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 informational asymmetry is minimal and (b) the demands for “unconstrained” and “constrained” privacy diverge substantially. Significantly, 86% of respondents express no willingness to pay for additional privacy when interacting with Google. Among the remaining 14%, the average expressed willingness to pay is low.

Keywords: digital privacy, survey, market failure, privacy paradox

JEL-Classification: D23, D62, K24, L86

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 2013)? Does information collection and sale to third parties generate market failure?

Surveys provide some evidence that markets provide less privacy than digital browsers would prefer (Acquisti and Gross 2006; Turow et al. 2009; Acquisti et al. 2013; Madden and Rainie 2015; Acquisti et al. 2016, pp. 476-478). For instance, Turow et al. (2009) show that 66% of Americans do not want marketers to target their offerings, but that most Americans use search engines that track consumers for that purpose. Furthermore, Turow et al. (2009) also report that 86% of "young adults" do not want to be shown ads resulting from them being tracked across websites. Why then do so many firms rely on that way of monetizing information about consumers rather than charging online visitors a transparent, pecuniary fee? Alternatively, firms could rely on less effective, but more privacy protective, non-targeted advertising.

One hypothesis is that companies benefit from behavioral biases enabling them to collect more data from consumers than if they were perfectly rational (Acquisti 2004; Calo 2013; Hoofnagle and Whittington 2013; Acquisti et al. 2016, pp. 477-478; Brown 2016). Calo (2013, p. 1003), for instance, refers to "the exploitation of cognitive bias", which enables firms to identify "the specific ways each individual consumer deviates from rational decision-making ... and leverage that bias

[1] See here for more information on Google's decision to drop the phrase: https://gizmodo.com/google-removes-nearly-all-mentions-of-dont-be-evil-from-1826153393.

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-run interests (Acquisti 2004; John et al. 2011).

Other scholars, however, have questioned whether it is reasonable to assume that consumers *are*, in fact, well-informed (Hirsch 2010; Acquisti et al. 2016, pp. 447-448). The possibility of asymmetric information between consumers and producers offers a distinct, yet complementary, explanation for why digital firms rely so heavily on information collection. Information may be over-collected relative to the case of perfectly informed browsers (Hoofnagle 2005; Hirsch 2010). As Acquisti et al. (2016, p. 448, emphasis added) put it, "issues associated with individuals' *awareness* of privacy challenges, solutions, and trade-offs cast doubts over the ability of market outcomes to accurately capture and reveal, by themselves, individuals' true privacy valuations."

In addition to the problems posed by behavioral biases and information asymmetry, initial information collectors may disseminate consumer data to unknown third parties (Varian 2009; Brown 2016). To the extent that consumers disapprove of such activity, digital interactions once again fail to satisfy consumer preferences.

If those three-mentioned market failures indeed are pervasive, markets may not be satisfying consumers' *unbiased, fully informed* preferences. Gertz (2002), for instance, contends that the digital marketplace is a "classic example of a market failure" in need of government regulation, a position advanced by many other scholars (Solove 2004; Vila et al. 2004; Hui and Png 2005; Hermalin and Katz 2006; Sachs 2009; Turow et al. 2009; Ohm 2010; Hoofnagle et al. 2012;

Strandburg 2013; Newman 2013; Acquisti et al. 2016).[2] Those drawing that conclusion also often appeal to the well-documented "privacy paradox", wherein individuals routinely voice preferences for more privacy, but just as frequently forego low-cost methods of protecting the same (Berendt et al. 2005; Norberg et al. 2007; Acquisti et al. 2016). Foregoing privacy protection after stating one's demand for it usually is attributed to either behavioral bias, poor information awareness, or both.

This paper offers theoretical and empirical reasons to be skeptical of the claim that the market for digital privacy fails to achieve ideal results. In so doing, I also propose a straightforward resolution to the privacy paradox. That resolution relies neither on biases causing consumers to behave irrationally, nor on them being poorly informed. Instead, the paradox may not emerge at all: it simply may reflect only a positive preference for a higher quality economic good, *ceteris paribus.* That conclusion follows from the fact that consumers likely articulate preferences for more of a good of higher quality, provided that the relevant constraints for obtaining that good have been relaxed. In other words, there may be no more need to explain a "privacy paradox" than to explain why consumers might express preferences for higher quality "traditional" economic goods, but subsequently and regularly fail to demonstrate that preference in their purchasing behavior. Stating a verbal preference incurs no opportunity cost; action necessarily does.

Section 2 provides an in-depth discussion of the three major claims of market failure in the digital context. Section 3 examines those three claims, first evaluating them by way of theory and secondly by presenting new survey evidence. Section 4 argues that government activity may itself

[2] Less consensus exists regarding what specific policy interventions should be implemented. Some scholars favor outright bans on information collection, others call for a legally mandated opt-in, and still others argue that greater transparency be required of firms. The EU, Japan, Canada, Singapore and South Africa all have passed comprehensive digital privacy legislation. For an analysis of intervention, see Fuller (2018).

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's device information, geographic location, browsing history, and click-trail (Berman and Mulligan 1998). Probably no website collects more "mouse-droppings" than Google. In fact, most of Google's revenue (almost $70 billion in 2015 and almost $80 billion in 2016) is earned from third-party advertisers who pay to use the platform to track browser behavior (Statista 2017).

