

# Privacy and Information Avoidance: An Experiment on Data-Sharing Preferences

Dan Svirsky

## ABSTRACT

There is a widespread intuition that people are inconsistent about protecting their privacy. This paper presents an experiment that demonstrates that people engage in information avoidance when making privacy decisions. People who are willing to pay nearly an hour's worth of wages for privacy are also willing to give away their data for small monetary bonuses if given a chance to avoid seeing the consequences to privacy. Placebo tests confirm that the same behavior does not occur when people make choices between two monetary bonuses. The paper also presents evidence on how this pattern changed during the Cambridge Analytica scandal.

## 1. INTRODUCTION

There is a widespread intuition that when making decisions about privacy, people are inconsistent. People share lots of data; people are angry about increased data collection. This intuition that people are inconsistent is shared widely enough to have a name—the privacy paradox (Acquisti, Brandimarte, and Loewenstein 2015).<sup>1</sup> This paper contributes to this literature with a simple experiment documenting a novel mechanism—information avoidance—that can help explain the privacy paradox. In brief, many people have a revealed preference for privacy, but some of the same people also have a revealed preference for avoiding learning about whether they will lose their privacy. The paper's second

DAN SVIRSKY is a data scientist at Uber Technologies. I thank Jennifer Arlen, Oren Bar-Gill, Omri Ben-Shahar, Christine Exley, Oliver Hart, Louis Kaplow, David Laibson, Florencia Marotta-Wurgler, Matthew Rabin, Josh Schwartzstein, Charlotte Svirsky, Tess Wilkinson-Ryan, and participants at the 2018 Harvard Law and Economics Workshop, the 2018 Harvard Economics Games and Markets Seminar, the 2018 New York University Empirical Contracts Workshop, and the 2019 Privacy Law Scholars Conference for helpful comments.

1. This phrase was mentioned in roughly 5,000 scholarly articles between 2010 and 2018.

[*Journal of Legal Studies*, vol. 51 (January 2022)]

© 2022 by The University of Chicago. All rights reserved. 0047-2530/2022/5101-0003\$10.00

contribution is showing how these preferences evolved during the Cambridge Analytica scandal, which occurred shortly after initial rounds of the experiment.

In the experiment, participants who completed a survey decided whether to take it anonymously or after logging in with a Facebook account in exchange for a monetary bonus. When participants in a direct trade-off treatment faced a choice between a \$.50 bonus and privacy, 64 percent of them refused to share their Facebook profiles in exchange for \$.50. Indeed, when facing a standard price list tool to elicit preferences, the majority of participants in an elicitation treatment (who make close to minimum wage) were unwilling to share their Facebook data for \$2.50, and 40 percent refused offers of \$5.00.

However, when the privacy settings were veiled (but could be revealed costlessly and instantly with the click of a button, as in the moral-wiggle-room experiment in Dana, Weber, and Kuang [2007]), many participants kept themselves in the dark and opted for more money. Participants in a veiled trade-off treatment faced a choice between a \$.50 bonus and a \$0 bonus. They knew that one bonus would mean giving out their Facebook profiles, and they could click a button to check which option involved a loss in privacy. Most people (58 percent) did not click, and only 40 percent kept their Facebook profiles private. Hence, people who were willing to pay nearly an hour's worth of wages to stay private were also able to take a \$.50 bonus and hope for the best, even when learning about the privacy settings would have taken seconds.

Importantly, this same avoidance pattern did not hold when participants made a choice between two monetary bonuses rather than between money and privacy. In a placebo veiled trade-off treatment, participants faced the same experimental interface as in the veiled trade-off group, but the second option was also a monetary bonus. The size of the second monetary bonus was drawn from the distribution of willingness-to-pay (WTP) prices from participants in the elicitation treatment. When facing this choice, participants in the placebo veiled trade-off treatment clicked to reveal the second option 66 percent of the time, a rate significantly different from the reveal rate in the veiled trade-off group.

This paper also presents data on changes in privacy preferences before, during, and after the Cambridge Analytica scandal, which made privacy issues more salient for many Facebook users. By happenstance, an initial round of the experiment was run several weeks before the scandal became public. Once the story broke, privacy issues—and more specifi-

cally, privacy issues surrounding Facebook data and third-party applications—dominated the news, appearing in articles on the front page of the *New York Times* on most days for a month. The experiment was rerun twice with new participants, once at the peak of the scandal and again a month later.

I find that privacy preferences did not change during the scandal, but information-avoidance behavior diminished. When facing the direct trade-off treatment, 64 percent of participants chose to keep their Facebook profiles private instead of getting \$.50, compared with 67 percent before the scandal (a slight and statistically insignificant drop). However, participants in the veiled trade-off treatment were more likely to click to learn the privacy setting before making their choice, which ultimately resulted in 58 percent opting for privacy (compared with 40 percent before the scandal). But this effect was short-lived. Forty days later, 46 percent opted for privacy over \$.50—a proportion statistically indistinct from the prescandal level and significantly lower than the peak-of-scandal level.

The results of the experiment make people's inconsistency over privacy choices more mysterious. Until now in the literature, the dominant explanations of the privacy paradox were revealed preference (Athey, Catalini, and Tucker 2017), behavioral or cognitive biases (Acquisti, Brandimarte, and Loewenstein 2015), and resignation (Turow, Hennessy, and Draper 2015). According to the revealed-preference explanation, people give away their data because they value the services they get in return. Under the bias explanation, people do not realize they are giving away their data or the implications of that decision, or they have difficulty optimizing like a rational actor. According to the resignation explanation, people feel that they have no realistic way of increasing their privacy and are resigned to losing it. In contrast, my experiment finds stark evidence of departures from rational decision-making, since participants avoid low-cost information that can help them make a better choice. Similarly, the experiment shows that the privacy paradox persists, even in a setting where many—though not all—of the behavioral biases that could explain the paradox are unlikely to affect the treatment group and control group differently. At the same time, the experiment also shows that instead of being resigned to losing privacy, people are willing to pay significant sums for privacy when given a chance.

These results have practical, important implications for privacy-law doctrine in the United States, which relies on giving consumers better notice before they make privacy decisions. A required notice makes sense if

people's inconsistency regarding privacy is explained by revealed preferences or ignorance, since either way better disclosure helps people make better choices. This experiment shows that such a policy will be difficult to execute, because even when, as in this experiment, a privacy disclosure is two words long ("high privacy" versus "low privacy"), many people are willing to avoid the disclosure and give away their data.

## **2. BACKGROUND**

This section gives a brief background on privacy law that is relevant to this experiment and on the literature on privacy and economics. It has three parts. First, it describes how existing privacy law in the United States relies on giving people information about data collection. Second, it describes the research on privacy that has led scholars to conclude that people are inconsistent about privacy choices because they are ignorant or boundedly rational. Third, it shows that information avoidance—a phenomenon well documented in other domains—is an alternative explanation for people's inconsistency.

