

TARGETED ADVERTISING AND ADVERTISING AVOIDANCE

JUSTIN P. JOHNSON[†]

October 29, 2009

ABSTRACT. I examine how the increasing ability of firms to target their advertisements to particular individuals influences market outcomes when consumers have access to advertising avoidance tools.

Today firms possess an unprecedented and rapidly improving ability to discover details about individuals and reach them with advertising based on this information. For instance, as many readers know, an individual's web browsing activities may be tracked via information stored in his browser as long as he "participates" by not regularly deleting it, allowing ads to be customized to him upon visiting a web site.¹ Less well known may be that new technology allows Internet Service Providers (ISPs) to target ads by directly tracking which websites consumers visit, meaning that not participating is not an option for consumers.² Similarly, advances allow cable television operators to deliver highly customized ads to viewers.³ For those who don't utilize such services, it may be surprising to learn that the many millions of users of social networking tools such as Facebook and MySpace,⁴ and of the popular email services of firms such as Yahoo and Google, allow ads to be served based on data generated by using the networking tools or gleaned from the contents of emails sent or received. Dunkin' Donuts is testing new technology that will change the ad displayed at the counter based on the appearance of the customer.⁵ Supermarkets and other retail outlets are beginning to offer coupons to consumers at the checkout aisle based on their previously recorded shopping habits.⁶ There are many more examples.⁷

I thank Nuffield College, Oxford University, for graciously allowing me to visit; part of this paper was written there. I also thank participants at various seminars, including Simon Anderson.

[†]Johnson Graduate School of Management, Cornell University. Email: jjp25@cornell.edu.

¹Data such as which websites were visited and what content was viewed are trackable via small pieces of information called "cookies" that are stored in browsers.

²"Watching What You See on the Web," *Wall Street Journal*, December 6, 2007. The technology being deployed by ISPs is potentially more invasive than that using cookies. The reason is that, heuristically, the use of cookies is most effective when many websites work together to share information, and also requires that users don't delete their cookies too often, whereas the new technology requires neither.

³"Cable TV Puts New Spin on Ads," *Wall Street Journal*, May 16, 2007. This article reports, for example, that "During a commercial break for 'Lost,' a young couple watching TV might see an ad for the latest cellphone, while at the same time their nextdoor neighbors with children may see a diaper commercial."

⁴Facebook claims 100 million active accounts.

⁵"The Ad Changes with the Shopper in Front of it," *Wall Street Journal*, August 21, 2008.

⁶"Personalized Store Ads Take Off," *Wall Street Journal*, October 23, 2008.

⁷Advertising sent to mobile phones is on the rise and will be based on the consumers' precise location as determined by embedded Global Positioning System devices. Ironically, the digital video recorder (DVR) manufacturer TiVo is exploring ways of tailoring advertisements to customers shown while these customers attempt to skip advertisements.

Because advertising is not free and consumers are heterogenous in their responses to it, real firms focus their ads on a subset of the great multitude of outlets, suggesting that an enhanced ability to target advertisements may be highly valuable. This especially may be true for niche firms who otherwise find themselves effectively locked out of traditional ad channels, unable cost-effectively to reach their target market. The idea that increased targeting and search capabilities helps niche firms and consumers is often referred to as the “long tail of the internet.”⁸

While firms may relish additional information about consumers, some consumers consider advertising annoying on average and see increased targeting as a further violation of their privacy. Consequently, a variety of techniques exist that allow consumers to avoid advertising, from the time-tested ignoring of ads or changing of the channel during commercials to recent innovations such as skipping past adverts,⁹ blocking online advertisements, filtering email, or subscribing to do-not-call, do-not-mail or do-not-track programs.

In this paper I construct and analyze a pair of models that captures the details of real-world advertising and privacy issues described above, namely targeted advertising by firms, advertising avoidance by consumers, and the long tail of the internet. The main conclusions are as follows. First, improved targeting raises the profits of all firms. This is despite the facts that consumers endogenously adjust their advertising avoidance decisions and that competition may increase for some ad inventory.

Consumers do not fare as well. While they may gain by witnessing more-relevant ads, there are two negative consequences of improved targeting that may lead consumers to prefer less of it. One is that the number of advertisements received may increase, which is bad for consumers since in equilibrium they tend to disdain marginal advertisements although they may appreciate inframarginal ones. The other reason is that, rather counterintuitively, improved information accuracy may actually lead consumers to receive ads that they prefer less than those received when targeting is less precise.

Second, consumers may underutilize advertising avoidance technologies from a (second-best) social standpoint. This is somewhat unexpected as there is also a force pushing for overuse of such technologies, which is that consumers ignore firms’ gains from trade. This externality implies that, even if ad revenues pay for content, there is no *a priori* reason to think such technologies cause social harm.¹⁰

Third, allowing for graduated information quality (as opposed to assuming either perfect secrecy or complete lack of secrecy over consumer information) provides useful insights. For

⁸The heightened ability to target consumers is also often listed as a main driver of an ongoing shift of advertising budgets away from traditional outlets such as television and towards new media.

⁹When content has been recorded on a DVR.

¹⁰Even the Federal Trade Commission (FTC), in its charge of investigating consumer privacy, has emphasized that consumers overlook the many benefits of targeted advertising, such as ad revenue sponsoring content.

instance, consumers may not have monotone preferences for increased targeting; indeed their utility is quasiconvex (or “U-shaped”) in the accuracy of targeting.

That utility may be quasiconvex highlights the intriguing possibility that consumer attitudes towards advertising in general and targeted advertising in particular might change dramatically as technology advances. In particular, for technology that only allows somewhat weak targeting, incremental improvements in it may only further frustrate consumers, encouraging them to block ads, complain about ads, or withhold information that could be used to target them. However, it may be that as targeting technologies pass some threshold and become fairly sophisticated, more and more consumers will welcome such targeting and improvements in it, and see advertising on average as less of a nuisance (or even as something to be appreciated). Since currently rapid advancements are being made in targeting technologies, this theoretical point is potentially of immediate relevance.

The fourth main contribution is the identification of a connection between two notions of privacy discussed in the literature, secrecy and autonomy (see Hirshleifer (1980)). Secrecy refers to the suppression of data that informs about a consumer, while autonomy refers to the right not to be disturbed.

I show that the demand for autonomy (or, the level of ad avoidance) is tied to the demand for secrecy (or, the accuracy of information). In particular, since utility is quasiconvex in accuracy, a low preference for secrecy is associated with relatively high utility and hence a low demand for autonomy. On the other hand, a low demand for autonomy may also be consistent with a high demand for secrecy, again due to the quasiconvexity of utility.

1. LITERATURE REVIEW

My work is related to the literatures on privacy, advertising, and targeted price offers. The early work in Economics on privacy, namely Posner (1981) and Stigler (1980), emphasized that a free flow of information rather than privacy would enhance economic efficiency.¹¹ Recently, a number of authors have identified important circumstances not covered by this work, for which the earlier conclusions are not valid or the analysis simply insufficient to answer the newly posed questions. For instance, the externalities that are critical in my model (and certainly in some work prior to mine) are not present in this earliest work, nor is the interplay between firm decisions on one side of the market (how much to advertise) and consumer decisions on the other (whether to block ads or not). See Hermalin and Katz (2006) and Hui and Png (2006) for in depth discussion of the literature.

Hirshleifer (1971) observed that parties may overinvest in information acquisition, where such information serves no socially productive role but instead redistributes wealth. While

¹¹Privacy is viewed as distortive because, for example, while a worker may gain from concealing a condition relevant to his productivity, this may lower total productivity if he is inefficiently assigned to a task.

Hirshleifer emphasized circumstances of individuals or firms predicting future states of nature for private gain, the basic idea applies to firms seeking information regarding consumers so as to more profitably price discriminate, which may be purely redistributive. A number of papers have explored this idea of price targeting in either monopolistic or oligopolistic settings, such as Bester and Petrakis (1996), Fudenberg and Tirole (2000), Taylor (2004), and Acquisti and Varian (2005).

The advertising literature has long focused on the idea that price competition is influenced by advertising, beginning with Butters (1977) and Grossman and Shapiro (1984). A number of recent papers look at targeted advertising more particularly, again in the context of how it influences strategic pricing behavior; see Akçura and Srinivasan (2005), Iyer, Soberman, and Villas-Boas (2005), Gal-Or, Gal-Or, May, and Spangler (2006), Esteban and Hernandez (2007) and Galeotti and Moraga-González (2008).¹² Unlike these papers, firms in my model do not compete in the product market, although they may compete for ad inventory.

Chen and He (September, 2006) and Athey and Ellison (May, 2008) examine a phenomenon related to targeted advertising, which is position auctions for advertisements that appear on online search engines. An important consideration is modeling how consumers behave when presented a list of potential links upon completing an online search.

The work that is closest to mine is that by Van Zandt (2004), Anderson and de Palma (2007), and Anderson and Gans (2006). Van Zandt (2004) provides the seminal model of information overload, carefully modeling the decision of firms about whom to send messages (for example, ads). An important observation is that firms may send too many messages, thereby generating a negative externality by crowding out other senders' messages. A big difference between his work and mine is that his is primarily focused on the decision problem facing senders, so that receivers are passive. Anderson and de Palma (2007) consider the consumer side more carefully, and also examine a number of potential corrective measures to move equilibrium closer to efficiency. Anderson and Gans (2006) focus on the interplay between advertising avoidance by consumers and advertising by firms, as I do. However, their main point is that in the case of broadcast television, increased use of TiVo (an advertising avoidance technology) by some consumers may lead to more advertising for groups not using TiVo. A major conceptual difference between my work and these others is that I am interested in the consequences of changing targeting opportunities, whereas this is mostly absent from the other work (targeting is most central in Van Zandt (2004)). Hann, Hui, Lee, and Png (2008) and Falkinger (2008) are also related to my work in that they both incorporate consumer defensive responses to advertising.