It has been argued that assessing whether information collection is a market failure is a difficult task. That is because some scholars refer to market failure in the technical sense implying failure to achieve Pareto optimality; others use the term casually to refer to market outcomes of which they disapprove (Acquisti 2012). Scholars arguing for market failure in the technical sense generally identify three sources of such failure. Brown (2016) provides a useful summary of them. First is information asymmetry between firms and consumers. Second, individuals fall prey to behavioral biases that cause them to act inconsistently with their "true" preferences. Third, data resale generates a negative externality.

2.2 Information asymmetry

Even if consumers were perfectly rational, some scholars contend that inefficiency can arise if information between firms and consumers is distributed unequally. As Acquisti et al. (2016, p. 448)

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 "attitudes") and their digital behavior. While claiming to value privacy highly, consumers' subsequent behavior—distorted as it is by behavioral biases—exhibits risky behaviors that would compromise the privacy they claim to value highly (Acquisti 2004). More recently, Acquisti et al. (2016, p. 477) emphasize that the observed dichotomy between attitudes and behavior is the result of "many coexisting, and not mutually exclusive factors." Those factors include behavioral biases, bounded rationality and asymmetric information. The result is that firms collect more consumer information than they would if consumers exhibited perfect behavioral rationality.

2.4 Data resale externalities

Suppose that a digital platform collects consumer data. Later, it discovers that other parties also

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

to 6,864.

The pool of respondents, all over the age of eighteen, mirrored 2010 US Census population demographics in percentage terms on the following dimensions: ethnicity, gender and religious affiliation. The survey was administered online intermittently between September 11, 2018, and September 26, 2018, to Internet users across the United States (the survey's full text can be found in Appendix A) and the data are available on request. The questionnaire was programmed and administered by Haven Insights, LLC, and hosted at SurveyGizmo.com. To ensure that all respondents were Google users, the respondents first answered a "screener question" ("Do you make searches on Google.com?"). That question eliminated 781 respondents, so that only 6,083 people were asked subsequent questions.

Respondents were directed to the survey by several online panel providers. A panel is a group of individuals who have expressed willingness to take online surveys, and who have been pre-screened according to a set of respondent criteria (to validate their identities and other basic characteristics). Lucid served as the panel aggregator.[3] Because the survey results were assembled from several different panels, many different methods were used to solicit respondents and reward them for participation. Whereas some respondents received invitations via email to participate in the survey, others logged into an online portal where they were subsequently invited to take a survey. Some respondents were compensated by receiving gift cards; others were compensated with small pecuniary rewards. Because of the survey's large scope, this paper's text does not discuss every result, but the full results are available in Appendix B. Sections 3.1.2, 3.2.2, and 3.3.2 describe the results of the survey. All percentages have been rounded to the nearest whole number.

[3] Haven Insights used Lucid's platform for academic research. See more details here: https://luc.id/lucid-for-academics/.

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-selection concerns may be mitigated. The objective of the survey at hand, however, is not to gain a fine-grained perspective of just how much consumers value privacy. Rather, the goal is to determine whether consumers offer different responses to "constrained" relative to "unconstrained" questions.

Third, surveys fail to uncover *revealed* preferences. Ultimately, they consist of "cheap talk." In other words, stated valuations in response to constrained questions may diverge from what an individual demonstrates in action. The present survey is concerned only with whether responses to constrained questions differ from those to unconstrained questions. If such a gap exists, it becomes more difficult to argue that unconstrained questions represent a person's "true" preferences, which they proceed to violate subsequently by their actions.

3.1 Asymmetric information

3.1.1. Theory

Since Akerlof's (1970) classic investigation of the market for "lemons", economists have been attuned to the possibility of information asymmetry generating market failure in a host of contexts,

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 exchanges to occur. Thus, adverse selection and, in other contexts, moral hazard, are the problems that stem from asymmetric information.

The foregoing reasoning raises the first theoretical difficulty with claiming that information asymmetry is causing failure in digital environments. To wit, asymmetric information does not appear to be generating significant retrenchment in online activity. As of July 2018, well over half of the world's population (roughly 4.1 billion) used the Internet, and the percentage of the world's internet-using population has been rising steadily for years (Statista 2018). Although those users are not perfectly informed regarding digital interactions, Internet activity certainly has not collapsed. Admittedly, such evidence may be unconvincing because many consumers remain poorly informed. Some commentators simply may be concerned with the risks that uninformed consumers bear.