### **2.1. Privacy Law Relies on Giving People Information**

Firms in the United States can legally harvest data from consumers so long as consumers receive proper notice and agree to the exchange. This framework, known as notice and choice, is the standard in US privacy law (Strahilevitz 2010). The notice-and-choice model was first outlined in a 1973 report by the US Department of Health, Education, and Welfare, and this legal framework was a departure from how privacy law developed. Before the rise in Internet commerce and telecommunications, privacy was governed by tort law (Warren and Brandeis 1890; Prosser 1960; Posner 1978). So long as it was not the government invading privacy—in which case constitutional protections would be relevant—a person could enforce various common-law rights to privacy under private causes of action (for example, a right to seclusion). As private data have become dominated by Internet transactions, privacy law has been increasingly governed by contract law principles.<sup>2</sup>

2. There is more stringent regulation for certain consumers and certain industries. Banks send annual privacy notices because of the Gramm-Leach-Bliley Act. Doctors require patients to sign an extra form because of the Health Insurance Portability and Accountability Act. Websites ask users if they are older than 13—not 18, not 12, not 16—because of the Children's Online Privacy Protection Act. Outside the United States, there is more stringent regulation still. The European Union has started enforcing the General

Since privacy is governed by free choice, it becomes important to understand when and why consumers sell their personal data. As a result, much of the empirical literature on privacy looks at how much consumers value keeping their data private in voluntary transactions.

## 2.2. Privacy Preferences Are Fickle

The question of how much people value privacy has been challenging to answer because people's privacy decisions are fickle. In the experiment reported in Acquisti, John, and Loewenstein (2013), the authors offered people gift cards in exchange for completing a survey. When endowed with a \$10 anonymous gift card, about half of participants chose to keep it rather than exchange it for a \$12 nonanonymous gift card. When endowed with the less private \$12 card, 90 percent of participants chose to keep it rather than exchange it for the \$10 anonymous card. John, Acquisti, and Loewenstein (2011) find that people volunteer more sensitive information when asked indirectly and when a website seems less professional. Similarly, an experiment by legal scholars testing different disclosure techniques finds that people's privacy behavior is not much affected by providing them more and better information about their privacy choices (Ben-Shahar and Chilton 2016; for more evidence, see Bakos, Marotta-Wurgler, and Trossen 2014). Along the same lines, Athey, Catalini, and Tucker (2017) report a field experiment in which Massachusetts Institute of Technology students were given Bitcoin and invited to start using one of four digital wallets with varying levels of privacy and convenience. Students' wallet choices were affected by the order in which the wallets were presented, and students' self-reported privacy preferences had no predictive power for their privacy choices. Hence, people's privacy decisions appear inconsistent.

There are several explanations for this inconsistency. Roughly, they can be broken down into three explanations: consumers are making rational decisions, consumers are struggling to optimize correctly, and consumers are resigned to not having meaningful choices.<sup>3</sup>

---

Data Protection Regulation, which imposes stronger consent requirements for data collection, forces firms to delete personal data at a consumer's request, and allows for fines up to 4 percent of a firm's global revenue.

3. These categories are meant as a helpful overview of the different classes of explanations, but this categorization has shortcomings that should be noted. First, the categories are broad and include many different explanations. Second, the boundaries between them are porous. Resignation may be rational or boundedly rational, for instance. Third, the categories are not mutually exclusive or exhaustive. All could be operating at the same time. Nonetheless, they do provide a sense of the dominant approaches in the literature.

The first class of explanations rely on the idea that rational models of consumer behavior are consistent with how people make privacy choices. It could be quite simple: people give up privacy simply because this maximizes their utility. People say that they do not like losing privacy, but people also say that they do not like losing \$5. That does not mean it is a paradox if lines of people in a Starbucks happily give away \$5 to a barista each morning—provided they get a fancy latte in return. Similarly, consumers may be giving up privacy because it is not worth the cost to protect it. If a consumer is choosing between two browsers to use, it might be difficult to figure out which one offers more privacy protection. That she values privacy does not necessarily mean that it is optimal to spend considerable time gathering information.

The second class of explanations rely on models of consumer behavior that allow for departures from rational optimization. Perhaps people are unaware of how much data they are emitting and simply do not realize the amount of privacy they are trading. Perhaps they struggle to value privacy because it is abstract or because privacy costs are inchoate and uncertain, both in scope and timing (Acquisti, John, and Loewenstein 2013). Perhaps the costs of privacy decisions are temporally distant, while the benefits of trading privacy are more proximate. This line of scholarship draws on classic findings from psychology and economics, like the endowment effect and framing effects, to explain people's fickle privacy preferences.

The third class of explanations argue that people may feel like they have no real choice. People may be resigned to the loss of privacy in society and feel that their decisions will not have any impact (Turow, Hennessy, and Draper 2015). Why spend time on one's Facebook privacy settings if Facebook will track consumers across websites in ways that cannot be monitored? Under this view, people's stated valuations of privacy are accurate. People really do value privacy, but they believe that changing their behavior makes no difference.

For the first two explanations—revealed preference or bounded rationality—more information is better. If it is costless, better information will help people make more informed choices in line with their preferences. Or, if people struggle to make consistent choices, better information can help dispel the cognitive biases or lack of awareness that might drive this inconsistency.

### 2.3. Information Avoidance Is an Important Driver of Behavior

Information avoidance is a well-documented phenomenon in other domains where people's stated preferences do not match their revealed preferences. There is a robust literature from psychology and economics on information avoidance (Golman, Hagmann, and Loewenstein 2017). While economists typically model information as an intermediate good (Posner 1978; Stigler 1961)—that is, valuable only because it helps people achieve ends—scholars in psychology and economics increasingly recognize that people sometimes behave as if information has emotional valence (Oster, Shoulson, and Dorsey 2013). More information is not always better. Consider a now widely replicated experiment on moral wiggle room (Dana, Weber, and Kuang 2007), which is the basis for the experimental design used in the present paper. In that experiment, a participant has to choose payoffs for herself and a partner whom she does not meet. In a baseline condition, she chooses between two options: \$6 for me and \$1 for my partner or \$5 for me and \$5 for my partner. Most people pick the second option. A treatment group faces a slightly modified choice: \$6 for me and \$ $X$  for my partner or \$5 for me and \$ $Y$  for my partner. In this case, either  $X$  is \$1 and  $Y$  is \$5 (as in the baseline group) or  $X$  is \$5 and  $Y$  is \$1. The person can costlessly click to find out the values of  $X$  and  $Y$ .

Consider what a typical economic model would predict. If in the baseline experiment I prefer \$5 and \$5 over \$6 and \$1, and this is a strong preference, then I should click to find out the value of  $X$  and  $Y$ . Either I will find that I am in the baseline case, and I can choose \$5 and \$5 again, or I will find that I am in the easier case and can choose \$6 and \$5.

But this is not how people act in the experiment. Instead, people avoid learning the values of  $X$  and  $Y$  and pick the option of \$6 for me and \$ $X$  for my partner. They exploit the wiggle room to act selfishly. Other experiments on altruism, lab and field based, have similar results (Exley 2016; Lazear, Malmendier, and Weber 2012; DellaVigna, List, and Malmendier 2012).

