

ECONOMIC ISSUES

A Brief Primer on the Economics of Targeted Advertising

by

Yan Lau



Bureau of Economics
Federal Trade Commission

January 2020

Federal Trade Commission

Joseph J. Simons	Chairman
Noah Joshua Phillips	Commissioner
Rohit Chopra	Commissioner
Rebecca Kelly Slaughter	Commissioner
Christine S. Wilson	Commissioner

Bureau of Economics

Andrew Sweeting	Director
Andrew E. Stivers	Deputy Director for Consumer Protection
Alison Oldale	Deputy Director for Antitrust
Michael G. Vita	Deputy Director for Research and Management
Janis K. Pappalardo	Assistant Director for Consumer Protection
David R. Schmidt	Assistant Director, Office of Applied Research and Outreach
Louis Silva, Jr.	Assistant Director for Antitrust
Aileen J. Thompson	Assistant Director for Antitrust

Yan Lau is an economist in the Division of Consumer Protection of the Bureau of Economics at the Federal Trade Commission. The views expressed are those of the author and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner.

Acknowledgments

I would like to thank Andrew Stivers and Jan Pappalardo for invaluable feedback on numerous revisions of the text, and the BE economists who contributed their thoughts and citations to this paper.

Table of Contents

1	Introduction	1
2	Search Costs and Match Quality	5
3	Marketing Costs and Ad Volume	6
4	Price Discrimination in Uncompetitive Settings	7
5	Market Segmentation in Competitive Setting	9
6	Consumer Concerns about Data Use	9
7	Conclusion	11
	References	13
	Appendix	16

1 Introduction

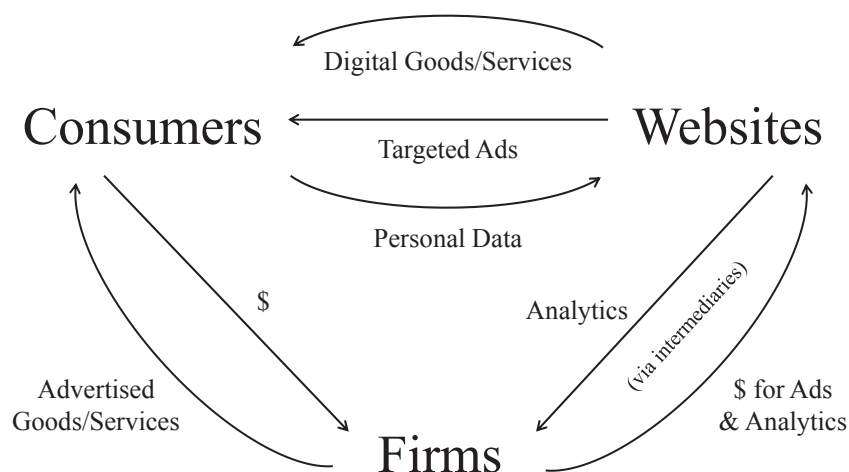
The internet has grown to touch a large part of our economic and social lives. This growth has transformed it into an important medium for marketers to serve advertising. Internet ad expenditure in the United States totaled \$69.2 billion in 2017, representing 35% of total ad expenditures across all media, compared to just 19% in 2012 (Zenith, 2018).

A crucial economic difference between internet advertising and more conventional forms of advertising is online technology's ability to more precisely target ads at consumers exhibiting specific characteristics based on browsing history and (derived) demographics (Goldfarb, 2014).¹ To harness its value, targeted online advertising necessitates the mass collection of consumers' data as they traverse cyberspace. These tracking activities have become more common over the years, but have also come under increasing public scrutiny since being introduced (Evans, 2009; Goldfarb and Tucker, 2011b). Most recently, policy makers in the European Union responded to broad consumer privacy concerns with the adoption of the General Data Protection Regulation (GDPR). One aim of the GDPR and similar policies is to limit the collection of consumer data by websites for the purposes of serving targeted ads online.

