Abstract

Conventional wisdom holds that the market for digital privacy fails owing to widespread informational asymmetry between digital firms and their customers, behavioral biases exhibited by those customers, and negative externalities from data resale. This paper supplies both theoretical and empirical reasons to question the standard market failure conclusion. On the theoretical side, I argue that digital markets are not qualitatively different from markets for other consumer goods. To wit, just as in traditional markets, it is costly to measure product attributes (such as "privacy") and, just as in more traditional settings, some firms offer credible commitments to reduce the threat of potential opportunism. On the empirical side, I conduct a survey of Google's users. The most important results of this survey suggest that, at least with respect to Google, (a) the extent of

1 Introduction

Between 2000 and 2018, Google’s unofficial motto was “Don’t Be Evil”, but the company’s surreptitious collection of information from more than one billion individuals annually has prompted some commentators to question whether its business model contradicts its famous dictum (Hoofnagle 2009).[1] Does information collection align with consumer preference, as argued by, e.g., Cooper (2012), or is there a disconnect between the two, as argued by others (Strandburg to the firm’s advantage.” Acquisti (2004, p. 7) concurs that behavioral biases contribute to the prevalence of information collection: “individuals who ... would like to protect their privacy may not do so because of psychological distortions well-documented in the behavioral economics literature.” For example, immediate-gratification and status-quo biases may cause even well-informed individuals to permit more information to be collected from them than is in their ultimate, long-

Strandburg 2013; Newman 2013; Acquisti et al. 2016).[2] Those drawing that conclusion also often appeal to the well-documented “privacy paradox”, contribute to privacy concerns. Section 5 concludes with a few implications.

2 Digital privacy and market failure

Web platforms collect “non-sensitive” information directly from visitors or allow third parties (advertisers) to use the site as a platform for information collection (Goldfarb and Tucker 2011; de Corniere and de Nijs 2016). Humorously referred to as “mouse droppings”, non-sensitive information usually consists of a user’ argue, “Information asymmetries regarding the usage and subsequent consequences of shared information raise questions regarding individuals’ abilities, as rational consumers, to optimally navigate privacy trade-offs.” Such a perspective grants perfect rationality to consumers for the sake of argument, but still concludes that an inefficient outcome emerges as a result of asymmetric information. Hirsch (2010, p. 455) claims that the primary objection to a *laissez-faire* approach to digital privacy is pervasive information asymmetry. Brown (2016, p. 5) concurs that consumers have “limited knowledge” of how digital platforms will use the information they collect. Tucker (2012, p. 328) emphasizes that consumers may not know what information is being collected and contends that “there is a need for empirical work that attempts to understand the extent of informational asymmetry between consumers and firms...about how much data are being collected...”

An overview of the economics of privacy in the *Journal of Economic Literature* sums up the consensus on information asymmetry in digital contexts by arguing that consumers are ignorant of when a firm is collecting information, what information it is collecting, or how the information will be used after collection (Acquisti et al. 2016).

2.3 Consumers’ behavioral biases

Surveys suggest that consumers value their privacy highly. Turow et al. (2009, p. 4) write that, “It is hard to escape the conclusion that our survey is tapping into a deep concern by Americans that marketers’ tailoring of ads for them and various forms of tracking that informs those personalizations are wrong.” Turow et al. (2009, p. 10) conclude from survey evidence that “it seems clear ... that Americans value the right to opt out from this sort of collection.” Acquisti (2004) cites older surveys that generate similar conclusions. For example, a 2002 Harris Interactive

Survey found that companies collecting personal information without prior consent was one of web consumers’ most significant concerns (Acquisti 2004).

Why then do so many consumers continue to patronize privacy-invasive services, such as Google, that track consumers? One possibility is that consumers are prone to myriad behavioral biases causing them to behave contrary to their true preferences, as elicited by surveys. Owing to bounded rationality, consumers rely on “simplified models” and “heuristics” that generate deviations from perfectly rational outcomes (Acquisti 2004; Brown 2016). That view emphasizes that consumers are poor judges of cumulative risk. They also tend to underestimate occurrences of low probability events (Acquisti 2004; Brown 2016). Consumers concurrently are plagued by immediate-gratification bias, which magnifies the rewards from engaging in risky privacy behaviors, while minimizing perceptions of potential threats (Acquisti 2004; Brown 2016).

The practical implication of widespread behavioral biases is that a gap exists between consumers’ *true* preferences (their “atti]TJ 05 (al) (t-2 (s)]TJ 0.001 Tc 0.0043.39.735 0 TTd (2016een)6 (

behar

value such information, so the initial platform sells it to them. Consumers may be both perfectly rational and perfectly informed about the initial act of collection. However, according to Brown (2016, p. 5), selling consumer data to a third party “imposes the cost of future invasive advertising on a data subject without compensation.” In the view of Acquisti et al. (2016, p. 452), such negative externalities may consist of “spam” and “adverse price discrimination.” For example, price discrimination might be facilitated when a digital merchant tracks a buyer’s browsing history or geographic location to better estimate the individual’s elasticity of demand.

