, so higher list prices imply larger—but not commensurately larger—discounts.

Then in an equilibrium with targeting, a consumer with value advantage y enjoys total surplus equivalent to paying $EP(y)$ for her favourite product, where $EP(y) = p^T$ if she is captive or $EP(y) = p^T - \Delta(y, p^T)$ if she is contested. Because the same consumer pays p^{NT} for her favourite product if targeting is banned, she is worse off with targeting if $EP(y) > p^{NT}$. We refer to $EP(y)$ as her expected "favourite-equivalent" price, since it accounts for the fact that a discount price p_2^d at her second-best firm is surplus-equivalent to a price $p_1^d = p_2^d + y$ at her favourite firm. Given the symmetry of demand across firms, aggregate consumer surplus is lower with targeting than without it if the average favourite-equivalent price $E_{y \geq 0}(EP(y))$ exceeds p^{NT}.

The consumer surplus ranking in Proposition 2 reflects two ways in which the targeting equilibrium looks relatively worse for consumers when captive demand is more convex. First, targeting inflates list prices more (that is, the gap $p^T - p^{NT}$ is larger) the more convex demand is—this makes the targeting case even less attractive to captive consumers. Second, the average value advantage, $E(y \mid y \geq 0)$, grows larger when demand is more convex—there are more consumers with relatively strong preferences for their favourite product. Since these are the consumers with the least to gain from discounting, this effect reduces the total gains from targeting that accrue to contested consumers.

The sufficient condition $\rho < A/p^{NT}$ reflects the fact that targeting costs are passed through to consumers, so targeting can be proven to hurt consumers for a broader range of demand systems when A is larger. However, it is a stronger than necessary condition, as demonstrated by the multinomial logit captive demand from Example 2. For this demand, banning targeting would improve consumer surplus as long as the targeting equilibrium involves some captive consumers who pay list prices.[27]

Proposition 2 concerns aggregate consumer surplus, but under a mild convexity condition ($g'(0) < 0$) it can be shown that *every* consumer is made worse off by targeting when the ad cost is sufficiently high. Details are in the Supplementary Appendix, but the logic is that targeting induces a first-order increase in the list price but only a second-order increase in discounts when $p^{NT} - A$ is sufficiently small. Thus, even the most fiercely contested consumers will pay more on average.

In contrast, as $A \to 0$, there are no captive consumers to pay higher list prices (since everyone is targeted), and discounts are no longer diminished by the passed-through costs of targeting—consequently consumer surplus is higher when targeting is permitted than when it is forbidden. We show this formally in the DP, but a sketch of the logic is straightforward. As the ad cost A

27. Let $1-G(y) = \frac{1}{1+(n-1)e^{y/\beta}}$. 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 2 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 $\overline{EP}^T > \overline{EP}^{NT}$.

tends to zero, the discounting strategies tend toward classic asymmetric Bertrand competition: each consumer purchases her most-preferred product (with probability tending to one) at a price tending to y_1—the highest price that her second-best firm cannot profitably undercut. (Section 3 gives further details.) This limiting case concurs with the standard price discrimination model in which ad costs are absent, so our results cover this model too. Consequently, the equilibrium of the targeting model in the $A=0$ limit is efficient: every consumer buys her first-best product, and the ad costs vanish. This matches the total welfare in the standard oligopoly equilibrium with no targeting, which is also efficient. But then, since we know that profits are lower under targeting, $\Pi^T_{A=0} < \Pi^{NT}$, it must be that consumer surplus is higher: $CS^T_{A=0} > CS^{NT}$. By continuity, the same ranking will hold for A sufficiently small.

6. OPT-IN AND CONSUMER PRIVACY

In reaction to concerns about the ubiquity of consumer data, its use in targeting, and the loss of privacy this entails, regulators have begun to consider policies to protect consumers. Most prominently, the General Data Protection Regulation (GDPR), which took force in the European Union in 2018, gives individuals the right to consent (or not) to the processing of their personal data. Inspired by the GDPR, we use our model to evaluate a policy under which a consumer cannot be targeted unless she opts in. We assume rational consumers with a taste for privacy that they trade off against the expected benefits from targeted discounts. This opt-in policy will be compared to the benchmarks of unrestricted, laissez-faire targeting and an outright ban on targeting (the T and NT cases from earlier). We prioritize the unrestricted targeting benchmark because consumer data use has already become widespread and there would be practical challenges with implementing a ban.[28]

We now assume each consumer chooses whether to opt into or out of data collection. A consumer who opts in suffers a lost-privacy cost $c \geq 0$ and can be targeted by any firm with a personalized discount. If she opts out, she suffers no privacy cost, cannot be targeted, and therefore will purchase at her best list price offer. Privacy costs are distributed across consumers according to cdf $H(c)$, independently of preferences over products.

A consumer opts in or out at the same time that firms set list prices, and before learning her preferences over products. This assumption reflects the idea that most people do not have a specific product or market in mind when they make decisions about privacy; rather, they have a more diffuse sense that their data could be used for or against them in some yet-to-be-determined future purchases. In the model, a consumer will weigh her privacy cost against the average targeted discount over all "locations" y. One interpretation is that she does not yet know which market her data will be used in; hence she does not know whether her relative preference for her top product will be strong or weak. An alternative interpretation is that she expects her data to be used in many different product markets, some where her y is small and others where it is large, and so she forms an expected benefit from discounts over all these markets.

Price competition among firms proceeds as described earlier, with two amendments. At Stage 2, only opt-ins may be sent a personalized discount, and firms set list prices in Stage 1 based on an expectation about the fraction of consumers λ who will opt in.

An equilibrium of the model with opt-in will require that (1) each firm sets a profit-maximizing list price with respect to correct beliefs about other firms' list prices and correct beliefs about λ and (2) consumers opt in if and only if the expected discount (based on correct

28. For example, it could be difficult to prevent an individual from sharing data with a firm in cases where it would be mutually beneficial.

beliefs about list prices) exceeds their privacy cost. As earlier, we focus on equilibria that are symmetric in list prices; in this case (1) can be summarized by a function $p(\lambda)$ that gives the equilibrium list price generated when fraction λ of consumers opt in. Meanwhile, (2) generates a correspondence $\lambda(p)$ identifying the fraction of consumers who opt in when a common list price p is expected. (The opt-in rate is single-valued except at any prices where a mass of consumers are indifferent.) Equilibria of the full model are intersections of these two curves. All claims in this section are proved in the Appendix.

Consumer opt-in decisions

If consumers anticipate symmetric list prices, an opt-in will enjoy expected discount $\Delta(y;p)$ if she turns out to be contestable by her top two firms ($y\in[0,p-A]$). If she turns out to be captive ($y>p-A$), she will get no discount; we extend the definition of $\Delta(y;p)$ to assign $\Delta(y;p)=0$ in this case. The *ex ante* expected discount is then $\bar{\Delta}(p)=E_G(\Delta(y;p))$, where the expectation is taken over locations y.

Consider price p and corresponding cost $c=\bar{\Delta}(p)$. Any consumer with cost $c'<c$ strictly prefers to opt in, whereas anyone with $c'>c$ will opt out. As long as there is not a mass of consumers at c, we simply have $\lambda(p)=H\left(\bar{\Delta}(p)\right)$. Consumers at $c'=c$ are indifferent at price p; if there is a mass of such consumers, we assign $\lambda(p)=\left[H_-,H\left(\bar{\Delta}(p)\right)\right]$, where $H_-=\lim_{p'\to p^-}H\left(\bar{\Delta}(p')\right)$. (An example where all consumers share the same privacy cost is illustrated in Figure 1 (b).) Because average discounts rise with the list price, the opt-in rate $\lambda(p)$ is increasing in p as well, strictly so if the privacy cost distribution has full support.

Price competition equilibrium among firms (at a given opt-in rate)

To distinguish it from an equilibrium of the full model, we say a *price competition equilibrium* (PCE) is a profile of list prices at which each firm maximizes its own profit, given the opt-in rate λ. Then a symmetric equilibrium of the full model is comprised of a symmetric PCE and an opt-in rate that are mutually consistent.

Without loss of generality, consider the marginal profit of Firm 1 when it expects all other firms to charge list price p:

$$\frac{d\Pi_1}{dp_1}=\lambda(1-G(p_1-A)-Ag(p_1-A))+(1-\lambda)(1-G(p_1-p)-p_1g(p_1-p)). \tag{9}$$

The expression is simply a weighted average of the targeting marginal profit on opt-ins and the no-targeting marginal profit on opt-outs. A symmetric PCE, if one exists, must therefore satisfy the equilibrium condition:

$$\Phi\left(p^*\right):=\lambda\left(1-G\left(p^*-A\right)-Ag\left(p^*-A\right)\right)+(1-\lambda)\left(1-G(0)-p^*g(0)\right)=0. \tag{10}$$

Of course this is just a weighted average of the equilibrium conditions for the cases of unrestricted targeting and no-targeting, respectively. In order for (10) to be not just necessary but sufficient for a PCE, we need demand mixtures of opt-in and opt-out consumers to be well-behaved. In the Appendix, we show the following. (The extra condition for concave demand ensures that list prices do not become irrelevant to opt-in consumers.[29])

29. The condition rules out boundary complications that arise because a concave captive demand function has finite support. It ensures that $p^*-A<\bar{y}$, implying there are some opt-in consumers who pay list prices because they turn out to be unprofitable to target.