However, information—in this context, information about relevant attributes of the digital environment—is costly to obtain (Stigler 1961). If digital markets fail because consumers are not perfectly informed about the digital environment, it is difficult to see how *every* market is not subject to asymmetric information to some extent. For example, what percentage of consumers would claim that they are perfectly informed regarding the wide array of laptops or motor vehicles, to name just two commonly purchased items? Like Internet use, the consumption of laptops and

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 "B" consists of a collection of vehicular attributes, including fuel efficiency, safety, comfort, aesthetics and countless other characteristics. Constrained as she is by a budget, a consumer wanting a little more fuel efficiency will sacrifice a little of something else, such as crash-worthiness. Using the Internet to search for information also consists of a bundled "experience" that includes certain privacy attributes, namely that some personal information will be collected. By contrast, using one's local public library to acquire information does not include nearly the same privacy attribute.[5]

Because product attributes are costly to measure, we should expect some degree of information asymmetry to be the norm, rather than the exception. As Barzel (1982) notes, buyer's "surprise" regarding the attributes of a good is inevitable.[6] When consumers purchase a car, they rarely are perfectly informed about its specific bundle of attributes. Yet, they still act with an *ex ante* expectation of gaining from exchange, having judged the acquisition of additional information to be costlier than it's worth. That a consumer may experience an *ex post* psychic loss (what Barzel calls an "unpleasant surprise") from certain digital activities does not itself define a market failure.

---

[4] It has been characteristic to describe privacy itself as an "economic good" (Farrell 2012; Acquisti et al. 2016, p. 446). This paper conceives of privacy, not as an economic good itself, but as an attribute of some *other* economic good.

[5] Of course, to get a library card, one usually provides name, physical address, email address, etc.…The information collected by digital firms tends to be a browser's location, browsing history, and (often) purchase history. If one wishes to avoid surrendering information to a library, it is possible to use the library without checking out any items.

[6] Expecting "perfect information" to describe the real world commits the "Nirvana Fallacy" (Demsetz 1969). Unsurprisingly, an orange grower will tend to be more informed about an orange's attributes than prospective fruit buyers (Barzel 1982).

Surprises, good or bad, are inherent in all actions because it never pays a consumer to be fully informed prior to purchase (Stigler 1961).

Given the ubiquity of asymmetric information, the question becomes whether the relevant market permits consumers who acquire new information about a product's unsatisfactory bundle of attributes to "punish" the seller by switching to one who offers an alternative mix of attributes. It certainly is possible to punish Google by way of unilateral boycott: refusal to use Google's services in the future, thus bringing the "discipline of continuous dealings" to bear on the company (Leeson 2014). Since the consumer never will use Google again, marketers will find themselves marginally less profitable and thus willing to pay marginally less to Google for an advertising slot. For a firm as large as Google, however, it is obvious that a single disgruntled consumer exiting the platform inflicts only little punishment.

However, a "one-person boycott" is not the only available option. Multilateral boycott harnesses the power of individuals who have not yet experienced harm themselves also to participate in punishing the offending party (Leeson 2008; Leeson and Coyne 2012). The Internet itself, having reduced the cost of disseminating information to anonymous others, facilitates the effectiveness of the multilateral boycott. Suppose that a consumer is disaffected by the way Google collects her information. She suspects that other users would feel similarly if only they knew the extent of the problem. Blogs provide a platform for the initial consumer to chronicle the privacy abuses she has experienced and to convey that information to millions of others. In calling for a boycott of the offending party, she can exact punishment far exceeding the threat posed by unilateral boycott.

If enough consumers agree with the initial dissatisfied consumer, the implication is that a demand exists for substitute services that differ from incumbent service providers in the bundle of

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' privacy preferences. Firms that refrain from such activity will lose market share continuously to rivals that refuse to respect privacy.

Given that the information collection process is alleged inherently to be opaque, how can a firm credibly commit to refrain from information collection? One way for a firm to overcome the prisoner's dilemma is by investing in a "hostage", which facilitates repeated interactions. Hostages are costly transaction-specific investments that pay off to the hostage-giver only when she cooperates, when she foregoes opportunistic behavior (Williamson 1983). As Benson (1998) notes, a good reputation makes an ideal hostage because the hostage-giver values it highly, while the hostage-recipient does not.

DuckDuckGo, a search engine that does not track its users, is an example of credibly committing to privacy protection by offering a hostage that consists of an investment in reputation.[7] Founded in 2008, DuckDuckGo advertised on a billboard in San Francisco (the

[7] The company earns revenue by displaying ads based merely on what search terms a browser enters but does not track the user.

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] Although DuckDuckGo has grown steadily, it averaged only a little more than 20 million queries daily as of early 2018, far less than 1% of Google's daily traffic. See https://duckduckgo.com/traffic.html for statistics on DuckDuckGo's traffic over time. See http://www.internetlivestats.com/google-search-statistics/ for a daily count of Google searches.

queried about their knowledge of Google's information-collection model (question three), respondents overwhelmingly are aware that the company gathers personal information about them as they use Google. At least regarding the *existence* of the practice, the extent of information asymmetry is low, with 89% ($n = 6{,}083$) of respondents indicating awareness of Google's collection of personal data.[9]

Mere knowledge that Google collects information is a relatively low standard to meet. Do consumers know what *types* of information firms collect? Following the initial question regarding awareness of data collection, respondents who acknowledged Google's information-collection practices ($n = 5{,}434$) were presented with 11 possible pieces of data (seven routinely collected by Google; four which it does not collect) and asked to select the items Google collects. "None of the above" was an additional (incorrect) possibility, so the respondents were presented with 12 total options.[10]

Here, too, the data suggest that most respondents possess a relatively high degree of awareness. Only 1% of "aware respondents" believe that Google collects "none" of the suggested pieces of information, 10% believe that the company collects driver's license information, 11% believe that Google collects social security information, while 13% believe that it may collect medical information. The most ignorance expressed related to financial information: as 21% responded that Google might collect "your credit card information." The possibility of ambiguity exists with that option, though, as many ecommerce portals save a consumer's credit card information in order to reduce the future costs of using their websites. While Google itself does not save a consumer's

[9] Respondents who indicated unawareness of Google's information-collection practices were not asked questions four, five, or six.