Varian (2009), Acquisti et al. (2016, p. 452), and Odlyzko (2003) concur regarding the threat of unauthorized third-party information use. Varian (2009) provides the example of a mailing list, collected initially by a single advertiser, who subsequently sells the list to other advertisers. Resale imposes a cost on anyone who is contacted in the future by advertisers who have gained unauthorized access to his or her home or business address.

3 Privacy market failure: Theory and evidence

If the foregoing arguments are backed by theory and evidence, the case for government regulation of digital privacy is bolstered. To the extent that the arguments are not easily supportable, the case becomes weaker. By examining the claims theoretically and empirically, this paper contributes to a debate in the economics of digital privacy literature: is the digital marketplace prone to failure (Acquisti 2004)?

On the empirical side, I conducted a survey of 6,883 Internet users. Nineteen of them were disqualified for the following reasons. Ten were removed for technical reasons. For instance, sometimes using an unusual browser can cause an answer not to be recorded. The other nine respondents were discarded for responding with highly unusual answers (see the discussion in section 3.2.2 for additional details). Those considerations reduced the number of valid respondents

Surveys are subject to criticism. The results may lack external validity for at least two reasons. First, the value of privacy differs across cultures and contexts (Milberg et al. 2000; Rose 2005). My results generate insight into a specific context (interactions with Google) at a specific time and in a specific place (the United States in the year 2018). Second, it is difficult to establish the randomness of the sample. As Turow et al. (2009) have noted, people who respond to an online survey may be less privacy-sensitive than those who do not. Respondents also may tend to be better informed about digital policies and practices, considering that they know how to participate in an online survey. Because a variety of panels were used, and because those panels solicit and reward respondents differently, at least some variation exists in the selection of respondents, and self-selec including used cars, health insurance, credit and labor. In each of those cases, asymmetric information precipitates a reduction in the number of mutually beneficial exchanges; in the limiting case, the market collapses altogether. For example, in the familiar Akerlof story about used cars, buyers continue to lower their bids as the average quality of cars on the lot falls, until no high-quality cars ("cream puffs") are offered for sale. The key to establishing the existence of market failure is that asymmetric information between buyer and seller causes fewer mutually beneficial motor vehicles also is risky. Laptops are vulnerable to damaging viruses; operating a motor vehicle may cause death.

In each of those cases, goods can be conceived of as bundles of attributes that are costly to measure (Barzel 1982).[4] Motor vehicles do not simply transport their occupants from “A” to “B.” The “experience” of getting from “A” to

attributes offered. Such substitutes might provide the ability to "search" as Google does without Google's privacy-invasive practices. Just as we observe vehicles comprised of alternative bundles of attributes, so we should expect the market for privacy to be characterized by firms occupying a spectrum of privacy policies, some of which cater specifically to privacy-sensitive users.

That argument has not gone uncontested, however. Some digital privacy scholars have argued that such a spectrum is unlikely to contain firms offering a bundle of attributes that prioritizes consumer privacy. The critics have argued that privacy on the Internet devolves into a "race to the bottom", a prisoner's dilemma (Hoofnagle 2003). In that view, collecting consumer information *always* is the profit-maximizing strategy and firms therefore will search for increasingly sophisticated techniques for acquiring that information, regardless of consumers' privacyw [illegible] location of Google's headquarters), proclaiming boldly that "Google tracks you. We don't." The proclamation serves as a hostage to potential DuckDuckGo users. Were DuckDuckGo to renege on its promise, privacy-sensitive users likely would abandon the service in droves. Destroying the investment's value would require only a single customer to discover DuckDuckGo's breach and to initiate a multilateral boycott by publishing that fact publicly.[8] Credible commitments also mitigate a potential "future-proofing" problem regarding privacy (Acquisti et al. 2016, p. 478). A consumer may be perfectly content with Google's *current* data practices, but might still prefer to conceal information from Google, if only to prevent future uses of which she disapproves. Investment in a reputational hostage thus mitigates future privacy invasions, as it raises DuckDuckGo's costs of reneging on its promise.

Perhaps in a world of fully-informed individuals, DuckDuckGo's traffic would dwarf Google's. If, however, consumers generally are well-informed about information-collection practices, yet *persist* in demonstrating a preference for browsing services that rely on that technique for monetizing such information, the case for a regulatory fix becomes even less compelling.

3.1.2 Empirical evidence

The empirical question remains as to how informed consumers are about online information collection. That question is relevant because the existence of many well-informed consumers improves the efficacy of the mechanisms described in Section 3.1.1.

The survey results suggest that many consumers indeed are relatively well-informed. When

[8] Alt

question five is 56%. By contrast, averaging across the incorrect options generates a selection rate of 43%. Most respondents select both correct and incorrect options, but correct answers are chosen more frequently.

Respondents clearly are far less well-informed about how Google uses their data than that personal information is collected. Is it then reasonable to conclude that consumer behavior would be different if they were better informed? If so, how might it differ? The most relevant question appears to be whether consumers are aware that Google unilaterally could enact—that is, without consumers' consent—a new privacy policy at any moment. A new policy hypothetically could permit data uses that the current policy prohibits. If consumers are unaware of that possibility, they may not be willing to pay as much for privacy because they fail to see the benefits inherent in "future-proofing" their information. On the other hand, if respondents are aware that Google could implement a new policy at any time, the fact that they are not particularly informed regarding Google's current policy becomes less important.