¹²See Meurer and Stahl (1994), Moscarini and Ottaviani (2001) and Johnson (2005) for work on how information influences perceived product differentiation and hence competition. Johnson and Myatt (2006) consider the effects of information release, with an emphasize on monopoly.

2. A MODEL OF TARGETED ADVERTISING AND ADVERTISING AVOIDANCE

There is a unit mass of consumers and a unit mass of firms. For any given consumer-firm pair there is some independently realized probability \tilde{q} distributed according to F that there is a match between the consumer and the firm's product, where F is continuously differentiable and strictly increasing on its support, which in turn is an interval contained in $[0, 1]$. Existence of such a match means that, were the consumer aware of the product, he and the firm would engage in trade generating surplus v_B for the consumer (buyer) and v_S for the firm (seller). Lack of a match means no gains are generated from trade. Note that firms are not competitors.

Consumers only become aware of a firm's product if they receive an advertisement from that firm. Firms in turn must decide to which potential buyers advertisements will be sent. Any given firm receives a signal q which is accurate with probability a , and is otherwise pure noise, so that if such a consumer were to receive an advertisement from this firm there would be trade with probability $aq + (1 - a)\hat{q}$, where $\hat{q} = E\tilde{q}$. All signals are conditionally independent across the different consumer-firm pairs. Further, there is a probability β that an advertisement that is sent is not actually observed by the consumer. Thus, β represents the probability that a given consumer will block or otherwise avoid seeing the advertisement. Sending an ad, regardless of whether it is received or not, costs the firm $c_S > 0$ (this assumption is discussed extensively at the end of Section 3).

A firm optimally advertises to a consumer of signal q if and only if

$$c_S \leq (1 - \beta)v_S[aq + (1 - a)\hat{q}].$$

If q is the smallest signal that is advertised to by this firm, total expected profits are

$$-c_S(1 - F(q)) + (1 - \beta)v_S \left[a \int_q^1 \tilde{q} dF(\tilde{q}) + \hat{q}(1 - a)(1 - F(q)) \right].$$

Note that it is equivalent and somewhat easier for the analysis to assume that a firm chooses the number or mass μ of advertisements, where it is understood these are sent to the consumers exhibiting the highest signals. Defining $q(\mu) = F^{-1}(1 - \mu)$ to be the signal of the marginal consumer receiving an advertisement, total firm profits are

$$\pi = -c_S\mu + (1 - \beta)v_S \left[a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1 - a)\hat{q}\mu \right], \quad (1)$$

while the optimal (if interior) μ satisfies $c_S = (1 - \beta)v_S[aq(\mu) + (1 - a)\hat{q}]$.

I now complete the description of the demand side. Suppose that a consumer receives all ads targeted towards him. Because of the independence assumptions imposed above on the distribution of matches and signals, all consumers receive the same number of ads. Each ad

causes a nuisance, summarized by the cost $c_B > 0$, so that the expected utility of a consumer receiving μ ads is

$$U = -c_B\mu + v_B \left[a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1-a)\hat{q}\mu \right]. \quad (2)$$

As discussed previously, there are many countermeasures by which consumers might avoid advertisements. I adopt this possibility into the model in a simple manner by supposing that avoiding advertisements is a zero-one decision, so that a consumer chooses to either avoid all or to receive all ads. A consumer who chooses to avoid ads receives a flat payoff of $\theta < 0$, where θ is distributed on $(-\infty, 0]$ according to G , a continuously differentiable and strictly increasing function. Thus there is a cost of avoiding ads that could capture a number of forces, for example the costs of changing channels (and remembering to come back before the original show starts again), of installing or correctly tuning blocking technologies, of subscribing to do-not-call style programs, missing incorrectly blocked emails, or of the mental energy associated with “tunnel vision” to ignore adverts.¹³

Taking the advertising volume μ as given, consumers choose to receive either U or θ , where U is as in Equation (2). Hence the proportion of consumers who block is

$$\beta = \Pr[\text{Block}] = 1 - G(U) < 1.$$

While β may equal zero for a range of μ , it is strictly increasing in μ once it is positive.

Consumers and firms simultaneously choose whether to block and the level of advertising, respectively, given what the other side of the market is doing. Note that because firms face the same objective function and this function has a unique maximizer given β , all firms choose the same advertising level μ . Hence, an equilibrium is composed of a pair (β^*, μ^*) .

Proposition 1 (Existence and Uniqueness). *There exists a unique Nash equilibrium of the game in which each firm chooses an advertising volume μ and consumers simultaneously decide whether to block advertisements.*

Proof. In a slight abuse of notation, let $\mu(\beta) \in [0, 1]$ be the (uniquely defined) best-response of firms given a level of blocking β and let $\beta(\mu) \in [0, 1]$ be the (uniquely defined) proportion of consumers who choose to block ads given μ . Define $h(\tilde{\mu}) = \mu[\beta(\tilde{\mu})] - \tilde{\mu}$. Then (β^*, μ^*) is an equilibrium if and only if $h(\mu^*) = 0$ and $\beta^* = \beta(\mu^*)$.

It is straightforward to show that h is continuous, and since $h(0) \geq 0$ and $h(1) \leq 0$, there must be some value μ^* such that $h(\mu^*) = 0$. Defining $\beta^* = \beta(\mu^*)$ therefore ensures that (β^*, μ^*) is an equilibrium; that is, there exists an equilibrium.

¹³Similarly, one might instead imagine that consumers have heterogeneous nuisance costs. Using the approach I do turns out to be simpler to analyze.

To verify uniqueness, suppose that both (β_1^*, μ_1^*) and (β_2^*, μ_2^*) are equilibria, with $\mu_1^* < \mu_2^*$ (if $\mu_1^* = \mu_2^*$ the equilibria coincide since $\beta(\mu)$ is a function). It must be that $\beta_2^* > \beta_1^*$ since $\beta(\mu)$ is non-decreasing, and if it were the case that $\beta_1^* = \beta_2^*$ then it would also be the case that $\mu_1^* = \mu_2^*$. But since $\mu(\beta)$ is non-increasing, $\beta_2^* > \beta_1^*$ implies $\mu_2^* \leq \mu_1^*$, a contradiction. ■

I will refer to the equilibrium (β^*, μ^*) whose existence is guaranteed by Proposition 1 as the “decentralized equilibrium.”

3. BASIC ANALYSIS: EXTERNALITIES AND WELFARE WHEN TARGETING IS FIXED

In this section I identify and explain the relevant externalities and forces at work in the context of welfare analysis, providing both first- and second-best welfare results. The two main results are that there may be too little advertising avoidance by consumers from a second-best perspective, although there is always too much from a first-best perspective. In this section I fix the ability to target consumers, summarized by the information accuracy parameter a , but allow both the levels of advertising and advertising avoidance to vary. Changes in a are covered thoroughly in Section 4.

Although my formal analysis below will proceed by identifying the relevant externalities by exploring first-best and then second-best analyses, there is another convenient way of categorizing these externalities, which is by whether they are direct or indirect. To see this observe that consumers impose direct externalities on firms when they choose to block ads or not, and similarly firms impose direct externalities on consumers when they choose to send advertisements or not. Furthermore, consumer actions have no direct impact on other consumers, and firm actions have no direct effect on other firms.

However, consumers indirectly impose externalities on other consumers and likewise firms indirectly impose externalities on other firms. The reason is that when consumers choose to block ads or not, they change the incentives of firms to advertise, which then influences other consumers. Similarly, when firms advertise or not, they change the incentives of consumers to block ads, which then influences other firms.

Ultimately, there are thus four externalities present in my model, which will be described in more detail below.

3.1. First-Best Analysis. In the first best, μ and β are chosen to maximize the sum of producer and consumer surplus. For notational simplicity, denote the expected volume of trade per consumer who doesn’t avoid ads by

$$w(\mu) = a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1 - a)\hat{q}\mu,$$

and recall that $\beta = 1 - G(\theta)$ where $\theta \leq 0$ is the marginal consumer blocking. Then, the first-best objective function is

$$S(\beta, \mu) = -c_S\mu + (1 - \beta)v_S w(\mu) + (1 - \beta)U(\mu) + \int_{\theta}^0 \tilde{\theta} dG(\tilde{\theta}).$$

The first two terms represent firm profits and the last two consumer utility.

Proposition 2. *In the first-best outcome, no consumers block advertisements ($\beta = 0$).*

Proof. If the proposition were false then at the optimum it would be the case that

$$\frac{\partial S}{\partial \beta} \geq 0 \iff -v_S w(\mu) - U(\mu) - \theta g(\theta) \frac{d\theta}{d\beta} = -v_S w(\mu) - U(\mu) + \theta \geq 0.$$

But since $\theta < 0$ this implies

$$0 > \theta \geq v_S w(\mu) + U(\mu) \Rightarrow 0 > (1 - \beta) [v_S w(\mu) + U(\mu)] \Rightarrow 0 > S(\beta, \mu).$$

Where the final inequality follows from the fact that $-c_S\mu$ and $\int_{\theta}^0 \tilde{\theta} dG(\tilde{\theta})$ are both negative. But if such were the case, then overall welfare could be raised by simply having no firms advertise and no consumers block, for in that case total welfare would equal zero. ■

Proposition 2 provides support for the (essentially) anti-privacy arguments of early writers such as Posner (1981) and Stigler (1980).¹⁴ For example, Posner argues that a worker who exercises privacy to conceal a condition relevant to productive efficiency makes himself better off but hurts his employer, at the ultimate loss of economic efficiency. This basic logic is at least superficially similar to that of Proposition 2, since consumers ignore the gains to trade for firms that are lost when they choose to avoid ads.

