



Marketing Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Consumer Privacy Choice in Online Advertising: Who Opt's Out and at What Cost to Industry?

Garrett A. Johnson, Scott K. Shriver, Shaoyin Du

To cite this article:

Garrett A. Johnson, Scott K. Shriver, Shaoyin Du (2020) Consumer Privacy Choice in Online Advertising: Who Opt's Out and at What Cost to Industry?. Marketing Science

Published online in Articles in Advance 09 Jan 2020

. <https://doi.org/10.1287/mksc.2019.1198>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2020, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Consumer Privacy Choice in Online Advertising: Who Opt Out and at What Cost to Industry?

Garrett A. Johnson,^a Scott K. Shriver,^b Shaoyin Du^c

^a Questrom School of Business, Boston University, Boston, Massachusetts 02215; ^b Leeds School of Business, University of Colorado Boulder, Boulder, Colorado 80309; ^c Simon Business School, University of Rochester, Rochester, New York 14620

Contact: garjoh@bu.edu,  <https://orcid.org/0000-0002-8202-5855> (GAJ); scott.shriver@colorado.edu,  <https://orcid.org/0000-0002-9711-3356> (SKS); shaoyin.du@simon.rochester.edu,  <https://orcid.org/0000-0002-4743-2535> (SD)

Received: August 1, 2017

Revised: November 17, 2018; February 20, 2019; June 19, 2019

Accepted: June 26, 2019

Published Online in Articles in Advance: January 9, 2020

<https://doi.org/10.1287/mksc.2019.1198>

Copyright: © 2020 INFORMS

Abstract. We study consumer privacy choice in the context of online display advertising, where advertisers track consumers' browsing to improve ad targeting. In 2010, the American ad industry self-regulated by implementing the AdChoices program: consumers could opt out of online behavioral advertising via a dedicated website, which can be reached by clicking the overlaid AdChoices icons on ads. We examine the real-world uptake of AdChoices using transaction data from an ad exchange. Though consumers express strong privacy concerns in surveys, we find that only 0.23% of American ad impressions arise from users who opted out of online behavioral advertising. We also find that opt-out user ads fetch 52% less revenue on the exchange than comparable ads for users who allow behavioral targeting. These findings are broadly consistent with evidence from the European Union and Canada, where industry subsequently implemented the AdChoices program. We calculate that the inability to behaviorally target opt-out users results in a loss of about \$8.58 in ad spending per American opt-out consumer, which is borne by publishers and the exchange. We find that opt-out users tend to be more technologically sophisticated, though opt-out rates are also higher in older and wealthier American cities. These results inform the privacy policy discussion by illuminating the real-world consequences of an opt-out privacy mechanism.

History: K. Sudhir served as the editor-in-chief and Catherine Tucker served as associate editor for this article. Accepted for the *Marketing Science* Special Issue on Consumer Protection.

Supplemental Material: Data and the online appendix are available at <https://doi.org/10.1287/mksc.2019.1198>.

Keywords: privacy • digital advertising • consumer protection • self-regulation

1. Introduction

Regulators wrestle with consumer privacy protection in the Internet age. The online display advertising industry relies on the controversial practice of behavioral advertising, where the industry tracks consumer browsing behavior to improve ad targeting. In the General Data Protection Regulation (GDPR), European regulators favor an opt-in policy, where firms must first obtain consumer consent. On the other hand, American regulators have favored an opt-out policy—where concerned consumers can choose to avoid behavioral advertising—in order to balance consumer privacy protection and growth of the \$11.4 billion display ad industry. The Federal Trade Commission (FTC 2012), the Obama administration, and several bills before both the Senate and Congress have all favored policies based on opt-out mechanisms for consumers. In 2010, the industry moved to self-regulate by launching the AdChoices program, which provides both notice and choice to consumers. AdChoices provides notice by superimposing the

AdChoices logo over the corner of online display ads. AdChoices enables consumer choice through a website that allows consumers to opt out of behaviorally targeted advertising. Consumers who opt out still see ads, just not ads that are targeted based on their previous browsing behavior.

We evaluate the American regulatory approach to privacy in online ads, which relies on self-regulation and a consumer opt-out mechanism. The AdChoices self-regulatory program provides a rare opportunity to both assess consumer-revealed preference for behavioral advertising and the consequences of an opt-out policy for the ad industry. A growing literature emphasizes that behavioral advertising is a key value generator for the display ad industry. However, the losses from an opt-out policy depend critically on who opts out and their value to advertisers. We provide the first evidence of the adoption rate of AdChoices: 0.23% of American impressions. We estimate the economic loss from opting out by obtaining a proprietary data set of ad transactions from an ad

exchange operating in the United States and internationally. Ad exchanges make behavioral targeting scalable by conducting real-time auctions where advertisers can buy ads across websites. We estimate that opt-out consumers fetch 52% lower prices on the exchange than the counterfactual prices that they would fetch with tracking cookies. Finally, we illuminate some characteristics of opt-out users and regions with greater opt-out rates, which improves our understanding of privacy-sensitive users.

Like Berendt et al. (2005) and Athey et al. (2019), we uncover a privacy paradox: consumers' stated preferences overstate revealed preference measures of privacy. Multiple surveys show that about two-thirds of US consumers oppose online behavioral advertising (Turow et al. 2009, McDonald and Cranor 2010b, Purcell et al. 2012), and 20% even claim to have opted out using AdChoices (Mandese 2018). Despite the survey evidence, we find that opt-out consumers represent a small share of the marketplace: only 0.23% of American ad impressions arise from opt-out consumers. We show that opt-out rates are similarly low in other countries that implemented the AdChoices program: 0.16% in Canada and 0.26% in the European Union (EU) before GDPR. We find empirical support for several contributing factors: lack of awareness, low consumer online search for and adoption of Internet privacy products, a fragile opt-out mechanism, and consumer confusion about the AdChoices icon.

Our ad exchange data are well suited to measure the cost of the AdChoices program. Whereas much of the past work on online behavioral advertising lacked transaction-level price data (Goldfarb and Tucker 2011, Rafieian and Yoganarasimhan 2016),¹ we observe the details of over 62 million ad impression transactions. Furthermore, our data feature rich covariation in prices and impression characteristics, including AdChoices adoption status. Price differences arise, because users with exchange-tracking cookies ("cookie users") can be targeted based on both the current site (contextually) and previously visited sites (behaviorally), whereas opt-out users can only be contextually targeted. As such, we estimate the lost revenue to the exchange, publishers, and sell-side intermediaries from the inability to behaviorally target opt-out users.

Our opt-out loss estimates account for opt-out user selection as it relates to advertisers' contextual and behavioral targeting. We balance competing concerns regarding opt-out user self-selection and representativeness by utilizing multiple cross-sectional and panel-based identification strategies. We model contextual targeting by including rich ad placement fixed effects in our regression of impression prices. We model behavioral targeting by adding the user's past browsing

history in the regression. To account for opt-out user selection, we predict the price of observed opt-out impressions both as opt-out impressions and in the counterfactual, where they instead have an exchange-tracking cookie. In order to approximate the browsing history of opt-out users, we leverage a fingerprinting technique (Eckersley 2010) to partially reidentify opt-out users. Specifically, we associate an opt-out user "fingerprint" identifier with the unique combination of user attributes—such as Internet Protocol (IP) address, operating system (OS), and browser type—which can successfully reidentify most cookie users. Using these methods, our preferred opt-out loss estimate indicates a 52% price reduction. This estimate is robust to an alternative difference-in-differences approach that leverages users who appear to switch between cookie and opt-out status. The 52% price difference chiefly arises from the upper tail of transaction prices for cookie users, which is consistent with online behavioral advertising. Because of low overall opt-out rates, we estimate that the total lost expenditure on behavioral targeting from self-regulation is a modest \$5.7 million (\$8.58 per opt-out consumer) for the US desktop display ad industry in 2015.

Though few users opt out, we note that certain types of users are more likely to opt out, which has distributional consequences for the industry. We find that opt-out rates are higher among users who install nondefault browsers (e.g., Firefox and Chrome), which suggests that opt-out users are more technologically sophisticated. We also note substantial variation in opt-out rates by US cities and states, which we correlate with region-level attributes. We find that opt-out rates are positively correlated with age and both high- and low-income extremes but negatively correlated with Asian- and African-American population shares as well as education levels.

Our paper contributes to four streams of literature. First, previous literature discusses the benefits of consumer-level data for marketing. In pricing applications, Rossi et al. (1996) show the value of customer purchase data for targeted marketing. Pancras and Sudhir (2007) further model such data's strategic value in a market between a data intermediary and competing manufacturers. Several researchers have also explored the benefits of behavioral data to advertisers. In theory, behavioral targeting could reduce ad prices by thinning markets (Chen and Stallaert 2014, Levin and Milgrom 2010). However, most empirical work suggests that behavioral targeting leads to higher prices. For example, Beales and Eisenach (2014) and Marotta et al. (2019) show that ad exchange prices are higher for users with cookies, and Johnson (2013) argues that niche advertisers would exit a market without cookies. Beales and Eisenach (2014) and Miller and Skiera (2017) further show

that older cookies are more valuable because they accumulate more information. Behavioral information is useful to advertisers because it predicts click-through rates and sales (Yan et al. 2009, Farahat and Bailey 2012, Aziz and Telang 2016, Rafieian and Yoganarasimhan 2016), even though this information can vary in quality (Neumann et al. 2019). Furthermore, Goldfarb and Tucker (2011) show that cookie restrictions reduced ad effectiveness in their study of the EU's 2002 e-Privacy Directive. The directive recommended a user opt-in policy for behavioral ads, though the authors do not observe opt-in choices.

Second, we contribute to the literature on privacy in marketing and economics surveyed by Acquisti et al. (2016). Much of this literature examines consumer privacy choices in surveys and laboratory experiments. Because privacy choices exhibit context dependence (Acquisti et al. 2013), a substream of literature examines real-world privacy choices using field data. Goldfarb and Tucker (2012) document that older people are less willing to answer privacy-sensitive survey questions. Athey et al. (2019) show that both small incentives and small costs generate large shifts in a consumer's supply of private information. In online settings, Kummer and Schulte (2019) show that the demand for mobile apps responds to privacy-sensitive permissions, and Gross and Acquisti (2005) find that few users (1.2%) alter the permissive default privacy settings on social media. We contribute the first observational study of consumer privacy choices in the domain of online advertising.

Third, our discussion of the AdChoices program contributes to the literature on industry self-regulation over the use and exchange of consumer data. In their privacy literature review, Acquisti et al. (2016) highlight trade-offs between regulation and self-regulation of consumer privacy. The rapid development and complexity of the online display ad industry represent a particular challenge for regulators to keep pace. Brill (2011) argues that privacy regulation can have important consequences for the progress and competitive structure of the industry. Self-regulation can have the speed and flexibility to keep pace with the sector's advances (Gunningham and Rees 1997). AdChoices exemplifies these characteristics because of its rapid rollout as well as its expansion to new countries and new platforms. Nonetheless, self-regulation may lead to limited constraints on firm behavior, for which AdChoices has been criticized (Brill 2011). Whereas past academic work evaluated the usability and consumer awareness of the AdChoices program using survey methods (Leon et al. 2012a, b), we provide the first observational data to assess the adoption and impact of AdChoices as a self-regulatory scheme.