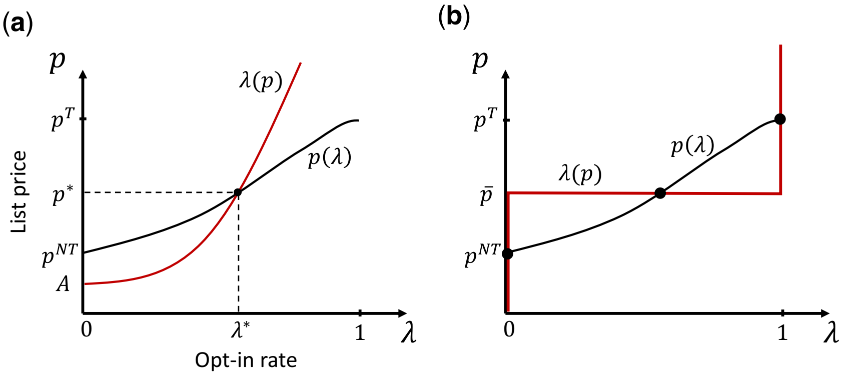

FIGURE 1

Examples of opt-in equilibrium when captive demand is convex.

There is a unique symmetric PCE, identified by $\Phi\left(p^*\right)=0$, *in either of these cases:*

Captive demand is convex, or

Captive demand is concave and $p^{NT} < \bar{y}+A$.

The unique symmetric PCE at each opt-in rate may be traced by a function $p(\lambda)$. We saw earlier that demand curvature dictates whether list prices are higher with unrestricted targeting or no targeting, and this logic extends to opt-in. If captive demand is strictly convex, then $p(\lambda)$ is strictly increasing, from $p(0)=p^{NT}$ up to $p(1)=p^T$. This reverses if demand is strictly concave: $p(0)=p^{NT}>p^T=p(1)$, and $p(\lambda)$ is strictly decreasing. Consumers impose spillovers on each other through the effect of their privacy choices on list prices, but the direction of that spillover depends on demand curvature: opting in hurts other consumers if demand is convex but helps other consumers if demand is concave.

6.1. *Equilibrium with consumer opt-in*

Because we regard convex captive demand as a more empirically plausible case, that will be where we focus most of our attention. All results in this section assume convex demand; Section 6.2 discusses which of these results extend or change if demand is concave. Figure 1 illustrates equilibria of the full model with opt-in. In Figure 1(a), privacy costs have full support on $[0,\infty)$. Consumers with privacy costs close to zero will opt in if there is any chance of a discount. This implies a vertical intercept $\lambda(A)=0$, as shown (since targeting becomes unprofitable for list prices below the ad cost). Meanwhile, the presence of consumers with arbitrarily large privacy costs ensures that even at high list prices the opt-in rate never reaches one. The figure depicts a unique equilibrium at $\left(\lambda^*,p^*\right)$. In Figure 1(b), all consumers have the same privacy cost $c>0$, so they shift in unison from opting out to in as list prices rise above a threshold price $\bar{p}$ (identified by $c=\bar{\Delta}(\bar{p})$). Here, there are three equilibria: full opt-out, full opt-in, and an interior equilibrium.

Because the empirical literature is only beginning to address the challenge of quantifying consumer tastes for privacy, we will focus on robust policy conclusions that are not sensitive to the details of the distribution $H(c)$. Notwithstanding the substantial differences between the panels of Figure 1, there are common threads to be found. A symmetric equilibrium always exists, and any symmetric equilibrium list price must lie in the interval between p^{NT} and p^T.[30] On this

30. Existence follows immediately from an application of the Kakutani fixed-point theorem to the mapping $\lambda(p(\lambda))$.

basis, we can evaluate the impact of imposing an opt-in requirement on *ex ante* consumer surplus (after learning one's privacy cost, but before learning y) and firm profits.[31]

Proposition 3. *Suppose some consumers opt out under policy OI. Then compared to unrestricted targeting:*
(i) The opt-in policy strictly improves profits.
(ii) The opt-in policy makes all consumers strictly better off.

Proof. (Part (i) is proved in the Appendix.) A consumer's expected payment net of discounts, $p-\bar{\Delta}(p)$, is strictly increasing in in the list price p. With convex demand, any equilibrium with some consumers opting out ($\lambda^*<1$) must have $p^* \leq p^T$. Because of the lower list price, a consumer who opts in when she has the choice to do so will enjoy a larger surplus than she would under unrestricted targeting: $E(r_1)-p^*+\bar{\Delta}\left(p^*\right)-c > E(r_1)-p^T+\bar{\Delta}\left(p^T\right)-c$. Any consumers opting out could have chosen to opt in, so they must be better off as well. □

In the convex demand case, consumers who wish to opt out benefit when they can do so, and by exercising that right they induce more competitive list prices, benefiting all consumers (including those who expect to receive a discount.) Meanwhile, firms benefit because they no longer need to discount to consumers who opt out. The fact that this gain outweighs the lower list price charged to all consumers is not obvious, but it can be demonstrated by using the equilibrium condition (10) to eliminate the opt-out rate $1-\lambda^*$ from the comparison.

Consumers do not internalize the spillover effect on list prices when they opt out. We next show that the spillover may be large enough that consumers would be better off if opt-out were mandatory—that is to say, if targeting were banned.

Proposition 4. *Suppose $A>0$ and some consumers opt in under policy OI. Then:*
(i) Banning targeting strictly improves profits ($\Pi^{NT} > \Pi^{OI}$).
(ii) There is some $\bar{\rho} \in (0,1)$ such that banning targeting strictly improves consumer surplus ($CS^{NT} > CS^{OI}$) if captive demand is $\bar{\rho}$-convex.

For consumers, the logic resembles that of Proposition 2. A targeting ban preserves the privacy of those who would otherwise opt in and brings list prices down for those who would otherwise opt out. For sufficiently convex demand, the latter effect is large enough on its own to compensate for the discounts that the opt-ins give up. For firms, the reasoning is similar to the previous result. Note that both results apply for any equilibrium under OI, regardless of the equilibrium opt-in rate λ^*, as long as it is strictly positive. As with Proposition 3, this is achieved by using (10) to eliminate λ^* from the comparison of countervailing effects.

Other consumer-friendly policies

Because an outright ban on targeting might be impractical to implement, we will mention a few possible ways to improve upon an opt-in policy. Because opt-in inflicts a negative spillover on other consumers when demand is convex, standard arguments show that consumer choice plus a Pigouvian tax on opt-in could improve aggregate consumer surplus relative to consumer choice alone, assuming the tax proceeds could be returned to consumers as a lump sum. In practice, implementing such a tax would be unwieldy, of course. More realistically, regulators could impose a nuisance cost on consumers choosing opt-in, perhaps by making opt-out the default and imposing a paperwork burden on those who opt in. If we write $c(\lambda)=H^{-1}(\lambda)$ for the

31. Since this is the information consumers have when opting in or out, it is the appropriate stage at which to evaluate consumer surplus.

privacy cost of the λth-percentile consumer, $c'\left(\lambda^*\right)$ gives a measure of how hard it is to nudge the opt-in rate downward with a nuisance cost. If $c'\left(\lambda^*\right)$ is large, privacy costs drop off quickly for inframarginal ($\lambda < \lambda^*$) opt-in consumers, so a relatively large nuisance cost would be required to convince them to opt out; conversely, if $c'\left(\lambda^*\right)$ is small, a small nuisance cost may reduce the opt-in rate substantially. In the Appendix, we show that a nuisance cost benefits *all* consumers if $p'\left(\lambda^*\right) > c'\left(\lambda^*\right)$. This condition ensures that the opt-in rate can be nudged downward with a nuisance cost smaller than the resulting drop in prices; thus even consumers who continue to opt in (and therefore bear the nuisance cost) benefit.

As seen in Figure 1(b), the benefits from an opt-in policy may also be hamstrung by coordination failure. With convex demand, multiple equilibria may arise because consumers choices are self-reinforcing: higher opt-in rates induce higher list prices, making opt-in and the prospect of discounts even more attractive. For consumers, the equilibrium in the figure where everyone opts out Pareto dominates the equilibrium where everyone opts in. However, if the *status quo ante* were unrestricted targeting at list price p^T, a new opt-in policy might fail unless accompanied by a coordinated campaign to shift consumer expectations to the low-price equilibrium. For example, policymakers might wish to emphasize the right to privacy so that low opt-in becomes focal.

6.2. *Caveats and extensions*

6.2.1. Concave captive demand. While we have discussed reasons to think that concave demand is not the norm, it can arise if consumers have polarized tastes. The condition $p^{NT} < \bar{y} + A$, which we imposed to ensure that a mixture of opt-in and opt-out consumers generates a well-behaved price competition equilibrium, implies that taste polarization is not too severe.[32] Under this condition, the privacy model has an equilibrium with $p^* \in \left[p^T, p^{NT}\right]$, and our conclusions about firms' profits extend: both Propositions 3.i and 4.i apply, so $\Pi^{OI} \in \left(\Pi^T, \Pi^{NT}\right)$. That is, firms are better off under an opt-in policy than with unrestricted targeting, but they would be better off yet if targeting were banned.

The implications for consumer surplus are less straightforward, because opting out imposes a negative externality on other consumers (through higher prices) when demand is concave. Switching from unrestricted targeting to an opt-in policy will generally make some consumers worse off; it can even make *all* consumers worse off, as illustrated by the example in Figure 2. The example involves two-firm Hotelling competition with consumers are clustered near the firms, leading to the density of relative preferences for Firm 1 that is seen in Figure 2(a). Compared to unrestricted targeting, an opt-in policy clearly benefits firms, as list prices rise from p^T to p^*, and there are fewer consumers to discount to. The impact on consumers depends on whether the preserved privacy of those who opt out outweighs the higher price level faced by everyone. For the privacy costs illustrated in Figure 2(b), all consumers would be made strictly better off if the right to opt out were rescinded!

