ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS
FELLOWS' DISCUSSION PAPER SERIES

WHY ARE PRIVACY PREFERENCES INCONSISTENT?

Dan Svirsky

Discussion Paper No. 81
Revision

01/2019

Harvard Law School
Cambridge, MA 02138

Why Are Privacy Preferences Inconsistent?

Dan Svirsky[†]

January 22, 2019

Abstract: There is a widespread intuition that people are inconsistent about protecting their privacy. People are angry about corporations collecting their data but often do not change simple default settings in their apps. This paper presents evidence from an experiment on data sharing. The results both document this paradox and provide evidence for a novel explanation: information avoidance. Even 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 money bonuses if given a chance to avoid seeing the privacy consequences of their choices. In an online experiment, participants must decide whether to share their Facebook profile data with a survey-taker in exchange for a higher payoff. When people make a *direct* tradeoff between 50 cents and privacy, roughly 64% refuse to share their Facebook data. However, when participants face a *veiled* tradeoff and must "click to reveal" to learn whether privacy is free or costs 50 cents, only 40% end up refusing to share their data for 50 cents. Consistent with information avoidance driving these differences in apparent valuations of privacy, 58% of participants did not click to reveal to learn which payment option was associated with privacy. In a Placebo Veiled Tradeoff treatment, participants choosing between two money bonuses did not engage in the same avoidance behavior. The paper also presents evidence on how this pattern changed before, during, and after the Cambridge Analytica Facebook scandal. While privacy valuations in the direct tradeoff treatment were unchanged, even at the height of the scandal, the veiled information treatment became less effective. One month later, the results returned to their pre-scandal values, and the treatment was as effective as before. The findings show that even people who would otherwise pay for privacy seem able to exploit strategic ignorance – keeping their head in the sand – and deal away their data for small amounts of money. The findings also suggest that privacy regulations aimed at giving people more information about data choices will be difficult to execute effectively.

Preliminary Draft

Please do not circulate

[†]Harvard Economics Department, dsvirsky@hbs.edu. I thank Jennifer Arlen, Oren Bar-Gill, Omri Ben-Shahar, Christine Exley, Louis Kaplow, David Laibson, Florencia Marotta-Wurgler, Oliver Hart, Matthew Rabin, Josh Schwartzstein, Charlotte Svirsky, Tess Wilkinson-Ryan, and participants at the Harvard Law and Economics Workshop, the Harvard Economics Games and Markets Seminar, and the NYU Empirical Contracts Workshop for helpful comments.

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 corporations collecting their data. This intuition that people are being inconsistent is shared widely enough to have a name – the privacy paradox (Acquisti et al. (2015)). This phrase has been mentioned in roughly 5,000 scholarly articles between 2010 and 2018.

This paper uses an experiment that heightens this paradox and provides evidence for a novel explanation: information avoidance. Even 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 money bonuses if given a chance to avoid seeing the privacy consequences of their choices.

In the experiment, participants who complete a survey decide whether to do the survey anonymously or after logging in with their Facebook account in exchange for a money bonus. When participants in a *Direct Tradeoff* Treatment face a choice between a 50 cent bonus and privacy, 64% of participants refuse to share their Facebook profile in exchange for 50 cents. 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) are unwilling to share their Facebook data for $2.50, and a plurality refuse offers of $5.00.

However, when the privacy settings are veiled (but revealed costlessly and instantly with the click of a button, as in a moral wiggle room experiment (Dana et al. (2007)), many participants keep themselves in the dark and opt for more money. Participants in a *Veiled Tradeoff* Treatment face a choice between a 50 cent bonus and a 0 cent bonus. They know that one of these bonuses will mean giving out their Faceboook profile, and they can click a button to check which option involves a loss in privacy. I find that most people (58%) do not click, and only 40% end up keeping their Facebook profile private. Hence, people who are willing to pay nearly an hour's worth of wages to stay private are also able to throw caution to the wind, take a 50 cent bonus, and hope for the best.

Importantly, this same avoidance pattern does not hold when participants make a choice between two money bonuses, rather than money versus privacy. In a *Placebo Veiled Tradeoff* Treatment, participants face the exact same experimental interface as in the Veiled Tradeoff group, but where the second column contains a second money bonus. The size of the second money bonus is drawn from the distribution of willingness-to-pay prices from participants in the Elicitation Treatment. When facing this choice, participants in the Placebo Veiled Tradeoff clicked to reveal the second column 66% of the time, a rate significantly different from the reveal rate in the Veiled Tradeoff group.

This paper also presents data on changes in privacy preferences before, during, and after the Cambridge Analytica / Facebook scandal, a major scandal that made privacy issues more salient for many Facebook users. If people are not aware of privacy issues,

scandals like this one might disabuse them of this lack of awareness, helping to resolve the inconsistency in people's privacy behavior. By happenstance, an initial round of the experiment was run several weeks before the scandal became public. Once the scandal broke, privacy issues – and more specifically, privacy issues surrounding Facebook data and third party apps – dominated the news, appearing on the front page of the New York Times on most days for a month. The experiment was re-run 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 Tradeoff treatment, 64% of participants chose to keep their Facebook profile private instead of getting 50 cents, compared to 67% before the scandal (a slight but statistically insignificant drop). However, participants in the Veiled Tradeoff treatment were more likely to click to learn the privacy setting before making their choice, ultimately resulting in 58% opting for privacy (compared to 40% before the scandal). But this effect was short-lived. Forty days later, 46% opted for privacy over 50 cents – a proportion statistically indistinct from the pre-scandal level, and significantly lower than the peak-of-the-scandal level.

The results of the experiment make people's inconsistency over privacy choices more mysterious. Until now in the literature, the two dominant explanations of the privacy paradox were revealed preference and ignorance. That is, maybe people give away their data because they value the services they get in return. Alternatively, maybe they do not realize they are giving away their data. In contrast, this experiment finds that the inconsistency persists, even in a setting where participants *know* the exact privacy loss at stake, and where ignorance is unlikely to affect the Direct Tradeoff and Veiled Tradeoff groups differently. At the same time, the experiment also shows that people *are* willing to pay for privacy. Therefore, the inconsistency can not be written off as mere talk.