This pattern of behavior is important across disparate domains. In health, one study finds that 27 percent of intravenous drug users at risk of human immunodeficiency virus (HIV) infection who were tested did not return to the clinic to see their results (Sullivan, Lansky, and Drake 2004), even though knowing one's HIV status can lengthen one's life. In family planning, 20 states have laws requiring a woman to see a picture

of the fetus before getting an abortion.<sup>4</sup> Presumably, women know what a fetus looks like, so the law was not passed because the increased information about the fetus's appearance will lead to more informed choices. In sum, people avoid information that upsets them, even if in theory a utility-maximizing agent would never reject free information. Given the central focus in privacy law on giving consumers better, cheaper information, and given the psychology and economics literature on how people avoid information, this paper focuses on testing an important open question: do people engage in information avoidance when making privacy decisions?

### 3. EXPERIMENTAL DESIGN

I conduct an experiment to test for information avoidance in privacy decisions. Participants are randomized to one of two treatments: a direct trade-off treatment and a veiled trade-off treatment. This section first discusses the overall timeline of the experiment and then describes the two treatments in detail.<sup>5</sup> The experiment was preregistered on AsPredicted under the title "Information Avoidance and Internet Privacy."<sup>6</sup>

A total of 795 participants were recruited on Amazon Mechanical Turk to take a short survey about health and financial status.<sup>7</sup> All participants were informed that before taking the survey, they would make decisions about the size of a bonus payment, to be received on completion, and the privacy settings of the survey.<sup>8</sup> The experiment was conducted on November 20, 2018.

After recruitment, the timeline of the experiment consisted of three

4. Guttmacher Institute, Requirements for Ultrasound (<https://www.guttmacher.org/state-policy/explore/requirements-ultrasound>).

5. The experiment was approved by Harvard's Committee on the Use of Human Subjects as protocol IRB18-0061.

6. For the preregistration document, see Dan Svirsky, Information Avoidance and Internet Privacy (#16702) (<https://aspredicted.org/gm9nv.pdf>).

7. Research increasingly suggests that, for the purpose of social science experiments, Mechanical Turk users are a reliable sample. Irvine, Hoffman, and Wilkinson-Ryan (2018) replicate three experiments using in-person labs, national online platforms, and Mechanical Turk and find that the results are constant across samples. The key difference is that Mechanical Turk users were significantly more attentive than the other sample groups. See also Hoffman et al. (2020), which replicates an experiment on Mechanical Turk users, college students in a physical lab, and college students in an online setting.

8. The median hourly wage for workers was \$14.96 (based on a median payment of \$1.52 for a median completion time of 8 minutes and 6 seconds).



stages: instructions and practice, privacy settings, and a survey.<sup>9</sup> First, participants were shown an introductory screen that gave an overview of their participation. Figure OA1 shows the first page of instructions, presented to all participants, and Figure OA2 shows the second page, which depended on the treatment group to which the participant was randomly assigned.<sup>10</sup> Participants were told that they would take a survey and that, while everyone would take the same survey, each participant would be given a choice between two privacy options. Participants could opt for high privacy, in which case their survey answers would be anonymous, or for low privacy, in which case they would click a Log in with Facebook button at the top of the survey. This meant that the survey administrator would see, in addition to a participant's survey answers, her public Facebook profile (including profile picture, name, and gender) and her email address.<sup>11</sup> Participants who chose low privacy would not be allowed to finish the survey until they logged in. Participants then completed two short practice rounds that looked identical to the privacy-settings task. After this, participants chose their privacy settings. Figure OA3 and Figure OA4 show screenshots of the module screen where participants made the privacy choices for the direct trade-off and veiled trade-off treatments, respectively. After completing the privacy-settings stage, participants completed the survey stage, shown in Figure OA5.

In the direct trade-off treatment, participants made only one decision: a direct choice between a \$.02 bonus and privacy option A or a \$.52 bonus and privacy option B. The privacy options were randomized so that half the time participants faced a degenerate choice between {more money, more privacy} and {less money, less privacy}. The other half of the time participants faced a true trade-off between money and privacy.

In the veiled trade-off treatment, participants faced the same decision as in the direct trade-off treatment, but the privacy setting was initially hidden. Participants had to click a button to reveal the column describing

9. Online Appendix OA presents the instructions for the experiment.

10. For the randomization process, participants were assigned a random integer between 0 and 100,000. The modulus with respect to 4 was taken, and participants were assigned to one of the four treatments accordingly.

11. Note that, by default, Facebook makes these profile details searchable on Google. It is not clear that participants were aware of this, so some participants may have believed that they were revealing information that was shared only with friends instead of simply making it easier for a stranger to identify them in an otherwise anonymous setting. In any case, this does not affect the experiment's results, since participants across all treatments would be just as likely to have this misconception.

the privacy settings, and there was a 50 percent chance that the higher monetary bonus would mean losing their anonymity.<sup>12</sup>

How should a participant behave in the veiled trade-off treatment? Online Appendix OB models the participant's decision-making problem. In brief, a participant is one of two types: she values privacy at either less than or equal to \$.50 or more than \$.50. If she values privacy very little, it never makes sense to click to reveal, since no matter what she learns, she will choose the option that gives her an extra \$.50. As Online Appendix OB demonstrates, if she values privacy at more than \$.50, and if clicking costs—which could be costs in terms of time spent or psychic costs or any nonmonetary cost—are 0, then she will always click to reveal the privacy settings. Online Appendix OB also provides a formula that relates the decision to reveal privacy settings to clicking costs.

The user interface for the experiment was coded using HTML and Javascript, which ensured that the reveal button would work instantaneously, without refreshing a page. When a user clicked the reveal button, Javascript code changed the visibility setting of the hidden column to visible. The hidden column therefore became visible immediately. The users' choices and data were sent to a MySQL database using PHP code.

### 3.1. Placebo Test

Any difference between the direct trade-off and veiled trade-off groups might be driven by clicking costs rather than information avoidance. Suppose many people value privacy at \$.51 but the reveal button imposes a few cents of costs from extra time or mental effort. Then I would observe a treatment effect because participants rationally conclude that it is not worth spending a few cents of effort for a \$.01 gain.

I test this alternative explanation in two ways. First, I use an elicitation treatment to gather the full distribution of WTP prices for privacy.

12. Note that for both groups, there was a 50 percent chance of facing a degenerate choice between {more money, more privacy} and {less money, less privacy}. These decisions cannot tell much about how a person values privacy, so they are omitted from the main analysis below. The resulting sample size is 535 participants: 117 in the direct trade-off group, 130 in the veiled trade-off group, 164 in the placebo veiled trade-off group, and 124 in the elicitation treatment. While the degenerate choices are not part of the main analysis, they can confirm how rationally the participants behaved. In the direct trade-off group, 5.6 percent of participants chose less money and less privacy. The veiled trade-off group clicked to reveal the privacy settings at similar rates in the degenerate-choice group (43 percent versus 42 percent for those facing a real choice), and on clicking 4 percent chose less money and less privacy. Hence, the degenerate-choice situations suggest that participants by and large behaved as one would expect.