This paper provides an economic approach to thinking about targeted online advertising, specifically with regard to the collection and use of personal data to enable targeting. Restricting the collection of personal data necessarily limits the use of targeted ads by firms to reach potential customers. I consider the effects of moving away from data collection and targeted ads, with particular focus towards the effects on consumers, and discuss the theoretical and empirical research conducted to address this issue.

¹While conventional forms of advertising do target to some extent (e.g. direct mailers), targeting of online ads is much more pervasive and accelerated.

Figure 1: Ad-supported Business Model with Targeted Ads



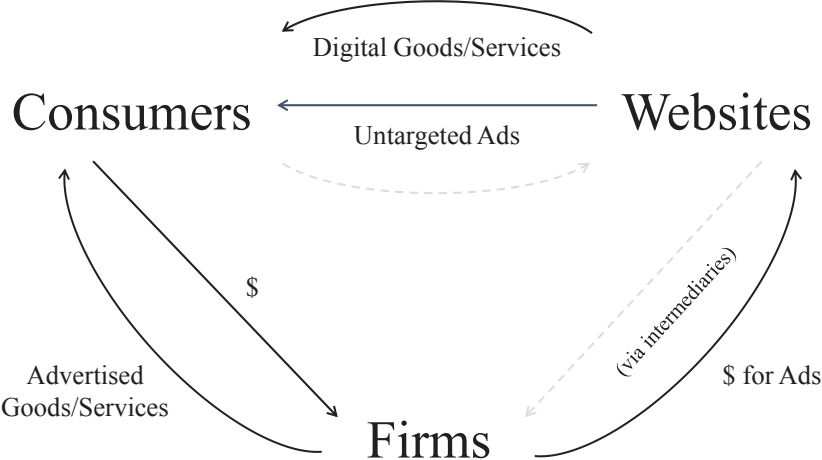
To properly consider these policy implications, it is important to first understand the current ad-supported business model that sustains much of the online economy.² The diagram in Figure 1 depicts interactions between the economic actors within (a simplified version of) this business model: consumers, websites, and firms producing the advertised product. Websites provide consumers with digital goods and services (e.g. news articles, videos, email, search), alongside which targeted ads appear. Websites receive money (via intermediaries such as ad networks) from firms that make the advertised products, and that specify which consumers they want targeted. In order to serve the targeted ads, websites use consumers’ personal data (e.g. IP addresses, browsing habits tracked using cookies, purchase history, behavioral patterns) as a basis for targeting. These data are then aggregated, packaged and provided by websites to firms (again, via intermediaries) as analytics for marketing purposes (e.g. to better target future ads or understand customers). Lastly, consumers—having viewed the targeted ads—purchase advertised goods and services from firms. An important aspect of this business model is that consumers are, in part, “paying for” the websites’ “free” digital

²For an in-depth economic perspective on the structure of the online advertising industry, see Evans (2009).

goods and services with their personal data. The presence of advertising implies they also pay with their time and attention (Bagwell, 2007).

In a world where the transfer of personal data from consumers to websites is disallowed—an extreme case no doubt, but one which serves the purposes of this exercise—it becomes infeasible for websites to serve targeted ads to consumers (they no longer have information on which to base targeting) or to provide analytics to firms. This policy would likely give rise to two alternatives against which I compare the effects of targeted advertising.

Figure 2: Ad-supported Business Model with Untargeted Ads




The first alternative is an ad-supported business model, except one with *untargeted* ads, depicted in Figure 2. Websites still receive money from firms to serve ads to consumers, albeit untargeted ones.³ However, websites do not provide analytics services to firms anymore because they no longer receive personal data from consumers.⁴ Even without targeting, ad revenues still finance the website’s digital goods and services,

³Of course, there may still be some macro-level targeting, in that firms can choose which websites to advertise on (e.g. car dealers on auto review websites), but definitely not the micro-level individualized targeting that requires the collection of consumers’ personal data.