The results also cast doubt on current privacy law doctrine in the United States, which relies on giving consumers better notice before they make privacy decisions. Such a policy makes sense if people's privacy inconsistency is explained by revealed preference 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" vs "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 and scholarship in 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 United States privacy law (Strahilevitz (2010)). This Notice and Choice model was first outlined in a 1973 report by the U.S. Department of Health, Education and Welfare, and at the time, this legal framework was a departure from how privacy law originally developed. Before the rise in internet commerce and telecommunications, privacy was governed by tort law (Warren & 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 (e.g., a right to seclusion). As private data has become dominated by internet transactions, privacy law has been increasingly governed by contract law principles.[1]

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 "how much do people value privacy" has been challenging to answer because people's privacy decisions are fickle. Acquisti et al. (2013) offer 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, non-anonymous gift card. When endowed with the less private $12 card, 90% of participants chose to keep it rather than exchange it for the $10, anonymous card. John et al. (2011) find that people volunteer more sensitive information when asked indirectly, and also when a website seems *less* professional. Similarly, an experiment by legal scholars testing different disclosure techniques finds that people's privacy

[1] 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 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% of a firm's global revenue.

behavior is not much affected by providing them more and better information about their privacy choices (Ben-Shahar & Chilton (2016); *c.f.* Bakos et al. (2014)). Along the same lines, Athey et al. (2017) conduct a field experiment where MIT students are given Bitcoin and are 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 two simple explanations for this inconsistency: ignorance and revealed preference.

Under the ignorance explanation, people are unaware of how much data they are emitting, or they struggle to value privacy, because it is abstract or because privacy costs are inchoate and uncertain, both in scope and timing (Acquisti et al. (2013)). Either way, they do not fully understand what is at stake. As a result, when deciding whether to exchange privacy for something more easily quantifiable, like money or convenience, small frictions may play an outsized role in decision-making. 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.

Under the revealed preference explanation, people give up privacy simply because this maximizes their utility. People say 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.

For either explanation – revealed preference or ignorance – more information is better. If it's 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 Fickle privacy preferences can be explained by bounded rationality *or* by information avoidance

The privacy literature points to bounded rationality or revealed preference to explain the privacy paradox, but information avoidance can just as easily explain the same pattern. There is a robust literature from psychology and economics on information avoidance (Golman et al. (2017)). While economists typically model information as an intermediate good (Posner (1978), Stigler (1961)) – i.e., valuable only because it helps us achieve ends – scholars in psychology and economics increasingly recognize that people sometimes behave as if information has emotional valence (Oster et al. (2013)). More information is not always better.

Consider a now widely-replicated experiment on moral wiggle room (Dana et al.

(2007)), which is the basis for the experimental design used in this paper. In the experiment, a participant has to choose payoffs for herself ("me") and a partner that she does not meet ("my partner"). 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 preferred \$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 find out that I am in the baseline case, in which case I can choose \$5 and \$5 again, or I will find out that I am in the easier case and 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 \$6 for me, \$X for my partner option. They exploit the wiggle room to act selfishly. Other experiments on altruism, lab- and field-based, find similar results (Exley (2016), Lazear et al. (2012), DellaVigna et al. (2012)).

This pattern of behavior is important across disparate domains. In health, one study found that 27% of intravenous drug users at risk of HIV who got tested did not return to the clinic to see their results (Sullivan et al. (2004)), even though knowing one's HIV positive status can lengthen one's life. In family planning, twenty states have laws requiring women to see a picture of the fetus before getting an abortion (Guttmacher Institute (2018)). Presumably, women know what a fetus looks like, so the law was not passed because the increased information of 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 Tradeoff Treatment and a Veiled Tradeoff Treatment. This section first discusses the overall timeline of the experiment, then describes the two treatments in detail.[2] The experiment was preregistered on AsPredicted under the title "Information Avoidance and Internet Privacy" (#16702).

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

795 participants were recruited on Amazon Mechanical Turk to take a short survey about health and financial status.[3] All participants were informed that before doing the survey, they would make decisions about the size of a bonus payment, to be received upon completion, and the privacy settings of the survey.[4] The experiment was conducted on November 20, 2018.

After recruitment, the timeline of the experiment consists of three stages: instructions and practice, privacy settings, and a survey.[5] First, participants were shown an initial introductory screen that gives an overview of their participation. Participants were told that they would take a survey, but while everyone would take the same exact survey, each participant would be given a choice between two privacy options. They could opt for high privacy, in which case their survey answers would be anonymous. Or, they could opt instead 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-taker would see, in addition to the participant's survey answers, her public Facebook profile (including profile picture, name, and gender) and her email address. Participants who chose low privacy would not be allowed to finish the survey until they logged in. Participants then completed two short practice rounds which looked identical to the privacy settings task.

After the instructions stage, participants chose their privacy settings. After completing the privacy settings stage, participants completed the survey stage.

The privacy measure in the experiment – whether to share Facebook information – has three advantages: it is a real decision, it is a realistic one, and it is an important one. First, participants who give up their privacy in this experiment must actually give over their profile data, so the choice is not a hypothetical one. Nor is it a behavior that can be faked: unlike other privacy experiments, which measure privacy as a person's willingness to answer an intrusive question, a participant in this experiment cannot pretend to give up privacy without actually giving anything up.[6] Second, the decision is a realistic one. The "Log In with Facebook" button is a ubiquitous part of the internet – many websites allow people to log in with their Facebook (or Google) account rather than with the website itself. Hence, it is a choice people routinely make: should I engage in online activity in a way that is linked to my Facebook profile or not? Third, the decision has important public policy implications, as suggested by the

[3]Research increasingly suggests that, for the purpose of social science experiments, Mechanical Turk users are a reliable sample. Irvine et al. (2018) replicates three experiments using in-person labs, national online platforms, and Mechanical Turk, and finds that the results are constant across samples. The key difference was that that Mechanical Turk users were significantly more attentive than the other samples. See also Hoffman et al. (2017) (replicating an experiment on Mechanical Turk, on college students in a physical lab, and college students in an online setting).