In the elicitation treatment, instead of making just one choice between privacy and \$.50, participants made 10 choices, with the bonus varying between \$.25 and \$5.00. Participants were told that one of their choices would be enforced. This is a standard technique in applied microeconomics to elicit a WTP price, in this case for staying anonymous. Participants faced a table of choices as in Figure OA3: either a \$.02 bonus and high privacy or a \$X.YY bonus and low privacy, with \$X.YY ranging from \$.27 to \$5.02. Hence, if someone opted to stay anonymous when offered \$.50, \$1.00, and \$1.50 but not \$2.00, then one can infer that her WTP for staying anonymous is between \$1.50 and \$2.00.

Second, I conduct a placebo veiled trade-off treatment. This treatment is identical to the veiled trade-off treatment, but instead of making a choice between one monetary bonus and privacy, participants make a choice between one monetary bonus and a second monetary bonus. The first bonus is \$.50, as in the main experiment, but the second bonus is randomly drawn from the distribution of WTP prices from the elicitation treatment. If clicking costs alone are driving results in the main experiment, where people have some distribution of WTP prices for privacy, then one would observe the same-sized treatment effect if the second column is instead a monetary bonus drawn from the same distribution of WTP prices.

### 3.2. Privacy Preferences and the Cambridge Analytica Scandal

On March 18, 2018, the *Guardian* first reported that Cambridge Analytica, a political consulting firm, had harvested data from nearly 90 million Facebook accounts to help conservative political candidates.<sup>13</sup> Most of the data were obtained without consent, and responses to the report quickly escalated until there was a public scandal. Cambridge Analytica largely relied on Mechanical Turk to construct its illicit data set. Mechanical Turk users were invited to share their Facebook data in exchange for bonuses between \$2 and \$4, but in addition the users gave permission to Cambridge Analytica (under false pretenses) to access their friends' profile data as well. The option to share friends' data was discontinued in 2016.

The pilot round of my experiment was initially run on February 23, 2018—23 days before news about the data mining broke. A second round

13. The story was first broken by the *New York Times* in November 2016 (Funk 2016) but did not capture the public's attention until the *Guardian's* story (Cadwalladr and Graham-Harrison 2018).

was conducted 11 days after the report became public. A third round was conducted 41 days later.

Figure OA7 uses Google trends data to show how often people searched for the phrase “Facebook privacy settings.” The graph shows a spike in such searches in the immediate aftermath of the scandal, coinciding with the second round of the experiment. This spike in search interest diminished by the time the third round was conducted.

It is important to note that this analysis is exploratory. The experiment was not designed to track how people’s privacy preferences evolved in response to a major scandal. It was merely rerun in response to the timing of the scandal to shed light on important questions about how people’s privacy preferences change in response to major scandals. These analyses were not preregistered.

#### 4. RESULTS

While the survey questions are not the focus of this paper, Table A1 presents summary statistics on the survey answers and a balance check providing evidence that the randomization effectively balanced the sample across the treatments. Of note, nearly all—94.6 percent—participants reported having a Facebook account. This is important, as it is not clear how a person without a Facebook account would make a valuation decision in this experiment (though the balance check confirms that, however this would affect results, the lack of a Facebook account was similar across treatments). All analyses are substantively unchanged whether these participants are excluded or included, but in the data below they are included. Across participants, Facebook use was common. The median participant reported using Facebook four or more times per week.

The analyses below are restricted to participants who completed both the privacy valuation task and the survey, but attrition from the study may be of substantive interest in its own right, for example if people drop out of the study when they see that they have to share Facebook information. Attrition was quite low. In the direct trade-off and veiled trade-off treatments, attrition (defined as people who read the instructions but quit before the survey round) was 5 percent and 2 percent, respectively—a statistically insignificant difference.

The treatment effect is robust even if I exclude participants who failed comprehension and attention checks. During the survey (and after completing their privacy choices), one question asked, “How old were you

when you were 10?” with several options in a drop-down menu. Roughly 84 percent of participants correctly answered. Another question in the survey asked, “How carefully did you make your choices?” with three options: not carefully at all, a little carefully, and very carefully. Roughly 75 percent of participants said they answered the questions very carefully, 21 percent said a little carefully, and 3 percent said not carefully at all. Note that by default “not carefully at all” was selected. The main results are substantively unchanged if I exclude participants who did not pay very careful attention or who answered the comprehension question wrong.

Another robustness concern is confusion—did participants in the veiled trade-off treatment mistakenly assume that a low bonus meant they would keep their privacy? That is, participants in the veiled trade-off treatment could have made (incorrect) guesses about the privacy settings, even though the instructions explicitly told them that the privacy settings were randomized. For example, a person could assume that the lower payoff always meant more privacy. In that case, one would expect that people would choose to never click to reveal the privacy setting but then nonetheless choose the lower payoff. Such behavior occurred in 4 percent of participants in the veiled trade-off treatment. The results discussed here categorize these participants as having chosen privacy over \$.50, but the results do not change if these participants are instead dropped. Table 1 reports the results of these robustness checks. The main results hold throughout.

#### **4.1. Direct Trade-off Treatment versus Veiled Trade-off Treatment**

I find a treatment effect from putting a costless veil on privacy settings. A total of 64 percent of people in the direct trade-off treatment refused to sell their Facebook data for \$.50.<sup>14</sup> In contrast, in the veiled trade-off treatment, when the privacy consequences of their actions were initially hidden, only 40 percent sold their Facebook data for \$.50. Table 1 shows that this result is statistically significant ( $p < .001$ ) in a linear regression predicting whether a participant opts for \$.50.<sup>15</sup> A majority of partici-

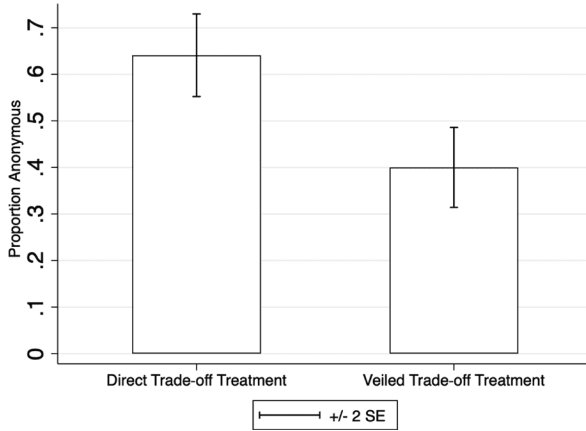
14. This is in line with the results from the elicitation treatment group described below. In that group, 59 percent rejected an offer of \$.50 to share their Facebook profiles, a slightly lower but statistically insignificant difference.

15. The results are significant using a Fisher’s exact test comparing the proportion of those who sold their data ( $p < .001$ ). For all reports of statistical significance in this paper, the significance level is tested using both a linear regression and a Fisher’s exact test, with no differences in the significance level.