To address those issues, the survey next asked a question regarding consumer awareness of how privacy policies work. The empirical results show that consumers are quite aware that Google's privacy policy is, at best, tentative. Those respondents who know Google collects information (J= 5,434) were asked: "Do you believe that Google could change its privacy policy to allow new uses for user data?" A large majority—85%—answer "yes." Thus, most consumers know that they are writing Google a "blank check" when they visit the site. That evidence suggests that concerns regarding future contingencies should be captured in their stated willingness to pay (WTP) for privacy.

The survey's results do not rule out the existence of information asymmetry, but nor should we expect them to. Costly as information about goods' attributes is to obtain, perfect information never is possible in the real world. Nonetheless, the results reveal the existence of many highly informed consumers. And because those consumers have both substitutes available to them and low-conceal personal information in the context of smartphone usage. They find relatively small one-time willingness to pay to conceal such information as browser histories ($2.28), cell phones identification numbers ($1.75), text messages ($3.58), locations ($1.19) and contact lists ($4.05).

Other surveys do not ask consumers to put a price on privacy. For example, Turow et al.'s (2009) survey asks questions like: "Please tell me whether or not you want the websites you visit to show you ads that are tailored to your interests." Finding that a significant percentage respond negatively to queries like that one, the authors conclude that governments should impose opt-in default options or set time limits on data preservation.[15] Turow et al. (2009) likewise find that 66% of respondents are "uncomfortable" with targeted ads, while a 2015 Pew Research Report says that 93% of Americans believe that being in control of who can access their information is important (Madden and Rainie 2015).

That type of query—one that reveals preferences for a higher quality good, *ceteris paribus*—is what might be called an "unconstrained approach" to privacy valuation. Unconstrained survey questions fail to remind consumers that acquiring a good with a more satisfactory bundle of attributes imposes an opportunity cost that they necessarily bear. Such an approach thus is not strictly "economic" because no tradeoffs are involved. We therefore should *expect* to see a difference between "talk" and "action" with those kinds of surveys, and we should expect to see a gap even in the complete absence of any behavioral biases. One likewise might expect individuals to articulate preferences for higher incomes, lower buying prices, higher selling prices, better working conditions and nicer friends, *ceteris paribus*.[16]

[15] Tucker and Goldfarb (2011) examine the economic impact of the EU's switch to an opt-in rather than an opt-out default option. They find that the switch reduced the effectiveness of the average digital ad dramatically because of the inability to target advertisements. Lerner (2012) finds that the EU's rules have lowered investments in ad-supported European firms.

[16] Unconstrained surveys also are common in other contexts. For example, see Clark and Powell's (2013) analysis of "non-economic" or "unconstrained"

The economic approach insists on using “constrained questions.”[17] That approach is superior to unconstrained ones because only tradeoffs, not solutions, are open to individuals choosing in the face of constraints. For example, a seller asking a low money price thereby is enabled to ask for more non-pecuniary equalizing differentials (Alchian 1967). In the case of Google, the firm asks a zero-money price, enabling it to collect a positive quantity of consumer information.[18]

The constrained approach suggests a straightforward resolution to the differences between what consumers say and what they do. Whereas Acquisti et al. (2016) argue that the gap between “privacy attitudes and privacy behaviors” arises because of “many, coexisting, and not mutually exclusive factors”, such as “asymmetric information, bounded rationality, and various heuristics”, my approach suggests that it can be explained by the difference between “constrained” and “unconstrained” survey questions. Unconstrained questions present achievement of privacy as a costless endeavor; constrained questions remind respondents that something must be sacrificed to attain privacy. In my questionnaire, the “something” is money, but in the real world it might be the convenience of searching online, the time invested in discovering privacy-protective services (such as DuckDuckGo or Adblock Plus), or even the benefits (for some consumers) of receiving targeted ads.

If a gap exists between s*tated* responses to “constrained” and “unconstrained” surveys, we would have evidence (though not conclusive evidence) that the difference between what consumers “say” and “do” can be explained without recourse to behavioral biases. A large stated WTP is evidence for divergence between “true” preferences (verbally expressed) and the

[17] Acquisti et al. (2016, pp. 44-445) affirm that both costs and benefits are associated with disclosure of personal information.

[18] Non-money differentials may include preferenceu f

between “constrained” and “unconstrained” preferences. Of those respondents who both voluntarily use Google and prefer not to be tracked, the overwhelming majority are unwilling to pay anything at all to achieve privacy. Indeed, 86% of Google users are unwilling to pay for privacy on Google’s search engine (J= 6,083).[20]

Just how intense are the stated demands for privacy on the part of the 14% of respondents in the minority? On average, their WTP for privacy is small. Nine respondents who entered an annual value of $10,000 or greater were dropped from the survey’s results on the basis that such stated amounts likely were errors or represented unserious responses. Among those respondents kept in the sample and indicating a positive WTP (J= 824), the average annual WTP was $59.59. Since all respondents in the sample report that they use Google at least once daily, it makes sense to convert that figure into daily WTP terms. The average daily WTP equals about 16 cents.