[4]Median hourly wages for workers was $14.96 (based on a median payment of $1.52 for a median completion of 8 minutes 6 seconds).

[5]Appendix A presents the entire experimental instructions.

[6]Even if participants have a fake account they can use – Facebook works hard to limit such behavior, but is not 100% successful – even handing over a fake account involves some cost. Doing so means the experimenter can link a fake Facebook account to a Mechanical Turk account (and the answers in the survey), which makes the fake account less effective.

Cambridge Analytica scandal.

Each person was randomized into one of two treatments during the privacy settings stage: the Direct Tradeoff Treatment and the Veiled Tradeoff Treatment. Figure 2 shows the exact format of the privacy choice made in each of the treatments.

In the Direct Tradeoff Treatment, participants only made one decision: a direct choice between a \$0.02 bonus and Privacy Option A or a \$0.52 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 tradeoff between money and privacy.

In the Veiled Tradeoff Treatment, participants faced the same decision as in the Direct Tradeoff Treatment, but the privacy setting was initially hidden. Participants had to click to reveal the column describing the privacy settings, and there was a 50% chance that the higher money bonus would mean losing their anonymity.[7]

After completing the privacy task, all participants completed a nine-question survey, shown in Figure 3. Five questions covered demographics, health, and financial topics. These questions asked about the person's age, the number of times they exercise in a week, the number of times they have attempted to diet in their life, their annual income, and their credit card debt. The survey also included two questions to check comprehension. One asked "How old were you when you were 10 years old" with a dropdown menu with several options, including 10. Another directly asked "How carefully did you make your choices?" with three options: "Not carefully at all", "I thought about it a little", and "I was very careful". Two questions asked whether participants had a Facebook profile and how often they used Facebook. After submitting the survey, participants were finished.

The user interface for the experiment was coded using HTML and Javascript, which ensured that the "reveal button" would work instantaneously – without a page refresh. When a user clicked the reveal button, Javascript code changed the visibility setting of the hidden column from hidden to visible. The hidden column would therefore become visible immediately. The users' choices and data were sent to a MySQL database using PHP code. All code is available on request from the author and includes survey instructions, experimental module coding, and the raw data.[8]

[7] Note that for both groups, there was a 50% chance of facing a degenerate choice between { more money, more privacy } and { less money, less privacy }. These decisions cannot tell us about how much a person values privacy, so they are omitted from the main analyses below. The resulting sample size is 535 participants: 117 in the Direct Tradeoff, 130 in the Veiled Tradeoff, 164 in the Placebo Veiled Tradeoff (described below), and 124 in the Elicitation Treatment (described below).

[8] Contact the author for the ZIP file: dsvirsky@hbs.edu .

3.1 Placebo Test

Any difference between the Direct Tradeoff and Veiled Tradeoff groups might be driven by clicking costs, rather than information avoidance. Suppose many people value privacy at 51 cents, but the "click to reveal" button imposes a few cents of effort costs. Then we would observe a treatment effect, but because participants rationally conclude that it's not worth spending a few cents of effort for a 1 cent gain.

I test this alternative explanation in two ways. First, I use an Elicitation Treatment to gather the full distribution of willingness-to-pay (WTP) prices for privacy. In the Elicitation Treatment, instead of making just one choice between privacy and 50 cents, participants made 10 choices, with the money bonus varying between 25 cents 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 as in Figure 2 in which they chose between two rows of a table. The top row meant a $0.02 bonus and "High Privacy", and the bottom row meant a $X.YY bonus and "Low Privacy", with $X.YY ranging from $0.27 to $5.02. Hence, if someone opted to stay anonymous when offered $0.50, $1.00, and $1.50, but not at $2.00, then we can infer that her WTP for staying anonymous is between $1.50 and $2.00.

Second, I conduct a Placebo Veiled Tradeoff treatment. This treatment is identical to the Veiled Tradeoff treatment, but instead of making a choice between one money bonus and privacy, participants make a choice between one money bonus and a second money bonus. The first money bonus is 50 cents, as in the main experiment, but the second money 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 we would observe the same size treatment effect if the second column is instead a money bonus drawn from the same distribution of WTP prices.

3.2 Privacy Preferences Before, During, and After the Cambridge Analytica / Facebook 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 in order to help conservative political candidates. Most of the data was obtained without consent, and the report quickly escalated into a public scandal. Cambridge Analytica largely relied on Mechanical Turk to construct its illicit dataset. Mechanical Turk users were invited to share their Facebook data in exchange for monetary 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 specific nature of the scandal could not have been better-suited to the dependent

variable for privacy used in this experiment. Specifically, the scandal dealt with people's willingness to share their Facebook data as part of an unrelated survey, which is precisely the dependent variable measured in this paper. Further, Cambridge Analytica targeted Mechanical Turk users, so the experiment in this paper was run on the same sample of people targeted in the scandal – though most likely not the exact same people, given natural turnover rates in Mechanical Turk's worker base.

The experiment was run three times, and the timing was chosen to measure whether privacy preferences changed during and after the scandal. The pilot round of the experiment was initially run on February 23, 2018 – 23 days before the scandal broke. A second round was conducted 11 days after the scandal became public. A third round was conducted 41 days later.

Figure 4 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.

The main results presented in this paper are from an experimental round run on November 20, 2018, whereas the Facebook results are from three earlier rounds of the experiment. These initial three rounds did not include an Elicitation Treatment, only a Direct Tradeoff and Veiled Tradeoff treatment. Importantly, the three initial rounds were all identical to each other, which ensures that comparisons across these three rounds are valid.

4 Results

Table 1 presents summary statistics on the survey answers, as well as a balance check. Nearly all – 94.6% – 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 4 or more times per week.