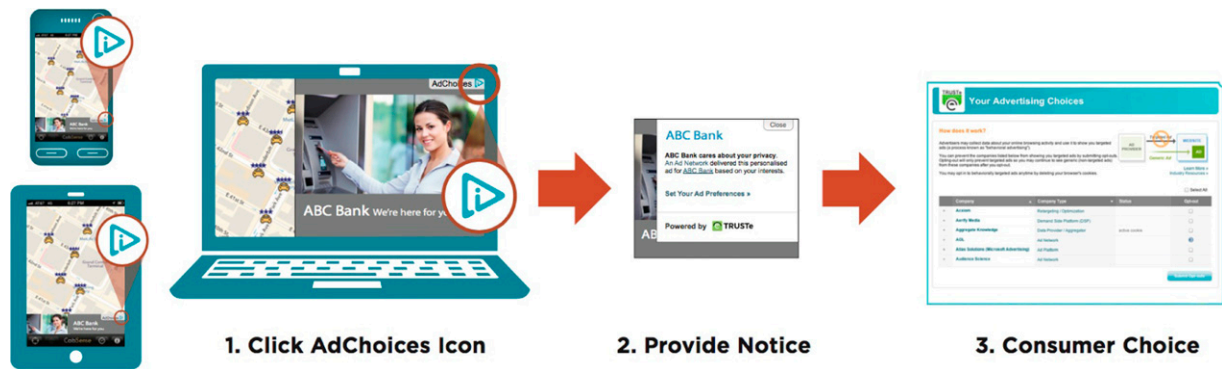
Fourth, we contribute to a growing literature on consumers' choice over the quantity and type of advertising they consume. Consumers have long passively consumed advertising that was bundled with content. This has complicated the task of measuring consumer demand for advertising, a theoretical idea put forward by Becker and Murphy (1993) and supported by empirical research showing that consumers make deliberate choices whether to watch television ads (Wilbur 2008, Tuchman et al. 2018). Meanwhile, the growth of ad blocking threatens the online display advertising industry and the publishers that depend on it (Shiller et al. 2018). Consumers who block ads cite privacy concerns arising from behavioral advertising as the top rationale for blocking ads (PageFair & Adobe 2015). As such, the AdChoices opt-out mechanism represents a middle ground where advertising can exist while mitigating consumer privacy concerns. Moreover, Tucker (2014) shows that providing users with privacy choices can make users more responsive to personalized ads. AdChoices represents a compromise option that is vital for the ad industry in an era of consumer ad choice, though the adoption of AdChoices lags far behind that of ad blocking.

The paper is organized as follows. In Section 2, we describe the AdChoices program in detail. In Section 3, we describe our data. We report our results in Section 4, and Section 5 concludes.

2. Digital Advertising Alliance AdChoices Program

In February 2009, the controversy over online behavioral advertising culminated in the Federal Trade Commission (2009) report that proposed opt-out choice as a policy direction. In response, the major marketing associations moved to self-regulate and formed the Digital Advertising Alliance (DAA) consortium. The DAA agreed on a set of self-regulatory principles in line with the Federal Trade Commission (2009) report. In October 2010, the DAA launched the centerpiece of its self-regulatory efforts: the AdChoices program.

Figure 1 illustrates how the AdChoices program provides choice and notice to Internet users. At its core, AdChoices offers the consumer choice website, where users can click a blanket—or selective—opt-out option for a list of online behavioral ad vendors. As such, the consumer choice page resembles the Federal Trade Commission's Do Not Call website, which offers consumers a blanket opt-out for direct marketing phone calls. However, desktop computers lack a permanent identifier analogous to registering a phone number: AdChoices instead records the opt-out with a browser cookie. Users can navigate to the

Figure 1. (Color online) How AdChoices Works

Source. TRUSTe (2016).

choice page directly or via small icons on display ads themselves. AdChoices notifies users through a small icon overlaid over the corner of online display ads. Figure 1 zooms in on this icon, which is a “forward i” logo that is sometimes accompanied by the word “AdChoices.” Users who click on the AdChoices icon arrive at an informational website that explains online behavioral advertising and links to the consumer choice page. Versions of these notification and choice webpages are offered by a patchwork of approved vendors. For instance, the DAA administers the informational site youradchoices.com and the consumer choice page aboutads.info/choices. Similar sites are offered by another industry group, the Network Advertising Initiative, as well as by approved vendors, such as TrustArc and Evidon. Some companies such as Google and Facebook also offer informational webpages that are linked to the DAA’s consumer choice page.

Self-regulation can be efficient because it allows industry to “get its house in order” without resorting to the coercive power of the state. The dynamism and complexity of modern ad technology pose a challenge for regulators, but the AdChoices program received government praise for its rapid rollout and broad coverage (Mastria 2014). Two years after it began, the AdChoices program extended to almost 90% of publishers in the US online behavioral market (Federal Trade Commission 2012). The DAA has kept pace with technological change by developing self-regulatory principles for mobile ads in addition to principles for both cross-application and multi-website data. By the end of 2014, AdChoices had been implemented in 33 countries (Mastria 2015), with the European Interactive Digital Advertising Alliance (EDAA) and the Digital Advertising Alliance of Canada administering AdChoices in their respective regions. The program has since grown to the point where the icon is shown over 1 trillion times per month worldwide. By the end of 2014, 5.4 million of the 39

million visitors to the DAA website had chosen to opt out (Mastria 2015). However, the total number of consumers who have opted out is unknown because the opt-out is implemented by multiple organizations.

The drawback of self-regulation is the potential underprovision of regulation, and the DAA’s AdChoices program has been criticized along several dimensions. The opt-out mechanism originally prevented online behavioral targeting but did not prevent tracking. The AdChoices program suffers from low consumer awareness, with surveys indicating American awareness levels of between 6% and 37% (Kelly Scott Madison & ORC International 2015, TRUSTe 2016). A third critique is that the cookie mechanism that preserves the opt-out choice is fragile. In particular, when users clear the cookies on their browser, users would erase their opt-out choice.

The DAA has worked to address these criticisms to varying degrees. The DAA revised the scope of the opt-out mechanism to also include tracking for the purposes of online behavioral advertising, but it still makes exceptions for ad measurement and frequency capping. To raise awareness, the DAA ran an online display ad campaign in support of youradchoices.com using over four billion donated impressions from 2012 to 2013. In 2016, the DAA launched a second awareness campaign. In response to the fragility issue, the DAA introduced browser extensions that serve as a persistent opt-out mechanism for Chrome, Firefox, and Internet Explorer. Users who read the opt-out page carefully can find a link to these extensions. Given that any cookie deletion would erase the default AdChoices program opt-out, the persistence of the 5.4 million opt-out decisions recorded by the DAA and other AdChoices-affiliated sites is unknown. Our paper provides a first look at the proportion of opt-out impressions that remain.

A fundamental concern is that users may find aspects of AdChoices confusing. Users usually do not notice the small AdChoices icon while browsing

(Leon et al. 2012b): the icons are “hidden in plain view” (Tugend 2015, section B, p. 6). Furthermore, most users do not understand what will happen when they click on the icon (Leon et al. 2012b, TRUSTe & European Interactive Digital Advertising Alliance 2014). This undermines the notification function of the AdChoices icon and contributes to its low click-through rate, which Creamer (2011) reports as being between 0.002% and 0.005%. However, AdChoices succeeds in that the majority of participants understand the informational landing page about online behavioral advertising and understand that the opt-out choice relates to behavioral advertising (Leon et al. 2012b). Nonetheless, a usability study by Leon et al. (2012a) gives a failing grade to *all* considered Internet privacy options, including AdChoices, browser-level privacy settings, and privacy-enhancing browser extensions developed by private companies. Unlike a self-regulatory program, private companies developing privacy products are fully incentivized to develop user products. If even those products fail, this would suggest that the general public’s poor technical understanding of Internet browsers is a fundamental challenge.

As we will see, these criticisms help explain the low adoption of the AdChoices program. Still, the underlying problems (low awareness of opt-out options, low technical understanding by users, and poor usability) are hard to resolve as long as users must take action to opt out of online behavioral ads. For instance, another proposed opt-out mechanism would have consumers change a browser setting so that browsers would communicate a “Do Not Track” message to websites. However, Leon et al. (2012b) show that participants also struggle to change browser settings as well. Consequently, the observed adoption of AdChoices ought to be informative for many regulatory opt-out solutions requiring direct consumer action.

3. Data

To examine consumer privacy choices and their economic consequences for the ad industry, we obtain a proprietary data set of online display ad transactions processed over one week in June 2015. In the subsections that follow, we first describe our data provider and its relation to the display ad industry. Next, we discuss the sampling scheme and its implications for inference. We conclude with a summary of our estimation data set.

3.1. Source and Related Industry Structure

Online display ads sell through two channels: guaranteed contracts and real-time auctions. Guaranteed contracts are prenegotiated bulk ad purchases between advertisers and publishers (websites), whereas

real-time auctions sell individual display ad impressions in a split second while the user loads a web page. In 2015, real-time auctions represented 31% of US online display ad spending (eMarketer 2015, 2016a, b). Online behavioral advertising is concentrated on the real-time auction channel for two reasons. First, advertisers can define narrow segments using online behavioral targeting, and real-time auctions allow advertisers to scale their target audience across many publishers. Second, real-time auctions provide flexibility for advertisers to assess users in real time based on the user’s browsing behavior. We exclude mobile impressions and focus on desktop impressions for three reasons: online behavioral advertising is most developed on desktop, the DAA was still refining its mobile opt-out tools in 2015, and mobile was a small share of our data provider’s business.

Our data provider operates an ad exchange, which is a clearinghouse for real-time auctions. As the market maker, the ad exchange uniquely observes both ad transactions and their revenue implications for the buyer, seller, and exchange. In particular, ad exchanges observe the supply of impressions regardless of whether they are sold and record the sale price.² Note that our data omit ad-blocking users because their browser extensions prevent the ad request from being sent to the ad exchange. Our agreement with the ad exchange prevents us from sharing information that would reveal the exchange’s identity or its customers’ identities. All identifying information in the data is anonymized, including publisher, advertiser, and user identifiers. This ad exchange operates in the United States and internationally, including Canada and the European Union, and it works with thousands of advertisers and publishers. The ad exchange has significant market share because it processes billions of transactions per day.

3.2. Sampling Scheme

With billions of ad exchange transactions daily, we sample the ad impressions in order to make our analysis computationally tractable. The AdChoices program operates through a cookie stored on the user’s browser that has implications for our sampling design. The opt-out cookie informs DAA participants not to collect or use the user’s browsing data for behavioral advertising. To differentiate remaining users, the exchange installs its own cookie identifier on the user’s browser whenever possible. When a user generates an impression, the exchange passes that identifier to advertisers, who may use it for behavioral advertising. From the perspective of the ad exchange, ad impressions thus fall into one of three cookie types: cookie (exchange cookie present),

opt-out (opt-out cookie present), and cookieless (both cookies absent).