One must be cautious however, as there is an important sense in which the reasoning behind Proposition 2 differs from that of Posner, which is that firms also exert negative externalities on consumers here. What this second externality implies is that while the first-best outcome has $\beta = 0$, it's also true that firm advertising levels are not set too high. So, it certainly does not follow that removing consumers' option to avoid ads and letting firms advertise as they please would lead to an efficient outcome.

The intuition for Proposition 2 is as follows. If it were socially optimal for any consumer to avoid ads and receive payoff $\theta < 0$, then the total surplus generated by ads would be negative, contradicting the supposed optimality of that level of advertisements. Finally, note that this proposition does not ensure that consumers receive positive utility—it is perfectly possible that in the first-best outcome consumers receive negative utility. For example, this would tend to be the case if v_S were quite large and v_B were quite small compared to c_B .

¹⁴At least conceptually: those authors primarily were concerned about secrecy of personal information, whereas the decision to avoid ads is one of autonomy.

3.2. Decentralized Welfare Analysis. Here I consider the effect of small changes in advertising levels away from the decentralized equilibrium. This analysis highlights a third externality, which is a negative one imposed by firms on other firms.

The following proposition provides a simple condition for whether there is too much or too little advertising local to the decentralized equilibrium.¹⁵

Proposition 3. *Suppose that $\mu^* \in (0, 1)$. Then, beginning from the decentralized equilibrium, a small (industry-wide) decrease in μ raises social welfare if and only if $c_B/v_B > c_S/v_S$. Furthermore, if $c_B/v_B > c_S/v_S$, then both the payoffs to consumers U and firms π increase.*

Proof. Firm optimization implies

$$\frac{c_S}{(1 - \beta^*)v_S} = aq(\mu^*) + (1 - a)\hat{q}, \quad (3)$$

while the utility effect from the marginal ad is

$$-c_B + v_B[aq(\mu^*) + (1 - a)\hat{q}] \geq 0 \iff \frac{c_B}{v_B} \leq aq(\mu^*) + (1 - a)\hat{q} = \frac{c_S}{(1 - \beta^*)v_S}. \quad (4)$$

The final equality follows from Equation (3).

Now, suppose that $c_B/v_B > c_S/v_S$. Then it must also be that $c_B/v_B > c_S/[(1 - \beta^*)v_S]$. To see this, suppose not. Then by Equation (4), consumers would receive positive surplus from marginal ads, implying $U > 0$ and hence $\beta = 0$. But then I would have

$$\frac{c_B}{v_B} \leq \frac{c_S}{[(1 - \beta^*)v_S]} = \frac{c_S}{v_S} < \frac{c_B}{v_B},$$

a contradiction. Hence, as claimed $c_B/v_B > c_S/[(1 - \beta^*)v_S]$ so that by Equation (4) the marginal return to ads from a consumer's perspective is negative. It follows that a local decrease in μ from μ^* raises the payoffs to consumers. But, since firms are (privately) indifferent to placing the marginal ad or not, and since $\beta(\mu)$ is increasing, a decrease in μ lowers β and raises the payouts of firms, too.

The remainder of the proof involves the case where $c_B/v_B < c_S/v_S$. Using arguments similar to those above, this implies that consumers always place positive value on the marginal ads, and hence that $U > 0$ and $\beta = 0$. Hence, a small increase in the volume of ads has no effect on β , but increases consumer surplus and hence raises total welfare. ■

The first part of Proposition 3 shows that whether there is excessive advertising is completely determined by a simple comparison of consumers' value for trade (deflated by the nuisance cost of ads) and firms' value of trade (deflated by the cost of placing an ad). The intuition builds on the fact that the benefit from the marginal ad to either a firm or a consumer is

¹⁵Remember that the decentralized equilibrium is the one whose existence is guaranteed by Proposition 1. Also remember that (β^*, μ^*) denote the equilibrium values.

proportional to the marginal probability of trade $aq(\mu) + (1 - a)\hat{q}$; there is disagreement only over the benefit of trade and the costs, where this disagreement is summarized by differences between c_B/v_B and c_S/v_S . Since on the margin firms are indifferent about placing an additional advertisement, heuristically the overall welfare effect is determined by whether this disagreement has consumers undervaluing ads relative to firms or not.¹⁶

There are two further noteworthy implications of Proposition 3. The first clarifies the last part of the proposition, which is that when social welfare increases from a reduction in μ , it is driven by increases not only in consumer utility but also in firm profits. The reason for this is that a forced decrease in μ ameliorates the externality firms impose on other firms in the form of advertising “clutter.” More specifically, when firms choose their advertising levels they ignore the impact this has on advertising avoidance decisions. Since (when $c_B/v_B > c_S/v_S$) higher advertising levels raise the number of consumers who block ads, and since this lowers the profits of all firms, the externality imposed by firms on other firms is a negative one.

This externality is very closely related to one identified in Van Zandt (2004) and Anderson and de Palma (2007). In their analyses, it’s shown that firms may send too many messages to consumers, which is bad for social welfare because consumers only view a portion of messages sent and this means that some higher value messages are crowded out by lower value ones. This observation of theirs is an important one since it suggests that some regulation of or intervention in the market for advertising may benefit firms. In other words, the market for advertising typically isn’t efficiently priced. While my end result is identical, the logic differs somewhat. In particular, in my analysis the benefit to firms from a reduction in μ is that consumers respond by avoiding fewer advertisements. In other words, consumer response is the critical element of my result, which is not the case for the results of Van Zandt (2004) and Anderson and de Palma (2007).

The second further implication of Proposition 3 is that for welfare to improve with an increase in advertising requires $c_B/v_B < c_S/v_S$ and this implies that consumers benefit from additional advertising. Equivalently, it implies $U > 0$ and so no consumers wish to avoid advertisements. To the extent that one believes that consumers typically do not favor increased advertisements, one might feel it unlikely that realistic parameters satisfy this condition. Were that the case, an implication would be that both sides of the market would be better off with a reduction in advertising levels.

3.3. Second-Best Advertising Avoidance. Here I return to the question of whether consumer autonomy in being able to avoid ads is socially beneficial or not. As discussed following Proposition 2, the arguments of Posner (1981) and Stigler (1980) opposing consumer privacy have merit in the current setting, but as I cautioned earlier that proposition involves the first

¹⁶This is heuristic because it ignores the change in β that occurs, but as made clear in the proof that effect actually works to raise firm profits.

best and so is not clearly the most appropriate benchmark. Here I look at the second best in which advertising avoidance levels are set by a planner, given that firms will freely respond with their privately optimal advertising. I also discuss the fourth and final externality, which is one that consumers impose on other consumers. My main result is that in this second best, consumers may avoid too few advertisements.

Suppose that the number of consumers who block ads is set to maximize total welfare, given that each firm will adjust μ to maximize its individual profits. Denote this level of blocking by β^{**} , and recall that the fully decentralized values of β and μ are denoted by (β^*, μ^*) .¹⁷ Also recall that the probability of trade conditional on an ad being viewed is

$$w(\mu) = a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1 - a)\hat{q}\mu,$$

and that $\beta = 1 - G(\theta)$ where $\theta \leq 0$ is the marginal consumer blocking. Then, suppressing the dependence of μ and θ on β , the objective function is

$$S(\beta) = -c_S\mu + (1 - \beta)v_S w(\mu) + (1 - \beta)U(\mu) + \int_{\theta}^0 \tilde{\theta} dG(\tilde{\theta}).$$

(Note that although this objective function might appear the same as that from the first-best analysis, this is an artifact of my suppressing the dependence of μ on β .)

While the following proposition admittedly provides a condition on endogenous variables, I feel it is nonetheless informative in identifying the relevant tradeoffs. See Corollary 1 below for a result in terms of exogenous variables.

Proposition 4. *Suppose that the level of blocking β^{**} is set to maximize total welfare, given that firms choose their individually optimal level of advertising in response to β^{**} . Assume the decentralized equilibrium exhibits $\beta^* > 0$ and $\mu^* \in (0, 1)$, and that*

$$\frac{v_B}{v_S} \geq \frac{q''(\mu)}{q'(\mu)}.$$

*Then, $\beta^{**} > \beta^*$ if and only if*

$$(1 - \beta^*)\mu'(\beta^*) \frac{\partial U(\mu^*)}{\partial \mu} > v_S w(\mu^*). \quad (5)$$

Proof. Using the envelope theorem on the terms involving profits, and the fact that $\mu'(\beta) = w(\mu)/[(1 - \beta)aq'(\mu)]$, it follows that

$$S'(\beta) = -v_S w(\mu) + \frac{w(\mu)}{aq'(\mu)} \frac{\partial U(\mu)}{\partial \mu} + [\theta - U(\mu)].$$

¹⁷To ensure that it is possible to have a non-degenerate result, I assume that $\beta^* > 0$ in Proposition 4. An implication of this is that $c_B/v_B > c_S/v_S$.

It is useful to decompose $S'(\beta)$ into two functions. Define

$$h(\beta) = -v_S w(\mu) + \frac{w(\mu)}{a q'(\mu)} \frac{\partial U(\mu)}{\partial \mu},$$

so that

$$S'(\beta) = h(\beta) + [\theta - U(\mu)].$$

Now, suppose that it could be shown that the following two statements were true: (i), $\theta - U(\mu)$ is strictly decreasing in β and equal to zero at β^* ; (ii), $h(\beta) = 0$ implies $h'(\beta) < 0$. Then, keeping in mind that $h(\beta) = -v_S w(\mu) + (1 - \beta)\mu'(\beta) \frac{\partial U(\mu)}{\partial \mu}$, the proposition would follow. The reason is that if $h(\beta^*) > 0$, then by (ii) for all $\beta \leq \beta^*$ it must be that $h(\beta) > 0$ and then using (i) it must be that $S'(\beta) > 0$ for all $\beta \leq \beta^*$; this implies $\beta^{**} > \beta^*$. Similarly, if $h(\beta^*) < 0$ then by (i) it must be that $h(\beta) < 0$ for all $\beta \geq \beta^*$ and then, using (i) it would be the case that $S'(\beta) < 0$ for all $\beta \geq \beta^*$; this implies $\beta^{**} < \beta^*$.