**Table 1.** Privacy Decisions in the Direct Trade-off and Veiled Trade-off Groups

	Full Sample		Passed Comprehension Check (3)	Answered Carefully (4)	Intersection of Robustness Checks (5)
	(1)	(2)	(3)	(4)	(5)
Veiled trade-off group	-.24 (.06)	-.23 (.07)	-.21 (.07)	-.22 (.06)	-.19 (.07)
Age		.01 (.004)			
Diet attempts in lifetime		-.01 (.02)			
Exercise sessions in a typical week		-.04 (.02)			
Annual income		.02 (.03)			
Credit card debt		.03 (.03)			
Constant	.64 (.05)	.34 (.17)	.63 (.05)	.63 (.05)	.62 (.05)
N	247	213	207	239	201
Adjusted R <sup>2</sup>	.05	.07	.04	.04	.03

**Note.** Results are from a linear probability regression of a binary variable for whether the participant stayed anonymous on a binary variable for whether the participant was in the veiled trade-off treatment. The omitted group is the direct trade-off treatment. Robust standard errors are in parentheses.



**Figure 1.** Treatment effect



**Figure 2.** Decision tree

pants in the veiled trade-off treatment (58 percent) chose not to look at the privacy setting before deciding to take the \$.50.

Figure 1 shows the proportion of participants who remained anonymous in the direct trade-off treatment and the veiled trade-off treatment. Figure 2 breaks down participants’ decisions in the treatments, including their privacy choices and decisions whether to click. In Table 1, where the unit of observation is an individual, the dependent variable is whether the participant remained anonymous, and the independent variable is an indicator variable for being in the veiled trade-off treatment, the treatment effect is robust even if I exclude participants who failed comprehension and attention checks. Columns 2–5 report robustness checks using different samples. Column 2 includes controls for survey answers, while columns 3–5 exclude participants on the basis of com-

prehension<sup>16</sup> and attention<sup>17</sup> and the intersection of the two. The main results hold throughout.

#### 4.2. Placebo Veiled Trade-off Treatment and Elicitation Treatment

The results from the placebo veiled trade-off treatment and elicitation treatment give strong evidence that clicking costs are not driving the treatment effect in the main experiment. An alternative explanation of the results is that clicking to reveal the privacy settings is costly. It is possible that many participants value privacy at only slightly more than \$.50, so when faced with the reveal button, they rationally decide that the costs of clicking and deciding are not worth the small gain in utility of potentially getting privacy over money.

One way to rule this out is by directly eliciting people's WTP for sharing their data. In doing so, I find that the majority of people value privacy at \$2.50 or more. Table 2 shows people's WTP price for staying anonymous in the elicitation treatment. The values indicate the proportion of participants who switched from high privacy to low privacy at the price offered. The results show that a plurality of participants—42.1 percent—refuse to share their Facebook profiles at all prices up to \$5.00. Note that the average hourly wage on Mechanical Turk is roughly \$5 per hour (Hara et al. 2018), so these participants would rather spend an hour of time completing mundane computer tasks than share their public Facebook profiles with a survey taker. Nonetheless, the second most common WTP price was at the lower end, with 20.6 percent choosing to share their Facebook profiles at \$.25. The remaining 37 percent evinced a WTP price between \$.25 and \$5.00.<sup>18</sup>

16. As described above, the comprehension check asked how old the participant was when she was 10 years old; the default response was 5 years old. Of the 247 participants in the direct and veiled trade-off groups, 40 failed this check.

17. For the attention question described above, the default option was "not carefully at all." Of the 247 participants in the direct and veiled trade-off groups, eight reported paying little attention.

18. Irrational behavior, defined as having multiple switching points, was rare. It is hard to interpret someone giving up her privacy for \$.50 ( $\text{privacy} > \$.50$ ) but not for \$1.00 ( $\$.50 > \text{privacy} \geq \$1.00$ ), assuming that she also values more money over less money. In the elicitation treatment, 84 percent gave rational answers in the sense of having at most one switching point. This is a relatively low level of multiple switching behavior compared with other experiments that use multiple price lists, which typically find levels of multiple switching behavior ranging from 10 percent to as high as 50 percent. See, for example, Jacobson and Petrie (2009), which gives thoughtful consideration to such behavior and asks whether it explains something substantive about participants. See also Andreoni and Sprenger (2012) and Meier and Sprenger (2016). This finding also



**Table 2.** Distribution of Willingness-to-Pay Prices in the Elicitation Treatment

Price	Percentage	Cumulative Percentage
.25	20.6	20.6
.50	20.6	41.1
1.00	1.9	43.0
2.00	1.9	44.9
3.00	3.7	48.6
4.00	9.3	57.9
5.00 <sup>a</sup>	42.1	100.0
Total	100.0	

**Note.** Data are the breakdown of prices for switching from high privacy to low privacy. People with multiple switching points (18 of 124 respondents) are omitted. *N* = 106.

<sup>a</sup> Includes respondents who refused to sell data at all prices.

Using the results of the elicitation treatment, I can say with more precision how high clicking costs would have to be to explain the main treatment effect. Using the model in Online Appendix OB, I find that clicking costs would have to be nearly \$2.00 to explain the treatment effect in this experiment. I consider this unlikely in this context, especially given that the median participant clicks a mouse 31 times during the experiment and is paid \$1.52 for her participation. If clicking to reveal the two-word privacy settings really imposed a cost of \$2.00, participants would be making a massive mistake by finishing the experiment.

The results of the placebo veiled trade-off treatment give more direct evidence that the results are not driven by clicking costs or confusion about experimental design. Recall that in the placebo veiled trade-off treatment, participants chose between two monetary bonuses, with the value of the second bonus drawn from the distribution of WTP prices in the elicitation treatment. Participants knew the size of the second bonus and had to click to reveal it. Among this group, the proportion of

---

suggests that Mechanical Turk workers evinced similar levels of this type of irrationality when compared with college students and people with moderate incomes in tax-filing centers, among other samples. In calculating the distribution of willingness-to-pay (WTP) prices, I exclude participants with multiple switches, but the results are similar if I instead include them and define their switching point as the lowest switch, the highest switch, or the average of the two.

**Table 3.** Behavior in the Placebo Veiled Trade-off Treatment

Second Bonus	Did Not Click	Clicked
.25	39	29
.50	2	1
1.00	1	2
2.00	2	8
3.00	4	8
4.00	0	9
5.00	8	51

**Note.**  $N = 164$ .

participants clicking to reveal the bonus was .66.<sup>19</sup> This is higher than the click rate of .42 in the main experiment, when participants chose between money and privacy, and the difference is statistically significant ( $p < .001$ , Fisher's exact test). Table 3 shows the proportion of clicks in the placebo veiled trade-off group by the size of the second bonus. These results suggest that people are capable of clicking to reveal the second bonus, and do so in a roughly rational way, when money is at stake instead of privacy. If the main treatment effect were driven by people misunderstanding the experiment, I find no evidence of such misunderstanding when two monetary bonuses were at stake.

#### 4.3. Treatment Effect and Time Spent Reading Instructions

One interpretation of the treatment effect is that people want to finish the experiment as quickly as possible and are willing to speed through the tasks to maximize earnings on the Mechanical Turk platform, even at the cost of privacy. If this is true, one might expect to observe heterogeneous treatment effects: people who click through the instructions fastest might be more likely to engage in information avoidance. I can test this by seeing whether there is a relationship between the speed with which people read through instructions and the treatment effect. I do not find such evidence. Figure OA8 graphs the proportion of participants who opted to

19. Note that for many of the 33 percent of participants who chose not to click to reveal the second bonus, this was a rational choice. The values in the second column were drawn from the distribution of WTP prices from the elicitation treatment. In many cases, the second bonus was less than \$.50. Participants knew the size of the bonus but not its row. For these participants, clicking to reveal the second column imposes extra time but should not affect their choice, so not clicking it makes sense.

remain private versus the time they spent reading instructions and shows no relationship.