⁴While I highlight the situation as one where no consumer data are collected for targeted advertising purposes, in reality, websites may still collect data internally to improve its own offerings or diagnose technical problems. This exception also applies to the business model discussed later in Figure 3.

where consumers (like before) “pay for” content with their time and attention. But since consumers are no longer also “paying for” the digital goods and services with their personal data, the volume of untargeted ads and the amount of ad revenue from firms will most likely differ from the volume and amount under the business model with targeted ads.⁵

Figure 3: Payment-supported Business Model without Ads



A second alternative is a payment-supported business model without *any* ads, depicted in Figure 3. Here, websites eschew ads and the associated revenue from firms, and instead charge consumers directly for providing them with digital goods and services. They no longer receive personal data from consumers for the purposes of targeting ads and providing analytics to firms in exchange for revenue.⁶

When considering a policy limiting the transfer of consumers’ personal data (in effect shutting down targeted ads), it is important to stipulate which alternative is being considered as the basis for comparison. Moreover, the business models presented are not mutually exclusive; indeed, “freemium” websites such as Pandora and Hulu

⁵In a sense, this alternative represents an early version of the internet, when expensive bandwidth limited the collection of big data for tracking purposes, and where random banner ads were pervasive. This is also the business model of most print media.

⁶Nonetheless, the exception mentioned in footnote 3 applies.

offer examples of cases that combine ad-supported free tiers (Figure 2) with costly but ad-less premium tiers (Figure 3).

In the sections to follow, I summarize various frameworks developed by economists to think about targeted advertising,⁷ and explain the benefits and costs to consumers that would arise when moving from an ad-supported business model with targeted ads, to one of the two alternatives presented above.

2 Search Costs and Match Quality

Advertising creates value through facilitating welfare-enhancing matches between consumers and firms. Consumers with imperfect information search for firms making products that meet their needs. Firms in turn search for consumers interested in their products. (“Search” here is a term used by economists to describe a costly process through which a match is sought, and does not refer to search in the sense of “search engines.”) Consumers viewing ads placed by a firm will gain information about the product (e.g. its price or quality), and the ones interested in the product will “match” with the firm and buy the product (Akerberg, 2003; Ippolito and Mathios, 1990).

Targeted ads are directed at a specific subset of consumers who are believed to be particularly interested in the firm’s product. Consumers receiving targeted ads will on average find them more “relevant” compared to untargeted ones, and spend less time searching for firms and their products to match their interests. Thus, targeting benefits the consumer because it effectively reduces their search costs.

These models also take product heterogeneity into consideration, in that firms produce products with different characteristics. In such models, average match quality—the “distance” between product characteristics desired by the consumer and the actual characteristics of the matched firm’s product—improves with targeted ads (de

⁷See Bagwell (2007) for a more general survey on the economics of advertising.

Corniere, 2016; Burguet and Petrikaite, 2017). Higher match quality means a greater increase in the consumer’s utility when the better-matched product is consumed. Empirical evidence confirms that match quality (in terms of ad effectiveness) declines when new privacy regulations limit the use of targeted ads (Goldfarb and Tucker, 2011b). Overall, search efficiency increases when advertising is targeted compared to the alternative with untargeted ads (Figure 2) because consumers search for less time to obtain better matches (de Corniere, 2016).

A secondary consequence of search cost reduction and improved match quality often identified in theoretical models is greater price competition between firms in equilibrium (de Corniere, 2016; Burguet and Petrikaite, 2017).⁸ “This result stems from the fact that targeting endogenously reduces the perceived cost of additional search because consumers know that with targeting they draw firms from a better pool. The intensification of price competition thus lowers firms’ mark-up. . . ” (de Corniere, 2016 p.157).

3 Marketing Costs and Ad Volume

Precise targeting can lead to fewer ads being served in equilibrium; this could lower marketing costs for firms by reducing wastage of ads served to disinterested consumers vis-à-vis the untargeted alternative (Esteban, Gil and Hernandez, 2001; Iyer, Soberman and Villas-Boas, 2005).⁹ Some of these cost savings could translate into lower prices for consumers, especially under competitive market conditions (Evans, 2009). Moreover, if consumers experience disutility from viewing ads, especially irrelevant untargeted ones, then having to view fewer of them is a benefit.