Thus, to complete the proof I merely must show that (i) and (ii) are true. The first claim follows from the facts that θ is decreasing in β by definition, that $\theta = U$ at β^* by the definition of β^* , and that U is increasing in β in the relevant region. To show this last part, note that it cannot be optimal to choose β^{**} such that $c_B/v_B < c_S/[(1 - \beta^{**})v_S]$ for to do so would imply that consumers prefer more advertisements (as shown in the proof of Proposition 3) so that $U > 0$ so that both firms and consumers would prefer less blocking. (Note that the assumption $\beta^* > 0$ guarantees that $c_B/v_B < c_S/v_S$.) Hence, the region of relevance is where β is such that consumers place negative value on marginal ads, which means that U is increasing in β since $\mu'(\beta) < 0$.

To prove (ii), suppose that $h(\beta) = 0$ for some β . Then

$$h'(\beta) = \mu'(\beta) w(\mu) \left[v_B - v_S \frac{q''(\mu)}{q'(\mu)} \right],$$

where I have eliminated some terms by using the facts that $h(\beta) = 0$ and $w''(\mu) = a q'(\mu)$. Now, since $\mu' < 0$ and $w > 0$, $h'(\beta) < 0$ at this point from an assumption of the proposition. Hence (ii) holds, and the proposition follows. ■

The intuition behind Proposition 4 involves two competing externalities. The first is a negative one imposed on firms by consumers, equaling the lost surplus when an ad is not viewed, $v_S w(\mu^*)$. The second is imposed by consumers on other consumers and is a positive externality. Namely, when an additional consumer elects not to view advertisements, he benefits all consumers who do view ads by reducing the volume of ads. The reason ad volume falls is that by blocking ads the cost-effectiveness of advertising falls, and this is beneficial to consumers since it is the advertisements that are least likely to be relevant that are removed from circulation first. Moreover, such marginal ads are always valued negatively

by consumers.¹⁸ Hence the increase in the utility of consumers overall from a small increase in β is $(1 - \beta^*)\mu'(\beta^*)\partial U(\mu^*)/\partial\mu > 0$.

Of course, Proposition 4 is a statement not only about whether a local change in β would be beneficial, but rather about the globally optimal second-best level β . Also, as a technical aside, the proof does not involve showing concavity of the objective function.¹⁹

3.3.1. *Discussion.* The positive externality associated with avoiding advertisements constitutes an important element absent from the debate over the overall impact of avoidance-enabling technologies. While the nuisance costs to consumers from ads are often noted, there at times appears to be a stronger emphasis on the benefits provided by ad revenue, for example in the form of content. Indeed, even the FTC emphasizes the unrecognized (to consumers) positives, in its document “Online Behavioral Advertising: Moving the Discussion Forward to Possible Self-Regulatory Principles”, writing

“First, while behavioral advertising provides benefits to consumers in the form of free web content and personalized ads that many consumers value, the practice itself is largely invisible and unknown to consumers. The benefits include, for example, access to newspapers and information from around the world, provided free because it is subsidized by online advertising; tailored ads that facilitate comparison shopping for the specific products that consumers want; and, potentially, a reduction in ads that are irrelevant to consumers interests and that may therefore be unwelcome. Although many consumers value these benefits, few appear to understand the role that data collection plays in providing them.”

But casual inspection of Equation (5) suggests that too little blocking is a plausible outcome, suggesting that ensuring consumer access to such technologies may be a valid mandate for regulators. Of course, I do not formally model content provision. Nonetheless, Proposition 4 indicates that the overall picture is more complicated than simply asserting that consumers don’t recognize “all the good things that advertising brings them.”

The following Corollary provides precise conditions under which there is too little avoidance in equilibrium. There is under-blocking of ads when the gains to sellers v_S is small compared to the nuisance costs facing consumers c_B .

¹⁸Marginal advertisements are always of negative utility for consumers so long as $c_B/v_B > c_S/v_S$, which is an implication of the assumption that $\beta^* > 0$.

¹⁹Along those lines, the optimal value β^{**} does not perfectly balance these two forces, that is does not equate their sum to zero. The reason is that there is a cost borne by consumers who choose to block, where this cost is given by $\theta - U$ —however, this cost is equal to zero at the decentralized level β^* and hence (heuristically) can be ignored.

Corollary 1. *Suppose that for a given set of parameters v'_S, c'_S, v'_B and c'_B that $c'_B/v'_B > c'_S/v'_S$, and that $q(\mu) = (1 - \mu)^\eta$ for some $\eta > 1$. Fix $v_B/c_B = v'_B/c'_B$, and $v_S/c_S = v'_S/c'_S$. Then there exists a $\rho > 0$ such that $\beta^{**} > \beta^*$ if and only if $v_S/c_B < \rho$.*

Proof. Fixing the parameters as specified ensures that neither μ^* nor β^* moves as v_S does. Moreover, $q(\mu)$ is decreasing and convex, meaning that $q''/q' < 0$, so that the appropriate assumption from Proposition 4 is satisfied. The condition for $\beta^{**} > \beta^*$ from Proposition 4 can be written as

$$\frac{-1 + \frac{v_B}{c_B}[aq(\mu^*) + (1 - a)\hat{q}]}{aq'(\mu^*)} > \frac{v_S}{c_B}.$$

Note that the left-hand side is always positive since $c'_B/v'_B > c'_S/v'_S$ (and from firm optimization). The result then follows directly. ■

There is an analogy to Lojack, a system installed on cars that allows them to be tracked should they be stolen. There are two important features of Lojack. First, such systems are unobservable to potential thieves, and second, the presence of Lojack-equipped cars in the population theoretically should lower the expected returns to theft because the cars are more likely to be recovered. Because of these two factors, there would seem to be a positive externality imposed on other consumers (car owners) by those who adopt Lojack, suggesting that there may be suboptimally low adoption. Indeed this is exactly what Ayres and Levitt (1998) find in their empirical analysis of Lojack.

I've already mentioned one important caveat, which is that my analysis neglects the fact that advertising revenue supports content. If an increase in blocking would lower such revenue to the extent that there would be a decline in content quality or availability, then even if Equation (5) held there might still be too much blocking in the decentralized equilibrium. However, even in this case I believe that the externality I identify is important, because it indicates the existence of a countervailing force.

A second important caveat is that the logic of Proposition 4 requires that the cost-effectiveness of advertising declines as more consumers block ads. This is clearly the case whenever consumer ad avoidance essentially is a mental decision to ignore ads in print, online, or on television. It's also appropriate for automated blocking tools like those that skip over television ads or block email. In other cases, such as with some (but by no means all) online ads, advertisers pay only when their ads are acted on in some way, usually by being clicked on. If consumers click ads if and only if they eventually purchase from that company, then heightened blocking may not raise costs, and the logic of Proposition 4 would fail.

I believe however that even in the case where payment is on a per-click basis, increased use of ad-avoidance technologies may very well reduce the overall effectiveness of advertising, suggesting that Proposition 4 will still hold. The reason is that the overall effect of ads is

likely more subtle than modeled here, meaning that viewing an ad could cause the consumer in the future to recall the company and then buy from it without ever clicking, for example by going directly to the company’s website. In fact, this is argued by some firms, such as Microsoft, prominent in online display advertising.²⁰

There is yet another reason why increased avoidance might reduce the effectiveness of advertising even if payment is on a per-click basis. In reality there are fixed costs associated with advertisements, and if consumer blocking of ads reduces the number of consumers who see the ad, it may be less worthwhile to create ads in the first place.

4. TARGETED ADVERTISING

In the analysis above I fixed the accuracy of firms’ information about consumers, given by *a*. Yet as discussed in the introduction, improvements in consumer targeting constitute a major recent development. Accordingly, here I analyze the effect of changes in information accuracy (or the privacy of consumer information).

While the ability to target particular consumer segments is typically highly valued by firms in the real world, as apparent from the huge investments being made in targeting technologies, consumers are less enthusiastic towards advertising in general and more ambivalent about targeted advertising in particular. Casual observation suggests that many if not most consumers dislike the typical advertisement and would prefer to avoid it. There are also concrete indications of privacy concerns related to targeted advertising. First, some consumers take active steps to make it more difficult to target ads towards them, for example by being sure their browser does not share information about which websites they have visited. Second, consumer privacy groups have submitted a Do-Not-Track proposal to the FTC. Third, firms that ask consumers to provide them with personal information sometimes pledge not to share the information, indirectly suggesting consumer concern. Fourth, as noted in a recent survey by TNS Global,²¹

“57 percent of the people surveyed said they are not comfortable with advertisers using their browsing history to serve ads even when they believe their names and other personal information is not being revealed. Fifty-four percent said they delete their cookies at least twice monthly.”

Complicating the picture is contradictory evidence that instead suggests consumers may be receptive to targeted advertising, at least in some circumstances. First, the TNS Global

²⁰See “Rallying Cry for Display Ads,” *Wall Street Journal*, September 19, 2008.

²¹See www.clickz.com, “Survey: Consumers Familiar with Behavioral Targeting, Don’t Love It,” March 31, 2008.

survey reports that many consumers are annoyed when they view ads they consider irrelevant.²² Second, some consumers willingly choose to receive advertisements from certain firms.²³ Third, a study by the Ponemon Institute found that consumers prefer relevant banner ads and would be more likely to click on them, and in addition

“45% of respondents said they would be willing to provide additional personal information if it meant they would receive more ads targeted to their individual interests.”²⁴

Therefore, the key questions are how an increased information flow influences firm and consumer decision making, and whether the final effect is beneficial for firms, consumers, or both. Answering these questions in turn suggests when or if consumers would wish to share information about themselves, and, if not, what could encourage such sharing.