#### 4.4. Privacy Valuations and Price List Framing

One interpretation of the results is that privacy valuations are highly sensitive to experimental design choices. Putting a costless veil on the privacy settings might have an outsized impact, but any variety of framing choices might also have an impact on people's willingness to pay for privacy.

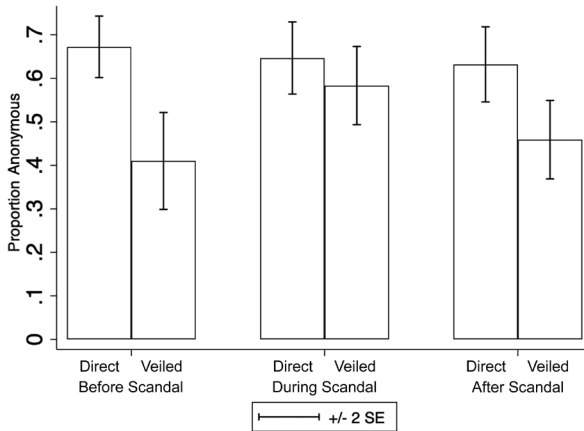
During the pilot phase of this study, a version of the elicitation treatment was run in which participants made a WTP decision for a price menu ranging from \$.06 to \$.54 in increments of \$.06 rather than from \$.25 to \$5.00. If this hypothesis is correct, one might expect that people's WTP price for privacy will be affected by the price list, just as people's risk preferences are (Beauchamp et al. 2015). I can test this hypothesis by looking at overlapping choices between these two groups. Though the groups faced different price menus, both groups made a decision of whether to sell data for roughly \$.25 and \$.50.

I find no evidence of price list effects. Even when one group faced a price menu of \$.06 to \$.54 and the other faced a menu of \$.25 to \$5.00, their decisions in the overlapping price range were almost identical, as shown in Figure OA9.<sup>20</sup>

#### 4.5. Privacy Preferences and the Cambridge Analytica Scandal

Did privacy preferences in either the direct or veiled trade-off groups change during the Cambridge Analytica scandal? Importantly, any changes are not necessarily attributable to the scandal, nor is the direction of any effect obvious *ex ante*. The experiment is limited in the sense that results could be driven by changes in the underlying sample of participants or trends that affect people's WTP for keeping their Facebook profiles private from a third party but that were unrelated to the Cambridge Analytica scandal. To get a sense of these issues, Table A2 presents a bal-

20. This is in sharp contrast to Benndorf and Normann (2014), who elicit privacy valuations using a Becker-DeGroot-Marschak (BDM) mechanism. Unlike in the present experiment, in which participants made a series of binary choices, Benndorf and Normann (2014) find that when the upper threshold in a BDM mechanism changes, people's privacy valuations change accordingly. Taken together, these findings suggest that people's valuations of privacy are sensitive to the elicitation techniques used but not more so than in other domains (Beauchamp et al. 2015). Open questions are whether the elicitation technique in the present paper dampens volatility in valuations, whether a BDM mechanism heightens volatility, or both.



**Figure 3.** Privacy and the Cambridge Analytica scandal

ance check to see whether the three samples of participants (or the sample from the main experiment in November) are significantly different in any of the survey responses. I find balance across all four groups, which suggests that in terms of reported age, credit card debt, income, and exercise patterns, the sample did not measurably change before, during, or after the scandal. Though there is evidence that Mechanical Turk participants can vary across time of day and day of week (Casey et al. 2017), these findings suggest that to the extent that is true here, it does not show up in participants' survey answers. The sample sizes for each experimental round were also roughly consistent.

Figure 3 shows the proportion of people who chose to keep their Facebook data private during the survey instead of getting a \$.50 bonus, by treatment and across the three experiment dates. At the height of the Cambridge Analytica scandal, people's behavior in the direct trade-off treatment was unchanged. Before the scandal, 66 percent opted for privacy over \$.50 in the direct trade-off treatment. At the height of the scandal, this number was 64 percent, and 1 month later the proportion was 63 percent. None of these changes were statistically significant.

However, the treatment became less effective. Before the scandal, the veiled trade-off treatment caused a 26-percentage-point drop ( $p < 0.001$ ) in the proportion of people opting to keep their Facebook profiles private. At the height of the scandal, the veiled trade-off treatment caused a 9-percentage-point drop ( $p = .06$ ). One month after the scandal, the

treatment was effective again, causing a 17-percentage-point drop ( $p = .003$ ). The treatment effect at the height of the scandal was significantly different from the treatment effects before ( $p = .01$ ) and after the scandal ( $p = .03$ ). Consistent with this result, the proportion of participants who clicked to reveal the privacy settings rose from .42 to .52 ( $p = .02$ ) during the scandal and then fell to .40 ( $p = .007$ , compared with clicking levels during the scandal).

Table A3 presents regression results and robustness checks. In the regression specification,  $p$  is an indicator variable for whether an individual remained anonymous,  $T$  is an indicator variable for whether the participant was in the veiled trade-off treatment,  $CA$  is an indicator variable for whether the experiment date occurred shortly after the Cambridge Analytica scandal, and  $Post$  is an indicator variable for whether the experiment occurred 40 days after the scandal:

$$p = \beta_0 + \beta_1 \times (T) + \beta_2 \times (CA) + \beta_3 \times (CA \times T) + \beta_4 \times (Post) + \beta_5 \times (Post \times T).$$

In the regression,  $\beta_1$  measures the treatment effect before the scandal,  $\beta_2$  measures the change in privacy preferences in the direct trade-off treatment group at the height of the Cambridge Analytica scandal,  $\beta_3$  measures the change in the treatment effect at the height of the Cambridge Analytica scandal,  $\beta_4$  measures the change in privacy preferences in the direct trade-off treatment group after the scandal, and  $\beta_5$  measures the change in the treatment effect after the scandal.

In sum, I find no measurable change in survey responses before, during, or after the Cambridge Analytica scandal, nor do I find any change in behavior in the direct trade-off treatment. I do, however, observe that the experimental treatment became significantly less effective, and this was driven by people in the veiled trade-off group opting for privacy more often.

## 5. DISCUSSION

The results of the experiment in this paper provide evidence for two conclusions. First, people do in fact behave inconsistently around privacy decisions. Second, this inconsistency can be explained in part by information avoidance. Because of unique timing, the paper also sheds some light on the effect of public privacy scandals on privacy behavior. The

treatment effect dissipated at the height of one of the biggest, most salient privacy scandals of the past decade, but not because people valued privacy more when directly asked. Rather, when the scandal arose, people's ability to take advantage of the costless veil seems to have weakened. But this change did not signal a new normal: privacy behavior returned to prescandal levels within 2 months of the news breaking.