The analyses below are restricted to participants who answered 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 Tradeoff and Veiled Tradeoff treatments, attrition (defined as people who read the instructions but quit before the survey round) was 5% and 2% respectively.

4.1 Results: Direct Tradeoff Treatment vs Veiled Tradeoff Treatment

I find a treatment effect from putting a costless veil on privacy settings. 64% of people in the Direct Tradeoff Treatment refuse to sell their Facebook data for 50 cents.[9] In contrast, in the Veiled Tradeoff Treatment, when the privacy consequences of their actions are initially hidden, only 40% refuse to sell their Facebook data for 50 cents. A majority in the Veiled Tradeoff Treatment (58%) chose *not* to look at the privacy setting before deciding to take the 50 cents.

Figure 5 shows the proportion of participants who remained anonymous in the Direct Tradeoff Treatment and the Veiled Tradeoff Treatment. Figure 6 breaks down participants' decisions in both treatments, including their privacy choice as well as their decision whether to click. Table 2 reports various regressions where the unit of observation is an individual, the dependent variable is whether the participant ended up remaining anonymous, and the independent variable is an indicator variable for being in the Veiled Tradeoff Treatment.

The treatment effect is robust even if we 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 dropdown menu. Roughly 84% of participants correctly answered. Another question in the survey asked "How carefully did you make your choices?", with three options: not carefully, a little carefully, and very carefully. Roughly 75% of participants said they answered the questions "very carefully", 21% said "a little carefully" and 3% said not at all carefully. Note that by default, "Not Carefully At All" was selected. The main results are substantively unchanged if we 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 Tradeoff Treatment mistakenly assume that a low bonus meant they would keep their privacy? That is, participants in the Veiled Tradeoff 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 monetary payoff always meant higher privacy. In that case, we 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% of participants in the Veiled Tradeoff treatment. The results discussed here categorize these participants as having chosen privacy over 50 cents, but the results do not change if these participants are instead dropped.

Table 2 reports the results of these robustness checks. Columns 2 - 5 report the result of the main regression, described above, but using different samples. Column 2 includes controls for survey answers, while columns 3 - 5 exclude participants based

[9] This is in line with the results from the Elicitation Treatment group, described below. In that group, 59% rejected an offer of 50 cents to share their Facebook profile, a slightly lower but statistically insignificant difference.

on comprehension, attention, and confusion (defined as opting for less money without clicking to reveal the privacy setting). The main results hold throughout.

4.2 Results: Placebo Veiled Tradeoff and Elicitation Treatment

The results from the Elicitation Treatment and Placebo Veiled Tradeoff 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 cents, so when faced with the "click to 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 is out is by directly eliciting people's WTP price for sharing their data, and in doing so, I find that the majority of people value privacy at $2.50 or above. Table 3 shows people's WTP for staying anonymous in the Elicitation Treatment. Each row shows the proportion of participants who switched from High Privacy to Low Privacy at the prices offered. The results show that a plurality of participants – 41.5% – refuse to share their Facebook profile at all prices, even 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 profile with a survey-taker. Nonetheless, the second most-common WTP price was at the lower end, with 20.8% choosing to sell their Facebook profile at 25 cents. The remaining 38% evinced a WTP between 25 cents and $5.00.[10]

Can the main experimental findings be explained by simple clicking costs? For example, if a user has a WTP for privacy of 50.1 cents, then it might not make sense to take a few seconds to reveal the privacy settings, even if she would have opted for privacy in the direct tradeoff treatment. Using the results of the Elicitation Treatment, I can say with more precision how high clicking costs would have to be to support such an explanation. Appendix B presents a more detailed mathematical approach to this question. It demonstrates three key points. First, anyone who values privacy at less than 50 cents should *never* click to reveal the privacy settings. Second, if clicking costs are zero, anyone who values privacy at more than 50 cents should *always* click to

[10]Irrational behavior, defined as having multiple switching points, was rare. It is hard to interpret someone giving up her privacy for 50 cents ($Privacy \succ \$0.50$) but not for $1.00 ($\$0.50 \succ Privacy \succ \$1.00$), assuming that she also values more money over less money. In the Elicitation Treatment, 84% gave rational answers in the sense of having at most one switching point. This is a relatively low level of multiple switch behavior compared to other experiments that use multiple price lists, which typically find levels of multiple switch behavior ranging from 10% to as high as 50% (Andreoni & Sprenger (2012), Meier et al. (2016)). This finding also suggests that Mechanical Turk workers evinced similar levels of this type of irrationality when compared to college students and people with moderate incomes in tax filing centers, among other samples. In calculating the distribution of WTP prices, I exclude participants with multiple switches, but the results are similar if I instead include them and define their switching point as either the lowest switch, the highest switch, or the average of the two.

reveal the privacy settings. Third, and most relevant here, 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 their mouse 31 during the experiment times 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 doing this experiment and finishing it.

The results of the Placebo Veiled Tradeoff give more direct evidence that the results are not driven by clicking costs. Recall that in the Placebo Veiled Tradeoff, participants chose between two money bonuses, with the value of the second money 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 which row the bonus was in. Among this group, the proportion of participants clicking to reveal the second column was 0.66. This is higher than the click rate of 0.42 in the main experiment, when participants chose between money and privacy, and the difference is statistically significant (Fisher's exact $p < 0.001$). Table 4 shows the click proportion in the Placebo Veiled Tradeoff group, broken down by the size of the second money 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.

In sum, the results of the Elicitation Treatment and Placebo Veiled Tradeoff Treatment suggest that the findings of the main experiment are not driven by clicking costs or confusion about the experimental design.

4.3 Privacy Preferences Before, During, and After the Facebook Cambridge Analytica Scandal

Roughly a month after a pilot round of the experiment was run, there was a controversial privacy scandal that directly involved people's willingness to share their Facebook data with third parties. A second and third round of the experiment were therefore run, one in the immediate aftermath of the scandal, and another roughly one month after.