[10] Google may collect any of the data listed in question four of the survey's text (Appendix A) except: "Your driver's license", "Your social security number", "Your medical information" and "Your credit card information."

financial information, some respondents may interpret that ecommerce capability as Google doing just that. When it comes to data that Google *does* collect, 87% of "aware respondents" know that Google harvests information pertaining to the sites a browser has visited. Similarly, 87% of the same group know that Google keeps a record of searches and 80% know that Google registers a browser's physical location.[11]

Lastly, are consumers aware of potential uses of their data, once it has been collected? The survey responses suggest that consumers are significantly less well-informed about such business practices, but not completely uninformed about them. Respondents who had indicated general awareness ($n = 5{,}434$) were presented with six possible ways that Google might use their data (three that Google's privacy policy permits and three that it does not).[12] While 81% of these respondents correctly identify that Google collects information "to target ads based on your search history and location", many of them consistently overestimate the number of uses to which Google puts their data. For example, 44% believe that Google might "sell your browsing history to potential employers or insurers who are hoping to learn about you", but Google's privacy policy forbids such usage. Likewise, 40% think that Google could "link your search history with your race, gender, religious preferences, or sexual orientation", but such activity also is expressly forbidden by Google's privacy policy. However, consumers are still somewhat more likely to correctly identify the purposes for which Google does use data than they are to mistakenly believe Google uses data in ways that it does not. The average selection rate for the correct options to

[11] The results show that individuals are least aware of the fact that Google gathers information about their devices. Still, 51% of "aware respondents" know that device information is collected. Arguably, for most users, device information is the least "sensitive" or "important" piece of information that Google collects. It also is possible that some consumers are unfamiliar with the term "device information."

[12] Google may use collected information to "target ads based on your search history and location", to "aggregate large quantities of anonymized data", and to "store your data indefinitely", but its privacy policy does not permit any of the other uses listed in question five of the survey (see the survey's text in Appendix A).

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 ($n = 5{,}434$) 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-cost means of invoking multilateral punishment, a conclusion that information asymmetry causes the market for privacy to fail is tenuous at best.

3.2 Behavioral biases

3.2.1 Theory

According to the conventional wisdom, even the existence of perfectly informed consumers is not enough to guarantee the absence of market failure. Consumers may be well-informed, but irrational. Claims that consumers are irrational in digital contexts typically derive from the gap between what consumers "say" and what they "do", a dichotomy that has been termed the "privacy paradox" (Norberg et al. 2007). In seeking to explain that gap, behavioral economics offers one reason why consumers behave in ways at odds with their expressed preferences.[13]

Surveys often have been used to elicit information regarding consumers' notional preferences. As Acquisti et al. (2013) observe, most empirical studies of consumers' privacy values focus on individuals' reservation prices for disclosing some piece of otherwise private information. Tsai et al. (2011) and Savage and Waldman (2015) are exceptions in that they investigate what individuals are willing to sacrifice in order to make otherwise public information private.[14] Tsai et al. (2011) find that, when a company makes its privacy-protective policies prominent, consumers are willing to pay a small premium for those features. Savage and Waldman (2015) investigate willingness to

[13] As noted by Acquisti et al. (2016), information asymmetry provides another explanation, but I am ruling that out for a moment so as to isolate the purported effect of behavioral biases.
[14] "[S]tudies in which consumers are ... asked to consider paying ... to protect their privacy are ... scarcer" (Acquisti et al. 2013, p. 254).

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" survey approaches in the literature on sweatshops.

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 preferences for beauty, love, discrimination and so on (Boettke and Candela 2017), but those differentials come in the form of personal information in the case of digital privacy.

respondent's behavior observed in digital environments, which seemingly disregards verbal preferences. By contrast, if the typical respondent voices a low WTP, it would suggest that "true" preferences and revealed preferences are relatively well-aligned. As such, online behavior suggesting little regard for privacy does not diverge significantly from consumers' stated, "constrained" preferences.

We can appeal to variations in costs to explain variation in behavior. In fact, we should expect that the gap between stated responses to "unconstrained" surveys and actual behavior is even larger than the difference between the two question types. That is the case even when consumers behave consistently with their true preferences. The reasoning is straightforward: talk is cheap. Action is not.

3.2.2 Empirical evidence

Like existing research, such as Turow et al. (2009) and the 2015 Pew survey, my questionnaire also finds that most consumers, in the absence of any constraints, would prefer to use Google without its information-collection practices. In fact, a large majority of Google users—76%—say they would prefer for Google *not* to collect their information ($n = 6{,}083$).