The experiment documents that information avoidance can play a role in people's privacy decisions, and this in turn adds to the debate on what drives the privacy paradox. The behavior documented in this experiment cannot be easily explained by some behavioral economics theories or by resignation. If the privacy paradox occurs because privacy costs are ephemeral, come in the distant future, or are hard to quantify, then this would affect the direct trade-off and veiled trade-off groups equally. Similarly, the resignation explanation—that people are resigned to losing privacy—is also problematic, since here people are willing to pay significant sums for privacy. If people are resigned to losing privacy—if they believe their choices make no difference—then it is not clear why participants in the elicitation treatment would forgo \$5.00 in exchange for giving up their data.

A related and unresolved question is why people avoid information. There are several plausible mechanisms. One is signaling. People care about privacy, but they also care about being the type of person who cares about privacy. This drives a wedge between the direct trade-off group and the veiled trade-off group, because members of the veiled trade-off group can take the monetary bonus without explicitly choosing to give away their data. A second mechanism is that thinking about a probabilistic chance of losing privacy is itself upsetting, as in the model of anxiety in Kőszegi (2003). A third mechanism is that people do care about privacy but are also able to ignore privacy losses that are not directly in front of them. Another mechanism is choosing costs, from either the weight of a difficult decision (Sunstein 2014) or aversion to regret (Hayashi 2008). It takes effort to make a decision between money and privacy, especially if privacy costs are inchoate or hard to measure. And making a direct choice risks future regret. Perhaps the direct trade-off group has no choice but to make this effort, but the veiled trade-off group might rationally decide that it is better to avoid doing the calculations, exploiting the veil either as a cognitive shortcut or as a way to avoid regret from decisions for which they are responsible. Similarly, it could be that the veiled trade-off scenario is better than having to actively

face the privacy versus money trade-off. In this case, the experimental findings would say more about the desire to avoid a difficult trade-off than about the privacy paradox *per se*. Under this view, the privacy paradox would suggest that in the direct trade-off treatment, more people ought to be rejecting money, given their stated preferences for privacy, whereas people's behavior in the veiled trade-off setting is not about a discrepancy between their stated and revealed preferences. Another explanation is salience—the privacy option is more immediately salient in the direct trade-off treatment, which could change people's focus as they make their decision. Another explanation is a preference for randomization (Agranov and Ortoleva 2017). Still another explanation of the results is that all of these mechanisms are valid to greater or lesser degrees depending on the person and the context. Future research can explore this, for example by changing whether the decision to hide information is active or passive or by changing the probabilities in the experiment. But the bottom line is that regardless of the exact combination of mechanisms, information avoidance might play a role in privacy preferences, which suggests different policy interventions than standard explanations of the privacy paradox.

While the primary goal of this paper is as a sort of existence proof of information avoidance in privacy decisions, teasing out its causes is important, since the policy response will depend on the mechanism. Further, the results of the experiment are arguably more consistent with some explanations than others, and they differ from information avoidance behavior in other domains in important ways. For example, in Dana, Weber, and Kuang (2007), participants avoid information about how much money they share with a partner. But when they avoid information, they never see the outcome of their choice. Participants in this experiment see that they will share their Facebook data after choosing to initially avoid information. This might be more consistent with choosing costs, anxiety, or a preference for randomization rather than a desire to not think about privacy, since privacy is front and center before, during, and even after making the choice to avoid information in the experimental module. This is also less consistent with signaling explanations. That is, a participant may want to think of herself as someone who takes privacy seriously or have an experimenter believe that about her. Avoiding information means avoiding having to demonstrate to herself or another that she would trade privacy away if given the chance. Here, the participant

cannot keep herself in the dark indefinitely, so self-signaling is a more strained explanation.

More broadly, the experiment suggests reasons for skepticism about policy interventions aimed at improving consumer decision-making with better information. Under the simpler explanations of privacy inconsistency—revealed preference and ignorance—policy makers agree that more and simpler information is better (Federal Trade Commission 2012; Kelley et al. 2010). Better notice means better choices, provided the notice is at low cost. Given this, there have been extensive efforts to improve privacy disclosures, for example with a privacy nutrition label. However, this experiment shows that such efforts will be a steep climb. While there is important literature documenting the shortcomings of such disclosures—Ben-Shahar and Chilton (2016) is especially stark—this experiment provides a mechanism that shows why these shortcomings might exist. The results presented here show that even when the privacy settings could be revealed instantly, and even when the settings are a mere two words long (“low privacy” and “high privacy”), most participants still opted not to click. Even when, or especially when, a privacy disclosure is salient, clear, and easily accessible, people may have struthious preferences.



**APPENDIX: ADDITIONAL TABLES**

**Table A1.** Summary Statistics for the Survey Answers

	Direct Trade-off	Veiled Trade-off	Placebo Veiled Trade-off	Elicitation	<i>p</i> -Value
Age	30.97 (7.255)	32.54 (8.646)	32.01 (9.791)	31.56 (9.554)	.56
Diet attempts in lifetime	2.272 (1.564)	2.254 (1.600)	2.085 (1.381)	2.306 (1.302)	.56
Exercise sessions in a typical week	2.429 (1.354)	2.305 (1.406)	2.427 (1.352)	2.234 (1.374)	.59
Annual income	1.333 (1.102)	1.112 (1.193)	1.366 (1.157)	1.137 (1.136)	.16
Credit card debt	.732 (1.099)	.795 (1.136)	.793 (.981)	.734 (1.029)	.93
Has Facebook account	.915 (.281)	.954 (.211)	.963 (.188)	.968 (.177)	.20
Weekly Facebook use	2.829 (1.481)	3.015 (1.381)	2.506 (1.537)	2.911 (1.437)	.02

**Note.** Standard deviations are in parentheses. Income and credit card debt variables are categorical (range of 1–4). Diet, exercise, and Facebook variables can be 0, 1, 2, 3, or 4 or more. Reported *p*-values are from *t*-tests comparing the means of the four groups.

**Table A2.** Summary Statistics: Before, during, and after the Cambridge Analytica Scandal

	Before Scandal	During Scandal	After Scandal	Main Experiment	<i>p</i> -Value
Age	33.54 (10.67)	32.40 (9.711)	33.06 (8.356)	31.75 (8.099)	.19
Diet attempts in lifetime	2.32 (1.58)	2.26 (1.49)	2.31 (1.56)	2.3 (7.79)	.82
Exercise sessions in a typical week	2.48 (1.33)	2.36 (1.34)	2.40 (1.44)	2.35 (1.40)	.51
Annual income	1.33 (1.11)	1.33 (1.15)	1.23 (1.12)	1.18 (1.17)	.27
Credit card debt	.64 (.99)	.71 (.95)	.64 (1.01)	.77 (1.13)	.91
<i>N</i>	311	309	306	535	

**Note.** The *p*-values are for an *F*-test, or the joint hypothesis that all the coefficients are insignificant; *p*-values are calculated by regressing the response to the survey on indicator variables for two of the three treatments.