Note that I am interested in exploring the equilibrium effects of changes in information flow, as the partial equilibrium effects are immediate. In particular, as is expected, if μ and β are fixed then an increase in a makes all consumers and firms better off.²⁵ Therefore, below I analyze the cases when μ and β equilibrate as a changes.

4.1. Equilibrium Effects of Increased Accuracy. Although both sides of the market do better as a increases when all else is fixed, there are tensions. For example, since an increase in a provides an incentive for firms to advertise more (when $q(\mu^*) > \hat{q}$), one might wonder whether this could raise nuisance costs sufficiently much that consumers would be worse off.²⁶ If this were to happen, consumers would actually have higher not lower incentives to avoid advertisements in equilibrium. Might this increased blocking overwhelm the positive effect of an increase in a , causing profits to fall? The answer is no, as I show just below.

However, it is indeed possible in general equilibrium for an increase in a to both increase advertising levels and the propensity of consumers to avoid ads—both μ and β can increase.²⁷

²²Arstechnica reports that “72 percent of those surveyed said that they find online advertising annoying when the ads are not relevant to their needs, and 87 percent of the group said that under a quarter of ads are well-targeted towards them,” and “64 percent said that they would willingly view ads from a list of online stores and brands that they know.”

²³See for example “Gucci-Stalking: New Sites Make Shopping a Game” (*Wall Street Journal*, May 1, 2008) and “Advertisers to Consumers: We’ll Text You” (*Wall Street Journal*, May 27, 2008).

²⁴“Top 5 privacy issues for 2005,” *Computerworld* December 28, 2004.

²⁵For example, consider consumers. Using Equation (2), $\frac{\partial U}{\partial a} = v_B \left[\int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) - \hat{q}\mu \right] > 0$.

²⁶When does an increase in a encourage more advertising? Recall that the marginal ad satisfies $c_S = (1-\beta)v_S[aq(\mu) + (1-a)\hat{q}]$, so that for fixed β an increase in a influences the marginal incentives to advertise thusly:

$$(1-\beta)v_S [q(\mu) - \hat{q}] > 0 \iff q(\mu) > \hat{q}.$$

In this partial equilibrium setting, firms wish to increase advertising only if the current level is such that the marginal ad is more effective than a response from a random consumer, and otherwise to decrease their advertising.

²⁷The reason β and μ may increase with a in equilibrium is that blocking influences the profits from all advertising, not merely marginal ads.

For the moment, I take this as a given (a formal demonstration is below, in Proposition 6). Despite this prospect, firms always prefer improved targeting opportunities.

Proposition 5. *Suppose that firms and consumers simultaneously choose advertising levels and whether to block ads, respectively. Then an increase in information accuracy raises the profits of firms.*

Proof. Using the envelope theorem, an increase in a influences profits (given by Equation (1)) via the direct effect of a and the indirect effect on profits through a change in β :

$$\frac{d\pi}{da} = v_S \left\{ (1 - \beta) \left[\int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) - \hat{q}\mu \right] - \frac{d\beta}{da} \left[a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1 - a)\hat{q}\mu \right] \right\}. \quad (6)$$

The two terms in brackets are positive, so that the only way profits could potentially decline is if $d\beta/da > 0$. In turn, this requires that $d\mu/da > 0$, which implies that the return from the marginal ad increases with a :

$$0 < v_S \left\{ (1 - \beta)[q(\mu) - \hat{q}] - \frac{d\beta}{da} [aq(\mu) + (1 - a)\hat{q}] \right\} \implies 0 < \frac{d\beta}{da} < \frac{(1 - \beta)[q(\mu) - \hat{q}]}{aq(\mu) + (1 - a)\hat{q}}.$$

Note that since $d\beta/da > 0$ by assumption this implies $q(\mu) - \hat{q} > 0$. Applying this inequality to Equation (6) it follows that a sufficient condition for $d\pi/da > 0$ is that

$$v_S(1 - \beta) \left\{ \left[\int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) - \hat{q}\mu \right] - \left[\frac{q(\mu) - \hat{q}}{aq(\mu) + (1 - a)\hat{q}} \right] \left[a \int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) + (1 - a)\hat{q}\mu \right] \right\} > 0.$$

Ignoring the common term $v_S(1 - \beta) > 0$, multiplying through by $aq(\mu) + (1 - a)\hat{q} > 0$ and eliminating like terms shows that this condition reduces to one that is known to hold:

$$\int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) > \hat{q}\mu. \quad \blacksquare$$

Proposition 5 is unambiguous and consistent with the advertising industry's emphasis on the benefits generated from targeted advertising. The intuition, specifically for why profits increase with a even if β does, is as follows. The only possible way β might increase is if μ does. However, since the marginal incentive to place ads decreases with β , there is an upper bound on how much β can increase and still be consistent with an increase in μ . Moreover, this bound involves the elasticity of the response probability of the marginal ad $\partial w(\mu, a)/\partial \mu$ with respect to a . Heuristically, the remaining effect on profits is determined by the elasticity of the overall response probability $w(\mu, a)$ with respect to a , and since the elasticity of this overall response is greater than the elasticity of this margin, the result follows.

I now turn to consumer attitudes towards targeted advertising, beginning with a formal defense of my claim above that it's possible for an increase in a to harm consumers.²⁸

Proposition 6. *Suppose that firms and consumers simultaneously choose advertising levels and whether to block ads, respectively. Also suppose that*

$$\min \left[\frac{c_B}{v_B}, q(0) \right] > \frac{c_S}{v_S} > \hat{q}.$$

Then there exists some $a^ \in (0, 1)$ such that when $a = a^*$ it is the case that $\mu = 0$, and a small increase in a from a^* strictly raises μ and β , and strictly lowers U .*

Proof. Let $a^* \in (0, 1)$ be the smallest value such that a firm is indifferent between placing the first ad or not, where such a value exists since the value of the initial ad is $-c_S + v_S \hat{q} < 0$ when $a = 0$ and $-c_S + v_S q(0) > 0$ when $a = 1$, by the assumption on parameters. Since there is a zero measure of advertising at a^* , $U = \beta = 0$. Since $c_S/v_S > \hat{q}$, advertising is increasing. Since $c_B/v_B > c_S/v_S$ by assumption, consumers value marginal ads negatively and hence they value all ads negatively in a neighborhood of a^* . Hence, U falls, meaning β rises at a^* . ■

Typically, consumers face two effects (and a potential tradeoff) from an increase in targeted advertising. First, there is a “mix effect,” meaning that the relevance of the existing volume of ads increases. Second, there is a “volume effect,” meaning the number of ads viewed changes. The mix effect is positive for consumers, whereas the volume effect is negative whenever the volume is increasing and $c_B/v_B > c_S/v_S$.

It's possible that the end result of higher a is higher consumer utility. This is always the case in the region of a where firms would reduce their advertising as a increases, everything else fixed (that is, where $q(\mu^*) < \hat{q}$). The following proposition provides conditions under which utility increases across the entire range of a .

Proposition 7. *Suppose that firms and consumers simultaneously choose advertising levels and whether to block ads, respectively. Fix all parameters save v_S and c_S . Then there exists a $\rho > 0$ such that for $c_S/v_S < \rho$, U is increasing in a .*

Proof. Let $\beta' < 1$ equal the level of blocking that would exist if $a = 0$ and $\mu = 1$, and choose ρ such that

$$\rho < \max \left[(1 - \beta') \hat{q}, \frac{c_B}{v_B} \right]$$

which implies that for $a = 0$ it is the case that $q(\mu^*) < \hat{q}$.

²⁸Note that in what follows below I use the fact that if $q(\mu^*) > \hat{q}$, then an increase in a leads firms to advertise more, even when consumers adjust their actions. Similarly, if $c_S/v_S > \hat{q}$, then μ is nondecreasing for any $a \in (0, 1)$. As the proof is straightforward, I omit it. Briefly, to sketch a proof of first claim, note that if an increase in a lowered μ when $q(\mu^*) > \hat{q}$, then it must be that β rose, which is impossible if a increased and μ fell.

Let subscripts denote values of a , and suppose $c_S/v_S < \rho$. Consider any $a > 0$ and suppose that μ_a^* is such that $q(\mu_a^*) \geq \hat{q}$. Then (by revealed preference of firms) $(1 - \beta_a^*)\hat{q} < c_S/v_S < (1 - \beta')\hat{q} \leq (1 - \beta_0^*)\hat{q}$, meaning $\beta_a^* > \beta_0^*$, which means that U has fallen from its value when $a = 0$. But that's not possible, since a has increased and μ has fallen (and $\beta_a^* > 0$). Hence for all a it is the case that $q(\mu_a^*) < \hat{q}$.

Now, suppose for the sake of contradiction that an increase in a lowered U . Then β must have increased, and, since the marginal return to firm advertising falls with a even for a fixed β when $q(\mu^*) < \hat{q}$, μ must have fallen. But since $c_S/v_S < c_B/v_B$ implies that consumers always dislike marginal ads, consumers must in fact be better off. ■

While Proposition 7 is intuitive and encouraging from the viewpoint of consumer welfare, it's unclear how realistic its restrictions are. That is, it seems possible that in the real world the typical firm is choosing μ such that $q(\mu) > \hat{q}$, meaning this firm would increase its advertising as targeting opportunities increase. To accommodate this concern I must impose a bit more structure. Doing so leads to a very intriguing result, however.