To see why, note that the example assumes all consumers have the same privacy cost, as illustrated by the flat $\lambda(p)$ curve. In the equilibrium at $\left(\lambda^*, p^*\right)$, all consumers are indifferent between opting in or out, so every consumer enjoys surplus equivalent to opting in at list price p^*. But rescinding the opt-in policy and reverting to unrestricted targeting would give each consumer surplus equivalent to opting in at the lower list price p^T. Homogeneous privacy tastes make the argument simple, but the conclusion that all consumers benefit from banning opt-out is robust

32. Because $p^{NT} = \frac{1-G(0)}{g(0)}$, the condition is satisfied $g(0)$ is not too small—that is, when number of consumers who like two products equally well is not too sparse.

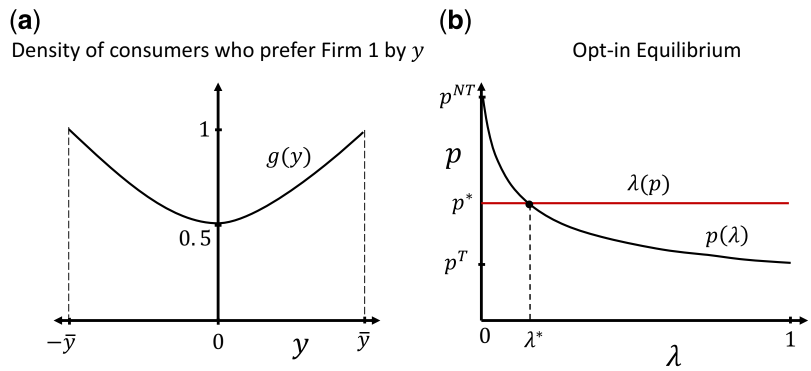

FIGURE 2

Concave demand: example where an opt-in policy makes all consumers worse off, relative to unrestricted targeting. (See the Appendix for details.)

to some heterogeneity in those tastes. While we do not suggest that consumers will typically be harmed by having more autonomy over their own data, the example does illustrate that demand curvature can cause privacy policies to have unexpected effects on prices.

6.2.2. Timing. In the Discussion Paper, we analyse the model when consumers make their privacy choices after observing their own tastes y and the price level p. This timing might make sense for sophisticated consumers making a major purchase. For example, after getting high initial quotes on a new car, a consumer with flexible preferences might wish to communicate her contestability by permitting her data to be collected on automotive websites. Our main results carry through—under a slightly stronger demand convexity condition, all consumers benefit from an opt-in policy relative to unrestricted targeting.

7. CONCLUDING REMARKS

In finding that competitive price discrimination shifts surplus from inframarginal to marginal consumers, we reinforce a consistent theme from the targeting literature. We advance that literature by showing how demand curvature determines the size of that shift and, consequently, targeting's impact on aggregate consumer welfare. Under plausible demand conditions, targeted discounting hurts consumers on average, and banning it would leave them better off.

Like much, but not all, of the preceding literature, we find that targeting also depresses profits. In those papers where firms benefit from targeting, the targeting technology usually has some imperfection or limitation that softens price competition over those targeted.[33] While slightly imperfect targeting would not change our model's conclusions about profits, large imperfections would require further study. Competitive price discrimination can also improve profits through a

33. Often (Iyer *et al.*, 2005; Galeotti and Moraga-González, 2008; 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.

market-expansion effect when the market is not fully covered, as illustrated by Thisse and Vives (1988).[34] Targeting would have a similar market-expanding effect in our model if the assumption that every consumer purchases were relaxed; presumably, this would make the impact of targeting look somewhat rosier for firms, and perhaps for consumers as well.

Our findings on privacy in Section 6 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, putting the choice to be targeted (or not) in individual consumers' hands makes all consumers better off. This conclusion hinges on the positive pricing externality bestowed on all consumers by the privacy-conscious individuals who opt out (when demand is convex). Indeed, due to this externality, consumers would be still better off if the right to opt out were supplemented with a tax, or even a dissipative nuisance cost, on those opting in. Furthermore, firms should not object to a policy that empowers consumers to opt in or out since, as we showed, such a policy will improve profits by committing firms to discount less.

Acknowledgments. We thank D. J. Campbell, 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 PSE/Paris 1 Panthéon-Sorbonne, 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 (January 2019). The first and third authors gratefully acknowledge support from the NSF and the first and second from the Bankard Fund for Political Economy.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

A. APPENDIX

Section 2 proofs

Proof of Lemma 1

Part (i) We appeal to known properties of log-concave distributions; see Bagnoli and Bergstrom (2005). 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 apply 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

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

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

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

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

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

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

But $f'(r)f(r)=\frac{1}{2}d\left(f(r)^2\right)$, so if $n=2$ we have $\frac{g'(0)}{n-1}=-\frac{1}{2}\left(f(\bar{r})^2+f\left(\underline{r}\right)^2\right)\leq 0$. Otherwise, integrate by parts to get

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

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

Part (iii) As the paper only uses this result in passing, we refer the reader to our Discussion Paper for the proof.

Section 3 analysis and proofs

(Targeting subgame in Stage 2)

As in the text, we consider Stage 2 competition for a consumer with tastes $r_1>r_2>...>r_n$, with $y_1=r_1-r_2$, given list prices p_1^l and $p_{j\neq 1}^l=p$. For the purpose of the paper, analysing this "semi-symmetric" subgame (where all of Firm 1's rivals have set the same list price) will suffice, since our interest is in the incentive to deviate from a symmetric list price profile. For completeness, the Discussion Paper gives an analysis of the Stage 2 subgame for arbitrary profiles of list prices; aside from the heavier notational burden, the logic is quite similar.

Define $P_1=\min\left(p_1^l,A\right)$, $P_{-1}=\min(p,A)$, and $y_1^*=p_1^l-P_{-1}$, and $\hat{y}_1=P_1-P_{-1}$. It is useful to partition the possible values of y_1 into three intervals: Region I is $y_1>y_1^*$, Region II is $y_1\in\left(\hat{y}_1,y_1^*\right)$, and Region III is $y_1<\hat{y}_1$. If Firm 1 cannot advertise ($p_1^l\leq A$), then $\hat{y}_1=y_1^*$, and so Region II vanishes. Proposition 5 is focused on Firm 1's equilibrium expected profit on this consumer. The proposition also gives other results (on who advertises to the consumer and the profits of other firms) that are used in the article.

Proposition 5. *(Firm 1's Stage 2 equilibrium profit)*
(I) A Region I consumer is captive to Firm 1, with profits $\pi_1=p_1^l$ and $\pi_{j\neq 1}=0$.
(II) Region II is non-empty if $p>A$, in which case profits on a Region II consumer are $\pi_1=y_1+P_{-1}-A$ and $\pi_{j\neq 1}=0$. In particular:

a. If $p\leq A$, Region II is $\left(A-p,p_1^l-p\right)$, the consumer is conceded to Firm 1, and $\pi_1=y_1+p-A$.
b. If $p>A$, Region II is $\left(0,p_1^l-A\right)$, the consumer is contested by Firms 1 and 2 only, and $\pi_1=y_1$.

(III) On a Region III consumer, Firm 1 earns $\pi_1=0$.

Proof. Part (I) If $p\leq A$, then $y_1^*=p_1^l-p$. Then, a consumer with $y_1>y_1^*$ prefers Firm 1's list price over other list price offers, and no other firm can afford to advertise a discount below its list price. If $p>A$, then a consumer type $y_1>y_1^*=p_1^l-A$ would require a discount offer $p_2^d\leq p_1^l-y_1<A$ to buy from Firm 2, but Firm 2 cannot profitably advertise a price this low. *A fortiori*, no lower-ranked firm can profitably target the consumer either.

Part (II.a) When $y_1<y_1^*$, the consumer's default is Firm 2, but Firm 1 can poach her with an "undercutting" offer $p_1^d\leq p+y_1$. As no other firm can afford to discount, it will target her with the minimal discount $p_1^d=p+y_1$ as long as the net profit $\pi_1=p_1^d-A$ from doing so is positive. Thus, it poaches her if $y_1>A-p$ and refrains from advertising if $y_1<A-p$.

Part (II.b) Proved in Lemmas 2–4 below.

Part (III) If $P_{-1}=p$, then $\hat{y}_1=\min\left(p_1^l,A\right)-p$. Then for $y_1<\hat{y}_1$, neither Firm 1's list price nor its lowest conceivable discount price beats Firm 2's list price. If $P_{-1}=A$ and $P_1=p_1^l$, then $\hat{y}_1=p_1^l-A$, and Firm 1 cannot afford to advertise. A consumer $y_1<\hat{y}_1$ either prefers Firm 2's list price (if $y_1<p_1^l-p$) or can be profitably won by Firm 2 with some discount $p_2^d\geq A$. If $P_{-1}=P_1=A$, then all firms can afford to discount, and $\hat{y}_1=0$. Then for $y_1<\hat{y}_1=0$, Firm 2 makes the consumer's favourite product, and $\pi_1=0$ follows by applying Part (II) to the re-ordered ranking of firms. □

Lemmas 2–4 establish Part II.b of the proposition. They presume that p_1^l and p strictly exceed A and that the consumer in question satisfies $y_1\in\left(\hat{y}_1,y_1^*\right)=\left(0,p_1^l-A\right)$.[35]

35. We omit arguments for the boundary cases $y_1=0$ and $y_1=y_1^*$. For the former, arguments similar to those here establish zero profits for all firms. For the latter, it can be shown there is a range of equilibria (depending on how often

Lemma 2. *Only the consumer's favourite firm earns a positive profit on her: $\pi_1 \geq y_1 > 0$ and $\pi_{j>1} = 0$.*