⁸Brahim, Lahmandi-Ayed and Laussel (2011) derive similar effects using a horizontal differentiation model.

⁹In contrast, Johnson (2013) presents a model without firm-side competition in which the volume of ads increases in equilibrium with targeting.

On the other hand, websites (or rather their intermediaries, the ad networks) may charge firms more for serving targeted ads and providing analytics, relative to untargeted ads. This would lead to higher advertising costs for firms, which may then be passed through to consumers in the form of higher prices (de Corniere and de Nijs, 2016; de Corniere, 2016).

4 Price Discrimination in Uncompetitive Settings

Targeted advertising may also lead to targeted pricing.¹⁰ Such price discrimination occurs in models where a single firm with market power (e.g. barriers to entry) separates a market into segments and charges consumers in each segment different prices.¹¹ Price discrimination in the context of targeted online advertising is relevant only when the price offered to consumers seeing targeted ads is different from the price offered to other consumers. Examples of this include providing coupon codes or tracking users' click-through from a "special offer" ad to the merchant's web store.

In price discrimination models with more and less price-sensitive market segments, the firm maximizes profits by offering the segment containing more (less) price-elastic (i.e. price-sensitive) consumers a low (high) price. In the context of an online advertising campaign, the firm would ask websites to use personal data to identify these price-elastic consumers and serve them targeted ads offering a discount. Alternatively, they could identify price-inelastic consumers and serve them targeted ads offering an inflated price.¹²

¹⁰All the models cited so far assume an undifferentiated price among consumers, except Iyer, Soberman and Villas-Boas (2005), who incorporate targeted pricing in an extension to their model.

¹¹More specifically, this practice of segmenting the market is referred to as third degree price discrimination. Shiller (2014) considers first degree price discrimination with targeted ads. Acquisti and Varian (2005) consider intertemporal price discrimination in online settings where firms track consumers' purchase history. In economics, price discrimination occurs when firms charge different consumers different prices despite there being no difference in the cost of producing the (identical) good or service.

¹²Whether the price difference is framed as "discounted" or "inflated" to which segment of consumers

Targeting technology enables websites to segment the market and charge consumers different prices; losing access to targeting (e.g. because of arbitrage or limits on its ability to collect information about consumers) would force the website to revert to a one-size-fits-all untargeted ad campaign (Figure 2), the basis for comparison. In this case, the price charged will be between the high and low prices under price discrimination. Thus, with price discrimination and targeted ads, price-inelastic consumers are worse off because they pay a relatively high price, while price-elastic consumers are better off because they pay a relatively low price (Aguirre, Cowan and Vickers, 2010). In an audit study, Hannak, Soeller, Lazer, Mislove and Wilson (2014) experimentally vary user features in simulated user profiles and empirically observe price differences based on data gathered by websites.

In the aggregate though, welfare analysis under price discrimination becomes more complex because total consumer surplus depends not only on the price individual consumers pay, but also on the quantity of consumers in each segment paying those prices. For instance, if the low price attracts a deluge of price-elastic consumers, then the aggregate gain in consumer surplus from that segment may surpass the aggregate loss in consumer surplus from the price-inelastic segment being charged the higher price. (The firm is necessarily better off because if it did not profit from segmenting the market, it would not have done so to begin with.) While total welfare depends on whether quantity demanded increases / stays constant / decreases, a necessary condition for an increase in total welfare is that aggregate quantity demanded increases (Schmalensee, 1981; Varian, 1985). More generally, economic models show that when moving from a situation without price discrimination (untargeted ads) to one with (targeted ads), changes in total welfare overall, and consumer surplus in particular, depend on the shape (in particular, the relative concavity) of demand functions in the two market

is irrelevant. In equilibrium, the pair of high and low prices would be the same.

segments (Aguirre, Cowan and Vickers, 2010; Cowan, 2012).