There is no evidence that the sample of participants was observably different across time. Importantly, any changes we see 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 profile private from a third party but that were unrelated to the Facebook scandal. To get a sense of these issues, Table 5 presents a balance check to see whether the three samples of participants are significantly different in any of the survey responses. I find balance across all three groups, suggesting that in terms of reported age, credit card debt, income, and exercise patterns, the sample did not measurably change before, during, and after the scandal.

At the height of the Facebook / Cambridge Analytica scandal, people's behavior in the Direct Tradeoff Treatment was unchanged. Before the scandal, 66% opted for privacy over 50 cents in the Direct Tradeoff Treatment. At the height of the scandal, this number was 64%, and one month later, the proportion was 63%. None of these changes were statistically significant.

However, the treatment became less effective. Before the scandal, the Veiled Tradeoff Treatment caused a 26 percentage point drop ($p < 0.001$) in the proportion of people opting to keep their Facebook profile private. At the height of the scandal, the Veiled Tradeoff Treatment caused a 9 percentage point drop ($p = 0.06$). One month after the scandal, the treatment was effective again, causing a 17 percentage point drop ($p = 0.003$). The treatment effect at the height of the scandal was significantly different from the treatment effects before (p = 0.01) and after the scandal (p = 0.03).

Figure 7 shows the proportion of people who chose to keep their Facebook data private during the survey instead of getting a fifty cent bonus, by treatment and across the three experiment dates.

Table 6 presents regression results and robustness checks. The regression specification is as follows, letting p be an indicator variable for whether an individual ended up remaining anonymous, T be an indicator variable for whether the participant was in the Veiled Tradeoff Treatment, FB be an indicator for whether the experiment date occurred shortly after the Facebook scandal, and $Post$ be an indicator for whether the experiment occurred forty days after the scandal.

$$\text{p} = \beta_0 + \beta_1 \cdot (T) + \beta_2 \cdot (FB) + \beta_3 \cdot (FB * T) + \beta_4 \cdot (Post) + \beta_5 \cdot (Post * T)$$

In the regression, β_1 measures the treatment effect before the scandal, β_2 measures the change in privacy preferences in the Direct Tradeoff Treatment group at the height of the Facebook scandal, β_3 measures the change in the treatment effect at the height of the Facebook scandal, β_4 measures the change in privacy preferences in the Direct Tradeoff Treatment group after the scandal, and β_5 measures the change in the treatment effect after the scandal. Column 1 includes the entire sample. Column 2 excludes participants who failed the comprehension check. Column 3 excludes participants who reported not answering carefully. Column 4 excludes participants who did not click to reveal the privacy setting but chose the lower money option.

In sum, I find no measurable change in survey responses before, during, and after the Facebook scandal, nor do I find any chance in behavior in the Direct Tradeoff treatment. I *do*, however, observe that the experimental treatment became significantly less effective, and this was driven by people in the Veiled Tradeoff 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 limited 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 hit, 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 pre-scandal levels within two months of the scandal breaking.

The results also suggest two directions for future scholarship on internet privacy. First, more research is needed to understand *why* people avoid information about privacy. Second, given that information avoidance can explain privacy inconsistency, more thought should be given to existing policy interventions in internet privacy.

An important 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 tradeoff group and the veiled tradeoff group, because members of the veiled tradeoff group can take the monetary bonus without explicitly choosing to give away their data. In this view, people who tap "No" when a browser asks them to share their location may be evincing a sort of phatic preference: their action helps express righteous anger as much as underlying preferences in a small decision where the stakes are low. A second mechanism is that thinking about a probabilistic chance of losing privacy is itself upsetting, as in the model of anxiety in Koszegi (2003). A third mechanism is that people do care about privacy, but are also able to turn off their minds to privacy losses that are not directly in front of their face. Another mechanism is choosing costs (Sunstein (2014)). It takes effort to make a decision between money and privacy, especially if privacy costs are inchoate or hard to measure. Perhaps the direct tradeoff group has no choice but to make this effort, but the veiled tradeoff group might rationally decide that it is better to avoid doing the calculations, exploiting the veil as a cognitive shortcut. Still another explanation of the results is that all these mechanisms are true, to greater or lesser degress 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.

More broadly, the experiment suggests reason 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)). Specifically, 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. The results presented here show that even when the privacy settings could be revealed instantly, and even when the settings were 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 and clear and easily accesible, people may have struthious preferences.

References

Acquisti, A., Brandimarte, L. & Loewenstein, G. (2015), 'Age of Information,', *Science* **347**(6221), 509–515.

Acquisti, A., John, L. K. & Loewenstein, G. (2013), 'What Is Privacy Worth?', *The Journal of Legal Studies* **42**(2), 249–274.

Andreoni, J. & Sprenger, C. (2012), 'Estimating Time Prefernces from Convex Budgets', *American Economic Review* **102**(7), 3333–3356.

Athey, S., Catalini, C. & Tucker, C. (2017), The Digital Privacy Paradox: Small Money, Small Costs, Small Talk.

Bakos, Y., Marotta-Wurgler, F. & Trossen, D. R. (2014), 'Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts', *The Journal of Legal Studies* **43**(1), 1–35.

Ben-Shahar, O. & Chilton, A. (2016), 'Simplification of Privacy Disclosures: An Experimental Test', *The Journal of Legal Studies* **45**(S2), S41–S67.

Dana, J., Weber, R. A. & Kuang, J. X. (2007), 'Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness', *Economic Theory* **33**(1), 67–80.

DellaVigna, S., List, J. A. & Malmendier, U. (2012), 'Testing for Altruism and Social Pressure in Charitable Giving', *Quarterly Journal of Economics* **127**(1), 1–56.

Exley, C. L. (2016), 'Excusing selfishness in charitable giving: The role of risk', *Review of Economic Studies* **83**(2), 587–628.

Federal Trade Commission, . (2012), Protecting Consumer in an Era of Rapid Change: recommendations for businesses and policymakers, Technical Report March, Federal Trade Commission.