Proof. Because any competitor's discount will satisfy $p_j^d \geq A$, Firm 1 can guarantee winning the consumer by advertising $p_1^d = A + y_1 - \varepsilon$, for $\varepsilon > 0$, thereby earning $\pi_1 = y_1 - \varepsilon$. Since $\pi_1 \geq y_1 - \varepsilon$ for all $\varepsilon > 0$, we have $\pi_1 \geq y_1 > 0$. Suppose toward a contradiction that $\pi_2 > 0$. This implies that both Firms 1 and 2 must target the consumer with probability one. (Of the two firms, the non-default firm strictly prefers to advertise a discount, since it would earn zero otherwise. But in this case the default firm will earn zero without discounting, so it strictly prefers to advertise too.) But then standard results for Bertrand competition preclude an outcome where both firms cover the ad cost A, contradicting the strict positivity of both profits. The same argument rules out $\pi_j > 0$ for any $j > 2$. □

Lemma 3. *Firm 1's profit on the consumer is $\pi_1 = y_1$.*

Proof. Let $\underline{p}_1$ be the infimum over Firm 1's support of discount offers to this consumer. If $\pi_1 > y_1$, then $\underline{p}_1$ must satisfy $\underline{p}_1 > A + y_1$. But then Firm 2 could earn a strictly positive profit with an undercutting discount $p_2^d = \underline{p}_1 - y_1 - \varepsilon > A$, contradicting Lemma 2. □

Lemma 4. *Only Firms 1 and 2 target the consumer with positive probability: $a_1 > 0$, $a_2 > 0$, and $a_{j>2} = 0$.*

Proof. If a lower-ranked firm $j > 2$ did advertise, it would need to earn a weakly positive profit on its lowest advertised price $\underline{p}_j$. But then by advertising the consumer-surplus-equivalent discount $p_2^d = \underline{p}_j + (r_2 - r_j)$, Firm 2 could win the consumer equally often but at a higher price, thereby earning a strictly positive profit (and contradicting Lemma 2). Next to establish $a_1 > 0$, note that Firm 2 could earn a strictly positive profit if Firm 1 never advertised (either as the consumer's default or by sending the undercutting offer $p_2^d = p_1^l - y_1 > A$), contradicting Lemma 2. Similarly, if Firm 2 never advertised, then Firm 1 could earn the profit $\pi_1 = p_1^l > y_1$ (if it is the default) or $\pi_1 = p + y_1 - A > y_1$ (by undercutting Firm 2's list price). Either case contradicts Lemma 3, establishing $a_2 > 0$. □

Proposition 6 below derives the discounting strategies for a contested consumer that appear in the text. It is assumed that list prices are symmetric and exceed the ad cost. We consider a consumer $y_1 \in (0, y_1^*) = (0, p - A)$ who (by Proposition 5) is contested by Firms 1 and 2 only.

Proposition 6. *Equilibrium discounting strategies with respect to a contested consumer are as described in the text.*

Proof. Advertised surplus offers have supports $(s_1^l, r_2 - A]$ and $[s_1^l, r_2 - A]$. Let $\bar{s}_j$ ($\underline{s}_j$) be Firm j's supremum (infimum) over advertised offers. The suprema satisfy $\bar{s}_1 = \bar{s}_2 = r_2 - A$. (Firm 2 cannot profitably offer $s_2 > r_2 - A$, and so Firm 1 need not offer $s_1 > r_2 - A$ either. And if either supremum were strictly below $r_2 - A$, the other firm could strictly exceed its equilibrium payoff by "overcutting" slightly, contradicting Lemmas 2 and 3.) Firm 1 will not make a discount offer $s_1 < \underline{s}_2$ that wins only if Firm 2 does not advertise; it will win just as often (and save A) by not discounting. So $\underline{s}_1 \geq \underline{s}_2$. Next we have $\underline{s}_2 = s_1^l$: since Firm 2's lowest advertised surplus offer wins only against Firm 1's list price, it should do so by no more than necessary. Finally, we cannot have $\underline{s}_1 > \underline{s}_2$, since 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 $\underline{s}_1 = s_1^l$ as well. The arguments against gaps and atoms on $(s_1^l, r_2 - A]$ are standard. We defer showing that Firm 2 makes an atom of advertised offers at $s_2 = s_1^l$ until the next step.
Mixed strategies over $s \in (s_1^l, r_2 - A]$ are given by $B_1(s)$ and $B_2(s)$. Firm 1's net expected profit from advertising surplus s is $\pi_1(s) = B_2(s)(r_1 - s) - A$. Then, use $r_1 = r_2 + y_1$ and the indifference condition $\pi_1(s) = y_1$ to obtain $B_2(s)$. Similarly, Firm 2's expected profit is $\pi_2(s) = B_1(s)(r_2 - s) - A$; the indifference condition $\pi_2(s) = 0$ delivers $B_1(s)$. Then $B_1(s_1^l) = \frac{A}{p - y_1} = 1 - a_1 > 0$ delivers the probability that Firm 1 sends no ad. In this case, Firm 1 wins only if Firm 2 does not advertise; hence $\pi_1 = (1 - a_2)p = y_1$ (where the last equality follows by indifference). We conclude $B_2(s_2^l) = 1 - a_2 = y_1/p$. It is established that Firm 2 makes no advertised offers on $s \in (s_2^l, s_1^l)$, and we also have $B_2(s_1^l) = \lim_{s \searrow s_1^l} B_2(s) = (A + y_1)/p$. The difference, $B_2(s_1^l) - B_2(s_2^l) = A/p$, must be an atom of offers at $s_2 = s_1^l$, undercutting Firm 1's list price. □

Section 4 analysis and proofs

(List prices in Stage 1)

Firm 2 advertises) yielding profits $\pi_1 \in [y_1^*, p_1^l]$ for Firm 1. As profits on zero measure sets have no impact on Stage 1 incentives, neither case is critical to the main results of the article.

Characterization of a unique symmetric equilibrium in list prices

Suppose that $A < p^{NT}$. (The case of $A \geq p^{NT}$ is covered in the text.) Firm 1's first-order condition (4) is equivalent to $\frac{1-G\left(p_1^l - P_{-1}\right)}{g\left(p_1^l - P_{-1}\right)} - P_1 = 0$. Define a function $\Theta(p)$ equal to the left-hand side of this expression, evaluated at the strategy profile in which *all* list prices are equal to p:

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

Note $\Theta(A) = p^{NT} - A > 0$. Furthermore, $\Theta(p)$ is strictly decreasing (as monotonicity of $\frac{1-G(y)}{g(y)}$ follows from Condition 1) and tends toward $m - A$ as $p - A \to \bar{y}$ (where $m = \frac{1-G(\bar{y})}{g(\bar{y})}$).[36] Thus, if $A > m$, then $\Theta(p) = 0$ has a unique solution for some $p \in (A, \bar{y}+A)$. Alternatively, if $A < m$, then $\Theta(p)$ is strictly positive at any price level such that $p - A < \bar{y}$.

Proposition 7. *Under Condition 1, there is a unique symmetric equilibrium. This is the unique equilibrium of the game if there are two firms. If $A \geq p^{NT}$, the common list price is p^{NT} and targeted discounts are not used. If $A \in \left(m, p^{NT}\right)$, the list price solves $\Theta\left(p^l\right) = 0$, targeting is used, and all non-captive consumers are contested by their top two firms. If $A < m$, then $p^l = \bar{y} + A$, and all but the most captive consumers are contested with targeting by their top two firms.*

Proof of Proposition 7

Suppose $A > m$. Let p^* be the unique solution to $\Theta(p) = 0$; as noted above, this solution must satisfy $p^* > A$, so the second part of the piecewise definition of $\Theta(p)$ applies, and p^* must solve (5). By construction, setting $p_j^l = p^*$ solves Firm j's first-order condition (4) when all other firms charge p^*. Furthermore, marginal profit $\partial \Pi_j \left(p_j^l\right) / \partial p_j^l$ is strictly positive for $p_j^l < p^*$ and strictly negative for $p_j^l > p^*$ (using strict logconcavity of $1 - G(y)$ and the fact that P_j is weakly increasing in p_j^l), so setting $p_j^l = p^*$ uniquely maximizes Firm j's profit. This establishes the symmetric equilibrium at p^*. The features of equilibrium follow from arguments in the text.

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

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 $\chi(u,v) = \frac{1-G(u-\min(v,A))}{g(u-\min(v,A))} - \min(u,A)$ so the first-order conditions imply $\chi\left(p_1^l, p_2^l\right) = 0$ and $\chi\left(p_2^l, p_1^l\right) \geq 0$. But $\chi(u,v)$ 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 $\chi\left(p_1^l, p_2^l\right) > \chi\left(p_2^l, p_2^l\right) \geq \chi\left(p_2^l, p_1^l\right) \geq 0$, contradicting $\chi\left(p_1^l, p_2^l\right) = 0$.

Section 5 analysis and proofs

Proof of Proposition 1

Because $p^{NT} = p^l(A)\big|_{A=p^{NT}}$, it suffices to show that $p^l(A)$ is strictly decreasing (increasing) if captive demand is strictly convex (strictly concave). Given $A < p^{NT}$, the equilibrium condition is $\Theta\left(p^l; A\right) = \frac{1-G\left(p^l - A\right)}{g\left(p^l - A\right)} - A = 0$, making the dependence on the parameter A explicit. Differentiate this equilibrium condition implicitly to get

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

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

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

Proof of Proposition 2

The profit ranking is established by the argument in the text. That argument relies on Condition 1 to ensure that the two equilibria exist, but otherwise it does not depend at all on demand curvature.