As a result, only cookieless impressions may be tied to specific users, and both opt-out and cookieless impressions are anonymous. This fact frustrates a user-based sampling scheme to capture impressions of all cookie types. We therefore use a stratified sampling scheme, with one cookie type per subpopulation (stratum). We sample cookieless impressions at the user level—that is, we randomly select a subset of cookieless users (<1%) and collect their complete impression histories during the sample window. This process generates a subsample with 28.8 million cookieless impressions. The exchange sees fewer opt-out and cookieless impressions, and therefore, we sample impressions of both types at higher rates—collecting 15.4 million opt-out and 18.7 million cookieless impressions. In total, our estimation sample contains 62.9 million desktop impressions.

In using our stratified sample, we require weights to construct population-level statistics, such as the average ad sale price. Here “population level” refers to all transactions on the ad exchange during the sample window. We use sampling (probability) weights to adjust for different sampling rates across strata as well as different sampling rates with respect to key impression characteristics. To generate the weights, our data provider supplied population-level tabulation data for all impressions during the sample window by the following impression attributes: cookie type, country, publisher, and user browser type. For an impression with a given attribute combination, the sampling weight is the ratio of the number of impressions with those characteristics in the population relative to the sample.³

3.2.1. Fingerprinting Technique. An inherent limitation of our sampling scheme is that we cannot directly observe the browsing histories of opt-out and cookieless users, which, in turn, frustrates the counterfactual valuation of such information if the user allows behavioral targeting (one of our research objectives). To mitigate this issue, we introduce a fingerprinting method (Eckersley 2010) in an attempt to reconstruct opt-out user browser histories.⁴ Fingerprinting leverages characteristics of the user’s network and browser that in combination can identify unique users or small groups of users. Here we use a combination of user IP address, OS, language code, browser type, and cookie type as the basis of the fingerprint.

The relationship between fingerprints and (potentially unobserved) users can be many to many. We benchmark the performance of our fingerprinting method using the users that we observe via our cookieless user sample. We find that 76.2% of

fingerprint identifiers are associated with a unique user identification. The remaining 23.8% of fingerprint identifiers correspond to multiuser IP numbers, such as with firewalled corporations or home networks, where users share common browsers, language, and so forth. Conversely, 97.7% of user identifications are associated with a single fingerprint identifier, so many fingerprints per user is the lesser problem. The remaining 2.3% of users appear on multiple networks. We expect better match rates for opt-out users, because their rarity should reduce the likelihood that they share a fingerprint. When fingerprinting fails to differentiate users, combining multiple users may overstate the value of browsing history. However, our data understate opt-out histories because we sample at the impression rather than at the IP level. We still expect reasonable coverage of opt-out user histories because opt-out users are rare and because we oversample opt-out impressions at a 100-fold rate.

Henceforth, for expositional convenience and clarity, we distinguish between identified (cookieless) users, for whom we directly observe user identification, and fingerprinted users, where the identity of opt-out/cookieless impressions is linked to the (IP, OS, language, browser, or cookie type) fingerprint identifier.

3.3. Summary

In Table 1, we summarize key variables in our estimation data set. The first set of columns, grouped under the heading “Unweighted,” reports moments of sample data in raw (stratified) form. The final, “Weighted” column reports the weighted mean of the reported variables using the probability weights described in the preceding section. For example, opt-out impressions represent 24% of the sample (sample mean of 0.24) but only 0.2% of the population (weighted mean of 0.002). The weights thus adjust the population mean downward to account for oversampling on opt-out impressions. Cookieless impressions represent 29.7% of sample observations but only 17.6% of the exchange volume. Implicitly, cookieless transactions represent 45.9% of sample transactions and 82.2% of exchange transactions.

Ad impression transaction prices reflect the price paid per thousand impressions (CPM) by the advertiser or its agent. The exchange runs a second price auction, where the winning bidder pays the bid of the second-highest bidder or the binding reserve price. Two features of the price distribution are noteworthy. First, prices have a long right-hand tail such that the 99% quantile is 19 times the mean, implying that a small number of high-value impressions contributes significantly to total ad sales. Second, another feature is a high prevalence of zero price observations, which correspond to unsold inventory. The population

Table 1. Estimation Sample Summary Statistics and Population Mean Estimates

Variable	No. of observations	Mean	Unweighted				Weighted ^a mean
			Standard deviation	q50	q95	q99	
<i>Opt-out</i>	62,861,238	0.244	0.430	0	1	1	0.002
<i>Cookieless</i>	62,861,238	0.297	0.457	0	1	1	0.176
<i>Price^b</i>	62,861,238	0.135	0.707	0.010	0.750	2.575	0.196
<i>Sold</i>	62,861,238	0.626	0.484	1	1	1	0.674
<i>US indicator</i>	62,861,238	0.506	0.500	1	1	1	0.506
<i>EU indicator</i>	62,861,238	0.379	0.485	0	1	1	0.385
<i>Canada indicator</i>	62,861,238	0.115	0.319	0	1	1	0.109
<i>Users</i>							
<i>Identified^c</i>	5,682,476						
<i>Fingerprinted^d</i>	6,598,355						

Note. q, quantile.

^aWeights reflect relative frequency of impressions (by cookie type, US/EU/Canada region, publisher identification, and browser type) in the population (all ad exchange impressions during the sample window) compared with the estimation sample.

^bPrice (CPM) paid for the impression.

^cIdentified users have persistent browser cookies (placed by our ad exchange) that incorporate a unique identifier for each user.

^dThe number of fingerprinted users is the number of fingerprints (unique IP number, language, OS, and browser combinations) used to identify opt-out/cookieless impressions.

share of impressions sold on the exchange is 67.4%. Figure 2 plots probability-weighted kernel density estimates of prices by cookie type. Figure 2 includes sold inventory, where most prices are less than \$0.25 CPM. From this plot, we see that cookied impression prices have the largest mode and the greatest dispersion, whereas opt-out impressions have the lowest mode but are more dispersed than cookieless transactions.

Returning to Table 1, the final set of reported means indicates that US, EU, and Canadian users account for 50.6%, 37.9%, and 11.5% of sample impressions, respectively, which are about the same as the population-level shares. Finally, following the summary of key variables, Table 1 reports key information related to the identifiability of impressions. We see that cookied

impressions are associated with approximately 5.7 million directly identified users, whereas opt-out and cookieless impressions are associated with 6.6 million fingerprinted users.

4. Results


In the three subsections that follow, we report observed opt-out rates, the economic impact of AdChoices, and characteristics of opt-out users, respectively.

4.1. What Are AdChoices Opt-Out Rates?

We provide a first look at the real-world rates of privacy choice in online display advertising. These rates are a point of departure for anticipating the consequences of a regulatory opt-in policy and for evaluating the AdChoices self-regulatory program.

One way to predict consumer opt-out rates is to ask consumers. Several surveys find that about two-thirds of consumers disapprove of online behavioral advertising (Turow et al. 2009, McDonald and Cranor 2010b, Purcell et al. 2012); 31% of consumers report having clicked on the AdChoices icon in a survey of 15 European countries by TRUSTe & European Interactive Digital Advertising Alliance (2016). More recently, 20% of Americans claim to have opted out using AdChoices in a *Business Insider* survey (Mandese 2018). These numbers are large relative to the 5.4 million users that the DAA reports have ever opted out (Mastria 2015) because 5.4 million is only 1.7% of the US population. This figure matches the survey evidence from Kelly Scott Madison & ORC International (2015) that finds that 1.8% of all respondents opted out, which is the product

Figure 2. Kernel Density Plots of US Prices by Cookie Type Conditioned on Sale (Price > 0)



of the respondents who are aware of the AdChoices program (26%), understand it (35%), report clicking on the icon (33%), and report opting out (60%).

We report the impression-level rate of opt-out choice because opt-out cookies by their nature do not allow us to differentiate between users. The share of opt-out users will resemble the share of opt-out impressions to the extent that the browsing intensity of opt-out users resembles that of the general population. Nonetheless, the impression rate better reflects the economic impact of the opt-out mechanism because online display ads are sold at the impression level.

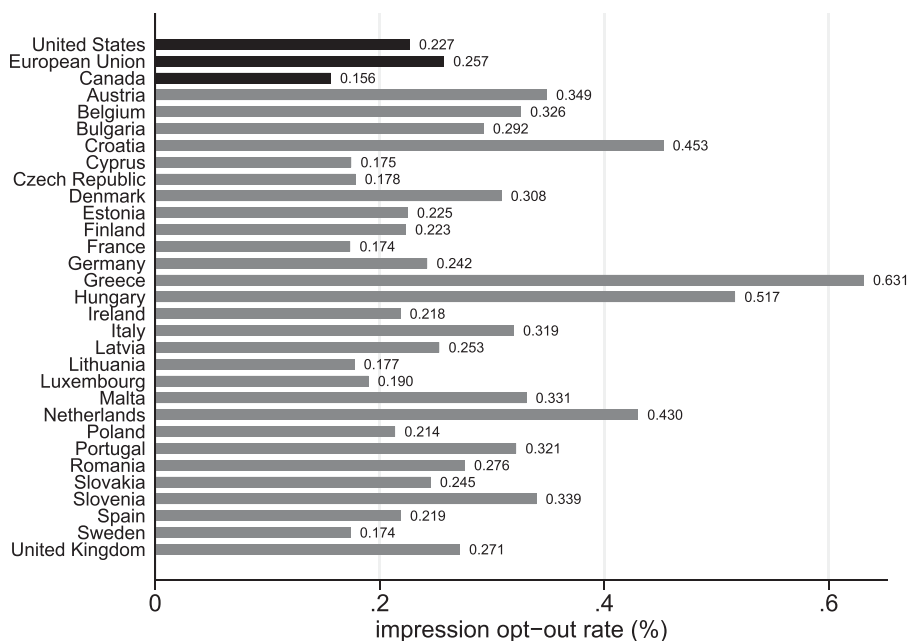
Despite consumer opposition to behavioral ads in surveys, we find much lower opt-out rates in the real world. Figure 3 reports the share of opt-out impressions rates by country using every impression on the exchange in our data period. The opt-out rate in the United States is 0.23% of impressions—meaning that 1 in 441 impressions arises from a user who opted out of online behavioral advertising. Opt-out rates are higher in the European Union at 0.26%, or 1 in 389 impressions. Opt-out rates are lower in Canada at 0.16%, or 1 in 640 impressions. Figure 3 also reports the opt-out rates for countries within the European Union. We see that Greece leads Europe with an opt-out rate of 0.63%. Hungary, Croatia, and the Netherlands are next at 0.52%, 0.45%, and 0.43%, respectively. The opt-out rates are 0.27% in the United Kingdom, 0.24% in Germany, and 0.32% in Italy. France and Sweden are both tied for the lowest opt-out rates at 0.17%.

4.1.1. Discussion: Why Are Opt-out Rates So Low? We find empirical support for several contributing explanations for the low rate of the AdChoices opt-out adoption using both our ad exchange data and external data. Note that the low opt-out rates accord with a pilot study of an AdChoices precursor, where less than 0.1% of site visitors opted out (Maier 2010).