Proposition 8. *Suppose that firms and consumers simultaneously choose advertising levels and whether to block ads, respectively. Also suppose that*

$$\frac{c_B}{v_B} > \frac{c_S}{v_S} \quad \text{and that} \quad \frac{\partial^2 \mu}{\partial a \partial a} < 0.$$

Then U is quasiconvex in information accuracy a (and β^ is quasiconcave).*

Proof. There may be an initial region of a in which $\mu = \beta = U = 0$. But, it's enough to show that in any region where $\mu > 0$, U is quasiconvex in a , which I will do by showing that U is convex at any critical point.

$$\frac{dU}{da} = v_B \left[\int_{q(\mu)}^1 \tilde{q} dF(\tilde{q}) - \hat{q}\mu \right] + \frac{d\mu}{da} [-c_B + v_B(aq(\mu) + (1-a)\hat{q})].$$

This implies

$$\frac{d^2U}{dada} = v_B \frac{d\mu}{da} (q(\mu) - \hat{q}) + v_B \frac{d\mu}{da} \left[q(\mu) - \hat{q} + \frac{d\mu}{da} aq'(\mu) \right] + \frac{d^2\mu}{dada} [-c_B + v_B(aq(\mu) + (1-a)\hat{q})].$$

Noting that $dU/da = 0$ implies $d\beta/da = 0$ gives

$$\frac{d\mu}{da} = \frac{\partial \mu}{\partial a} + \frac{\partial \mu}{\partial \beta} \frac{d\beta}{da} = \frac{\partial \mu}{\partial a} = -\frac{(q(\mu) - \hat{q})}{aq'(\mu)}, \quad (7)$$

allowing me to write

$$\frac{d^2U}{dada} = \frac{-v_B[q(\mu) - \hat{q}]^2}{aq'(\mu)} + \frac{d^2\mu}{dada} [-c_B + v_B(aq(\mu) + (1-a)\hat{q})]. \quad (8)$$

For the sake of contradiction, suppose that $d^2U/dada \leq 0$. As the first term in Equation (8) is positive and the portion of the second term in brackets is negative (by the assumption that $c_B/v_B > c_S/v_S$ and from firm optimization), this requires $d^2\mu/dada > 0$. Building on the Equation (7),

$$\frac{d^2\mu}{dada} = \frac{\partial^2\mu}{\partial a\partial a} + \frac{\partial^2\mu}{\partial a\partial\beta} \frac{d\beta}{da} + \frac{\partial\mu}{\partial\beta} \frac{d^2\beta}{dada} = \frac{\partial^2\mu}{\partial a\partial a} + \frac{\partial\mu}{\partial\beta} \frac{d^2\beta}{dada}.$$

By assumption $\partial^2\mu/\partial a\partial a < 0$, and $\partial\mu/\partial\beta < 0$. Furthermore, using the identity $\beta = 1 - G(U)$, it can be seen that $\text{sign}[d^2\beta/dada] = -\text{sign}[d^2U/dada] \geq 0$ (given that $dU/da = 0$). Since I'm assuming (for the sake of contradiction) that $d^2U/dada \leq 0$, this means that $d^2\beta/dada \geq 0$. But this means $d^2\mu/dada < 0$, which contradicts $d^2\mu/dada \leq 0$. ■

Proposition 8 shows that consumers either prefer extremely inaccurate or extremely accurate targeting, that is $\arg \max_a U \in \{0, 1\}$.²⁹ I provide a brief example and then discussion.

4.1.1. *Example.* As an important (sufficient but not necessary) assumption of Proposition 8 is that μ is concave in a , fixing everything else, it's natural to ask what might lead to this. One might also wish to confirm that it is possible for utility to be in fact strictly decreasing and increasing over different ranges of a .

Suppose that $q(\mu) = (1 - \mu)^\eta$ for some $\eta \geq 1$, which implies that q is decreasing and convex. This is so if $F(q) = q^{1/\eta}$. More importantly, $\partial^2\mu/\partial a\partial a < 0$, so this specification satisfies the requirements of Proposition 8.³⁰ Let us also suppose that

$$\min \left[1, \frac{1}{2} \left(1 + \frac{c_s}{v_s} \right) \right] > \frac{c_B}{v_B} > \frac{c_S}{v_S} > \frac{1}{2}.$$

Result 1. *In the example described just above, there exists an $\bar{\eta} > 1$ such that for $\eta \in [1, \bar{\eta}]$, the following is true. There exists an \hat{a} such that $U(a)$ is decreasing for $a \in [0, \hat{a}]$ and increasing for $a \in (\hat{a}, 1]$. Moreover, $U(1) > U(0)$.*

The proof involves two steps, which I sketch. First, one can verify that for η not too large, Proposition 6 in addition to Proposition 8 is satisfied under the additional parametric assumptions above, ensuring quasiconvexity and also that U is strictly decreasing in some region. Second, set $a = 1$ and $\eta = 1$, assume that $U(1) > 0$ so that $\beta = 0$, compute the level of advertising by firms, and then verify that $U(1)$ is indeed strictly positive. By continuity, the result then holds.

²⁹It is worth noting that this result also holds if consumer advertising avoidance decisions are fixed. So, if consumers either don't have access to blocking technologies, or for some reason do not adjust their usage levels in response to changes in a , their utility nonetheless is quasiconvex.

³⁰Indeed, it also satisfies the requirements of Proposition 4, and these are the two main propositions that lean on particular properties of $q(\mu)$ or $\mu(a)$.

4.2. **Discussion.** In the real world firms appear to be in favor of targeted advertising while consumers are more ambiguous. These facts are reflected in Propositions 5-8.

A key implication of Proposition 8's prediction that consumers have preferences for extreme levels of information accuracy is that consumer attitudes towards targeted advertising may depend crucially on the state of targeting technology. For example, suppose that the hypotheses of that proposition are satisfied, and that the set of feasible technologies is indexed by \bar{a} , meaning that for a given technology level the information accuracy is constrained to $[0, \bar{a}]$. Then unless \bar{a} is sufficiently large, consumers will always oppose increased targeted advertising. In contrast, if \bar{a} is very large so that potential targeting is very precise, consumer utility may be in the region where it increases with targeting, meaning they would favor further improvements.

For similar reasons, Proposition 8 also suggests that consumer attitude towards sharing information about themselves may change as technology improves. In particular, it's reasonable to assume that the actual ability to target a is a function of several factors, one of which is how much information consumers share in addition to what might otherwise be available. Thus, it may be that when the underlying ability to target is not very high, consumers will be hesitant to share additional information, but this attitude may change as the underlying ability to target passes some threshold.

Given the rapidly advancing ability to target ads in the real world, such a watershed time may be near. At the least, there is anecdotal evidence that the targeted advertising of today is not always very accurate, and sometimes quite bad, but appears to be improving. For example, I am a member of the social networking tool Facebook. Facebook knows, because I provided information about myself at its request when signing up for it, that I am male, thirty six, a runner, and interested in (only) women. But in the past Facebook has delivered many ads seemingly based on the belief that I, as a thirty-six year old male, am quite overweight and looking to lose substantial weight, and that I am looking for romantic male companionship. That is, despite the fact that Facebook had the information to know that I likely wasn't a good match for those ads, I received them frequently. Yet, recently, Facebook is doing a better job of using the information supplied to them by me, for now am more likely to get ads geared towards runners, and no longer receive ads geared towards gay men.

Unfortunately, since quasiconvexity admits the possibility of strictly monotone decreasing utility, it's also possible that further technological advances will only make consumers worse off with an increase in undesired advertisements. Therefore, to the extent that consumer preferences are placed before firm profits, providing consumers with means to take counter-measures remains critical as targeting technology improves.

Providing consumers with the means to avoid targeting if they choose, but not ruling out the possibility of their reaping the benefits of tracking by prohibiting it, leaves open the

prospect of an entirely new world in which consumers generally value advertisements and are happy with tracking, while protecting them from the worst-possible outcomes. Moreover, such measures potentially allow consumers to tailor how much information they release in different circumstances. For example, a consumer receiving an ad while using a search engine to find a certain type of product may be very happy to divulge further information, but be less willing to do so while randomly surfing.³¹

The work above also emphasizes the connection between two different notions of privacy identified in the literature, secrecy and autonomy. Secrecy refers to the suppression of data that informs about a consumer, while autonomy refers to the right of consumers not to be disturbed. Here, this means that secrecy refers to a while autonomy refers to ad blocking.

Changes in the level of secrecy alter the equilibrium payoffs of viewing ads, and the equilibrium response is a change in the portion of consumers who exercise autonomy in the form of ad blocking. Under the conditions of Proposition 8, a decrease in secrecy available to consumers may initially raise the demand for autonomy, while further decreases in secrecy may lower the demand for autonomy. Another viewpoint is that since utility is quasiconvex in accuracy, a low preference for secrecy is associated with relatively high utility and hence a low demand for autonomy, while a low demand for autonomy may also be consistent with a high demand for secrecy, again due to the quasiconvexity of utility.

Hermalin and Katz (2006) also emphasize the importance of considering graded information structures, unlike the earliest work in the economics of privacy. They demonstrate the existence of cases in which a perfect ability to recognize consumer types leads to efficiency but an intermediate ability can lower social welfare compared to the case where no information is shared. My results in this section, and also in Section 5 below, strongly support the key takeaway from their analysis, albeit in a very different setting.

5. “THE LONG TAIL OF THE INTERNET”

In the analysis above I assumed that firms were symmetric, that the cost of advertisements was fixed, and that the number of ads fluctuated with a . Here I change all of these assumptions, allowing me to identify additional reasons consumers may favor targeted advertising or not, and revealing a richer view of how firms benefit (or not) from increased targeting.