Golman, R., Hagmann, D. & Loewenstein, G. (2017), 'Information Avoidance', *Journal of Economic Literature2* **55**(1), 96–135.

Guttmacher Institute, . (2018), Requirements for Ultrasound, Technical report.
URL: *https://www.guttmacher.org/state-policy/explore/requirements-ultrasound*

Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C. & Bigham, J. P. (2018), 'A Data-Driven Analysis of Workers ' Earnings on Amazon Mechanical Turk', *Conference on Human Factors in Computing Systems (CHI 2018)* pp. 1–14.

Hoffman, E., Schwartz, D., Spitzer, M. & Talley, E. (2017), Patently Risky : Framing , Innovation and Entrepreneurial Preferences.

Irvine, K., Hoffman, D. A. & Wilkinson-Ryan, T. (2018), 'Law and Psychology Grows Up, Goes Online, and Replicates', *Journal of Empirical Legal Studies* **15**(2), 320–355.

John, L. K., Acquisti, A. & Loewenstein, G. (2011), 'Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information', *Journal of Consumer Research* **37**(5), 858–873.

Kelley, P., Cesca, L., Bresee, J. & Cranor, L. (2010), 'Standardizing Privacy Notices : An Online Study of the Nutrition Label Approach', *Proc. 28th International Conference on Human Factors in Computing Systems* pp. 1573 – 1582.

Koszegi, B. (2003), 'Health anxiety and patient behavior', *Journal of Health Economics* **22**(6), 1073–1084.

Lazear, E. P., Malmendier, U. & Weber, R. A. (2012), 'Sorting in Experiments with Application to Social preferences', *American Economic Journal: Applied Economics* **4**(1), 136–163.

Meier, S., Sprenger, C., Meier, B. S. & Sprenger, C. (2016), 'Present-Biased Preferences and Credit Card Borrowing * revolving with debt at one credit card report', **2**(1), 193–210.

Oster, E., Shoulson, I. & Dorsey, E. R. (2013), 'Optimal expectations and limited medical testing: Evidence from huntington disease', *American Economic Review* **103**(2), 804–830.

Posner, R. A. (1978), 'The right of privacy', *Georgia Law Review* **12**(3), 393.

Prosser, W. L. (1960), 'Privacy', *California Law Review* **48**(3).

Stigler, G. J. (1961), 'The Economics of Information', *Journal of Political Economy* **69**(3), 213–225.

Strahilevitz, L. J. (2010), 'Reunifying Privacy Law', *California Law Review* **98**(6), 2007–2048.

Sullivan, P. S., Lansky, A. & Drake, A. (2004), 'Failure to Return for HIV Test Results Among Persons at High Risk for HIV Infection', *JAIDS Journal of Acquired Immune Deficiency Syndromes* **35**(5), 511–518.

Sunstein, C. R. (2014), 'Choosing Not to Choose', *SSRN Electronic Journal* **5**(1999), 1–52.
URL: *http://www.ssrn.com/abstract=2377364*

Warren, S. D. & Brandeis, L. D. (1890), 'The Right to Privacy', *Harvard Law Review* **4**(5), 193–220.

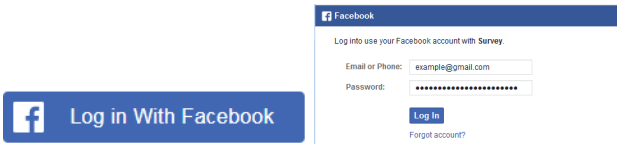

Figure 1: In a low privacy setting, participants logged into their Facebook account before completing the survey by clicking the login button shown above. Once they clicked on the login button, the window on the right appeared, allowing them to login.

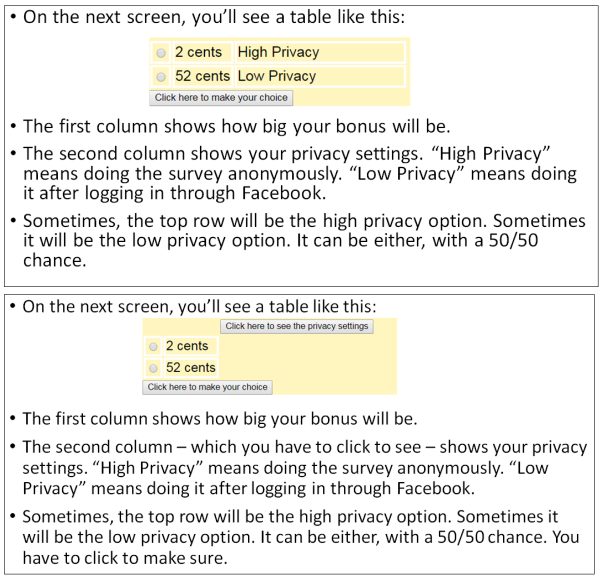

Figure 2: This figure shows the instructions page for each of the two treatments. The Direct Tradeoff Treatment group was shown the instructions in the top panel. The Veiled Tradeoff Treatment group was shown the instructions in the bottom panel.

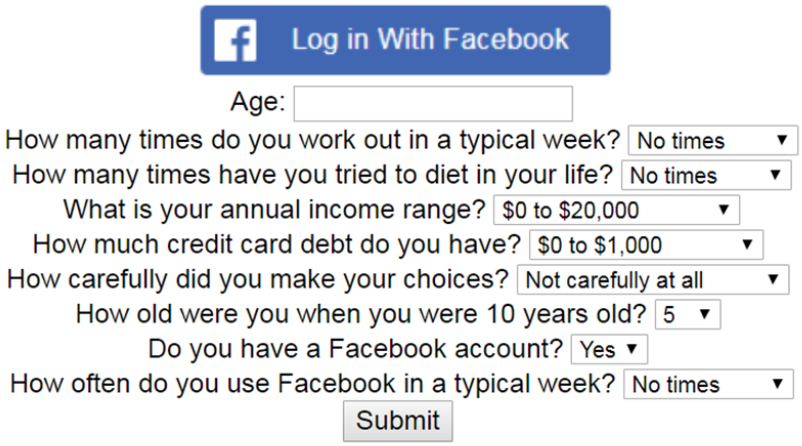

Figure 3: After making their privacy choices, all participants completed the survey above. Those participants who opted for the anonymous survey were not shown the Facebook login button. Those that opted for the low privacy setting saw the login button, as in the picture above.