Awareness contributes to lower opt-out rates. We correlate country-level opt-out rates with an unaided recognition measure of the AdChoices icon from 15 EU countries in 2015 (TRUSTe & European Interactive Digital Advertising Alliance 2016). We find that the two are positively correlated with a correlation coefficient of 0.59. Nonetheless, low awareness rates do not fully explain the opt-out rates. Even using the conservative 6% awareness rate found by Parks Associates (2014), the observed US opt-out rates would only be 3.8% of aware consumers. We also examined data from the EDAA on awareness-raising campaigns that it ran in 10 of 28 EU countries between 2013 and the first half of 2015. We see no significant relationship in univariate regressions on EU opt-out rates with campaign reach per capita, campaign recency, or the interaction of the two.

The low observed opt-out rate could also be in part because of the fragility of the opt-out cookie mechanism. The gap between the DAA's 5.4 million opt-outs (1.7% of the US population) and the 0.23% impression opt-out rate must result from some combination of opt-out cookie deletion, multiple desktop device-browser combinations per person, and opt-out choices by non-Americans.

Figure 3. Country-Level Opt-Out Rates



Note. The black bars represent the three regions that we analyze, and the gray bars break apart the EU by country.

We expect more technologically sophisticated users to have a lower effort cost of opting out. Table 2 shows the opt-out rates by browser. We see that Firefox users have the highest opt-out rate at 0.357% followed by Chrome at 0.256%. Because Firefox and Chrome users are sophisticated enough to install and use a non-default browser, we take this as evidence that technologically sophisticated users have lower costs of undergoing the opt-out process.⁵

Low opt-out rates can also be explained by the alternative privacy choices available to consumers in the marketplace or little demand for these options. To examine this, we first looked up the number of users for top privacy-focused browser extensions. Table 7 in Online Appendix A reports this information for Chrome and Firefox, which reveal daily active global extension usage numbers. Focusing on Chrome, we see that the DAA's "Protect My Choices" extension has about 68,000 daily active users. By comparison, the two most popular privacy extensions that block tracking—Ghostery and Privacy Badger—have 2.7 and 0.5 million daily active users, respectively. The corresponding Firefox numbers follow the same pattern. Ad blockers provide limited privacy protection, but their broad use serves as a benchmark. The two main ad-blocking extensions each have over 20 million daily active users on Chrome, and the largest has 19 million daily active users on Firefox. In the United States, ad impression blocking rates are 16% (PageFair & Adobe 2015), which is 70 times that of AdChoices opt-out rates.

Search of volume data from Google Trends confirms this pattern. Table 8 in Online Appendix A shows that US search traffic for "AdChoices," "Internet Privacy (topic)," and "Do Not Track (topic)" are within a few percent of each other. "Privacy (topic)" receives 58 times more search volume, and "Ad Blocking" receives 21 times more volume. The browser extension Ghostery receives 2.8 times more search volume than AdChoices, but Privacy Badger receives 0.33 times as much volume. By comparison, the candy "Swedish Fish," the Star Wars character "Jar Jar Binks," and the 2003 Tommy Wiseau film "The Room" receive 2.1, 4.1, and 5.3 times, respectively, more search volume than AdChoices. What emerges from the search volume and browser extension

data is that consumer surveys of privacy concern overstate consumer adoption of and search for any of these privacy technologies.

The Google Trends data also indicate evidence of consumer confusion about the AdChoices icon. We find that eight of the top nine AdChoices-related queries—89.2% of total index weighting—contained terms like "remove AdChoices." The associated search links show that consumers are confusing the AdChoices with adware because AdChoices is the only identifying characteristic of the resulting pop-up ads. The top associated links show instructions for deleting suspicious programs that may be malware. Thus, most of the searches for AdChoices do not reflect an interest in the program or avoiding behavioral advertising.

4.2. Economic Impact of AdChoices

We wish to understand the cost of the AdChoices program to industry by estimating the lost revenue from the inability to behaviorally target opt-out users. Our primary motivation is to inform the discussion around self-regulatory and legislative approaches to online privacy among policymakers and industry participants. Our secondary motivation is to help both regulators and industry understand the value of online behavioral advertising, particularly among self-selected opt-out users.

We analyze the economic impact of the AdChoices program in three steps. The first step formulates an empirical approach to measuring AdChoices' influence on ad industry revenues using regression models of impression prices. Next, we report estimates and key predictions from our regression models. We conclude by discussing the economic implications of our empirical findings.

4.2.1. Empirical Approach. Our objective is to estimate the unobserved counterfactual price that opt-out users would fetch as cookie users. We propose multiple cross-sectional and panel-based strategies to identify the associated opt-out effect. We begin by modeling cross-sectional differences in impression prices by inventory and user characteristics. However, opt-out users are selected, and the cross-sectional approach will miss unobservable differences in those users. We then reconstruct our data as a

Table 2. Browser-Level Opt-Out Rates

Browser	Opt-out rate (impression level), %	Impression share, %
Internet Explorer	0.20	47.24
Chrome	0.25	34.61
Firefox	0.34	9.52
Safari	0.19	6.05
Opera	0.23	1.93
Other	0.16	0.66

panel of impressions by user. Because price differences also arise from user browsing history, we augment our model with observed user-level browsing history. To address lingering selection concerns, we apply a difference-in-differences strategy to measure price differences among users who appear to switch between cookie and opt-out status. This strategy, in turn, has a representativeness problem because it arises from a minority of opt-out users. Nonetheless, these identification strategies all yield similar results despite their respective problems of selection on unobservables and representativeness.

Advertisers purchase ads on the exchange using two main criteria: contextual targeting and behavioral targeting. Because of contextual targeting, advertiser bids can vary as a function of inventory type, such as the site's content category and domain quality—regardless of cookie type. Because of behavioral targeting, advertiser bids can vary as a function of a user's browsing history, which requires the exchange's cookie identifier. We propose the following reduced-form regression for the price of impression t from user i to model opt-out user price difference and selection:

$$\ln(1 + P_{it}) = \alpha + \beta I_i^{\text{Opt-out}} + \lambda I_i^{\text{Cookieless}} + \gamma_t^{\text{Inventory}} + f(H_{it}) I_i^{\text{Cookieless}} + \theta X_i + \varepsilon_{it}. \quad (1)$$

User index i incorporates both identified (cookie) users and fingerprinted (opt-out, cookieless) users, as discussed in Section 3.2.1. Thus, cookie-type indicators are user specific, and they are indexed by i .

The exchange uses a second price auction so that sale price P_{it} reflects the maximum of the second-highest buyer's bid and the seller's reserve price. We transform the price variable into $\ln(1 + P_{it})$ because the multiplicative model improves fit, and the price is zero whenever the impression is unsold. The coefficient β on the opt-out indicator $I_i^{\text{Opt-out}}$ estimates the average difference between opt-out and cookie prices; the coefficient λ does the same for cookieless impressions. As we elaborate next, we model contextual targeting with inventory fixed effects $\gamma_t^{\text{Inventory}}$ and behavioral targeting—for cookie users only—with the function $f(\cdot)$ on user browsing history to date H_{it} . We allow advertisers to value user attributes X_i , such as browser and country, through θ , which can vary by cookie type. We estimate (1) separately for the United States, Canada, and the European Union.

We model contextual targeting with inventory fixed effects $\gamma_t^{\text{Inventory}}$ for content category, site domain, ad format, ad placement tag, day, and hour of day. Because of high-level contextual targeting, advertiser valuations vary by content categories, such as “finance” and “sports.” Web domains also command different prices; for instance, premium domains such

as yahoo.com command higher prices.⁶ Advertisers also favor certain ad formats, such as 728×90 or 300×250 pixels. Web publishers may categorize their impression inventory using ad tags, for instance, to indicate ads above or below the fold. Because a domain can contain multiple content categories, ad formats, and tags, we include fixed effects for the interaction of these four variables. To allow market conditions to vary over time, we also include both day and hour-of-day fixed effects.

We model behavioral targeting by relating price to the cookie user's browsing history through $f(H_{it})$. Advertisers may value users who previously browsed finance or shopping content more than those who visited games content. We construct a user's browsing history using the sequence of previous ad auctions for his or her cookie or fingerprint identifier as available, which captures that part of the history in our exchange data. To simplify user histories, we express H_{it} as a vector capturing cumulative browsing activity (impression counts) in 10 high-level content categories. We consider two specifications of $f(\cdot)$: indicators for previous browsing in each category and the log of one plus the number of previous impressions in each category. The indicator $I_i^{\text{Cookieless}}$ clarifies that advertisers can only learn a user's behavioral profile from the exchange's cookie identifier.

To interpret our nonlinear model coefficients, we calculate marginal effects as the relative difference between the model's predicted opt-out impression prices $\hat{P}_{it}^{\text{Opt-out}}$ and the *counterfactual* predicted cookie prices $\hat{P}_{it}^{\text{Cookieless}}$,⁷ holding all other predictors constant:

$$\frac{\sum_{it} w_{it} (\hat{P}_{it}^{\text{Opt-out}} - \hat{P}_{it}^{\text{Cookieless}}) I_i^{\text{Opt-out}}}{\sum_{it} w_{it} \hat{P}_{it}^{\text{Cookieless}} I_i^{\text{Opt-out}}}. \quad (2)$$

Note that the indicator $I_i^{\text{Opt-out}}$ appears in both the numerator and the denominator because the marginal effect calculation is taken over the set of opt-out impressions. We use sampling weights w_{it} to deliver population-level estimates. We refer to (2) as our marginal effect estimate for the inability to behaviorally target opt-out users.⁸

We also report an analogous market-level opt-out effect where we replace the denominator with the total predicted revenue across all three cookie types:

$$\frac{\sum_{it} w_{it} (\hat{P}_{it}^{\text{Opt-out}} - \hat{P}_{it}^{\text{Cookieless}}) I_i^{\text{Opt-out}}}{\sum_{it} w_{it} (\hat{P}_{it}^{\text{Cookieless}} I_i^{\text{Cookieless}} + \hat{P}_{it}^{\text{Opt-out}} I_i^{\text{Opt-out}} + \hat{P}_{it}^{\text{Cookieless}} I_i^{\text{Cookieless}})}. \quad (3)$$

Note that the opt-out rates in Section 4.1 place a low ceiling on the market-level effect.

We emphasize that our marginal effect estimates account for selection of opt-out users in terms of both

their browsing preference and their history. Opt-out users may select into browsing more or less valuable inventory. By predicting prices only for opt-out inventory and including rich inventory fixed effects, we isolate the role of selection in inventory. To account for browsing history in the cookieless counterfactual, prior impressions observed for an opt-out user's fingerprint identifier are allowed to influence predicted prices through the function $f(H_{it})$ in Equation (1).

We expect that the market will treat cookieless⁹ and opt-out users similarly in that neither can be behaviorally targeted. Although both coefficients should have the same sign, selection by each type of user into different inventory types differentiates the two estimates.