Individuals expressing a verbal preference for a higher quality good, however, does not imply that markets are failing by instead providing a lower-quality good. Tellingly, of the respondents who would prefer not to be tracked, only 18% of these verbalize willingness to pay anything to retain their privacy ($n = 4{,}621$).[19] That finding is strong evidence in favor of a large difference

[19] Note that 149 respondents indicated a willingness to pay for privacy on Google, but when they subsequently were prompted to state the amount they would be willing to pay, they entered $0. Those 149 respondents were re-categorized as being *unwilling* to pay for privacy and thus included amongst the 86% of all respondents not willing to pay for privacy.

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 ($n = 6{,}083$).[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 $10,000 or greater were dropped from the survey’s results on the basis that such stated amounts likely were errors or represented unserious responses. Among those respondents kept in the sample and indicating a positive WTP ($n = 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 day, privacy on Google does not seem to be an overwhelming concern (Classroom.com 2017).

Because individuals might have difficulty calculating what a year of privacy is worth to them,

[20] The survey began with a sample of 6,864 respondents, but 781 were eliminated because they did not use Google. It is unclear how those non-users would respond to the remainder of the survey. At one extreme, it is possible that 100% of them refrain from using Google because of privacy concerns and all of them would also be willing to pay for privacy on Google. If that were the case, 23% of the Internet-using population would be willing to pay for privacy on Google. At the other extreme, 100% of them could also be unwilling to pay for privacy on Google because they never use Google (for reasons other than privacy concerns). If that were the case, only 12% of the Internet-using population would be willing to pay for privacy. The truth probably lies somewhere between the extremes.

the same respondents also were asked about their willingness to purchase privacy on a "per-search" basis. On that issue, respondents were asked to select one of the following per-search measures of WTP: "less than 1 cent", "1 cent to 99 cents", "$1 to $5", or "more than $5." In response ($n = 824$), 54% indicate a per-search WTP of "less than 1 cent", 28% select "between 1 and 99 cents", with the remaining 18% roughly split evenly between the final two options. Such small per-search valuations are consistent with the small annual WTPs.[21]

If respondents' verbal preferences accurately reflect their demonstrated preferences, the results can inform speculation about what would happen to Google's revenue if it switched from information collection to charging a use fee. Google has about one billion users annually and earned roughly $70 billion in 2015 from information collection. Multiplying the number of respondents with positive WTPs by the mean annual WTP (14% of one billion multiplied by $59.59) yields annual revenue of $8.3 billion. That sum amounts to about 12% of Google's 2015 revenue.

While it serves as an interesting thought experiment, charging the survey respondents' mean annual WTP runs into a fundamental problem. Were Google to collect data on users' average WTP, and then set a price based on that average, the company immediately would price many Google users out of the market. Pricing at the average WTP would generate a "lemons market" of sorts, since only people expressing above-average WTPs would be willing to pay the fee (Akerlof 1970).

An alternative way for Google to earn $70 billion annually would be to charge $70 per year to every user. Thus, as a final measure of WTP—and as a check on the preceding results—the survey asked respondents who had indicated positive WTPs ($n = 824$) a simple "yes/no" query regarding

[21] It is impossible to determine whether respondents are perfectly consistent between their annual and "per-search" valuations. For example, someone selecting "$1 to $5" may have had $1 in mind, whereas another had $5 in mind. Nonetheless, the answers are "generally consistent" in that both the annual and "per-search" prompts elicit relatively low WTPs.

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 3.1.2's results indicate minimal information asymmetry between consumers and Google. If largely uninformed respondents place a low value on their privacy, little WTP might be implied because those respondents also are ignorant. To the contrary, the results suggest the existence of relatively well-informed consumers who, on average, express slight WTP for privacy.

Among respondents who know that Google unilaterally could alter its privacy policy (question six), 19% are willing to pay, whereas only 13% of uninformed consumers indicate such willingness ($n = 4{,}083$).[22] By itself, that result seemingly would suggest that informed consumers have a greater demand for "future-proofing" themselves against the possible policy changes Google might introduce. However, such a perspective is undermined by the average WTP of the two groups: $53.32 for the informed group and $93.61 for the uninformed group ($n = 742$).[23] Despite the apparently large difference in average WTP, it is not statistically significant at the 5% level (t-statistic $= -1.04$).

The notion of a strictly positive relationship between privacy awareness and WTP is undermined further by survey question five—the question about possible data uses by Google. I examined whether respondents' knowledge, as judged by question five, systematically was

[22] Respondents who believe that Google does not collect information were excluded from the question about whether Google can change its privacy policy unilaterally. Thus, the relevant sample comprises users who are aware of Google's information-collection practices and who express a desire for Google not to collect their data.

[23] The respondents are comprised of those users who were aware that Google engages in information collection (question three) and expressed a willingness to pay for privacy (question nine).

correlated with WTP. The first possibility is that being uninformed tends to depress WTP. Someone might reason that, if only a respondent were more informed about privacy harms, she necessarily would be willing to pay more. That perspective corresponds to the idea that a world of fully informed individuals would see digital firms collecting less consumer information (Hirsch 2010, p. 455). On the other hand, becoming more informed might reduce stated WTP if potential privacy harms are viewed as negligible. Were that true, the idea that informing people would tend to make them more privacy-conscious would be undermined.