Even after having removed the nine values exceeding $10,000, the mean is still driven by several outliers, as evidenced by a standard deviation of 150.11; the median is thus a more representative measure. The median annual WTP is $25 annually. In other words, of the roughly 14% of respondents willing to pay to protect their information, only half are unwilling to pay more than $25 per year. This annual WTP converts to between six and seven cents daily. Seeing as how the average American household spends, on average, 34 times as much on soft drinks per dayin.1 (e)- -19..3

the same respondents also were asked about their willingness to purchase privacy on a “per-search” basis. Oarch their willingness to pay $70 annually to protect their privacy on Google’s search engine. Roughly 45% of those willing to pay for privacy indicate willingness to pay the $70 fee. That result translates into about 6% of all Google users in the survey. If Google charged members of that group $70 per year, total revenue would amount to around $4.2 billion annually.

Low WTP for privacy is significant given that Section

protective technologies. Some respondents already might have "paid" for privacy by investing in the search for a complementary browsing technology enabling them to consume Google without unwanted privacy intrusions. Such respondents have purchased a higher-quality Internet experience, but their purchase comes at the expense of time invested in search, as many adblockers can be installed at no charge. At the same time, users who have installed a privacy-protective technology likely are to be more privacy-sensitive than the average respondent. Thus, it is possible that the most privacy-sensitive users indicate *little* WTP, already having satisfied their demands for privacy by way of adblocking technologies.

To understand the magnitude of that potential issue, I asked respondents whether they use a privacy-protecting technology—such as Adblock Plus—while browsing. Of the total number of respondents (J = 6,083), 39% do so. Among those respondents who *do* employ a means of privacy protection, 21% say they are willing to pay for privacy, whereas only 16% of non-ad-block users are willing to pay (J = 4,621).[25] While a larger percentage of respondents using privacy-protective options are willing to pay, their average WTP is smaller at $52.48, in comparison with the $64.81 average WTP of respondents who do not use a privacy-protective technology (J = 824). That difference is not statistically significant at the 5% level, however (t-statistic = F1.28). Taken together, the foregoing results seem ambiguous and do not suggest strongly that privacy-protective technology users are biasing the WTP questions systematically. On the one hand, the larger percentage of users expressing willingness to pay would suggest that members of the ad-block-using group are more privacy-sensitive, even after having invested in a technology to protect themselves. On the other hand, the dollar figures suggest that that group is willing to pay less,

[25] The respondents are comprised of those who prefer not to have their information collected (including those both willing and unwilling to pay for privacy).

having already secured their privacy by alternative means.

Suppose that the latter possibility—that the privacy-sensitive respondents express a lower WTP because of already having secured their own privacy—is the dominant effect. Far from being evidence of market failure, however, such a possibility would serve merely to highlight the wide array of services that permit the privacy-sensitive to alter the attributes of the Internet good they are consuming. The situation would parallel the ability of the most safety-conscious car consumers to pay for an add-on option that strengthens vehicle safety. That the car lacked such a characteristic before the consumer purchased the add-on is a feature, not a bug. The absence of the safety feature permits car buyers with trivial demands for safety to purchase vehicles at a lower price. Meanwhile, buyers with high demands for safety can pay extra if they value the added features sufficiently. That steering wheels almost always come packaged with automobiles, but additional safety features do not, is a function of the costs associated with various components and diversity in customers' demands for them. Everyone wants to buy cars with steering wheels. It thus needlessly raises transaction costs for steering wheels to be a separately priced option. By contrast, many car buyers may not value an extra safety feature; it thus makes sense for that feature to be purchased separately. The same transaction cost logic can be applied to the purchase (either by money or time) of additional privacy features in digital contexts. Not everyone values such features more than their cost, but those who do can purchase them.

That reasoning and the empirical results reported above do not demonstrate that every consumer is unwilling to pay for privacy. However, the analysis of survey responses does reveal a significant difference between unconstrained and constrained preferences for privacy. The significance of that finding is that while behavioral biases cannot be ruled out conclusively, they may be superfluous for explaining the well-documented dichotomy between stated preferences and actual behavior. To explain behavior in digital environments, appeal to immediate-gratification bias need not be necessary or even helpful. Instead, consumers simply may be unwilling to bear the cost of obtaining a higher-quality search engine. The results also hint that more awareness generally may be inversely related to WTP, suggesting that many search engine users evaluate privacy harms as being negligible.

3.3 Resale externalities

3.3.1 Theory

Although it has received less attention than information asymmetry and behavioral biases, it is worth discussing a final market failure, namely third parties accessing personal information, thereby imposing a negative externality on the consumers initially relinquishing it (Hermalin and Katz 2004; Hui and Png 2005; Varian 2009; Acquisti et al. 2016). As Acquisti et al. (2016, p. 452) argue, "The firm may sell the consumer's data to third parties, which may lead to spam and adverse price discrimination, among other concerns.… Such negative externalities may not be internalized by the consumer nor by the firm that distributes the information."

First, price discrimination should not be categorized as a negative externality. Externalities refer to third-party effects that are not captured in the prices at which parties exchange. Price discrimination merely moves the price closer to a consumer's reservation price, but it does not impose uncompensated costs involuntarily on third parties. The price-discrimination claim also is puzzling because it is a consumer' negative externality on a digital user, the logical conclusion seems to be that *every* mutually beneficial exchange—in fact, every social interaction—is rife with the possibility of generating negative externalities.