4.2.2. Regression Analysis. We begin by examining ad exchange prices for American users. In Table 3, we present the regression estimates for the price differences between cookieless impressions and both opt-out and cookieless impressions. Column (1) of Table 3 provides the two raw differences in means using only opt-out and cookieless indicators as covariates. At this level, our coefficient on opt-out impressions is -0.063 (standard error of 0.002), and our coefficient on cookieless impressions is -0.105 (0.001). We present marginal effect and market effect estimates for the opt-out policy at the bottom of Table 3. Using

Equation (2), we compute the marginal effect to be -37.5% , implying that opt-out impressions on average fetch 37.5% lower prices than comparable cookieless impressions. The market effect estimate (from Equation (3)) implies that total ad revenue on the exchange falls by 0.10% when opt-out impressions cannot be behaviorally targeted.

We proceed to include ad inventory fixed effects ($\gamma_t^{\text{Inventory}}$) and user attributes (X_i) to our regression model in order to account for the selection of opt-out impressions. In column (2) of Table 3, we include $1,638,157$ fixed effects that correspond to the interactions of site domain, ad tag, ad format, and content category as well as fixed effects for the user's browser type and country. In this specification, the opt-out coefficient is -0.115 , which drops the marginal effect estimate to -50.0% . Column (3) of Table 3 retains the controls of column (2) and includes our first implementation of the user history function $f(H_{it})$. Here our specification of $f(H_{it})$ includes 10 regressors for the log of one plus the number of previous impressions in each of 10 high-level content categories. Compared with column (2) of Table 3, including these history regressors has the modest effect of increasing the marginal effect estimate to -48.0% . Column (4) of Table 3 introduces our preferred regressors for browsing history $f(H_{it})$: simple indicators for previous browsing history in each of these top 10

Table 3. Regressions of Ad Sale Prices on Cookie Types

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$	$\log(1 + \text{Price})$
Region	United States	United States	United States	United States	United States	European Union	Canada
Opt-out	-0.063 (0.002)	-0.115 (0.003)	-0.098 (0.003)	-0.104 (0.002)	-0.116 (0.005)	-0.116 (0.013)	-0.102 (0.012)
Cookieless	-0.105 (0.001)	-0.045 (0.001)	-0.037 (0.001)	-0.041 (0.001)	-0.116 (0.014)	-0.139 (0.014)	-0.168 (0.019)
Fixed effects and controls							
Tag \times site \times size \times category		X	X	X	X	X	X
Day, hour of day		X	X	X	X	X	X
Browser, country		X	X	X	X	X	X
Browser, country \times cookie type					X	X	X
Lagged no. category impressions			X				
Lagged category indicators				X	X	X	X
Observations	31,828,871	31,828,871	31,828,871	31,828,871	31,828,871	23,795,924	7,236,443
R ²	0.016	0.669	0.670	0.669	0.670	0.623	0.534
No. of controls	1	1,638,157	1,638,167	1,638,167	1,638,177	1,682,054	483,226
Identified users	3,711,987	3,711,987	3,711,987	3,711,987	3,711,987	1,314,768	669,981
Fingerprinted users	3,474,740	3,474,740	3,474,740	3,474,740	3,474,740	2,529,783	593,964
Opt-out effects							
Marginal effect	-37.45%	-50.04%	-48.08%	-50.48%	-52.04%	-75.03%	-58.84%
Market effect	-0.10%	-0.15%	-0.14%	-0.15%	-0.16%	-0.32%	-0.14%

Notes. Regressions are probability weighted to reflect the relative frequency of impressions (by cookie type, US/EU/Canada region, publisher identification, and browser type) in the population. Standard errors are clustered by user identification for cookieless impressions and by fingerprint identification (unique IP number, OS, browser, language, and cookie-type combination) for opt-out and cookieless impressions. "No. of controls" indicates the number of regressors, exempting opt-out and cookieless main effects. Marginal effect computes the (weighted) average proportional change in opt-out impression prices compared with a baseline counterfactual with cookies. Market effect computes the (weighted) average proportional change in total ad exchange revenue compared with a baseline counterfactual where opt-out impressions have cookies.

content categories. The parameter estimates in column (4) of Table 3 are statistically consistent with column (3), and the marginal effects are also similar. Column (5) of Table 3 presents our overall preferred US specification, which adds interactions of browser type and country code with cookie type to the controls used in column (4). Our preferred marginal effect estimate implies that opt-out impressions fetch 52.0% lower prices than cookied impressions. As a portion of total exchange expenditure, our market effect estimate reveals that AdChoices reduces ad expenditure by 0.16%.

We note that the selection of opt-out impressions is unfavorable to industry because it is correlated with higher-value inventory. If we ignore selection, we could instead compute the market effect by multiplying our preferred marginal effect estimate (52.0%) and the US opt-out rate (0.23%), which yields a market effect of 0.12%. This is 25% lower than our preferred 0.16% market effect estimate earlier so that selection into more valuable inventory accounts for a quarter of the loss due to AdChoices. We further note the near equivalence of the opt-out and cookieless coefficients in our preferred US specification (column (5) of Table 3), implying that—with sufficiently rich controls for selection—opt-out and cookieless impressions fetch comparable prices.

Our analysis of exchange transactions in the European Union and Canada yields broadly consistent results. Columns (6) and (7) of Table 3 present our preferred specifications for EU and Canadian impressions, respectively. We see that the marginal effect of opt-out on impression prices in the European Union (−75.0%) is higher than that in the United States, whereas the effect in Canada (−58.8%) is comparable with that in the United States. In terms of market effects, the loss to industry is larger in the European Union (−0.32%), owing to both the larger marginal opt-out effect and the higher frequency of opt-out impressions. In Canada, the lower frequency of opt-out impressions compared with the United States dominates the higher marginal opt-out effect, yielding a lower loss to industry of −0.14%. As with the United States, the 95% confidence intervals for EU opt-out and cookieless coefficients overlap, whereas the Canadian coefficients overlap within the 99% confidence intervals.

4.2.2.1. Robustness: A Difference-in-Differences Estimator. After accounting for inventory quality and browsing history in Equation (1), opt-out users could still differ in some persistent user-level unobservable: for example, their responsiveness to behavioral advertising. If their distaste for behavioral advertising means that they are less responsive to the tactic, then our estimates will overstate the value of

opt-out users. To account for the residual opt-out selection, we apply a difference-in-differences estimator to the set of *switcher* users who move between cookied and opt-out impressions. Our strategy is to identify the opt-out effect from the (pooled) within-subject variation in prices by cookie type. To implement the strategy, we first modify our fingerprinting technique to identify switchers, that is, users who change from cookied to opt-out (or vice versa). Because switchers are not directly identified in the data, we impute their identity using a set of criteria. We assume that a (combination of IP, OS, language, and browser) fingerprint is associated with a switcher if the fingerprint is associated with (1) only cookied and opt-out impressions, (2) only one cookie identified user, (3) no cookieless impressions, (4) a single switching point in time, and (5) a minimum of five cookied and five opt-out impressions. This filtering criterion reduces—but does not eliminate—the chance that apparent switchers are two separate users.¹⁰ We then restrict the estimation sample to impressions from switchers and all cookied and cookieless impressions.

Thus, our difference-in-differences model associates switchers with the treatment condition and the interaction of switchers and opt-out with the treatment effect. The market prices of other cookied and cookieless impressions serve as the control. Formally, we estimate the following model:

$$\ln(1 + P_{it}) = \alpha + \omega I_i^{Switcher} + \beta I_{it}^{Opt-out} * I_i^{Switcher} + \lambda I_i^{Cookieless} + \gamma_t^{Inventory} + f(H_{it}) I_{it}^{Cookied} + \theta X_i + \varepsilon_{it}. \quad (4)$$

Unlike model (1), model (4) excludes opt-out impressions from nonswitchers.

Table 4 summarizes the results of this analysis. For each region, the table presents a simple difference-in-differences estimator without controls and Equation (4) specification that retains the controls from our preferred regressions in columns (5)–(7) of Table 3. We first note that the coefficient on switchers is consistently negative and significant without controls and weakly significant when rich controls are included—which suggests a negative selection bias among switchers that is captured by the controls in the richer specifications. Turning to the opt-out coefficients and marginal effects, the results are broadly consistent with those reported in Table 4. Comparing specifications with controls, the US marginal effect in column (2) of Table 4, −52.7%, mirrors our preferred estimate from column (5) of Table 3 of −52.0%.¹¹ For the European Union and Canada, we find larger in magnitude marginal effects using the difference-in-differences estimator than the comparable effects in Table 3: the European Union effect

Table 4. Difference-in-Difference Regressions of Ad Sale Prices on Cookie Types

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable	log(1 + Price)	log(1 + Price)	log(1 + Price)	log(1 + Price)	log(1 + Price)	log(1 + Price)
Region	United States	United States	European Union	European Union	Canada	Canada
Opt-out × Switcher	−0.063 (0.007)	−0.062 (0.013)	−0.025 (0.002)	−0.176 (0.019)	−0.018 (0.002)	−0.087 (0.022)
Switcher	−0.056 (0.007)	0.007 (0.004)	−0.079 (0.003)	0.001 (0.001)	−0.052 (0.003)	0.001 (0.002)
Cookieless	−0.106 (0.001)	−0.116 (0.014)	−0.101 (0.002)	−0.139 (0.014)	−0.065 (0.002)	−0.169 (0.019)
Fixed effects and controls						
Tag × site × size × category		X		X		X
Day, hour of day		X		X		X
Browser, country		X		X		X
Browser, country × cookie type		X		X		X
Lagged category indicators		X		X		X
Observations	24,249,186	24,249,186	17,336,474	17,336,474	6,116,568	6,116,568
R ²	0.016	0.670	0.019	0.624	0.019	0.534
No. of controls	2	1,436,537	2	1,446,396	2	447,863
Identified users	3,705,940	3,705,940	1,302,729	1,302,729	665,919	665,919
Switching users	2,159	2,159	4,858	4,858	871	871
Fingerprinted users	2,656,719	2,656,719	1,511,143	1,511,143	489,994	489,994
Opt-out effects						
Marginal effect	−51.61%	−52.73%	−35.29%	−88.07%	−40.63%	−79.93%

Notes. Regressions are probability weighted to reflect the relative frequency of impressions (by cookie type, US/EU/Canada region, publisher identification, and browser type) in the population. Standard errors are clustered by user identification for cookieless impressions and by fingerprint identification (unique IP number, OS, browser, language, and cookie-type combination) for opt-out and cookieless impressions. No. of controls indicates the number of regressors, exempting opt-out and cookieless main effects. Marginal effect computes the (weighted) average proportional change in the switcher's opt-out impression prices compared with their baseline counterfactual with cookies.

drops from −75.0% to −88.1%, and the Canada effect drops from −58.8% to −79.9%. This may indicate more potential for selection on unobservables among EU and Canadian residents, or it may indicate that sample switchers are not generally representative of opt-out users.