To adjudicate between those views, respondents were first categorized as being either "relatively informed" or "relatively uninformed", as judged by their answer to question five. "Relatively informed" respondents must have *either* selected all three correct answers and no incorrect answers *or* selected all three correct answers and only one incorrect response. Every other possible response permutation was categorized as being "relatively uninformed."[24] The results provide some support for the idea that greater awareness is correlated with lower—not higher—WTP. Among respondents who possess a high level of awareness of what Google can do with consumer data, 14% are willing to pay, but 18% of the uninformed are willing to pay. Similarly, the average WTP of the informed respondents is $48.19, whereas the average WTP among the uninformed is $58.16. However, that difference is not statistically significant at the 5% level (t-statistic $= -0.86$).

Awareness thus may, in fact, depress stated WTP. The same may be true for the use of privacy-

[24] Of course, other ways of categorizing respondents as "relatively informed" or "relatively uninformed" with respect to question five are possible. My strategy for categorization was selected in the interest of generating a sufficiently large sample size for both "informed" and "uninformed" groups, given that most respondents are unwilling to pay. Respondents who selected only two correct answers, but no incorrect answers are categorized as "uninformed" because they seemingly exhibit less awareness of overall data collection practices than those who selected all three correct answers *and* exhibited some degree of misinformation by also selecting an incorrect response.

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 ($n = 6{,}083$), 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 ($n = 4{,}621$).[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 ($n = 824$). That difference is not statistically significant at the 5% level, however (t-statistic $= -1.28$). 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's own behavior that would be generating an externality that she herself would subsequently bear.

Furthermore, pushing the reasoning of "data-resale-as-externality" to its logical conclusion seems to generate an uncomfortable implication. If the possibility of information resale imposes a

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 ($n = 6{,}083$) 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 ($n = 4{,}621$) 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 ($n = 4{,}621$) 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 "a government agency forcing an internet entity that has collected your information to hand over the information" is a concern. For a point of reference, only 27% express any distaste for price discrimination, which has been suggested as a contributor to expressions of dislike for current online information-collection practices.

Admittedly, government overreach is the item of second-least concern of the options presented to consumers in the survey, but the fact that a significant minority of respondents who dislike information collection list it as a contributing factor should not be overlooked. The result suggests that government failure–defined here as government encroaching on private property rights by forcing companies to relinquish data–contributes to consumer mistrust of information collection.

The finding suggests that, rather than systematic "over-collection" of information by private

[26] Acquisti et al. (2016) also list "quantity discrimination in insurance and credit markets", but I did not present respondents with that option because it was the most technical of the possibilities suggested by those authors.

[27] In addition to the four possibilities listed by Acquisti et al. (2016), my survey added: "Advertisers being able to target you directly", "A government agency forcing an internet entity that has collected your information to hand over the information", and "Other (please specify)".

firms, it is at least theoretically possible that systematic "under-collection" occurs relative to a benchmark in which governments are perfectly constrained from accessing the information gathered by search engines and other online entities, including social media sites. Scholars have argued that if consumers were informed perfectly about privacy policies, firms would collect less information (Hirsch 2010).[28] But if a world of perfect information and perfectly enforced property rights is to be taken as the relevant benchmark, then the real world–in which governments often behave in predatory fashion (Leeson 2007)–may suffer from information *under-collection* in comparison to an ideal world.

That conclusion follows because consumer worries about government overreach functions as a constraint on the quantity of information firms may collect. For one, individuals engage in less (and different) Internet search activity than otherwise. Second, firms collect less (and less sensitive) information given that consumers fear the governmental threat. In other words, the existence of uninformed consumers may, indeed, push Google and other search engines toward over-collection. But the existence of predatory government may push toward under-collection, and it is not clear which effect dominates.

5 Conclusion

This paper has two primary implications. First, it is possible to explain the so-called "privacy paradox" by showing that individuals express greater demands for digital privacy when they are not forced to consider the opportunity cost of that choice. The question never has been whether people value privacy at all, but rather how strongly they value it. At least in the context of interacting with Google, the findings suggest that most individuals place relatively low values on

[28] Section 3.2.2's results provide evidence for doubting this argument.

privacy. A small expressed willingness to pay for privacy is consistent with behavior that seemingly disregards privacy threats. The results reported herein also may explain why so many digital firms engage in information collection rather than adopting alternative methods of earning revenue: consumers prefer exchanging information to exchanging money. Put differently, no privacy paradox may exist at all–online consumer behavior simply may reveal positive preferences, *ceteris paribus*, for what many see as a higher quality good. For a large majority of this paper's survey respondents, being able to use Google without relinquishing personal information is a higher quality good than the Google they currently use.

That survey respondents generally placed low value on their privacy is significant given that one of 2018's highest-profile news stories regarded privacy violations at Facebook. In early 2018, it became widely known that the consulting firm Cambridge Analytica used personal data from Facebook without users' consent in order to target political ads. Only a few months later, Europe's General Data Protection Regulation went into effect, prompting a flurry of emails from firms with a digital presence to their users to alert them about privacy policy changes. These events suggest the presence of well-informed populace, but one that still evaluated potential privacy harms as relatively minor.