5 Market Segmentation in Competitive Setting

Price discrimination models assume a single firm with market power splitting the market into segments. However, another form of market segmentation may occur in competitive settings with multiple firms selling differentiated products. Under the untargeted ads regime (Figure 2), firms serve ads randomly. When firms move to a regime with targeted ads (Figure 1), they may now find it more effective to concentrate advertising onto a small segment of the product space, generating an equilibrium which endogenously differentiates and divides the market. As firms become “mini-monopolists” in their respective targeted market segments, price competition declines, and consumers may end up paying higher prices relative to an unsegmented market (Roy, 2000; Iyer, Soberman and Villas-Boas, 2005; Galeotti and Moraga-Gonzalez, 2008; Brahim, Lahmandi-Ayed and Laussel, 2011).

6 Consumer Concerns about Data Use

Limiting the collection of personal data and its use for targeted advertising may be beneficial to consumers for privacy reasons. Concerns relating to privacy¹³ can enter a consumer’s utility function in several ways. First, consumers often underestimate the degree and consequence of the personal data collection websites carry out in exchange for providing free digital goods and services (McDonald and Cranor, 2010; Acquisti, Taylor and Wagman, 2016). Consumers may experience disutility from the uncertainty regarding how the data will be used—possibly in a manner that could result in higher

¹³Acquisti, Taylor and Wagman (2016) define privacy, in the informational economics sense, as concerns relating to “trade-offs arising from protecting [as opposed to] sharing of personal data” (p. 433).

prices or fewer choices when ads are targeted.¹⁴ Moreover, there is the potential adverse outcome regarding data security; should a website be breached, “consumers could incur the costs of having private information disclosed and potentially misused” (Evans, 2009). Limiting the collection of personal data mitigates these risks because there would no longer be stored personal data for websites to exploit or hackers to steal.

Second, consumer harm can “come directly from disutility that a customer might feel from the perceived intrusiveness of the advertising” (Tucker, 2012 p. 327). Indeed, a field experiment by Goldfarb and Tucker (2011a) finds that consumers are less responsive to obtrusive ads when they are targeted based on the content of the website on which they appear, possibly because they deem targeted ads manipulative. Others postulate that consumers respond to overly intrusive ads by resisting their appeal in a process known in social psychology as “reactance” (Tucker, 2012).

That consumers care about privacy is reflected in numerous studies that attempt to empirically quantify their value of privacy.¹⁵ Importantly, valuations depend on the way choices are framed, revealing a gap between a willingness to pay for privacy protection (starting from a no privacy standpoint) versus a willingness to accept a loss of privacy (starting from a fully protected standpoint) (Acquisti, John and Loewenstein, 2013). While most studies report privacy valuations to be low, they generally confirm that consumers do place a positive—though highly variable and difficult to establish—value on protecting their privacy (Acquisti, Taylor and Wagman, 2016).

¹⁴One approach to curing the information asymmetry is to require websites to make disclosures that fully inform consumers. Yet some research suggest that privacy disclosures may have a limited impact on consumers’ willingness to allow access to their data (Berendt, Gunther and Spiekermann, 2005).

¹⁵While this literature relies heavily on surveying stated preferences, there is research highlighting a disconnect between consumers’ stated and revealed preferences when it comes to privacy (Berendt, Gunther and Spiekermann, 2005).

7 Conclusion

In considering policies limiting personal data collection practices, it is important to remember that consumers highly value the free online goods and services provided by websites. In one study using large-scale discrete choice experiments, Brynjolfsson, Eggers and Gannamaneni (2018) estimate that the median consumer in 2017 values search engine services at \$17,530 annually and email at \$8,414 annually. While a shift to the payments-supported alternative (Figure 3) would not represent a total loss of this value, it would nonetheless reduce the consumer surplus somewhat, as websites start charging for their digital goods and services. (See the Appendix for an economic explanation of what these valuations represent.)

The advent of the targeted ad-supported business model essentially allows individual consumers to monetize their personal data in exchange for valuable digital goods and services, financed by revenues from targeted ads (Figure 1). If policies restrict websites from collecting consumer personal data, then consumers lose the opportunity to monetize their personal data, representing an income effect, which consumers pay for either in dollars (Figure 3) or in time and attention spent watching a (potentially) greater number of untargeted ads (Figure 2).¹⁶ Moreover, the loss of this monetization opportunity could disproportionately affect more wealth-constrained users, who may end up losing access to these free services. Policy makers must carefully consider the ramifications of altering the current targeted ad-supported regime, which generates significant consumer and economic value.