4.2.2.2. Sold vs. Price|Sold Analysis. The reduction in exchange expenditure on opt-out users could arise either because online behavioral targeting is needed to sell ads on the exchange or because online behavioral targeting lifts exchange prices. Because the latter relates to the match value between advertisers and users, we expect that it will be difficult to replace with a nonbehavioral form of targeting. To determine the relative importance of these forces, Table 5 decomposes the opt-out expenditure difference into the change in the price conditional on sale (intensive margin) and the probability of sale on the exchange (extensive margin). Column (1) of Table 5 repeats our preferred specification for American user expenditure differences from Table 3. Column (2) of Table 3 uses the same specification and uses price conditional on sale as the outcome. We find that most of the −52.0% price reduction comes from a −50.3% marginal effect on the exchange price conditional on sale. By contrast, the marginal effect on the probability of sale in column (3) of Table 5 is only −13.6%.

Note that the low extensive margin estimate reflects the general equilibrium nature of the evidence: contextual buying shifts to opt-out users where prices are lower. These findings imply that behavioral targeting is instrumental in fetching higher prices for sold inventory.

4.2.2.3. Quantile Regressions. When we compare the distributions of the prices for opt-out and cookieless impressions, we see that the differences arise from the upper tail of the cookieless price distribution. Recall from Figure 2 that the distribution of price conditional on sale for opt-out users and cookieless users does not exhibit the large differences that we expect from Table 3. However, cookieless impressions are more likely to fetch prices above \$1 CPM. We use quantile regressions, which like Figure 2, allow us to compare the distribution of prices by cookie type but also allow us to include inventory covariates to account for the selection of inventory by opt-out users. Figure 4 superimposes price quantile predictions by cookie type resulting from a series of quantile regressions. Because of computational constraints, these regressions only include fixed effects for 10 high-level content categories. Compared with opt-out and cookieless impressions, the cookieless price quantile predictions spike after the 85th quantile—meaning that cookieless impressions fetch significantly higher prices about


Table 5. Decomposition of Unconditional Ad Prices into Prices Conditional on Sale (Price > 0) and Probability of Sale

	(1)	(2)	(3)
Dependent variable	$\log(1 + \text{Price})$	$\log(\text{Price} \text{Sold})$	$\text{Pr}(\text{Sold})$
Region	United States	United States	United States
Opt-out	-0.116 (0.005)	-0.659 (0.015)	-0.075 (0.007)
Cookieless	-0.116 (0.014)	-0.599 (0.028)	-0.138 (0.024)
Fixed effects and controls			
Tag \times site \times size \times category	X	X	X
Day, hour of day	X	X	X
Browser, region	X	X	X
Lagged category indicators	X	X	X
Browser, region \times cookie type	X	X	X
Observations	31,828,871	20,265,411	31,828,871
R ²	0.670	0.826	0.726
No. of controls	1,638,177	1,346,419	1,638,177
Opt-out effects			
Marginal effect	-52.02%	-50.29%	-13.61%
Market effect	-0.16%	-0.11%	-0.03%

Notes. Regressions are probability weighted to reflect the relative frequency of impressions (by cookie type, US/EU/Canada region, publisher identification, and browser type) in the population. Standard errors are clustered by user identification for cookieless impressions and by fingerprint identification (unique IP number, OS, browser, language, and cookie-type combination) for opt-out and cookieless impressions. No. of controls indicates the number of regressors, exempting opt-out and cookieless main effects. Marginal effect computes the (weighted) average proportional change in opt-out impression prices compared with a baseline counterfactual with cookies. Market effect computes the (weighted) average proportional change in total ad exchange revenue compared with a baseline counterfactual where opt-out impressions have cookies.

15% of the time. This widening price gap implies that opt-out impressions are much less likely to fetch occasional high prices than their cookieless counterparts. Thus, online behavioral advertising enables some advertisers that have high willingness to pay (WTP) for narrow target audiences to bid up prices for some impressions.

4.2.2.4. Summary. US opt-out users fetch 52.0% lower prices on the exchange, though this loss represents

Figure 4. Price Quantile Prediction by Cookie Type

Note. Quantile regressions include 10 high-level content category fixed effects.

only 0.16% of total exchange revenue because opt-out impressions are so rare. These results are robust to alternative modeling and identification assumptions. We also show that opt-out impressions are selected such that the lost expenditure from AdChoices is about 25% higher than if opt-out were independent from inventory. We demonstrate that opt-out impressions tend to be sold at rates comparable with cookieless impressions but that they fetch lower prices when successfully auctioned. We further find that 80% of the difference in opt-out prices arises from higher prices conditional on sale: particularly, the top 15% of prices for cookieless impressions. Without online behavioral targeting, advertisers reach opt-out users with traditional, less valuable forms of targeting. Finally, we find evidence of larger magnitude marginal and market effects in the EU (-75.0% and -0.32%, respectively), whereas in Canada, a slightly higher marginal opt-out effect (-58.8%) coupled with lower opt-out base rates gives a lower market effect (-0.14%).

4.2.3. Discussion: Value of Online Behavioral Targeting. Our price estimates accord with most estimates of the value of a cookie. On two ad exchanges, Beales and Eisenach (2014) find that cookieless impressions fetch at least 66% lower prices than even

new cookie impressions. This accords with Goldfarb and Tucker (2011), who find that survey measures of ad effectiveness fall 65% after a cookie restriction policy in Europe. Marotta et al. (2019) report that one publisher's revenue falls only 4% for users without cookies. By contrast, we analyze opt-out effects by type of industry participant in Online Appendix C, and we show that publisher revenue falls 40% for opt-out users. Thus, the findings of Marotta et al. (2019) may indicate that some publishers are less dependent on behavioral targeting.

Based on our results, we calculate an approximate upper bound of \$5.7 million for the cost of opt-out users to the display ad industry. Our back-of-the-envelope calculation is based on the product of three terms. First, the US desktop display ad expenditure is \$11.42 billion in 2015 (eMarketer 2016a, b). Second, real-time auctions represent 31% of this expenditure in the same year (eMarketer 2015, 2016a, b). Third, the last term is the market effect estimate (−0.16%) from our preferred specification in column (5) of Table 3.

Our industry cost estimate can also speak to a propertization policy (Laudon 1996, Varian 2009) in which the ad industry could purchase online behavioral data from the consumer. We can recast the \$5.7 million cost estimate as a simple measure of the ad industry's WTP for online behavioral data. We obtain an annual per consumer industry WTP estimate of \$8.58 by dividing the industry cost estimate of \$5.7 million by the number of opt-out consumers in the United States: we estimate the latter as the product of the number of US Internet users (287 million) and the opt-out impression rate (0.23%).¹² Now that the United States allows Internet service providers to track consumers, such a market for privacy could take the form of a discount on Internet access. Nevertheless, this policy is predicated on consumers' willingness to bargain their privacy for at most \$8.58 per person annually, which is less than many households pay for Internet access monthly.

Although we can estimate industry willingness to pay, we are unable to compute consumer's valuation for privacy with our data. Acquisti et al. (2013) discuss some of the challenges in eliciting consumers' value of privacy, like differences by willingness to pay and willingness to accept framing. In our context, McDonald and Cranor (2010a) show marked differences in a survey: 69% of consumers are unwilling to accept \$1 monthly to be tracked, but only 11% are willing to pay \$1 monthly to stop tracking. Regardless, we expect the small minority who opt out to have greater valuation of privacy.

Our back-of-the-envelope calculation has several limitations beyond our aforementioned assumptions. Our estimates capture the lost value of behaviorally

targeting opt-out users in the real-time bidding exchange channel of the desktop display ad market but ignore other value-creating uses of this data. We do not know how the market would adjust to a large increase in opt-out users, though our findings are a useful starting point to consider this scenario. Our data arise from a single exchange operating globally, its affiliated data providers, thousands of advertisers, and thousands of publishers; still, these may not represent the whole industry. Our data contain a wide cross-section of users but capture a single week in June 2015 that may not reflect high-demand periods, like the holiday season. We do not observe the marginal costs of exchange participants: these costs are small but would decrease our measure of the industry's willingness to pay. Budak et al. (2016) discuss important differences in distributional consequences by type of publisher and advertiser that complement our analysis of the incidence of the opt-out effect by type of industry participant in Online Appendix C. Finally, our impact estimate ignores the compliance costs of administering the AdChoices program.

4.3. Who Are Opt-out Users?

We return to the consumer's choice to opt out by examining the characteristics of opt-out users. Regulators would benefit from this knowledge, because it informs the relative value of opt-out users in the marketplace and therefore, the cost of the policy to industry. However, the characteristics of opt-out users are difficult to study, because the opt-out eliminates behavioral data on these users. To address this issue, we aggregate the data by geographic region in order to correlate opt-out rates with regional characteristics. Although the opt-out users may not share their region's characteristics, this evidence represents a first step.

We use the Berry (1994) aggregate logit model to relate the opt-out market share to regional characteristics. As in Section 4.1, we use the impression-level opt-out rate as a measure of the consumer-level opt-out rate, though this introduces error whenever the relative browsing intensity of opt-out users varies by region. We infer city- and state-level opt-out rates using weighted subsamples of our estimation data set.¹³ Figure 5 in Online Appendix B maps the opt-out rate across states. Our city data are at the census place level and include 14,474 cities. Following Gart et al. (1985), we transform the opt-out share s using $\log\left(\frac{0.5+s}{1.5-s}\right)$, because the logit model otherwise struggles with s close to 0. In our baseline specification, we consider the roles of income, age, education, and nondefault browser share. We allow average age and per capita income to enter nonlinearly by including quadratic terms. We measure education by the share

of postsecondary graduates. We also consider the role of technological sophistication as proxied by the nondefault browser share in our data. In Table 6, we regress the opt-out rates on these normalized variables to ascertain their relative explanatory power.

The US city-level results in column (1) of Table 6 have the most statistical power and provide the clearest look at opt-out users. We find statistically

Table 6. Opt-Out User Characteristics

	(1)	(2)	(3)
Dependent variable	$\log[(0.5 + \text{opt} - \text{out rate}) / (1.5 - \text{opt} - \text{out rate})]$		
Geography	US cities	US states	EU countries ^a
<i>Per capita income</i>	−0.0103 (0.00954)	−0.00152*** (0.000489)	0.00118 (0.00118)
<i>Per capita income</i> ²	0.00401*** (0.00148)	0.000296 (0.000289)	−0.000424 (0.000403)
<i>Age</i>	0.0120** (0.00489)	0.000569* (0.000286)	0.000890 (0.000625)
<i>Age</i> ²	0.00660*** (0.00251)	0.000113 (0.000084)	0.000011 (0.000245)
<i>College graduate share</i>	−0.0344*** (0.00712)	0.00129* (0.000657)	−0.000440 (0.000865)
<i>Nondefault browser share</i>	0.0210*** (0.00572)	−0.000497* (0.000265)	0.000499 (0.000812)
<i>Partisan Voter Index</i>	0.00817** (0.00368)	0.000384 (0.000621)	
<i>Asian population share</i>	−0.0106*** (0.00189)	0.000220 (0.000342)	
<i>Black population share</i>	−0.0159*** (0.00366)	−0.000049 (0.000403)	
<i>Ad-block share</i>		0.000374 (0.000362)	0.000983 (0.000744)
<i>Do not call share</i>		−0.000158 (0.000290)	
<i>Internet Privacy search index</i>		−0.000378* (0.000195)	
<i>Privacy search index</i>		0.000316 (0.000400)	
<i>Change privacy setting</i>			−0.000798 (0.000778)
Constant	−1.001*** (0.00419)	−1.093*** (0.000277)	−1.091*** (0.00091)
Observations	14,474	50	28
R ²	0.017	0.423	0.326

Notes. Robust standard errors are in parentheses. All independent variables are normalized. The US income, age, education, and Asian- and African-American shares data are from the US Census (2015 American Community Survey 5-year estimates). Voting data are from the 2016 presidential election at the voting district level from the Cook Political Index. Nondefault browser rates are from our exchange data. Ad-block rates are reported in PageFair & Adobe (2015). Do Not Call shares are reported in Federal Trade Commission (2015). “Internet Privacy (topic)” and “Privacy (topic)” search volumes are indexed by Google Trends. Country-level income data are from the International Monetary Fund; age and education are from the United Nations. Proportion of respondents who affirm changing their browser privacy setting comes from European Commission’s Flash Eurobarometer 443: e-Privacy survey.