Second, this paper's results should add a dose of humility to the impulse to regulate digital privacy. That conclusion is particularly relevant given little consensus regarding how governments ought to regulate digital privacy (Hirsch 2010; Acquisti et al. 2016, pp. 452-453). The justification for regulating digital privacy rests on the pillars of widespread information asymmetry, pervasive behavioral biases, and negative externalities from data resale. Yet, if the extent of information asymmetry is little more than what can be expected in markets for other complex consumer goods, if consumers are not biased, but simply responding to constraints, and if every social interaction is

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

Acquisti, A. (2004). Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, 21–29. ACM.

Acquisti, A. (2012). Privacy and market failures: Three reasons for concern, and three reasons for hope. *Journal on Telecommunications. & High Technology Law* 10, 227-233.

Acquisti, A., and Gross, R. (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies*, 36–58. Springer.

Acquisti, A., John, L.K., and Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies* 42(2), 249–274.

Acquisti, A., Taylor, C., Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature* 54(2), 442-92.

Akerlof, G.A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics* 84(3), 488–500.

Alchian, A.A. (1967). *Pricing and society*. Institute of Economic Affairs.

Barzel, Y. (1982). Measurement cost and the organization of markets. *The Journal of Law and Economics* 25(1), 27-48.

Benson, B.L. (1998.) How to Secede in Business Without Really Leaving: Evidence of the Substitution of Arbitration for Litigation. *Secession, State, and Liberty.* New Brunswick, NJ: Transaction.

Berendt, B., Günther, O., and Spiekermann, S. (2005). Privacy in e-commerce: stated preferences vs. actual behavior. *Communications of the ACM* 48(4), 101–106.

Berman, J. and Mulligan, D. (1998). Privacy in the digital age: Work in progress. *Nova Law Review* 23, 551-582.

Boettke, P.J., and Candela, R.A. (2017). Price theory as prophylactic against popular fallacies. *Journal of Institutional Economics* 13(3), 725-752.

Brown, I. (2016). The economics of privacy, data protection and surveillance. *Handbook on the Economics of the Internet*. Available at SSRN 2358392 (2013).

Calo, R. (2013). Digital Market Manipulation. *George Washington Law Review* 82, 995-1051.

Clark, J.R., and Powell, B. (2013). Sweatshop working conditions and employee welfare: Say it ain't sew. *Comparative Economic Studies* 55(2), 343–357.

Classroom.com. How much do Americans spend on soft drinks? (September 29, 2017). https://classroom.synonym.com/how-much-do-americans-spend-on-soft-drinks-12081634.html

Cooper, J.C. (2012). Privacy and Antitrust: Underpants Gnomes, the First Amendment, and Subjectivity. *George Mason Law Review* 20, 1129-1146.

De Corniere, A., and De Nijs, R. (2016). Online advertising and privacy. *The RAND Journal of Economics* 47(1), 48–72.

Demsetz, H. 1969. Information and efficiency: another viewpoint. *The Journal of Law and Economics* 12(1), 1-22.

Farrell, J. (2012). Can privacy be just another good? *Journal on Telecommunications & High Technology Law* 10, 251-264.

Federal Trade Commission. (2012). Protecting Consumer Privacy in an Era of Rapid Change. FTC Report, 1-112. https://www.ftc.gov/sites/default/files/documents/reports/federal-trade-

commission-report-protecting-consumer-privacy-era-rapid-change-recommendations/120326privacyreport.pdf
Fuller, C.S. (2018). Privacy law as price control. *European Journal of Law and Economics*. 45(2), 225-250.
Gertz, J.D. (2002). The purloined personality: consumer profiling in financial services. *San Diego Law Review* 39, 943.
Hermalin, B.E. and Katz, M.L. (2004). Sender or receiver: Who should pay to receive an electronic message? *RAND Journal of Economics* 35(3), 423-448.
Hermalin, B.E., and Katz, M.L. (2006). Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics* 4(3), 209–239.
Hirsch, D.D. (2010). The law and policy of online privacy: Regulation, self-regulation, or co-regulation. *Seattle University Law Review* 34, 439-480.
Hoofnagle, C.J. (2003). Reflections on the NC JOLT Symposium: the privacy self-regulation race to the bottom. *NCJL & Tech.* 5, 213-217.
Hoofnagle, C.J. (2009). Beyond Google and evil: How policy makers, journalists and consumers should talk differently about Google and privacy. *First Monday* 14, 4-6.
Hoofnagle, C.J., and Whittington, J. (2013). Free: accounting for the costs of the internet's most popular price. *UCLA Law Review* 61, 606-670.
Hui, K.L., and Png, I. (2005). Economics of privacy. *Handbook of Information Systems and Economics.* Available at SSRN 786846.
John, L.K., Acquisti, A., and Loewenstein, G. (2011). Strangers on a plane: context-dependent willingness to divulge sensitive information. *Journal of Consumer Research* 37(5), 858–873.
Leeson, P.T. (2007). Better off stateless: Somalia before and after government collapse. *Journal of Comparative Economics* 35(4), 689-710.
Leeson, P.T. (2008). Social distance and self-enforcing exchange. *The Journal of Legal Studies* 37(1), 161-188.
Leeson, P.T. (2014). Pirates, prisoners, and preliterates: anarchic context and the private enforcement of law. *European Journal of Law and Economics* 37(3), 365-379.
Leeson, P.T., and Coyne, C.J. (2012). Conflict-inhibiting norms. *Oxford handbook of the economics of peace and conflict.* Oxford University Press.
Lerner, Josh. (2012). The impact of privacy policy changes on venture capital investment in online advertising companies. *Analysis Group*.
Madden, M., and Rainie, L. (2015). Americans' attitudes about privacy, security, and surveillance. http://www.pewinternet.org/2015/05/20/americans-attitudes-about-privacy-security-and-surveillance/
Marthews, A., and Tucker, C. (2017). Government surveillance and internet search behavior. Available at SSRN 2412564.
Milberg, S. Smith, J., Burke, S.J. (2000). Information privacy: corporate management and national regulation. *Organization Science* 11(1), 35–57.
Newman, N. (2013). The costs of lost privacy: consumer harm and rising economic inequality in the age of Google. *William Mitchell Law Review* 40, 849.
Norberg, P.A., Horne, D.R. and Horne, D.A. (2007). The privacy paradox: personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs* 41(1), 100–126.
Odlyzko, A. (2003). Privacy, economics, and price discrimination on the Internet. In *Proceedings*