**Table A3.** Privacy Decisions before, during, and after the Cambridge Analytica Scandal

	Full Sample	Passed Comprehension Check	Answered Carefully	Excludes Those Who Did Not Click
Privacy Setting Hidden	-.26 (.06)	-.21 (.08)	-.19 (.07)	-.30 (.07)
During Cambridge Analytica Scandal	-.03 (.05)	.00 (.06)	-.00 (.06)	-.03 (.05)
During Cambridge Analytica Scandal $\times$ Veiled Trade-off Treatment	.20 (.08)	.14 (.11)	.15 (.10)	.21 (.10)
After Cambridge Analytica Scandal	-.04 (.05)	-.01 (.06)	-.00 (.06)	-.04 (.05)
After Cambridge Analytica Scandal $\times$ Veiled Trade-off Treatment	.09 (.08)	-.01 (.10)	.02 (.10)	.09 (.09)
Constant	.67 (.03)	.68 (.04)	.68 (.04)	.67 (.04)
N	755	689	619	734
Adjusted $R^2$	.03	.03	.02	.04

**Note.** Results are from a linear probability regression of a binary variable for whether the participant stayed anonymous on the binary variables indicated. Treatment-level block bootstrap standard errors are in parentheses.

## REFERENCES

- Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. 2015. Privacy and Human Behavior in the Age of Information. *Science*, January 30, pp. 509–14.
- Acquisti, Alessandro, Leslie K. John, and George Loewenstein. 2013. What Is Privacy Worth? *Journal of Legal Studies* 42:249–74.
- Agranov, Marina, and Pietro Ortoleva. 2017. Stochastic Choice and Preferences for Randomization. *Journal of Political Economy* 125:40–68.
- Andreoni, James, and Charles Sprenger. 2012. Estimating Time Preferences from Convex Budgets. *American Economic Review* 102:3333–56.
- Athey, Susan, Christian Catalini, and Catherine Tucker. 2017. The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. Working Paper No. 23488. National Bureau of Economic Research, Cambridge, MA.
- Bakos, Yannis, Florencia Marotta-Wurgler, and David R. Trossen. 2014. Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts. *Journal of Legal Studies* 43:1–35.
- Beauchamp, Jonathan P., Daniel J. Benjamin, Christopher F. Chabris, and David I. Laibson. 2015. Controlling for Compromise Effects Debases Estimates of Preference Parameters. Working Paper No. 21792. National Bureau of Economic Research, Cambridge, MA.
- Benndorf, Volker, and Hans-Theo Normann. 2014. The Willingness to Sell Personal Data. Discussion Paper No. 143. Heinrich Heine University, Faculty of Business Administration and Economics, Düsseldorf.
- Ben-Shahar, Omri, and Adam Chilton. 2016. Simplification of Privacy Disclosures: An Experimental Test. *Journal of Legal Studies* 45:S41–S67.
- Cadwalladr, Carole, and Emma Graham-Harrison. 2018. Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach. *Guardian*, March 17.
- Casey, Logan S., Jesse Chandler, Adam Seth Levine, Andrew Proctor, and Dara Z. Strolovitch. 2017. Intertemporal Differences among MTurk Workers: Time-Based Sample Variations and Implications for Online Data Collection. *Sage Open* 7:1–15.
- Dana, Jason, Roberto A. Weber, and Jason Xi Kuang. 2007. Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness. *Economic Theory* 33:67–80.
- DellaVigna, Stefano, John A. List, and Ulrike Malmendier. 2012. Testing for Altruism and Social Pressure in Charitable Giving. *Quarterly Journal of Economics* 127:1–56.
- Exley, Christine L. 2016. Excusing Selfishness in Charitable Giving: The Role of Risk. *Review of Economic Studies* 83:587–628.
- Federal Trade Commission. 2012. *Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers*. Washing-

- ton, DC: Federal Trade Commission.
- Funk, McKenzie. 2016. The Secret Agenda of a Facebook Quiz. *New York Times*, November 19, p. SR5.
- Golman, Russell, David Hagmann, and George Loewenstein. 2017. Information Avoidance. *Journal of Economic Literature* 55:96–135.
- Hara, Kotaro, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. 2018. A Data-Driven Analysis of Workers' Earnings on Amazon Mechanical Turk. In *CHI '18: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, paper 449, 14 pp. New York: Association for Computing Machinery.
- Hayashi, Takashi. 2008. Regret Aversion and Opportunity Dependence. *Journal of Economic Theory* 139:242–68.
- Hoffman, Elizabeth, David Schwartz, Matthew Spitzer, and Eric Talley. 2020. Patently Risky: Framing, Innovation, and Entrepreneurial Preferences. *Harvard Journal of Law and Technology* 34:191–257.
- Irvine, Krin, David A. Hoffman, and Tess Wilkinson-Ryan. 2018. Law and Psychology Grows Up, Goes Online, and Replicates. *Journal of Empirical Legal Studies* 15:320–55.
- Jacobson, Sarah, and Ragan Petrie. 2009. Learning from Mistakes: What Do Inconsistent Choices over Risk Tell Us? *Journal of Risk and Uncertainty* 38:143–58.
- John, Leslie K., Alessandro Acquisti, and George Loewenstein. 2011. Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information. *Journal of Consumer Research* 37:858–73.
- Kelley, P., L. Cesca, J. Bresee, and L. Cranor. 2010. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach. Pp. 1573–82 in *CHI '10: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York: Association for Computing Machinery.
- Kőszegi, Botond. 2003. Health Anxiety and Patient Behavior. *Journal of Health Economics* 22:1073–84.
- Lazear, Edward P., Ulrike Malmendier, and Roberto A. Weber. 2012. Sorting in Experiments with Application to Social Preferences. *American Economic Journal: Applied Economics* 4:136–63.
- Meier, Stephan, and Charles Sprenger. 2016. Present-Biased Preferences and Credit Card Borrowing. *American Economic Journal: Applied Economics* 2:193–210.
- Oster, Emily, Ira Shoulson, and E. Ray Dorsey. 2013. Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease. *American Economic Review* 103:804–30.
- Posner, Richard A. 1978. The Right of Privacy. *Georgia Law Review* 12:393–422.
- Prosser, William L. 1960. Privacy. *California Law Review* 48:383–423.
- Stigler, George J. 1961. The Economics of Information. *Journal of Political Economy* 69:213–25.

- Strahilevitz, Lior Jacob. 2010. Reunifying Privacy Law. *California Law Review* 98:2007–48.
- Sullivan, Patrick S., Amy Lansky, and Amy Drake. 2004. Failure to Return for HIV Test Results among Persons at High Risk for HIV Infection: Results from a Multistate Interview Project. *Journal of Acquired Immune Deficiency Syndromes* 35:511–18.
- Sunstein, Cass R. 2014. Choosing Not to Choose. *Duke Law Journal* 64:1–52.
- Turow, Joseph, Michael Hennessy, and Nora Draper. 2015. The Tradeoff Fallacy: How Marketers Are Misrepresenting American Consumers and Opening Them Up to Exploitation. Report. University of Pennsylvania, Annenberg School for Communication, Philadelphia.
- Warren, Samuel D., and Louis D. Brandeis. 1890. The Right to Privacy. *Harvard Law Review* 4:193–220.