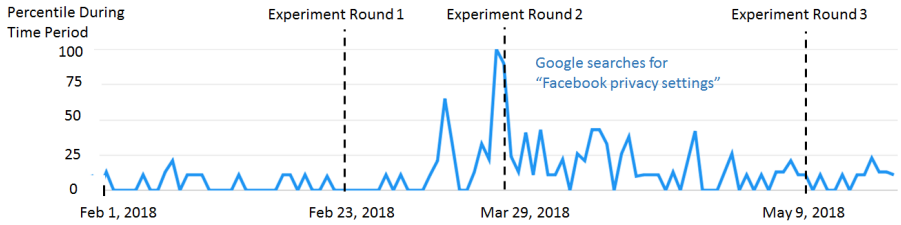

Figure 4: This figure shows the relative volume of Google searches for the phrase "Facebook privacy settings" over time, as well as the timing of the initial three rounds of the experiment. These three rounds, all identical in design, are used to measure changes in privacy valuations and information avoidance behavior during the Facebook Cambridge Analytica Scandal.

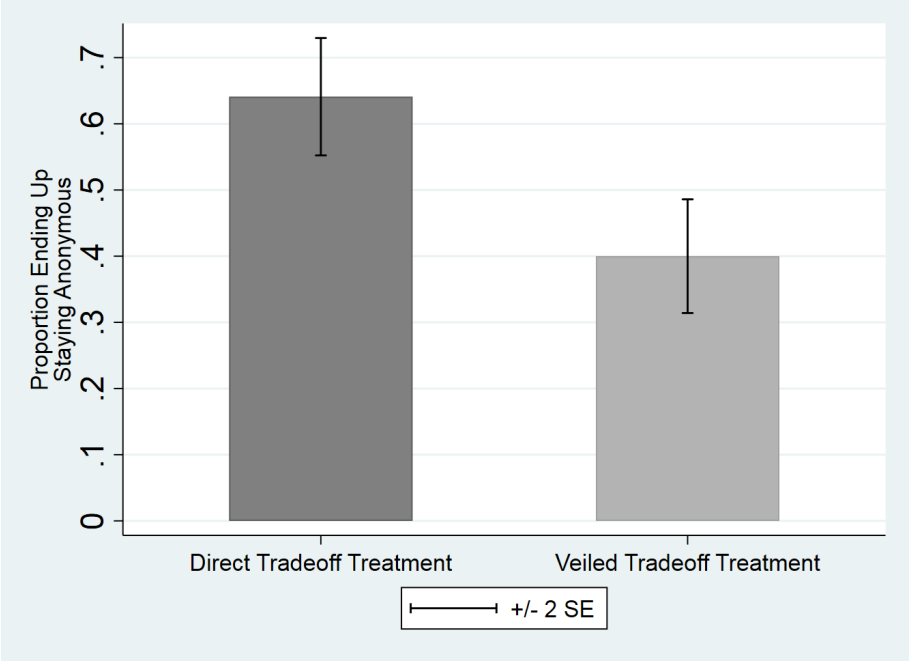

Figure 5: This figure shows the proportion of participants who ended up remaining anonymous instead of sharing their Facebook profile for 50 cents, for the Direct Tradeoff Treatment (N = 117) and the Veiled Tradeoff Treatment (N = 130). These results exclude all participants who, by randomization, faced a degenerate tradeoff of 50 cents and high privacy vs 0 cents and low privacy. Therefore, for the Veiled Tradeoff Treatment, anyone who chose the higher money option is counted as having chosen 50 cents over anonymity, regardless of whether they clicked to reveal the privacy setting before making their decision.

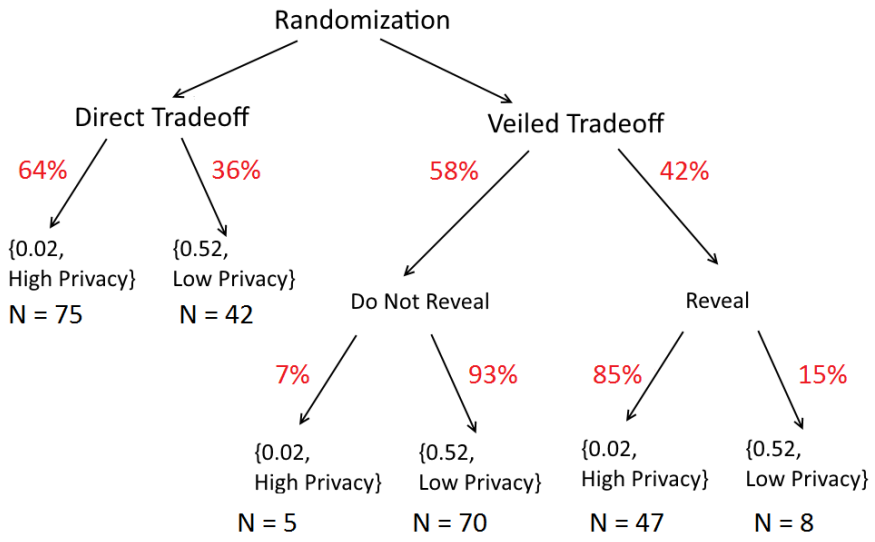

Figure 6: This figure shows the decisions made by participants across the Direct Tradeoff Treatment and Veiled Tradeoff Treatment. This figure excludes participants who were randomized into a degenerate choice between more money *and* high privacy vs less money and low privacy. In the Direct Tradeoff Treatment, participants made a choice between { \$0.02, High Privacy } versus { \$0.52, Low Privacy }. In the Veiled Tradeoff Treatment, participants first decide whether to reveal or not to reveal. If they do not reveal, then they choose between { \$0.02, Privacy Option A } and { \$0.52, Privacy Option B }. If they do reveal, then they face the same choice as in the Direct Tradeoff Treatment. Because I exclude all participants who face a degenerate choice, the lower monetary bonus always corresponds to high privacy, though participants in the Veiled Tradeoff Treatment who do not click to reveal cannot be certain of this, and only know that there is a 50% chance that low money corresponds to high privacy and a 50% chance that low money corresponds to low privacy.