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. *Would you prefer that Google collected no information about you when you use Google.com?*

    a. Those responding that they would prefer Google to collect personal information were disqualified from further queries.

    b. Possible responses:
        i. I would prefer Google collect information about me
        ii. I would prefer Google NOT collect information about me

9. *Would you prefer to pay to use Google.com in exchange for a guarantee that Google will NOT collect any information about you?*

    a. Those answering “No” to this question were disqualified from further queries.

    b. Possible responses:

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[31]

Table 1: Survey Results[32]

	n	Percent
How often do you make searches on Google.com?		
Dozens of times per day (or more)	2406	40%
A few times per day	2749	45%
Once a day (or less)	928	15%
Column Total	*6083*	*100%*
Do you believe that Google collects information about you as you use Google.com?		
Yes	5434	89%
No	649	11%
Column Total	*6083*	*100%*
What information do you believe Google collects and saves about you? Select all that apply.		
Your driver's license number	556	10%
Your social security number	638	12%
Videos you watch	3707	68%
Device information	3207	59%
Ads you click on or tap	3863	71%
Your credit card information	1165	21%
Websites you visit	4704	87%
Your location	4363	80%

[31] Appendix B contains the results from all survey questions except for question one (a screener question to determine whether respondents are Google users) and question 12 which asks about how much consumers would be willing to pay for privacy. The paper's text reports the results of question 12.
[32] As in the paper's text, percentages in the tables are rounded to the nearest whole number.

Things you search for	4709	87%
Your medical information	729	13%
IP address and cookie data	3943	73%
None of the above	60	1%
Column Total	*5434*	*N/A*
Which of the following do you think Google may use your information for? Select all that apply.		
To target ads based on your search history and location	4409	81%
To link your search history with your personal identity	2374	44%
To link your search history with your race, gender, religious preferences, or sexual orientation	2164	40%
To aggregate large quantities of anonymized data	2539	47%
To store your data indefinitely	2250	41%
To sell your browsing history to potential employers or insurers who are hoping to learn more about you	2379	44%
Column Total	*5434*	*N/A*
Do you believe that Google could change its privacy policy to allow new uses for user data?		
Yes	4636	85%
No	798	15%
Column Total	*5434*	*100%*

Do you use a tool to protect your privacy while browsing, such as Adblock Plus?		
Yes	2395	39%
No	3688	61%
Column Total	*6083*	*100%*
Would you prefer that Google collected no information about you when you use Google.com?		
I would prefer Google collect information about me	1462	24%

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%
Column Total	*4621*	*100%*
Why do you prefer that Google NOT collect information about you? Select all that apply.		
A government agency forcing an internet entity that has collected your information to hand over the information	1866	41%
Sellers offering different prices to buyers for the same good	1235	27%
Uneasiness just not knowing who knows what about you	3167	69%
The risk of identity theft	3453	76%
The threat of spam	2776	61%
Advertisers being able to target you directly	2379	52%
Other (please specify)	0	0%
Column Total	*4570*	
What do you think about the ads targeted to you based on the information Google collects about you?		
I like seeing the ads customized to my preferences	1088	24%
I don't like the ads and would rather not see them	3533	77%
Column Total	*4621*	*100%*
How much would you be willing to pay per search to use Google.com without Google collecting any personal information about you?		

[33] As described in the paper's text, some respondents indicated a positive WTP, but then subsequently entered a value of "zero" for question 12. Those respondents (totaling 149) were re-categorized in both the text and in Appendix B's table as being *unwilling* to pay.

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%
No	446	54%
Column Total	*824*	*100%*

Table 2: WTP Contingent on Responses to Questions Five, Six and Seven (Percentages)

	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?			
	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	349	21%	475	16%
No	1305	79%	2492	84%
Column Total	*1654*	*100%*	*2967*	*100%*
	Do you believe that Google could change its privacy policy to allow new uses for user data?			
	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	662	19%	80	13%
No	2821	81%	520	87%

[34] The paper's text describes how respondents were assigned to either the "informed" or "uninformed" categories. Respondents who *prefer* their information to be collected are excluded from this analysis.

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