Personal data collection and targeted advertising can be beneficial or detrimental to consumers depending on many factors. Targeted ads reduce search costs and improve

¹⁶This line of thinking is related to the allocation of property rights over personal data, first postulated by Posner (1981). Hermalin and Katz (2006) provide a summary of these arguments, but also argue that complete assignment of property rights may be insufficient for protecting consumer privacy.

match quality, which in turn may increase price competition; this increases the total value consumers derive from acquiring the products they match with. Targeting could mean fewer ads overall; consumers benefit directly from not having to view ads, but also indirectly from cost-savings passed on by firms. On the other hand, price discrimination by a single firm or market segmentation by previously-competing firms may lead to higher prices for some or all consumers involved, though the overall welfare effect depends on the shape of demand functions. There are also consumer privacy concerns that need to be addressed. Policy decisions in this arena must account for all these various aspects of economic analysis.

References

- Akerberg, Daniel A. (2003) “Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination,” *International Economic Review*, Vol. 44, No. 3, pp. 1007–1040.
- Acquisti, Alessandro and Hal R. Varian (2005) “Conditioning Prices on Purchase History,” *Marketing Science*, Vol. 24, No. 3, pp. 367–381.
- Acquisti, Alessandro, Leslie K. John, and George Loewenstein (2013) “What Is Privacy Worth?” *The Journal of Legal Studies*, Vol. 42, No. 2, pp. 249–274.
- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016) “The Economics of Privacy,” *Journal of Economic Literature*, Vol. 54, No. 2, pp. 442–492.
- Aguirre, Inaki, Simon Cowan, and John Vickers (2010) “Monopoly Price Discrimination and Demand Curvature,” *American Economic Review*, Vol. 100, No. 4, pp. 1601–1615.
- Bagwell, Kyle (2007) “The Economic Analysis of Advertising,” in M. Armstrong and R. Porter eds. *Handbook of Industrial Organization*, Vol. 3 of Handbook of Industrial Organization: Elsevier, Chap. 28, pp. 1701–1844.
- Berendt, Bettina, Oliver Gunther, and Sarah Spiekermann (2005) “Privacy in e-Commerce: Stated Preferences vs. Actual Behavior,” *Communications of the ACM*, Vol. 48, No. 4, pp. 101–106.
- Brahim, Nada Ben Elhadj-Ben, Rim Lahmandi-Ayed, and Didier Laussel (2011) “Is Targeted Advertising Always Beneficial?” *International Journal of Industrial Organization*, Vol. 29, No. 6, pp. 678–689.
- Brynjolfsson, Erik, Felix Eggers, and Avinash Gannamaneni (2018) “Using Massive Online Choice Experiments to Measure Changes in Well-being,” Working Paper 24514, National Bureau of Economic Research.
- Burguet, Roberto and Vaiva Petrikaite (2017) “Targeted Advertising and Costly Consumer Search,” Working Paper 971, Barcelona Graduate School of Economics.
- Cowan, Simon (2012) “Third-Degree Price Discrimination and Consumer Surplus,” *The Journal of Industrial Economics*, Vol. 60, No. 2, pp. 333–345.