³¹Another circumstance in which consumers would clearly benefit from increased targeting is when this increased ability to track is not shared by all firms, or more generally when there is an assurance that it will not lead to increased levels of annoying ads (as shown in Section ??). Since in some real-world circumstances consumers have control over how much information firms are able to gather about them, it follows that firms may be able to convince consumers to share more information with them if they can convince consumers that they will not share this information with others; simultaneously, firms may be tempted to renege on their promises.

While there are several ways in which firm heterogeneity could be incorporated, I adopt an approach that allows for “niche” firms vying for the attention of consumers from “mainstream” firms (but not competing in the product market). This makes it possible to consider the so-called “long tail of the internet,” which is the idea that the internet, especially search and targeted advertising, provides enhanced opportunities for niche firms.³²

I use a model related to but different from that above.

5.1. Model Details. There are two firms and two groups of consumers, indexed by $i = 1, 2$. The mass of consumer group i is $m^i > 0$. It is uncertain which consumer is in which group. For each customer, a signal $\sigma \in \{1, 2\}$ is realized that both firms observe. Let $\sigma(i|j)$ denote the probability that a consumer of type j generates signal i . Assume that $\sigma(j|j) = a \in [0.5, 1]$ for $j = 1, 2$.

A consumer of group i has a match with firm i with probability q_H and with firm $j \neq i$ with probability $q_L < q_H$. The existence of a match between a consumer of either group and firm j means that if this consumer has viewed an advertisement from j , then trade occurs, generating surplus v_B^j and v_S^j for the buyer and seller, respectively.

An important new element is that each consumer views only one advertisement in total. Which advertisement is viewed is determined by a second-price auction. In particular, each firm simultaneously bids for the right to advertise to any given consumer. Whoever bids the highest places an ad to that consumer, and pays the bid placed by the other firm.

On the consumer side, I ignore advertising avoidance, so all consumers view exactly one advertisement. As such, it is without further loss to ignore nuisance costs, so that consumer utility from viewing an ad of firm j is v_B^j if there is a match with that firm, and zero otherwise.

To reiterate and be entirely precise, the timing is as follows. First, for each given consumer, a single signal is generated, and both firms observe the signal. Second, firms simultaneously place bids for the right to advertise to the various consumers. Third, the winner for any given consumer pays the second highest bid and places an ad to that consumer. Fourth, the existence of a match is realized, and if a match exists trade occurs.

Let r^{ij} denote the gross profits, or the profits excluding advertising costs, that firm i would gain if its ads were viewed by $\sigma = j$ (that is, by consumers exhibiting signal j).

5.2. Analysis. As mentioned above, my main goal is to incorporate firm asymmetry. Nonetheless, as a starting point, I begin with the symmetric case where $m^1 = m^2 = m$, $v_S^1 = v_S^2 = v_S$, and $v_B^1 = v_B^2 = v_B$.

³²See Brynjolfsson, Hu, and Simester (2007) for a discussion and further references.

Proposition 9. *Consider the symmetric case: $m^1 = m^2 = m$, $v_S^1 = v_S^2 = v_S$, and $v_B^1 = v_B^2 = v_B$. An increase in information accuracy a raises consumer utility and firm profits.*

Proof. As the firms are symmetric, without loss consider firm 1. Advertising to $\sigma = 1$ (that is, to consumers exhibiting signal 1) generates gross profits from these consumers of

$$r^{11} = v_S m [q_H a + q_L (1 - a)] = v_S m [q_L + a(q_H - q_L)].$$

If firm 1 were to advertise to $\sigma = 2$ its gross profits from those consumers would be

$$r^{12} = v_S m [q_H (1 - a) + q_L a] = v_S m [q_H - a(q_H - q_L)].$$

Now, the gross profits of targeting $\sigma = 1$ exceed those of targeting $\sigma = 2$, since $a \geq 0.5$.

$$r^{11} - r^{12} = v_S m (q_H - q_L) (2a - 1) \geq 0.$$

Appealing to symmetry, this means that firm 1 always outbids firm 2 for $\sigma = 1$ and likewise firm 2 outbids 1 for $\sigma = 2$. Furthermore, the difference above is increasing in a , so that the profits of both firms increase as targeting accuracy does.

Now, the expected utility of any consumer is $v_B [q_H a + (1 - a)q_L]$, which increases in a . ■

The intuition is as follows. As a increases, the signals become more correlated with the preferences of consumers. Because both firms value trade with consumers the same, such increases raise the gross profits of firm 1 targeting $\sigma = 1$ and lower the gross profits of targeting $\sigma = 2$. By symmetry, this means that firm 1 is willing to bid more for access to $\sigma = 1$ while firm 2 is willing to bid less. Both forces push towards higher firm profits.

Consumers in turn are more likely to receive an advertisement for a product they are more likely to buy. Because they receive the same surplus from either firm conditional on trade (and because they only receive one ad no matter what) they value improved targeting.

Proposition 9 shows targeted advertising in a very positive light. To investigate its robustness, I impose three parametric assumptions that in essence make firm 1 a niche firm.

Assumption 1 (Firm 1 is a Niche Firm). *The following hold:*

- (1) $v_S^1 (q_H m^1 + q_L m^2) < v_S^2 (q_H m^2 + q_L m^1)$,
- (2) $v_S^2 q_H > v_S^1 q_L$, and
- (3) $v_S^1 q_H > v_S^2 q_L$.

As the following proposition makes precise, Assumption 1 means that firm 1 is a niche firm in the sense that for low levels of a it does not advertise at all, while for higher levels it advertises to $\sigma = 1$ with firm 2 continuing to advertise to $\sigma = 2$.

Proposition 10. *Suppose that Assumption 1 holds. Then there exists an $a^* \in (0.5, 1)$ such that for $a < a^*$, firm 2 advertises to $\sigma = 1$ and $\sigma = 2$ (and so firm 1 does not advertise), while for $a > a^*$ firm 1 advertises to $\sigma = 1$ while firm 2 advertises to $\sigma = 2$. Furthermore, the profits of both firms are nondecreasing in a , being strictly increasing for $a \in [a^*, 1]$.*

Proof. To avoid a tedious waste of space, I merely explain the main steps and omit most computations. Part (1) of Assumption 1 ensures that for $a = 0.5$, $r^{2j} > r^{1j}$ for $j \in \{1, 2\}$. For $a = 1$, parts (2) and (3) ensure, respectively, that $r^{11} > r^{21}$ and $r^{22} > r^{12}$. By linearity of the r^{ij} in a , the first part of the proposition follows, and all that remains is to show how profits change with a .

That firm 1's profits increase past a^* follows from what's already been shown, including linearity of the r^{ij} . So consider firm 2, which has profits of

$$\max [0, r^{21} - r^{11}] + \max [0, r^{22} - r^{12}] .$$

These are continuous, and for $a < a^*$ given by

$$(r^{21} - r^{11}) + (r^{22} - r^{12}) = (r^{21} + r^{22}) - (r^{11} + r^{12}) ,$$

which is invariant to a (since each term in parentheses on the right-hand side is the profits generated if all consumers receive an advert from the same firm). Since the threshold result (already proven) concerning a^* implies that $r^{21} - r^{11}$ is decreasing, $r^{22} - r^{12}$ must be increasing. Since for $a > a^*$ profits for firm 2 are $r^{22} - r^{12}$, the result follows. ■

What is interesting about Proposition 10 is that it shows that increased targeting may benefit all firms, reminiscent of the earlier analysis (Proposition 5). In particular, even those “mainstream” firms that face growing competition for ad inventory experience higher profits.

The reason the long tail of the internet story benefits mainstream firms in addition to niche ones is that advertising rates are set via a competitive bidding process. Thus, simultaneous to firm 1's bidding more to reach $\sigma = 1$, firm 1 also bids less to reach $\sigma = 2$. In other words, an increase in a softens competition for some ad inventory, allowing both firms to benefit even though one advertises more while another advertises less in total.

This raises the question of whether consumers also benefit from the emergence of new advertisers, and a complete long tail story must explain whether this is so. It turns out that, as in the earlier model, consumers certainly may gain from improved targeting, but need not.

The following proposition demonstrates the diversity of possible outcomes for consumers. Note that the three cases covered also exhaust all possible (non-knife edge) parameter configurations, so that the characterization is complete.³³

³³It cannot be that $\frac{q_H}{q_L} < \max \left[\frac{v_B^1}{v_B^2}, \frac{v_B^2}{v_B^1} \right]$, since $q_H/q_L > 1$ and one of the other two terms must be less than 1.

Proposition 11. *Suppose that Assumption 1 holds.*

- (1) *If $\frac{q_H}{q_L} > \max\left[\frac{v_B^1}{v_B^2}, \frac{v_B^2}{v_B^1}\right]$, then an increase in a from $a' \in (0.5, a^*)$ to $a'' \in (a^*, 1)$ strictly raises the utility of group 1 consumers and strictly lowers that of group 2 consumers.*
- (2) *If $\frac{v_B^1}{v_B^2} > \frac{q_H}{q_L} > \frac{v_B^2}{v_B^1}$, then an increase in a from $a' \in (0.5, a^*)$ to $a'' \in (a^*, 1)$ strictly raises the utility of all consumers.*
- (3) *If $\frac{v_B^2}{v_B^1} > \frac{q_H}{q_L} > \frac{v_B^1}{v_B^2}$, then an increase in a from $a' \in (0.5, a^*)$ to $a'' \in (a^*, 1)$ strictly lowers the utility of all consumers.*

Proof. The proof is little more than algebra, along the lines of the intuition presented just below. As such, to avoid a waste of space, I do not provide the relevant computations. ■

Proposition 11 follows from three observations. First, a consumer of group i would prefer to receive an advert from firm i if and only if $q_H v_B^i > q_L v_B^{-i} \iff q_H/q_L > v_B^{-i}/v_B^i$.³⁴ Second, the probability that a consumer of type i receives an advert from firm i is increasing in a , conditional on $a > a^*$. Third, the discrete shift in which firms place advertisements, occurring at a^* , leads to a discrete jump in the probability that group 1 receives ads from firm 1, and a discrete drop in the probability that group 2 receives ads from firm 2.