Suppose that a well-informed individual voluntarily relinquishes personal information in exchange for accessing a digital service. After the exchange, the collecting platform sells the information to a third party who uses it to target ads, solicit business via email, or engage in price discrimination. Suppose further that the initial consumer dislikes that market outcome. Now compare the same scenario to a simple and "traditional" market exchange: "A" exchanges cash for "B's" good. After the transaction is consummated, B uses the cash for some purpose that imposes a cost—perhaps only a psychic cost—on A. For example, B donates the cash to a non-profit organization advocating a cause that A detests. Alternatively, B uses the cash to purchase a weapon that he then uses to harm A physically. Can we conclude that the initial exchange between A and B generated a negative externality because following the exchange B used what he gained from it to harm A? That is not an externality because the risk that B might do something harmful to A after the exchange is captured in the price at which A and B first trade. Instead, we might say that A has suffered a psychic loss from engaging in the exchange; by his own estimation, he would have been better off had he refrained from trading with B. But psychic losses are not market failures; they are a possibility in every transaction.

3.3.2 Empirical evidence

One additional problem with claiming that personal data resale is a negative externality is that some individuals positively *prefer* to receive targeted advertisements or email solicitations. As Varian (2009) notes, one of the reasons people receive so much "junk mail" (both physical and digital) is because potential sellers lack information about prospective buyers. If sellers possessed more information about buyers' attributes, they could target their solicitations to individuals with larger probabilities of buying.

Data resale enables information to flow to sellers attempting to more closely tailor their offerings. That observation may explain why 24% of survey respondents (J= 6,083) indicate a preference for Google continuing to collect their information. For them, "Google-without-targeted-ads" is a *lower* quality good than "Google-with-targeted-ads", possibly because showing consumers targeted ads lowers their search costs for products. That possibility is supported further by the fact that 24% of those preferring not to be tracked (J= 4,621) also indicate that they "like seeing the ads customized to my preferences." That is perhaps a surprising result given that the same respondents wish that Google would refrain from tracking, yet still want to see targeted ads. Of course, enjoying the latter depends on the former. The bottom line is that with a significant minority of digital users indicating a preference for receiving targeted ads, it seems wrong to conclude that information resale *universally* is viewed as a cost.

4 Do governments contribute to privacy hysteria?

Recent studies find that government surveillance programs exert a "chilling" effect on Internet search activity (Penney 2016; Marthews and Tucker 2017). If the threat of government surveillance acts as a constraint on digital activities, then government failure, rather than (or, at least in addition to) market failure contributes to distaste for information collection. Private information collection itself may not be sufficient for generating the level of discomfort expressed in unconstrained consumer surveys.

My survey briefly investigated *why* respondents dislike information collection by asking those who had indicated a preference for Google to refrain from gathering data about what motivated that answer. Acquisti et al. (2016, p. 483) summarize possible reasons why consumers might express dislike of online information collection: “price discrimination … spam … risk of identity theft … [and] the disutility inherent in just not knowing who knows what.”[26] The findings of the survey at hand ($J = 4{,}621$) provide general support for the conjectures offered by Acquisti et al. (2016). For example, 75% of respondents indicate concern regarding “the risk of identity theft.”

However, the findings also suggest that the literature largely has ignored an important reason explaining why individuals express dislike of digital information collection. Respondents to question ten were presented with seven options and were instructed to select as many of them as they found applicable.[27] Of respondents who would prefer Google not to collect information, 41% indicate that “

subject to the same negative externality critique, the case for unique regulation of digital privacy is weakened.

In the United States, however, policymakers continue to debate the merits of implementing comprehensive, EU-style regulation. As a Federal Trade Commission (2012, p. i) report on the topic states: “Although companies use this information to deliver better products and services to consumers, they should not do so at the expense of consumer privacy.” Such a value judgment is not supported well by this paper’s results.

References

commission-report-protecting-

of the 5th international conference on Electronic commerce, ACM, 355-366.

Penney, J. (2016). Chilling effects: online surveillance and Wikipedia use. *Berkeley Technology Law Journal*.

Sachs, B.R. (2009). Consumerism and information privacy: how Upton Sinclair can again save us from ourselves. *Virginia Law Review* 95(1), 205–252.

Savage, S.J., and Waldman, D.M. (2015). Privacy tradeoffs in smartphone applications. *Economics Letters* 137, 171-175.

Solove, D.J. (2004). *The digital person: Technology and privacy in the information age*. NYU Press.

Statista. (2017). Google's ad revenue from 2001 to 2017 "in billion US dollars). https://www.statista.com/statistics/266249/advertising-revenue-of-google/

Statista. (2018). Global digital population as of July 2018 (in millions). https://www.statista.com/statistics/617136/digital-population-worldwide/

Stigler, G.J. (1961). The economics of information. *Journal of Political Economy* 69(3), 213-225.

Strandburg, K.J. (2013). Free fall: the online market's consumer preference disconnect. *University of Chicago Legal Forum* 5, 95-172.