^aEU countries are listed in Figure 3.

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

significant associations for all variables. Compared with the opt-out share at the median income, the opt-out share is higher at the extremes of the income distribution, though differences are more pronounced at higher-end incomes. This result could reflect the notion that the rich have more to lose and is in line with a Pew Research Center (2014) survey showing that higher-income individuals report most forms of personal data as being more sensitive. We find the same qualitative pattern for age: the opt-out share increases at a faster rate for cities with older populations. The positive relationship with age is consistent with an array of survey evidence, including recognition of the AdChoices icon (Kelly Scott Madison & ORC International 2015), dislike of online behavioral advertising (Turow et al. 2009), and response rates for more private survey questions (Goldfarb and Tucker 2012). We find a negative association between opt-out and the college graduate share. We found this surprising, though we speculate that more educated users may consider a wider set of privacy alternatives, like Ghostery and Privacy Badger. The positive relationship with nondefault browser share accords with the aggregate evidence in Table 2, lending credence to the city-level findings. Though we expect that privacy cuts across political lines, we find that cities with greater Republican vote share have higher opt-out rates. The opt-out rate falls with both the Asian- and African-American population shares, which may reflect different privacy preferences and resonance with AdChoices in those groups. However, the share of Asian (African) Americans is positively (negatively) related with ad blocking, which indicates a preference for alternatives to AdChoices among Asian Americans.¹⁴

We also consider US state-level and country-level analyses in columns (2) and (3) of Table 6, though these have low power and indicate few statistically significant relationships. In the state-level analysis, we see that opt-out again increases with age but that education and nondefault browser reverse sign. Opt-out rates fall with income, but the narrower income range in the state-level data excludes high average incomes observed in column (1) of Table 6. In addition to the variables in column (1) of Table 6, we include the adoption of a prominent alternative in the share of ad-block users (PageFair & Adobe 2015), willingness to opt out of marketing given by the Do Not Call list (Federal Trade Commission 2015), and privacy interest using Google Trends data for both the “Privacy” and “Internet Privacy” topics. Of these, only the “Internet Privacy” search index is weakly statistically significant ($p < 0.1$) but has the unexpected sign (negative). Column (3) of Table 6 analyzes the opt-out rates for the 28 EU countries given in Figure 3. We see that none of the coefficients are statistically

significant, including a survey measure for whether users change their browser privacy settings. Because an estimated 16% of US users blocked ads at the time (PageFair & Adobe 2015), the share of exchange opt-outs can be said to understate the population share of opt-out consumers. Nonetheless, the results in Table 6 are similar when we account for ad blocking by transforming the outcome to be $\frac{s}{1-b}$, where b is the country- or state-level ad-block rate, as well as when we instead use a linear probability model.¹⁵

In addition to demographic differences, opt-out users exhibit different tastes for online content. Figure 6 in Online Appendix B illustrates differences in opt-out rates by content category. “Games” websites have the largest share of opt-out impressions in the United States. This is followed by “Arts & Entertainment,” “Internet & Telecom,” “News,” and “Education.” On the opposite end, “Law & Government” websites have the lowest opt-out rates, followed by “World Localities,” “Recreation,” and “Books & Literature.” The high share among “Games” and “Internet & Telecom” is consistent with more sophisticated users opting out, though “Computers & Electronics” content sites have relatively low opt-out rates. These differences in opt-out user characteristics imply heterogeneous impacts of the opt-out policy among publishers and advertisers.

5. Conclusion

We present a novel assessment of the AdChoices self-regulatory program. We calculate a modest loss from the inability to behaviorally target opt-out users of \$5.7 million (\$8.58 per opt-out consumer) in the US desktop display ad market in 2015. We find that only 0.23% of exchange impressions arise from Americans who opt out of online behavioral ads, which reduces the scope for harm to industry. This contrasts with survey evidence where the majority of consumers oppose online behavioral ads.

Our study is limited in some aspects. For instance, we observe data for a single week from a single ad exchange, and thus, our results may not represent the market as a whole. Our fingerprinting technique is inexact because of incomplete data sampling and because of users sharing a fingerprint. Despite our inventory- and user-level covariates, our preferred specification may yet omit important unobservables. Our difference-in-differences check trades the user-level unobservable problem for one of representativeness of switching users. Nevertheless, our various strategies provide similar estimates of the effect of opting out on ad prices, which is reassuring. Finally, our analysis of opt-out user characteristics aggregates impressions at a coarse regional level.

Both stated and revealed preference measures can aid policymakers. Indeed, our evidence can be interpreted in two opposing ways. The first interpretation emphasizes the low opt-out rate as evidence that few consumers are concerned enough about their privacy to take this action. The industry has a relatively low willingness to pay for an individual user’s browsing information, but this scales to a large value over the entire population. As such, this first interpretation favors an opt-out policy where privacy-concerned users can choose to avoid tracking, but the industry benefits from the majority of consumers who remain.

A second interpretation favors consumer-stated preferences and concludes that the AdChoices program failed to allow most privacy-concerned consumers to avoid online behavioral advertising. In a related setting, Hoofnagle (2006) points out that the telemarketing self-regulatory body collected fewer than five million opt-outs, whereas the Federal Trade Commission’s more usable Do Not Call list has now enrolled 222 million phone numbers (Federal Trade Commission 2015). The second interpretation favors more aggressive policies, such as an opt-in policy, where industry cannot track consumers without their explicit consent. In the EU’s General Data Protection Regulation, Europe’s *de jure* approach includes opt-in consent as grounds for data processing, though the current *de facto* approach is site-level opt-out consent. Nonetheless, opt-out rates are reportedly still under 10%, even when users are required to make an active, salient, and frequent choices on their privacy settings (Quantcast 2018).

A third policy proposal envisions a market where consumers can sell the right to their personal information. In our setting, consumers could sell the right to their online browsing information to the ad industry. Our analysis suggests that the ad industry would only be willing to spend up to about \$8.58 annually for the right to behaviorally target the average opt-out consumer. Though we cannot estimate consumer’s valuation for their online browsing privacy with our data, the \$8.58 figure places a modest limit on the price that consumers could expect to receive for their browsing data.

Privacy policy here contrasts with the retirement savings problem, which also exhibits a gap between consumer’s stated and revealed preferences. Although changing saving defaults can ameliorate this gap (Madrian and Shea 2001), the long-run welfare implications of changing the default are largely beneficial for all. However, a no behavioral advertising default offers no such “win-win” scenario: losing online behavioral advertising for the majority of users would have severe implications for advertisers, publishers, and ad exchanges (Wilbur 2008, Anderson and Gans 2011).

Acknowledgments

The authors thank S. S., D. W., and C. W. from our data partner for their assistance. They also thank Eric Anderson, Brett Gordon, Li Jiang, Klaus Miller, Navdeep Sahni, Anna Tuchman, Florian Zettelmeyer, and the editorial team for helpful comments.

Endnotes

¹ Beales and Eisenach (2014) and recently, Marotta et al. (2019) use ad exchange transaction data to show that users without exchange cookies fetch lower prices than cookied users. However, these papers ignore the AdChoices program and user opt-out choices. These papers also do not model the selection of noncookied users in terms of their browsing history.

² In the case of unsold impressions, the user is typically served a public service announcement ad or the publisher's default ad. Neither generates revenue for ad industry participants.

³ We also note that because of sampling at the user level for cookied impressions, ours is a stratified cluster sampling design where users are clusters and cookie types define strata. As such, we cluster the standard errors for our regression coefficients at the user level for cookied impressions and at the fingerprint level (see Section 3.2.1) for cookieless and opt-out impressions. We found this approach to be the most conservative of the parameter variance estimators that we explored.

⁴ AdChoices permits user identifiers to facilitate ad measurement and frequency capping. Therefore, storing IP address information does not itself violate the program as long as IP addresses are not used for behavioral advertising.

⁵ One potential confound is that these two browsers have Protect My Choices extensions, which would reduce the chance of opt-out cookie deletion. However, the DAA also offers the Protect My Choices option for Internet Explorer, and its opt-out rate is the lowest of the top four browsers at 0.183%. Safari has no such extension, and its opt-out rate is 0.232%.

⁶ Yahoo.com is a hypothetical example: domains and publishers are anonymized in our data.

⁷ Ignoring inventory and history controls in Equation (1), the predicted price for a cookied impression is given by $\hat{p}^{\text{Cookied}} = \exp(\alpha + \frac{\sigma^2}{2}) - 1$ and that for opt-out is given by $\hat{p}^{\text{Opt-out}} = \exp(\alpha + \beta + \frac{\sigma^2}{2}) - 1$, where σ^2 is the variance of the residual ϵ .

⁸ These estimates can be thought of as worst-case revenue loss estimates. To see why, consider that eliminating the opt-out option would shift users to being cookied, cookieless, or ad blockers—in decreasing order of ad revenue. As such, the worst-case revenue loss corresponds to the case where all opt-out impressions switch from cookied to opt-out. In the best case, all opt-out users would have been ad blockers, in which case all opt-out ad revenue is incremental.

⁹ Cookieless impressions arise from browser version defaults—or user settings—that block third-party cookies or clear cookies when a browsing session ends. Also, the cookie identifier look-up process can time out on occasion. Thus, we cannot identify which cookieless impressions represent a deliberate user privacy choice, though we expect the majority result from browser-level default settings.

¹⁰ For candidate switchers, we consider the 166,000 fingerprints that contain at least one opt-out and cookied impression from the 8,424,431 fingerprints in the US impression data. From this baseline, the preceding steps filter out the following proportion of fingerprints: (2) 63.4%, (3) 6.7%, (4) 0.3%, and (5) 24.85%. Our inferred switchers represent only 4.75% of baseline fingerprints.

¹¹ Our marginal effect estimates are robust to instead requiring at least 2, 10, or 25 impressions of each type per switcher as well as further

restricting to switchers who visit at least one site both as cookied and opt-out types. Results available from the authors on request.