Consumer surplus. Given symmetry, it suffices to aggregate over consumers $y_1 \geq 0$ with favourite 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 $\lim_{y\to\infty} m(y) = m > A$). Because $L(y_1, p^T, A) \geq A$ (see (8)), 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 5(i) and the fact that $p^{NT} = \mu(0)$, we have $y^* \geq (p^{NT} - A)/\rho$. Then, use Lemma 5(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:

$$\rho < \alpha\left(\frac{1-\alpha^{1/\rho}}{1-\alpha}\right). \tag{A.1}$$

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 (A.1).

Lemma 5. *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 $1-G(y) = (1-G(0))\exp\left(-\int_0^y \nu(y')dy'\right)$. 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. □

Section 6 Analysis and proofs

EXISTENCE OF A UNIQUE SYMMETRIC PRICE COMPETITION EQUILIBRIUM

Recall our standing assumption that $1-G(y)$ is strictly logconcave, so the hazard rate $h(y) = \frac{g(y)}{1-G(y)}$ is strictly increasing. Suppose firms anticipate an opt-in rate λ. As noted in the text, any symmetric price competition equilibrium at list price p^* must satisfy the necessary condition $\Phi(p^*) = 0$. We will show that the condition $\Phi(p) = 0$ has a unique solution and is both necessary and sufficient for an equilibrium. As noted in the text, for the concave captive demand case, we impose the condition $p^{NT} < \bar{y} + A$.

Concave captive demand

Suppose captive demand is concave, so $g'(y) \geq 0$ on $[0, \bar{y}]$. Note that $\Phi(p) = \lambda\Phi_T(p) + (1-\lambda)\Phi_{NT}(p)$, where $\Phi_T(p) = 1-G(p-A) - Ag(p-A)$ and $\Phi_{NT}(p) = 1-G(0) - pg(0)$. Both $\Phi_T(p)$ and $\Phi_{NT}(p)$ are strictly decreasing, so $\Phi(p)$ is strictly decreasing as well. Furthermore, we have $\Phi_T(p^T) = 0$ and $\Phi_{NT}(p^{NT}) = 0$, with $p^T \leq p^{NT}$ by Proposition 1. This implies $\Phi(p)$ is strictly positive for $p < p^T$ and strictly negative for $p > p^{NT}$. Thus, $\Phi(p) = 0$ has a unique solution p^*, located on $[p^T, p^{NT}]$.

To show sufficiency, it suffices to show that $p_1 = p^*$ maximizes Firm 1's profit when all other firms charge p^*. Write Firm 1's marginal profit as

$$\Theta(p_1) := \left.\frac{d\Pi_1}{dp_1}\right|_{p_{-1}=p^*} = \lambda\Theta_T(p_1) + (1-\lambda)\Theta_{NT}(p_1),$$

where $\Theta_T(p_1) = 1-G(p_1-A) - Ag(p_1-A)$ and $\Theta_{NT}(p_1) = 1-G(p_1-p^*) - p_1 g(p_1-p^*)$ are the marginal profits associated with opt-in and opt-out consumers, respectively. By construction, $\Theta(p^*) = 0$, with $\Theta_T(p^*) < 0$, $\Theta_{NT}(p^*) > 0$,

$p^* \in \left[p^T, p^{NT}\right]$, and, by assumption, $p^{NT} < \bar{y}+A$. In the sequence of claims below, we establish that $\Theta(p_1)$ is strictly positive for $p_1 < p^T$, strictly decreasing for $p_1 \in \left[p^T, p^{NT}\right]$, and strictly negative for $p_1 > p^{NT}$. It follows that Firm 1's profit is maximized at the first-order condition solution, $p_1 = p^*$.

Claim: $\Theta(p_1)$ is strictly positive for $p_1 < p^T$ and strictly negative for $p_1 > p^{NT}$.

The fact that $\Theta_T\left(p^T\right)=0$ and strict logconcavity of $1-G$, implies that $\Theta_T(p_1)$ is strictly positive for $p < p^T$. A similar argument establishes that $\Theta_{NT}(p_1)$ is strictly positive for $p_1 \leq p^*$ (since $\Theta_{NT}(p^*)>0$). *A fortiori*, $\Theta(p_1)>0$ for $p_1 < p^T$.

Another logconcavity argument implies that $\Theta_{NT}(p_1)<0$ for $p_1 > p^{NT}$. The targeting marginal profit $\Theta_T(p_1)$ is strictly negative for $p_1 \in \left(p^T, \bar{y}+A\right)$ and jumps to zero for $p_1 > \bar{y}+A$. (Once all opt-in consumers are targeted, further list price increases are irrelevant to the opt-in portion of profits.) Since we have assumed $\bar{y}+A > p^{NT}$, $\Theta_T(p_1)$ is at least weakly negative for all $p_1 > p^{NT}$, and so the convex combination $\Theta(p_1)$ is strictly negative for all $p_1 > p^{NT}$.

Claim: $\Theta_T(p_1)$ is strictly decreasing for $p_1 \in \left[p^T, p^{NT}\right]$.

Note that $p^T - A > 0$ and $p^{NT} - A < \bar{y}$, so on this range, $p_1 - A$ is within the support of $g(y)$, with $g'(p_1-A)>0$. Then, the claim follows immediately from the derivative $\Theta_T'(p_1)$.

Claim: $\Theta_{NT}(p_1)$ is strictly decreasing for $p_1 \geq p^T$.

For $p_1 \geq p^*$, we have $g'(p_1-p^*) \geq 0$, and an argument similar to the previous one applies. For $p_1 \in [p^T, p^*)$, the argument $p_1 - p^*$ is negative and so (by the symmetry of the captive demand function), $g'(p_1-p^*)<0$. In this case, note that strict logconcavity of captive demand implies $g'(y) > -\frac{g(y)^2}{1-G(y)}$. (Differentiate the hazard rate.) Substituting this inequality into $\Theta_{NT}'(p_1) = -2g(p_1-p^*) - p_1 g_1'(p_1-p^*)$ produces

$$\Theta_{NT}'(p_1) < -g\left(p_1-p^*\right) - \frac{g(p_1-p^*)}{1-G(p_1-p^*)}\Theta_{NT}(p_1)$$

But we established above that $\Theta_{NT}(p_1)>0$ for $p \in [p^T, p^*)$, so it follows that $\Theta_{NT}(p_1)$ is strictly decreasing on this range as well.

Convex captive demand

Now suppose captive demand is strictly convex, so $g'(y)<0$ on $(0,\infty)$. As above, we seek to establish the existence of a unique symmetric PCE by showing (1) that $\Phi(p)=0$ has a unique solution and (2) a firm's profit function is single-peaked and maximized at p^* when all other firms charge p^*.

(1) A unique solution to $\Phi(p)=0$ exists.

By strict convexity, the no-targeting and unrestricted targeting list prices satisfy $p^{NT} < p^T$. Because $\Phi_T(p)$ and $\Phi_{NT}(p)$ are both positive below p^{NT} and both negative above p^T, $\Phi(p)=0$ has some solution on the interval $\left[p^{NT}, p^T\right]$ and no solutions outside this interval. Suppose $\Phi(p^*)=0$ is such a solution. To show uniqueness, it suffices to show $\Phi'(p^*)<0$. Since $\Phi_{NT}'(p^*)<0$ is immediate, showing $\Phi_T'(p^*)<0$ will suffice. For this, write $\Phi_T(p) = (1-G(p-A))(1-Ah(p-A))$ and differentiate:

$$\begin{aligned}\Phi_T'(p) &= -g(p-A)(1-Ah(p-A)) - A(1-G(p-A))h'(p-A) \\ &= -h(p-A)\Phi_T(p) - A(1-G(p-A))h'(p-A)\end{aligned}$$

Evaluated at p^*, the first term is weakly negative because $p^* \leq p^T$ implies $\Phi_T(p^*) \geq 0$, and the second term is strictly negative by the monotonicity of the hazard rate.

(2) Profit is uniquely maximized at $p_1 = p^*$ when all other firms charge p^*.

It suffices to consider Firm 1, whose marginal profit may be written $\Theta(p_1) = \lambda\Theta_T(p_1) + (1-\lambda)\Theta_{NT}(p_1)$ as above. Clearly $p_1 = p^*$ is one solution to the first-order condition $\Theta(p_1)=0$. We will show that any solution $\hat{p}_1$ to $\Theta(p_1)=0$ must satisfy $\Theta'\left(\hat{p}_1\right)<0$; this implies that Firm 1's profit is strictly quasiconcave and uniquely maximized at $p_1 = p^*$.

First, we claim that $\Theta\left(\hat{p}_1\right)=0$ implies $\hat{p}_1 < p^T$.

Proof Note that $\Theta_{NT}(p^*)<0$ (because p^{NT} is defined by $1-G(0)-p^{NT}g(0)=0$, and $p^* > p^{NT}$). Because $\Theta_{NT}(p_1)$ crosses zero once, from above, and $p^T > p^*$, $\Theta_{NT}(p_1)<0$ for all $p_1 \geq p^T$. Since $\Theta_T(p_1)$ is also negative above p^T, we have $\Theta(p_1)<0$ for all $p_1 \geq p^T$.