Tsai, J.Y., Egelman, S., Cranor, L., and Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research* 22(2), 254–268.

Tucker, C.E. (2012). The economics of advertising and privacy. *International Journal of Industrial Organization* 30(3), 326–329.

Turow, J., King, J., Hoofnagle, C.J., Bleakley, A., and Hennessy, M. (2009). Americans reject tailored advertising and three activities that enable it. *Available at SSRN 1478214*.

Varian, H.R. (2009). Economic aspects of personal privacy. In *Internet policy and economics*, 101–109. Springer.

Vila, T., Greenstadt, R., and Molnar, D. (2004). Why we can't be bothered to read privacy policies. In *Economics of Information Security*, 143–153. Springer.

Williamson, O.E. (1983). Credible commitments: using hostages to support exchange. *The American Economic Review* 73(4), 519-540.

Appendix A[29]

1. *Do you make web searches on Google.com?*

    a. If the respondent indicated they did not, they were disqualified from further questions.

    b. *Possible responses:*

        i. *Yes*

        ii. *No*

2. *How often do you make searches on Google.com?*

    a. Possible responses:

        i. Once a day

        ii. A few times per day

        iii. Dozens of times per day (or more)

3. *Do you believe that Google collects information about you as you use Google.com?*

    a. *Possible responses:*

        i. *Yes*

        ii. *No*

4. *What information do you believe Google collects and saves about you? Select all that apply.*

    a. This question was asked of those who answered "Yes" to question three.

    b. Possible responses:

        i. Your driver's license number

        ii. Your social security number

        iii. Videos you watch

        iv. Device information

        v. Ads you click on or tap

        vi. Your credit card information

        vii. Websites you visit

        viii. Your location

        ix. Things you search for

        x. Your medical information

        xi. IP address and cookie data

        xii. None of the above

[29] Questions four, five and ten randomized the response options to respondents. The other questions presented the response options in the order displayed in Appendix A.

5. *Which of the following do you believe Google may use your information for? Select all that apply.*

    a. This question was asked of those who answered “Yes” to question three.

    b. Possible responses:
        i. To target ads based on your search history and location
        ii. To link your search history with your personal identity
        iii. To link your search history with your race, gender, religious preferences, or sexual orientation
        iv. To aggregate large quantities of anonymized data
        v. To store your data indefinitely
        vi. To sell your browsing history to potential employers or insurers who are hoping to learn more about you

6. *Do you believe that Google could change its privacy policy to allow new uses for user data?*

    a. This question was asked of those who answered “Yes” to question three.

    b. Possible responses:
        i. Yes
        ii. No

7. *Do you use a tool to protect your privacy while browsing, such as Adblock Plus?*

    a. Possible responses:
        i. Yes
        ii. No

8.

[illegible]

        i. Yes
        ii. No

10. *Why do you prefer that Google not collect information about you? Select all that apply.*

    a. Possible responses:
        i. A government agency forcing an internet entity that has collected your information to hand over the information
        ii. Sellers offering different prices to buyers for the same good
        iii. Uneasiness just not knowing who knows what about you
        iv. The risk of identity theft
        v. The threat of spam
        vi. Advertisers being able to target you directly
        vii. Other (please specify)

11. *What do you think about the ads targeted to you based on the information Google collects about you?*

    a. Possible responses:
        i. I like seeing the ads customized to my preferences
        ii. I don't like the ads and would rather not see them

12. *How much would you be willing to pay per year to use Google.com without Google collecting any personal information about you? Enter a whole number in US dollars.*

13. *How much would you be willing to pay per search to use Google.com without Google collecting any personal information about you? Enter a whole number in US dollars.*[30]

    a. Possible responses:
        i. Less than 1 cent
        ii. 1 cent to ninety-nine cents
        iii. $1 to $5
        iv. More than $5

14. *Would you be willing to pay $70 per year for a guarantee that Google will NOT collect any information about you while using Google.com?*

    a. Possible responses:
        i. Yes
        ii. No

[30] Question 13 contains a wording error. The question should *not* have included the phrase: "Enter a whole number in US dollars" because respondents were not offered an open-ended response option.