¹² Alternately, we can compute this figure as \$6.42, which is the product of the per capita programmatic spend (\$12.34) and the marginal effect of opt-out (–52.0%). In the spirit of providing a worst-case estimate, we present the higher figure of \$8.58 because it accounts for the higher value of inventory among opt-out impressions.

¹³ As described in Section 3.2, we use probability weights to account for our stratified sampling design. Here we calculate impression-level opt-out rates as weighted averages by state and city, as identified in our estimation data set.

¹⁴ Results are available from the authors on request.

¹⁵ Results are available from the authors on request.

References

- Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? *J. Legal Stud.* 42(2):249–274.
- Acquisti A, Taylor C, Wagman L (2016) The economics of privacy. *J. Econom. Literature* 54(2):442–492.
- Anderson SP, Gans JS (2011) Platform siphoning: Ad-avoidance and media content. *Amer. Econom. J. Microeconomics* 3(4):1–34.
- Athey S, Catalini C, Tucker C (2019) The digital privacy paradox: Small money, small costs, small talk. Working paper, Stanford Graduate School of Business, Stanford, CA.
- Aziz A, Telang R (2016) What is a digital cookie worth? Preprint, submitted March 31, <http://dx.doi.org/10.2139/ssrn.2757325>.
- Beales JH, Eisenach JA (2014) An empirical analysis of the value of information sharing in the market for online content. Technical report, Navigant Economics, Chicago.
- Becker GS, Murphy KM (1993) A simple theory of advertising as a good or bad. *Quart. J. Econom.* 108(4):941–964.
- Berendt B, Günther O, Spiekermann S (2005) Privacy in e-commerce: Stated preferences vs. actual behavior. *Comm. ACM* 48(4): 101–106.
- Berry ST (1994) Estimating discrete-choice models of product differentiation. *RAND J. Econom.* 25(2):242–262.
- Brill J (2011) The intersection of consumer protection and competition in the new world of privacy. *Competition Policy Internat.* 7(1):7–23.
- Budak C, Goel S, Rao J, Zervas G (2016) Understanding emerging threats to online advertising. *Proc. 2016 ACM Conf. Econom. Comput.* (ACM, New York), 561–578.
- Chen J, Stallaert J (2014) An economic analysis of online advertising using behavioral targeting. *Management Inform. Systems Quart.* 38(2):429–449.
- Creamer M (2011) Despite digital privacy uproar, consumers are not opting out. *Ad Age* (May 30), <http://adage.com/article/digital/digital-privacy-uproar-consumers-opting/227828/>.
- Eckersley P (2010) How unique is your web browser? *Proc. Internat. Sympos. Privacy Enhancing Tech.* (Springer, Berlin, Heidelberg), 1–18.
- eMarketer (2015) US programmatic ad spending: eMarketer's forecast for 2015 and beyond. Accessed October 8, 2015, <https://www.emarketer.com/Report/US-Programmatic-Ad-Spending-eMarketers-Forecast-2015-Beyond/2001591>.
- eMarketer (2016a) US digital display ad spending to surpass search ad spending in 2016. Accessed July 17, 2017, <https://www.emarketer.com/Article/US-Digital-Display-Ad-Spending-Surpass-Search-Ad-Spending-2016/1013442>.
- eMarketer (2016b) US programmatic ad spending forecast: Most mobile display and video ad dollars to be automated by 2018. Accessed July 17, 2017, <https://www.emarketer.com/Report/US-Programmatic-Ad-Spending-Forecast-Most-Mobile-Display-Video-Ad-Dollars-Automated-by-2018/2001894>.

- Farahat A, Bailey MC (2012) How effective is targeted advertising? *Proc. 21st Internat. Conf. World Wide Web* (ACM, New York), 111–120.
- Federal Trade Commission (2009) Self-regulatory principles for online behavioral advertising. Technical report, Federal Trade Commission, Washington, DC.
- Federal Trade Commission (2012) Protecting consumer privacy in an era of rapid change. Technical report, Federal Trade Commission, Washington, DC.
- Federal Trade Commission (2015) National do not call registry data book FY 2015. Technical report, Federal Trade Commission, Washington, DC.
- Gart JJ, Pettigrew HM, Thomas DG (1985) The effect of bias, variance estimation, skewness and kurtosis of the empirical logit on weighted least squares analyses. *Biometrika* 72(1):179–190.
- Goldfarb A, Tucker C (2011) Privacy regulation and online advertising. *Management Sci.* 57(1):57–71.
- Goldfarb A, Tucker C (2012) Shifts in privacy concerns. *AER: Papers Proc.* 102(3):349–353.
- Gross R, Acquisti A (2005) Information revelation and privacy in online social networks. *Proc. 2005 ACM Workshop Privacy Electronic Soc.* (ACM, New York), 71–80.
- Gunningham N, Rees J (1997) Industry self-regulation: An institutional perspective. *Law Policy* 19(4):363–414.
- Hoofnagle CJ (2006) Privacy self-regulation: A decade of disappointment. Winn JK, ed. *Consumer Protection in the Age of the 'Information Economy'* (Routledge, London), 379–401.
- Johnson G (2013) The impact of privacy policy on the auction market for online display advertising. Preprint, submitted May 21, <http://dx.doi.org/10.2139/ssrn.2333193>.
- Kelly Scott Madison & ORC International (2015) Do consumers actually care about online privacy? Technical report, Kelly Scott Madison, Chicago.
- Kummer M, Schulte P (2019) When private information settles the bill: Money and privacy in Google's market for smartphone applications. *Management Sci.* 65(8):3470–3494.
- Laudon KC (1996) Markets and privacy. *Comm. ACM* 39(9):92–104.
- Leon P, Ur B, Shay R, Wang Y, Balebako R, Cranor L (2012a) Why Johnny can't opt out: A usability evaluation of tools to limit online behavioral advertising. *Proc. SIGCHI Conf. Human Factors Comput. Systems* (ACM, New York), 589–598.
- Leon PG, Cranshaw J, Cranor LF, Graves J, Hastak M, Ur B, Xu G (2012b) What do online behavioral advertising privacy disclosures communicate to users? *Proc. 2012 ACM Workshop Privacy Electronic Soc.* (ACM, New York), 19–30.
- Levin J, Milgrom P (2010) Online advertising: Heterogeneity and conflation in market design. *AER: Papers Proc.* 100(2):603–607.
- Madrian BC, Shea DF (2001) The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quart. J. Econom.* 116(4):1149–1187.
- Maier F (2010) Don't fear the opt-out. *TrustArc Blog*. Accessed November 16, 2010, <https://www.trustarc.com/blog/2010/11/16/dont-fear-the-opt-out/>.
- Mandese J (2018) Few Americans choose AdChoices, know it even exists research.
- Marotta V, Abhishek V, Acquisti A (2019) Online tracking and publishers' revenues: An empirical analysis. Working paper, University of Minnesota, Minneapolis.
- Mastria L (2014) Testimony. *Hearing on Online Advertising and Hidden Hazards to Consumer Security and Data Privacy* (US Senate Committee on Homeland Security & Governmental Affairs), Washington, DC.
- Mastria L (2015) DAA's 2014 year in review: Momentum, engagement & presence. Accessed January 12, 2015, <http://digitaladvertisingalliance.org/blog/daas-2014-year-review-momentum-engagement-presence>.
- McDonald AM, Cranor LF (2010a) Americans' attitudes about Internet behavioral advertising practices. *Proc. 9th Annual ACM Workshop Privacy Electronic Soc.* (ACM, New York), 63–72.
- McDonald AM, Cranor LF (2010b) Beliefs and behaviors: Internet users' understanding of behavioral advertising. *Proc. 38th Res. Conf. Comm. Inform. Internet Policy* (Citeseer).
- Miller KM, Skiera B (2017) Economic damage of cookie lifetime restrictions. Working paper, NET Institute, New York.
- Neumann N, Tucker CE, Whitfield T (2019) How effective is third-party consumer profiling and audience delivery? Evidence from field studies. *Marketing Sci.* 38(6):918–926.
- PageFair & Adobe (2015) The cost of ad blocking: PageFair and Adobe 2015 ad blocking report. Technical report, PageFair and Adobe, Dublin, Ireland.
- Pancras J, Sudhir K (2007) Optimal marketing strategies for a customer data intermediary. *J. Marketing Res.* 44(4):560–578.
- Parks Associates (2014) Harnessing the power of big data: New media and advertising. Technical report, Parks Associates, Addison, TX.
- Pew Research Center (2014) Public perceptions of privacy and security in the post-Snowden era. Technical report, Pew Research Center, Washington, DC.
- Purcell K, Brenner J, Rainie L (2012) Search engine use 2012. Technical report, Pew Research Center, Washington, DC.
- Quantcast (2018) Quantcast Choice powers one billion consumer consent choices in two months since GDPR. Press release, Quantcast, San Francisco.
- Rafieian O, Yoganarasimhan H (2018) Targeting and privacy in mobile advertising. Preprint, submitted May 3, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3163806.
- Rossi PE, McCulloch RE, Allenby GM (1996) The value of purchase history data in target marketing. *Marketing Sci.* 15(4):321–340.
- Shiller B, Waldfogel J, Ryan J (2018) The effect of ad blocking on website traffic and quality. *RAND J. Econom.* 49(1):43–63.
- TRUSTe (2016) Solutions brief: Ads privacy compliance. Technical report, TRUSTe.
- TRUSTe & European Interactive Digital Advertising Alliance (2014) European advertising consumer research report 2014. Technical report, TRUSTe & European Interactive Digital Advertising Alliance, San Francisco.
- TRUSTe & European Interactive Digital Advertising Alliance (2016) European advertising consumer research report 2015. Technical report, TRUSTe & European Interactive Digital Advertising Alliance, San Francisco.
- Tuchman AE, Nair HS, Gardete PM (2018) Television ad-skipping, consumption complementarities and the consumer demand for advertising. *Quant. Marketing Econom.* 16(2):1–64.
- Tucker C (2014) Social networks, personalized advertising, and privacy controls. *J. Marketing Res.* 51(5):546–562.
- Tugend A (2015) Key to opting out of personalized ads, hidden in plain view. *New York Times* (December 20), <https://www.nytimes.com/2015/12/21/business/media/key-to-opting-out-of-personalized-ads-hidden-in-plain-view.html>.
- Turow J, King J, Hoofnagle C, Bleakley A, Hennessy M (2009) Americans reject tailored advertising and three activities that enable it. Preprint, submitted September 29, <http://dx.doi.org/10.2139/ssrn.1478214>.
- Varian HR (2009) Economic aspects of personal privacy. Lehr W, Pupillo L, eds. *Internet Policy and Economics* (Springer, Boston), 101–109.
- Wilbur KC (2008) A two-sided, empirical model of television advertising and viewing markets. *Marketing Sci.* 27(3):356–378.
- Yan J, Liu N, Wang G, Zhang W, Jiang Y, Chen Z (2009) How much can behavioral targeting help online advertising? *Proc. 18th Internat. Conf. World Wide Web* (ACM, New York), 261–270.