Next, regroup the terms in Firm 1's marginal profit as:

$$\Theta(p_1) = \lambda(1-G(p_1-A))Z_T(p_1) + (1-\lambda)\left(1-G\left(p_1-p^*\right)\right)Z_{NT}(p_1),$$

where $Z_T(p_1) = 1-Ah(p_1-A)$ and $Z_{NT}(p_1) = 1-p_1h(p_1-p^*)$. Then,

$$\begin{aligned}\Theta'(p_1) &= \left[\lambda(1-G(p_1-A))Z_T'(p_1) + \lambda\left(1-G\left(p_1-p^*\right)\right)Z_{NT}'(p_1)\right] - X(p_1), \text{ where} \\ X(p_1) &= \lambda g(p_1-A)Z_T(p_1) + (1-\lambda)g\left(p_1-p^*\right)Z_{NT}(p_1).\end{aligned}$$

The first term is negative for any p_1 because $Z_T(p_1)$ and $Z_{NT}(p_1)$ are strictly decreasing, so it will suffice to show $X(\hat{p}_1) > 0$ holds whenever $\Theta(\hat{p}_1) = 0$. Manipulate $X(p_1)$ to get:

$$\begin{aligned} X(p_1) &= h(p_1 - A) \cdot \lambda \Theta_T(p_1) + h(p_1 - p^*) \cdot (1-\lambda)\Theta_{NT}(p_1) \\ &= h(p_1 - p^*)\Theta(p_1) + \lambda\big(h(p_1 - A) - h(p_1 - p^*)\big)\Theta_T(p_1). \end{aligned}$$

By conjecture, at $\hat{p}_1$ the first term drops out: $X(\hat{p}_1) = \lambda\big(h(\hat{p}_1 - A) - h(\hat{p}_1 - p^*)\big)\Theta_T(\hat{p}_1)$. Both terms in the remaining expression are strictly positive (by monotonicity of the hazard rate and because we showed that $\hat{p}_1 < p^T$). As claimed, this establishes that $\Theta'(\hat{p}_1) < 0$ at any $\hat{p}_1$ satisfying the first-order condition $\Theta(\hat{p}_1) = 0$.

Monotonicity of $p(\lambda)$

As elsewhere, we restrict attention to the cases where captive demand is either strictly convex or strictly concave. (If captive demand is linear, it is easily seen that $p(\lambda) = p^{NT} = p^T$.) Since $p(\lambda)$ is defined implicitly by the condition $\Phi(p) = 0$, we have $p'(\lambda) = -\Phi_\lambda / \Phi_p\big|_{p=p^*}$. In proving equilibrium uniqueness, we showed that $\Phi_p\big|_{p=p^*} < 0$. If demand is strictly convex, and $\lambda \in (0,1)$, the equilibrium price $p^* \in (p^{NT}, p^T)$ satisfies $\Phi_T(p^*) > 0 > \Phi_{NT}(p^*)$, so $\Phi_\lambda|_{p=p^*} > 0$, and thus $p'(\lambda) > 0$. The same argument applies with very slight adaptations at $\lambda = 0$ and $\lambda = 1$. If demand is strictly concave, $p^* \in (p^T, p^{NT})$, the argument above reverses, and so $p'(\lambda) < 0$.

All consumers can benefit from a nuisance cost on opt-in.

Suppose demand is strictly convex and there is an interior equilibrium (λ^*, p^*). To ensure this equilibrium is stable, we also suppose the $\lambda(p)$ curve crosses $p(\lambda)$ from below, as in Figure 1(a). Because $\lambda(p)$ curve can be written $p = \bar{\Delta}^{-1}(c(\lambda))$, this stability condition simplifies to $c'(\lambda^*) > p'(\lambda^*)\bar{\Delta}'(p^*)$.) By construction, $c(\lambda^*) = \bar{\Delta}(p^*)$, since the λ^* consumer is indifferent between opting in or out.

Suppose the government implements a lower opt-in rate $\hat{\lambda} < \lambda^*$ by imposing a nuisance cost τ on opt-in. The size of the nuisance cost must be such that $\hat{\lambda}$ consumers are indifferent: $c(\hat{\lambda}) + \tau = \bar{\Delta}(\hat{p})$, where $\hat{p} = p(\hat{\lambda})$. Thus, $\tau = \bar{\Delta}(\hat{p}) - c(\hat{\lambda})$. Consumers at $\lambda \geq \lambda^*$ opt out before and after the nudge; these consumers strictly benefit when list prices fall from p^* to $\hat{p}$. A consumer at $\lambda < \lambda^*$, with privacy cost $c(\lambda)$, opts in before and makes net payments (including privacy cost) $p^* - \bar{\Delta}(p^*) + c(\lambda) = (p^* - c(\lambda^*)) + c(\lambda)$, using the equilibrium condition at the "star" prices. If this consumer also opts in after the nudge, she makes net payments $\hat{p} - \bar{\Delta}(\hat{p}) + c(\lambda) + \tau = \big(\hat{p} - c(\hat{\lambda})\big) + c(\lambda)$, using the equilibrium at the new, "hat" prices. Since she also has the option to switch to opting out after the nudge, this consumer is made unambiguously better off if $\hat{p} - c(\hat{\lambda}) < p^* - c(\lambda^*)$. So, there is a nudge that benefits all consumers if $p(\hat{\lambda}) - c(\hat{\lambda}) < p(\lambda^*) - c(\lambda^*)$ for some $\hat{\lambda} < \lambda^*$. Since $p'(\lambda^*) > c'(\lambda^*)$ ensures this is true for all $\hat{\lambda}$ sufficiently close to λ^*, we have the claim in the text.

Note that $\bar{\Delta}'(p) < 1$, so the stability condition does not preclude the $p'(\lambda^*) > c'(\lambda^*)$ condition from being met.

Concave captive demand example

For the example in the text, the captive demand function is $1 - G(y) = \frac{1}{2} - \frac{1}{2}y - \frac{k}{3}y^3$ on support $[-\bar{y}, \bar{y}]$, with density $g(y) = \frac{1}{2} + ky^2$, upper limit $\bar{y} = 0.6$, and normalization constant $k = \frac{3-3\bar{y}}{2\bar{y}^2}$. The ad cost is $A = 0.4$. The captive demand function (slightly) violates logconcavity when $y < 0$, but it can be confirmed that profit functions (for no-targeting, unrestricted targeting, and a firm's blended profit function with opt-in) all remain quasiconcave, and this suffices for the equilibrium characterization. (One can recover logconcavity in this parametric family by reducing the curvature of $g(y)$; the price difference $p^{NT} - p^T$ becomes smaller, but otherwise the conclusions are similar.) By construction, the no-targeting list price is $p^{NT} = 1$; with unrestricted targeting, the list price is $p^T \approx 0.766$. Under opt-in, the equilibrium condition is $(1-\lambda)(1 - G(0) - (y+A)g(0)) + \lambda(1 - G(y) - Ag(y)) = 0$, using the relation $p - A = y$ for the marginal targeted consumer (among those who opt in). For $\lambda \in [0,1]$, the left-hand side reduces to the strictly decreasing cubic $\left(\frac{14}{9} - \frac{169}{150}\lambda\right) + \left(\frac{41}{10}\lambda - \frac{13}{3}\right)p + 5(1-\lambda)p^2 - \left(\frac{20}{9} - \frac{5}{3}\lambda\right)p^3$; the equilibrium list price function $p(\lambda)$ plotted in Figure 2(b) is given by the unique real root of this cubic.

Proofs of Propositions 3 and 4

Proof of Proposition 3

Part (ii) is proved in the text. The profit results in part (i) and Proposition 4 are proved together in Lemma 6.

Lemma 6. *A firm's profit in an equilibrium under the opt-in policy satisfies* $\Pi_{OI} \in [\Pi_T, \Pi_{NT}]$. *Furthermore,* $\Pi_{OI} \in (\Pi_T, \Pi_{NT})$ *if the equilibrium is interior* ($\lambda^* \in (0,1)$).

Proof. Start with the case of strictly convex captive demand.

If captive demand is strictly convex:

Let (λ^*, p^*) be an equilibrium under regime OI, with $y^* = p^* - A$. The result is immediate if $\lambda^* = 0$ or $\lambda^* = 1$, so we focus on the case $\lambda^* \in (0,1)$. Then, the price satisfies $p^* \in \left(p^{NT}, p^T\right)$ and solves the equilibrium condition

$$\left(1-\lambda^*\right)\left(1-G(0)-\left(y^*+A\right)g(0)\right)+\lambda^*\left(1-G\left(y^*\right)-Ag\left(y^*\right)\right)=0$$

which may be rearranged as:

$$p^* = p^{NT} + \frac{\lambda^*}{1-\lambda^*}\frac{1-G(y^*)-Ag(y^*)}{g(0)} \tag{A.2}$$

Profit in this equilibrium is $\Pi_{OI} = (1-\lambda^*)\Pi_O + \lambda^*\Pi_I$, where $\Pi_O = (1-G(0))p^*$ and $\Pi_I = (1-G(p^*-A))p^* + \int_0^{p^*-A} y\ g(y)dy$ are profits on opt-outs and opt-ins, respectively. Use (A.2) and $p^{NT} = (1-G(0))/g(0)$ to replace $(1-\lambda^*)\Pi_O$:

$$\Pi_{OI} = \left(1-\lambda^*\right)\Pi_{NT} + \lambda^*\left(p^{NT}\left(1-G\left(y^*\right)-Ag\left(y^*\right)\right)+\Pi_I\right.$$

With an eye toward showing $\Pi_{OI} < \Pi_{NT}$, define

$$Z(y) = p^{NT}(1-G(y)-Ag(y)) + (1-G(y))(y+A) + \int_0^y y'g\left(y'\right)dy' - \Pi_{NT}.$$

Then, we have $\Pi_{OI} = \Pi_{NT} + \lambda^* Z(y^*)$. To prove $\Pi_{OI} < \Pi_{NT}$, it suffices to show that $Z(y)$ is strictly negative for $y \in \left(0, p^T - A\right)$ (since $y^* \in \left(p^{NT}-A, p^T-A\right)$ and $\lambda^* > 0$). First, observe that $Z(0) = 0$, so it will suffice to show $Z'(y) < 0$ for all $y \in \left(0, p^T - A\right)\right)$, where $Z'(y) = (1-G(y)-Ag(y)) - p^{NT}\left(g(y)+Ag'(y)\right)$. Strict logconcavity of $1-G(y)$ implies $g'(y) > -\frac{g(y)^2}{1-G(y)}$, so

$$\begin{aligned} Z'(y) &< (1-G(y)-Ag(y)) - p^{NT}\left(g(y) - \frac{Ag(y)^2}{1-G(y)}\right) \\ &= (1-G(y)-Ag(y))\left(1 - \frac{1-G(0)}{g(0)}\frac{g(y)}{1-G(y)}\right). \end{aligned}$$

Because the hazard rate $g(y)/(1-G(y))$ is strictly increasing, the first term is strictly positive for $y < p^T - A$ and the second is strictly negative for $y > 0$, so $Z'(y) < 0$ holds on $\left(0, p^T - A\right)$ as claimed. This establishes $\Pi_{OI} < \Pi_{NT}$.