Appendix B

Things you search for

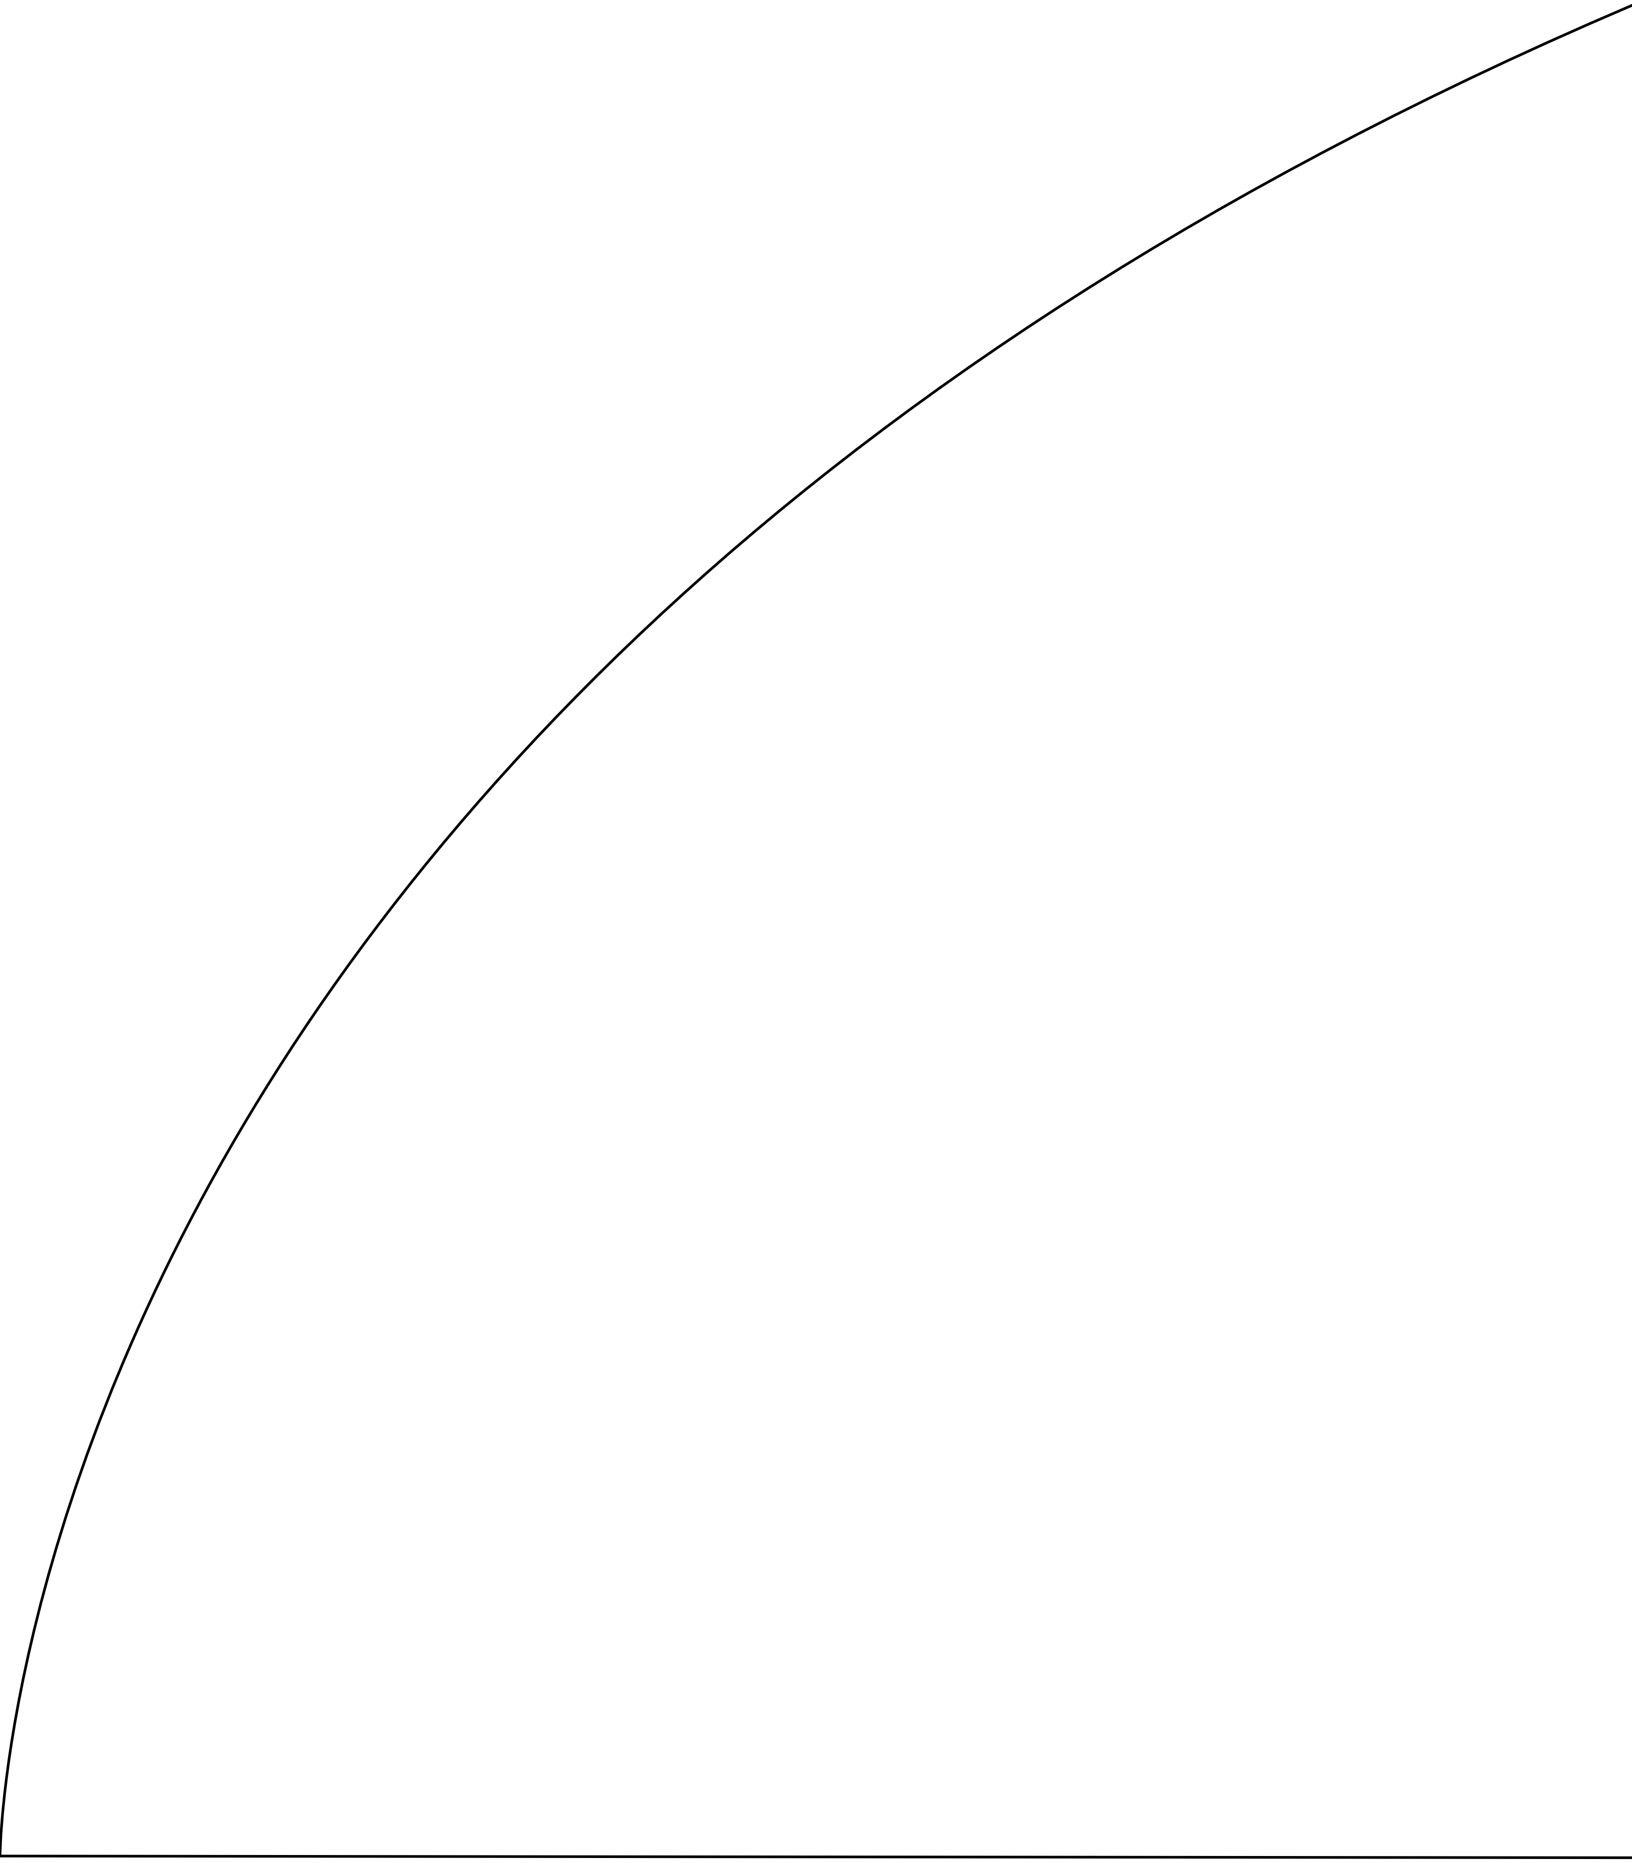

I would prefer Google NOT collect information about me	4621	76%
Column Total	*6083*	*100%*
Would you prefer to pay to use Google.com in exchange for a guarantee that Google will NOT collect any information about you?[33]		
Yes	824	18%
No	3797	82%

Less than $0.01	448	54%
$0.01 to $0.99	230	28%
$1 to $5	73	9%
More than $5	73	9%
Column Total	*824*	*100%*
Would you be willing to pay $70 per year for a guarantee that Google will NOT collect any information about you while using Google.com?		
Yes	378	46%

	Informed about data use?[34]			
	Yes		**No**	
	n	**Percent**	**n**	**Percent**
Would you prefer to pay to use Google.com in exchange for a guarantee that Google will NOT collect any information about you?				
Yes	37	14%	705	18%
No	220	86%	3121	82%
Column Total	257		3826	

Do you use a tool to protect your privacy while browsing, such as Adblock Plus?

Column Total	*3483*	*100%*	*600*	*100%*

Table 3: WTP Contingent on Responses to Questions Five, Six and Seven (Dollar Values)

Total		**Informed about data use?**		**Do you believe that Google could change its privacy policy to allow new uses for user data?**		**Do you use a tool to protect your privacy while browsing, such as Adblock Plus?**	
		Yes	**No**	**Yes**	**No**	**Yes**	**No**
		n = 37	n = 705	n = 662	n = 80	n = 349	n = 475
Mean WTP	$59.59	$48.19	$58.16	$53.32	$93.61	$52.48	$64.81