- de Corniere, Alexandre (2016) “Search Advertising,” *American Economic Journal: Microeconomics*, Vol. 8, No. 3, pp. 156–188.
- de Corniere, Alexandre and Romain de Nijs (2016) “Online Advertising and Privacy,” *The RAND Journal of Economics*, Vol. 47, No. 1, pp. 48–72.
- Esteban, Lola, Agustin Gil, and Jose M. Hernandez (2001) “Informative Advertising and Optimal Targeting in a Monopoly,” *The Journal of Industrial Economics*, Vol. 49, No. 2, pp. 161–180.
- Evans, David S. (2009) “The Online Advertising Industry: Economics, Evolution, and Privacy,” *Journal of Economic Perspectives*, Vol. 23, No. 3, pp. 37–60.
- Galeotti, Andrea and Jose Luis Moraga-Gonzalez (2008) “Segmentation, Advertising and Prices,” *International Journal of Industrial Organization*, Vol. 26, No. 5, pp. 1106–1119.
- Goldfarb, Avi (2014) “What is Different About Online Advertising?” *Review of Industrial Organization*, Vol. 44, No. 2, pp. 115–129.
- Goldfarb, Avi and Catherine E. Tucker (2011a) “Online Display Advertising: Targeting and Obtrusiveness,” *Marketing Science*, Vol. 30, No. 3, pp. 389–404.
- (2011b) “Privacy Regulation and Online Advertising,” *Management Science*, Vol. 57, No. 1, pp. 57–71.
- Hannak, Aniko, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson (2014) “Measuring Price Discrimination and Steering on E-commerce Web Sites,” in *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 305–318: ACM.
- Hermalin, Benjamin E. and Michael L. Katz (2006) “Privacy, Property Rights and Efficiency: The Economics of Privacy As Secrecy,” *Quantitative Marketing and Economics*, Vol. 4, No. 3, pp. 209–239.
- Ippolito, Pauline M. and Alan D. Mathios (1990) “Information, Advertising and Health Choices: A Study of the Cereal Market,” *The RAND Journal of Economics*, Vol. 21, No. 3, pp. 459–480.

- Iyer, Ganesh, David Soberman, and J. Miguel Villas-Boas (2005) “The Targeting of Advertising,” *Marketing Science*, Vol. 24, No. 3, pp. 461–476.
- Johnson, Justin P. (2013) “Targeted Advertising and Advertising Avoidance,” *The RAND Journal of Economics*, Vol. 44, No. 1, pp. 128–144.
- McDonald, Aleecia and Lorrie Faith Cranor (2010) “Beliefs and Behaviors: Internet Users’ Understanding of Behavioral Advertising,” in *2010 TPRC - 43rd Research Conference on Communication, Information and Internet Policy*, TPRC.
- Posner, Richard A. (1981) “The Economics of Privacy,” *The American Economic Review*, Vol. 71, No. 2, pp. 405–409.
- Roy, Santanu (2000) “Strategic Segmentation of a Market,” *International Journal of Industrial Organization*, Vol. 18, No. 8, pp. 1279–1290.
- Schmalensee, Richard (1981) “Output and Welfare Implications of Monopolistic Third-Degree Price Discrimination,” *The American Economic Review*, Vol. 71, No. 1, pp. 242–247.
- Shiller, Benjamin Reed (2014) “First Degree Price Discrimination Using Big Data,” Working Paper 58, Brandeis University, Department of Economics and International Business School.
- Tucker, Catherine E. (2012) “The Economics of Advertising and Privacy,” *International Journal of Industrial Organization*, Vol. 30, No. 3, pp. 326–329.
- Varian, Hal R. (1985) “Price Discrimination and Social Welfare,” *The American Economic Review*, Vol. 75, No. 4, pp. 870–875.
- Zenith (2018) “Adspend Forecast Live.” Retrieved from <http://adforecast.zenithmedia.com/> on June 1, 2018.

Appendix

Brynjolfsson, Eggers and Gannamaneni (2018) estimate the median payments consumers would be willing to accept to forgo various online services. When the services are offered free of charge, the valuation amounts represent the full consumer surplus of the median consumer. However, consumers would not lose the entire amount of consumer surplus even in the extreme event where all websites shift to a payments-supported business model. Price competition among websites and entry of new competitors would keep payments well below these willingness-to-accept amounts.

Figure 4 illustrates how these median valuations relate to overall consumer surplus in a simplified market setting. The downward sloping line represents market demand for a specific online service. The median valuation is a point on this demand curve representing the median consumer’s willingness to pay. When the website moves from charging a price of zero to a price of p , the shaded amount of consumer surplus is lost in aggregate. The higher the price p charged, the greater the loss in consumer surplus.

Figure 4: The Market for an Online Service