To show that $\Pi_{OI} > \Pi_T$, note that the latter may be written $\Pi_T = \Pi_{NT} + Z\left(p^T - A\right)$. Then, the following sequence of inequalities follows because Z is strictly negative and strictly decreasing, respectively.

$$\Pi_T - \Pi_{OI} = Z\left(p^T - A\right) - \lambda^* Z\left(y^*\right) < \lambda^*\left(Z\left(p^T - A\right) - Z\left(y^*\right)\right) < 0.$$

If captive demand is weakly concave:

We take an entirely different approach in this case. Write $\Pi(\lambda)$ for the PCE profit with opt-in rate λ. Since $\Pi_T = \Pi(1)$, $\Pi_{NT} = \Pi(0)$, and $\Pi_{NT} > \Pi_T$, it suffices to show that $\Pi(\lambda)$ is strictly decreasing in λ. Without loss of generality, express this profit from Firm 1's point of view as $\Pi(\lambda) = \Pi(p_1, P_{-1}, \lambda)|_{p_1 = p(\lambda),\ P_{-1} = p(\lambda)}$, where we explicitly separate Firm 1's list price from the common list price P_{-1} of its rivals, and $p(\lambda)$ is the PCE price. We have $\Pi(p_1, P_{-1}, \lambda) = (1-\lambda)\Pi_O(p_1, P_{-1}) + \lambda\Pi_I(p_1)$, where $\Pi_O(p_1, P_{-1}) = p_1(1-G(p_1 - P_{-1}))$ is the profit on an opt-out consumer, and $\Pi_I(p_1) = p_1(1-G(p_1 - A)) + \int_0^{p_1 - A} y\, dG(y)$ is the profit on an opt-in. Note that only $\Pi_O(p_1, P_{-1})$ depends on the rivals' list price. Then,

$$\frac{d\Pi(\lambda)}{d\lambda} = \left[\frac{\partial\Pi(p_1, P_{-1}, \lambda)}{\partial\lambda} + \frac{\partial\Pi(p_1, P_{-1}, \lambda)}{\partial p_1}p'(\lambda) + \frac{\partial\Pi(p_1, P_{-1}, \lambda)}{\partial P_{-1}}p'(\lambda)\right]_{p_1 = p(\lambda),\ P_{-1} = p(\lambda)}.$$

The middle term vanishes, since it includes Firm 1's first-order condition for its profit-maximizing price in the PCE. So,

$$\frac{d\Pi(\lambda)}{d\lambda} = -(\Pi_O(\lambda) - \Pi_I(\lambda)) + (1-\lambda)\left(g(0)p(\lambda)p'(\lambda)\right).$$

The first term is strictly negative, since opt-outs are more profitable than opt-ins, as confirmed below:

$$\begin{aligned} \Pi_O(\lambda) - \Pi_I(\lambda) &= p(\lambda)(G(p(\lambda)) - G(0)) - \int_0^{p(\lambda)-A} y\ dG(y) \\ &= \int_0^{p(\lambda)-A} (p(\lambda) - y)\ dG(y) > 0. \end{aligned}$$

Then because $p'(\lambda)$ is negative if captive demand is concave (and strictly negative if captive demand is strictly concave), we have $d\Pi(\lambda)/d\lambda < 0$, as claimed. □

Proof of Proposition 4

Part (i) is proved in Lemma 6. The proof of part (ii) follows.

Preliminaries

Define the following: let $\bar{y} = \max\left(p^{NT}\ln\left(\frac{2p^{NT}}{A}\right), p^{NT}-A\right)$. Let

$$\gamma_\rho(y) = \left(p^{NT}-A-\rho y\right)\left(1-\rho\frac{y}{p^{NT}}\right)^{\frac{1-\rho}{\rho}} + \frac{p^{NT}}{1+\rho} - \frac{p^{NT}}{1+\rho}\left(1-\rho\frac{y}{p^{NT}}\right)^{\frac{1+\rho}{\rho}} + A, \text{ and}$$

$$\gamma_0(y) = \lim_{\rho\to 0}\gamma_\rho(y) = p^{NT} + A\left(1-e^{-y/p^{NT}}\right)$$

Noting that $\gamma_\rho(y)$ converges uniformly to $\gamma_0(y)$ for $y \in \left[p^{NT}-A, \bar{y}\right]$, choose ρ_1 such that $\left|\gamma_\rho(y)-\gamma_0(y)\right| < \frac{A}{2}\left(1-e^{-\left(p^{NT}-A\right)/p^{NT}}\right)$ holds for all $\rho \le \rho_1$ and $y \in \left[p^{NT}-A, \bar{y}\right]$. (The right-hand side is strictly positive because $A < p^{NT}$.) Let $\rho_2 = \frac{1}{2}\frac{A}{p^{NT}-A}$, and set $\bar{\rho} = \min(\rho_1, \rho_2)$. Suppose that captive demand is $\bar{\rho}$-convex. Let (λ^*, p^*) be an equilibrium under regime OI, when opt-in is permitted.

Special cases: all consumers opt in, or all consumers opt out.

If $\lambda^* = 1$, then all consumers may be targeted, and the Proof of Proposition 2 applies (*a fortiori*, because we now have privacy costs that are avoided under regime NT). If $\lambda^* = 0$, then no consumers opt in, and the OI and NT outcomes are identical.

Interior equilibrium with opt-in

Henceforth, assume the regime OI equilibrium is interior, $\lambda^* \in (0,1)$. Then, the list price $p^* \in \left(p^{NT}, p^T\right)$ satisfies the equilibrium condition:

$$\left(1-\lambda^*\right)\left(1-G(0)-p^*g(0)\right) + \lambda\left(1-G\left(p^*-A\right)-Ag\left(p^*-A\right)\right) = 0.$$

Using $p^{NT} = (1-G(0))/g(0)$, this equilibrium condition may be rearranged as:

$$\left(1-\lambda^*\right)\left(p^*-p^{NT}\right) = \lambda\frac{p^{NT}}{1-G(0)}\left(1-G\left(y^*\right)-Ag\left(y^*\right)\right). \quad \text{(A.3)}$$

Recast the consumer surplus comparison in terms of favourite-equivalent prices

Some consumers bear privacy costs in the OI equilibrium, but none do under regime NT. Therefore, showing that consumers also face a higher average favourite-equivalent price under OI than under NT is sufficient to prove the claim of the proposition. So, the goal is to establish $\overline{EP}^{OI} > \overline{EP}^{NT} = p^{NT}$, where $\overline{EP}^{OI} = (1-\lambda^*)\overline{EP}^{O} + \lambda^*\overline{EP}^{I}$, where $\overline{EP}^{O} = p^*$ and $\overline{EP}^{I}$ are the average favourite-equivalent prices in the OI equilibrium for consumers who opt out or in, respectively.

Claim 1: $\overline{EP}^{OI} > \overline{EP}^{NT}$ holds if $y^* \ge \bar{y}$.

Because $\overline{EP}^{O} = p^* > \overline{EP}^{NT} = p^{NT}$, it suffices to establish that $\overline{EP}^{I} \ge p^{NT}$. By the same bounding argument as in Proposition 2, the expected favourite-equivalent price for consumers who opt in satisfies:

$$\begin{aligned}\overline{EP}^{I} \ge \int_0^{y^*}\frac{1-G(y)}{1-G(0)}dy + A &\ge \frac{p^{NT}}{1+\bar{\rho}} - \frac{p^{NT}}{1+\bar{\rho}}\left(1-\rho\frac{y^*}{p^{NT}}\right)^{\frac{1+\rho}{\rho}} + A \\ &\ge \frac{p^{NT}}{1+\bar{\rho}} - \frac{p^{NT}}{1+\bar{\rho}}e^{-y^*/p^{NT}} + A \\ &\ge \frac{p^{NT}-\frac{A}{2}}{1+\bar{\rho}} + A \ge p^{NT}.\end{aligned}$$

The first step applies Lemma 5.ii, then integrates. The second step uses the fact that $(1-\rho x)^{\frac{1+\rho}{\rho}} < (1-\rho x)^{\frac{1}{\rho}} < e^{-x}$ for $x > 0$. The third step uses $y^* \ge \bar{y} \ge p^{NT}\ln\left(\frac{2p^{NT}}{A}\right)$, and the final step uses $\bar{\rho} \le \rho_2$.

Claim 2: $\overline{EP}^{OI} > \overline{EP}^{NT}$ holds if $y^* \le \bar{y}$.