For instance, consider case (1) of Proposition 11. For these parameters, any consumer i would ideally receive an advert from firm i . As a moves from a point below a^* to one above it, group 1 consumers become more likely to receive their preferred ad while group 2 consumers become less likely to receive theirs (since they received it with probability one when $a < a^*$). Hence, group 1 is better off, while group 2 is worse off.³⁵

In case (2) of Proposition 11, the mainstream product doesn't supply much surplus to consumers. Hence, even group 2 consumers would prefer to receive an ad from the niche firm, despite having a lower probability of a match. The change in advertising as a passes a^* benefits all customers by making it more likely they receive ads from the niche firm.³⁶

In case (3) of Proposition 11, the niche product doesn't supply much surplus to consumers. Hence, even group 1 consumers would prefer not to receive an ad from the niche firm. Nonetheless, as a passes a^* , the niche firm begins reaching some consumers with its ads, displacing ads from firm 2. This makes all consumers worse off.³⁷

³⁴Note that this is also the condition for an increase in a to benefit type i conditional on $a > a^*$.

³⁵Further increases in a from a'' are valued positively by both groups. However, group 2 consumers are still worse off than when $a < a^*$.

³⁶Further increases in a from a'' are valued positively by group 1, but negatively by group 2 since such increases make it more likely that these groups receive ads from firm 1 and 2, respectively.

³⁷Further increases in a from a'' are valued positively by group 2, but negatively by group 1 since such increases make it more likely that these groups receive ads from firm 2 and 1, respectively.

5.3. **Discussion.** Here I further discuss the results of this section, with an emphasis on comparing them to those from earlier sections. First, consistent with earlier analysis, an increased ability to target consumers increases profits for all firms, where here the mechanism involves the strategic interaction between firms in the competition for ad inventory. Even though once-dominant firms face stiffened competition for access to some consumers, the softened competition for others allows overall profits to rise.

Second, also consistent with earlier analysis, consumers can but need not be made better off by improved targeting. Here the main downside of increased a is that, counterintuitively, this may cause consumers not to receive their preferred ads. There are two broad reasons for this. First, ads are placed by firms, while consumers are completely passive. Second, there is variation in the ranking of preferred matches across firms and consumers. For example, in case (2) of Proposition 11 even the mainstream consumer group would prefer to receive the niche ad, but for low values of a firm 2 nonetheless outbids firm 1 and dominates the ad channels. In other words, because firms and consumers may feel differently about potential matches, and since firms are the initiators of matches, consumers may prefer to see ads other than the ones they actually do.

Third, as in earlier analysis and reminiscent of Hermalin and Katz (2006), consumer utility may be non-monotone in the precision of the targeting technology. For instance, case (1) of Proposition 11 represents a situation in which all consumers prefer *perfect* targeting ($a = 1$) to any other level of targeting. However, in its absence the emergence of new advertisers makes mainstream consumers worse off. Similarly, in cases (1) and (3), mainstream consumers have changing attitudes about small increases in a . When $a < a^*$, these consumers are indifferent to small changes, but once $a > a^*$, they are receiving ads that are “irrelevant” to them (that is, ads from firm 1 that have a low match probability). At this stage, increases in a are preferred, since this reduces the likelihood of irrelevant ads.

In other words, in cases (1) and (3), mainstream consumer utility is quasiconvex in a .³⁸ Utility drops discretely at a^* , then begins climbing again, and finally at $a = 1$ reaches the level it was at for $a < a^*$. Given the earlier quasiconvexity result (Proposition 8), as well as the work of Hermalin and Katz (2006), it is clear that ignoring gradations in the quality of information available leads to an incomplete picture.

This leads to a fourth and final point, which is that different consumer groups may feel differently about improved targeting. For example, as case (1) of Proposition 11 makes clear, mainstream consumers may be worse off while niche consumers are better off.

³⁸Actually, niche consumer utility is also quasiconvex in cases (1), (2) and (3). In case (2), mainstream consumers have quasiconcave utility.

6. CONCLUSION

In this paper I investigate the ongoing transformation of the advertising market. The main conclusions are as follows. First, improved targeting raises the profits of all firms. Consumers need not fare as well, as they may be better or worse off when targeting improves. Second, consumers may underutilize advertising avoidance technologies from a (second-best) social standpoint. This externality implies that, even if ad revenues pay for content, there is no *a priori* reason to think such technologies cause social harm. Third, consumers may not have monotone preferences for increased targeting; indeed their utility is quasiconvex or “U-shaped.” That utility may be quasiconvex highlights the intriguing possibility that consumer attitudes towards advertising in general and targeted advertising in particular might change dramatically as technology advances. Finally, I identify a connection between two notions of privacy previously discussed in the literature, secrecy and autonomy.

REFERENCES

- ACQUISTI, A., AND H. VARIAN (2005): “Conditioning Prices on Purchase History,” *Marketing Science*, 24(3), 367–381.
- AKÇURA, M. T., AND K. SRINIVASAN (2005): “Research Note: Customer Intimacy and Cross-Selling Strategy,” *Management Science*, 51(6), 1007–1012.
- ANDERSON, S. P., AND A. DE PALMA (2007): “Information Congestion: Open Access in a Two-Sided Market,” *Working Paper*.
- ANDERSON, S. P., AND J. S. GANS (2006): “TiVoed: The Effects of Ad-Avoidance Technologies on Broadcaster Behavior,” *Working Paper*.
- ATHEY, S., AND G. ELLISON (May, 2008): “Position Auctions with Consumer Search,” *Working Paper*.
- AYRES, I., AND S. D. LEVITT (1998): “Measuring Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack,” *Quarterly Journal of Economics*, 113(1), 43–77.
- BESTER, H., AND E. PETRAKIS (1996): “Coupons and Oligopolistic Price Discrimination,” *International Journal of Industrial Organization*, 14(2), 227–242.
- BRYNJOLFSSON, E., Y. HU, AND D. SIMESTER (2007): “Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales,” *Working Paper*.
- BUTTERS, G. R. (1977): “Equilibrium Distribution of Sales and Advertising Prices,” *Review of Economic Studies*, 44(3), 465–491.
- CHEN, Y., AND C. HE (September, 2006): “Paid Placement: Advertising and Search on the Internet,” *Working Paper*.
- ESTEBAN, L., AND J. M. HERNANDEZ (2007): “Strategic Targeted Advertising and Market Fragmentation,” *Economics Bulletin*, 12(10), 1–12.

- FALKINGER, J. (2008): “A Welfare Analysis of ‘Junk’ Information and Spam Filters,” *University of Zurich Socioeconomic Institute Working Paper 0811*.
- FUDENBERG, D., AND J. TIROLE (2000): “Customer Poaching and Brand Switching,” *Rand Journal of Economics*, 31(4), 634–657.
- GAL-OR, E., M. GAL-OR, J. H. MAY, AND W. E. SPANGLER (2006): “Targeted Advertising Strategies on Television,” *Management Science*, 52(5), 713–725.
- GALEOTTI, A., AND J. L. MORAGA-GONZÁLEZ (2008): “Segmentation, Advertising and Prices,” *International Journal of Industrial Organization*, 26(5), 1106–1119.
- GROSSMAN, G. M., AND C. SHAPIRO (1984): “Informative Advertising with Differentiated Products,” *Review of Economic Studies*, 51(1), 63–81.
- HANN, I.-H., K.-L. HUI, S.-Y. T. LEE, AND I. P. PNG (2008): “Consumer Privacy and Marketing Avoidance: A Static Model,” *Management Science*, 54(6), 1094–1103.
- HERMALIN, B. E., AND M. L. KATZ (2006): “Privacy, Property Rights and Efficiency: THE Economics of Privacy as Secrecy,” *Quantitative Marketing and Economics*, 4(3), 209–239.
- HIRSHLEIFER, J. (1971): “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 61(4), 561–574.
- (1980): “Privacy: Its Origin, Function, and Future,” *Journal of Legal Studies*, 9(4), 649–664.
- HUI, K.-L., AND I. PNG (2006): “The Economics of Privacy,” in *Handbooks in Information Systems*, ed. by T. Hendershott, vol. 1, chap. 9, pp. 471–497. Elsevier, Amsterdam.
- IYER, G., D. SOBERMAN, AND J. M. VILLAS-BOAS (2005): “The Targeting of Advertising,” *Marketing Science*, 24(3), 461–476.
- JOHNSON, J. P. (2005): “Entry-Level Products with Consumer Learning,” *Contributions to Economic Analysis & Policy*, 4(1), Article 9.
- JOHNSON, J. P., AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96(3), 756–784.
- MEURER, M., AND D. STAHL (1994): “Informative Advertising and Product Match,” *International Journal of Industrial Organization*, 12, 1–19.
- MOSCARINI, G., AND M. OTTAVIANI (2001): “Price Competition for an Informed Buyer,” *Journal of Economic Theory*, 101(2), 457–493.
- POSNER, R. A. (1981): “The Economics of Privacy,” *American Economic Review (Papers and Proceedings)*, 71(2), 405–409.
- STIGLER, G. J. (1980): “An Introduction to Privacy in Economics and Politics,” *Journal of Legal Studies*, 9(4), 623–644.
- TAYLOR, C. R. (2004): “Consumer Privacy and the Market for Consumer Information,” *Rand Journal of Economics*, 35(4), 631–650.
- VAN ZANDT, T. (2004): “Information Overload in a Network of Targeted Communication,” *Rand Journal of Economics*, 35(3), 542–560.