Use (A.3) to write $(1-\lambda^*)\overline{EP}^{O} = \lambda\frac{p^{NT}}{1-G(0)}(1-G(y^*)-Ag(y^*)) + (1-\lambda)p^{NT}$. Substitute this into $\overline{EP}^{OI}$ to establish that $\overline{EP}^{OI} > \overline{EP}^{NT}$ is equivalent to the inequality

$$\frac{p^{NT}}{1-G(0)}\left(1-G\left(y^*\right)-Ag\left(y^*\right)\right) + \overline{EP}^{I} > p^{NT}$$

and (using the bound on $\overline{EP}^I$ from Claim 1) a sufficient condition is $\gamma(y^*) > p^{NT}$,

$$\text{where } \gamma(y^*) := \frac{p^{NT}}{1-G(0)}\left(1-G(y^*)-Ag(y^*)\right)+\int_0^{y^*}\frac{1-G(y)}{1-G(0)}dy+A.$$

But by applying Lemma 5, we have $\gamma(y^*) \geq \gamma_{\bar{\rho}}(y^*)$. (For the first term, note that $1-G(y^*)-Ag(y^*) = (1-G(y^*))\left(1-\frac{A}{\mu(y^*)}\right)$, apply both parts of the lemma, and simplify.) Then by the construction of $\bar{\rho}$, we have

$$\gamma(y^*) > \gamma_0(y^*)-\frac{A}{2}\left(1-e^{-\left(p^{NT}-A\right)/p^{NT}}\right) \geq \gamma_0\left(p^{NT}-A\right)-\frac{A}{2}\left(1-e^{-\left(p^{NT}-A\right)/p^{NT}}\right) > p^{NT}$$

as claimed, where the middle step follows because $p^* \geq p^{NT}$ and so $y^* \geq p^{NT}-A$.

Summary Claims 1 and 2 establish that if captive demand is $\bar{\rho}$-convex, then at any interior equilibrium under regime OI, $\overline{EP}^{OI} > \overline{EP}^{NT}$ holds, and therefore that consumer surplus is lower under regime OI than it would be under regime NT.

REFERENCES

ACEMOGLU, D., MAKHDOUMI, A., MALEKIAN, A. and OZDAGLAR, A. (2022), "Too Much Data: Prices and Inefficiencies in Data Markets", *American Economic Journal: Microeconomics*, **14**, 218–256.

ACQUISTI, A., TAYLOR, C. and WAGMAN, L., (2016), "The Economics of Privacy", *Journal of Economic Literature*, **54**, 442–492.

ALI, N., LEWIS, G. and VASSERMAN, S. (2021), "Voluntary Disclosure and Personalized Pricing" (NBER Working Paper 26592).

ANDERSON, S., BAIK, A. and LARSON, N. (2015), "Personalized Pricing and Advertising: An Asymmetric Equilibrium Analysis", *Games and Economic Behavior*, **92**, 53–73.

ANDERSON, S., BAIK, A. and LARSON, N. (2019), "Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts" (CEPR Working Paper 15169).

ANDERSON, S. and DE PALMA, A. (1988), "Spatial Price Discrimination with Heterogeneous Products", *Review of Economic Studies*, **55**, 573–592.

ANDERSON, S., DE PALMA, A. and THISSE, J.-F. (1989), "Spatial Price Policies Reconsidered", *The Journal of Industrial Economics*, **38**, 1–18.

ANDERSON, S., DE PALMA, A. and THISSE, J.-F. (1992), *Discrete Choice Theory of Product Differentiation* (Cambridge, MA, USA: MIT Press).

BAGNOLI, M. and BERGSTROM, T. (2005), "Log-Concave Probability and Its Applications", *Economic Theory*, **26**, 445–469.

BELLEFLAMME, P. and VERGOTE, W. (2016), "Monopoly Price Discrimination and Privacy: The Hidden Cost of Hiding", *Economics Letters*, **149**, 141–144.

BERGEMANN, D. and BONATTI, A. (2019), "Markets for Information: An Introduction", *Annual Review of Economics*, **11**, 85–107.

BERGEMANN, D., BONATTI, A. and GAN, T. (2022), "The Economics of Social Data", *RAND Journal of Economics*, **53**, 263–296.

BESANKO, D., DUBÉ, J.-P., AND GUPTA, S. (2003), "Competitive Price Discrimination Strategies in a Vertical Channel Using Aggregate Retail Data", *Management Science*, **49**, 1121–1138.

BESTER, H. and PETRAKIS, E. (1995), "Price Competition and Advertising in Oligopoly", *European Economic Review*, **39**, 1075–1088.

BESTER, H. and PETRAKIS, E. (1996), "Coupons and Oligopolistic Price Discrimination", *International Journal of Industrial Organization*, **14**, 227–242.

BOURREAU, M. and DE STREEL, A. (2018), "Personalised Pricing in the Digital Era–Note by Marc Bourreau and Alexandre de Streel, OECD" https://one.oecd.org/document/DAF/COMP/WD(2018)150/en/pdf.

BRAHIM, N., LAHMANDI-AYED, R. and LAUSSEL, D. (2011), "Is Targeted Advertising Always Beneficial?", *International Journal of Industrial Organization*, **29**, 678–689.

BUTTERS, G. 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.

CHOI, J. P., JEON, D.-S. and KIM, B.-C. (2019), "Privacy and Personal Data Collection with Information Externalities", *Journal of Public Economics*, **173**, 113–124.

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

[EC 2018] EUROPEAN COMMISSION (2018), "Consumer Market Study on Online Market Segmentation Through Personalised Pricing/Offers", https://ec.europa.eu/files/aid_development_cooperation_fundamental_rights/aid_and_development_by_topic/documents/synthesis_report_online_personalisation_study_final_0.pdf.

ESTEBAN, L., GIL, A. and HERNÁNDEZ, J. (2001), "Informative Advertising and Optimal Targeting in a Monopoly", *The Journal of Industrial Economics*, **49**, 161–180.
ESTEVES, R.-B. and RESENDE, J. (2016), "Competitive Targeted Advertising with Price Discrimination", *Marketing Science*, **35**, 576–587.
[EOP 2015] EXECUTIVE OFFICE OF THE PRESIDENT OF THE UNITED STATES (2015), "Big Data and Differential Pricing", https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf.
GALEOTTI, A. and MORAGA-GONZÁLEZ, J. (2008), "Segmentation, Advertising, and Prices", *International Journal of Industrial Organization*, **26**, 1106–1119.
GROSSMAN, G. and SHAPIRO, C. (1984), "Informative Advertising with Differentiated Products", *Review of Economic Studies*, **51**, 63–81.
HIDIR, S. and VELLODI, N. (2021), "Privacy, Personalization, and Price Discrimination", *Journal of the European Economic Association*, **19**, 1342–1363.
HILLMAN, A. L. and RILEY, J. G. (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.
ICHIHASHI, S. (2020), "Online Privacy and Information Disclosure by Consumers", *American Economic Review*, **110**, 569–595.
ICHIHASHI, S. (2021), "The Economics of Data Externalities", *Journal of Economic Theory*, **196**, 105316.
IYER, G., SOBERMAN, D. and VILLAS-BOAS, J. (2005), "The Targeting of Advertising", *Marketing Science*, **24**, 461–476.
KOÇAŞ, C. and KIYAK, T. (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. (2019), "The Value of Personal Information in Online Markets with Endogenous Privacy", *Management Science*, **65**, 1342–1362.
MYATT, D. and RONAYNE, D. (2019), "A Theory of Stable Price Dispersion" (University of Oxford Discussion Paper 873).
NARASIMHAN, C. (1988), "Competitive Promotional Strategies", *Journal of Business*, **61**, 427–449.
[OECD-SEC 2018] OECD (2018), "Personalised Pricing in the Digital Era–Background Note by the Secretariat" (28 November 2018). https://one.oecd.org/document/DAF/COMP(2018)13/en/pdf.
[OECD-UK 2018] OECD (2018), "Personalised Pricing in the Digital Era–Note by the United Kingdom" (28 November 2018). https://one.oecd.org/document/DAF/COMP/WD(2018)127/en/pdf.
PERLOFF, J. and SALOP, S. (1985), "Equilibrium with Product Differentiation", *Review of Economic Studies*, **52**, 107–120.
PRAT, A. and VALLETTI, T. (2021), "Attention Oligopoly", *American Economic Journal: Microeconomics*, forthcoming.
ROY, S. (2000), "Strategic Segmentation of a Market", *International Journal of Industrial Organization*, **18**, 1279–1290.
SHAFFER, G. and JOHN ZHANG, Z. (1995), "Competitive Coupon Targeting", *Marketing Science*, **14**, 395–416.
SHAFFER, G. and JOHN ZHANG, Z. (2002), "Competitive One-to-One Promotions", *Management Science*, **48**, 1143–1160.
SHARKEY, W. W. and SIBLEY, D. (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**, 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 (Amsterdam, Netherlands: North-Holland) 541–562.
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.
[UK-OFT 2013] U.K. OFFICE OF FAIR TRADING (2013), "Personalised Pricing: Increasing Transparency to Improve Trust" available at http://webarchive.nationalarchives.gov.uk/20140402142426/http:/www.oft.gov.uk/shared_oft/marketswork/personalised-pricing/oft1489.pdf.
[UK-CMA 2018] U.K. COMPETITION AND MARKETS AUTHORITY (2018), "Pricing Algorithms: Economic Working Paper on the Use of Algorithms to Facilitate Collusion and Personalised Pricing" (8 October 2018), available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/746353/Algorithms_econ_report.pdf.
VARIAN, H. R. (1980), "A Model of Sales", *American Economic Review*, **70**, 651–659.
WHITE, T. B., ZAHAY, D. L., THORBJØRNSEN, H. and SHAVITT, S. (2008), "Getting Too Personal: Reactance to Highly Personalized Email Solicitations", *Marketing Letters*, **19**, 39–50.