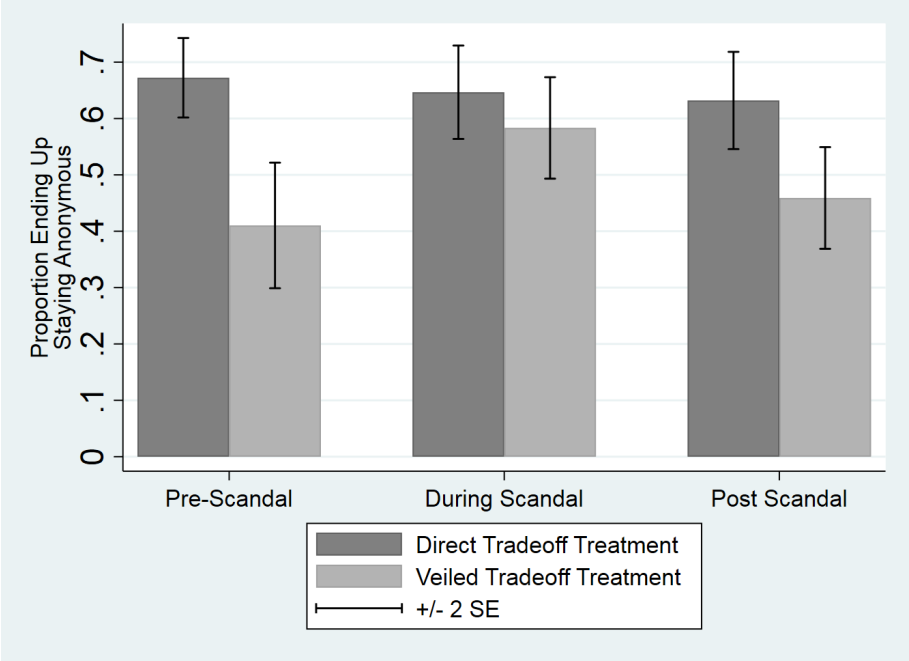

Figure 7: This figure shows the proportion of people in the Direct Tradeoff and Veiled Tradeoff treatments who ended up staying anonymous instead of getting a 50 cent bonus, across experiment dates. The Facebook scandal became public on March 18, 2018. The first round of the experiment occurred on February 23. The second round occurred on March 29. The third round occurred on May 9. Error bars are +/- two standard errors.

Table 1: Summary statistics

	Direct Tradeoff	Veiled Tradeoff	P-Value
Age (Years)	31.0	32.5	0.13
	(7.3)	(8.6)	
Diet Attempts in Lifetime (0 - 4+)	2.3	2.3	0.93
	(1.6)	(1.6)	
Exercise Workouts in a Typical Week (0 - 4+)	2.4	2.3	0.49
	(1.4)	(1.4)	
Annual Income (0 - 4)	1.3	1.1	0.14
	(1.1)	(1.2)	
Credit Card Debt (0 - 4)	0.73	0.80	0.66
	(1.1)	(1.1)	
Has Facebook (0,1)	0.92	0.95	0.21
	(0.28)	(0.21)	
Weekly Facebook Use (0 - 4+)	2.8	3.0	0.31
	(1.5)	(1.4)	

Summary statistics for the Direct Tradeoff and Veiled Tradeoff groups. Standard deviation reported in parenthesis. Each statistic is taken from the participants' survey answers. Income and credit card debt variables are categorical. Each category from 0 to 4 represents a different income or debt range. Diet attempts, Exercise and Facebook use can be 0, 1, 2, 3 or "4 or more." Reported p values taken from t-tests comparingthe means of the two groups.

Table 2: Privacy Decisions In Direct Tradeoff Group and Veiled Tradeoff Group, with Robustness Checks

Sample	(1) Full Sample	(2) Full Sample	(3) Passed Comprehension Check	(4) Answered Carefully	(5) Intersection of (3) and (4)
Veiled Tradeoff Group	-0.24***	-0.23***	-0.21**	-0.22***	-0.19**
	(0.06)	(0.07)	(0.07)	(0.06)	(0.07)
Age		0.01**			
		(0.004)			
Diet Attempts in Lifetime		-0.01			
		(0.02)			
Exercise in a Typical Week		-0.04			
		(0.02)			
Annual Income		0.02			
		(0.03)			
Credit Card Debt		0.03			
		(0.03)			
Constant	0.64***	0.34*	0.63***	0.63***	0.62***
	(0.05)	(0.17)	(0.05)	(0.05)	(0.05)
Observations	247	213	207	239	201
Adjusted R^2	0.05	0.07	0.04	0.04	0.03

Table reports a linear probability regression of a binary variable for whether the participant ended up staying anonymous on a binary variable for whether the participant was in the Veiled Tradeoff Treatment. Omitted group is the Direct Tradeoff Treatment. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Distribution of Willingness-to-Pay Prices for Remaining Anonymous, in Elicitation Treatment

WTP	Column %	Cumulative %
$0.25	20.8	20.8
$0.50	20.8	41.5
$1.00	1.9	43.4
$2.00	1.9	45.3
$3.00	3.8	49.1
$4.00	9.4	58.5
$5.00	41.5	100.0
Total	100.0	

N = 106

This table presents the breakdown of Willingness-to-Pay ("WTP") prices in the Elicitation Treatment. in the Elicitation Treatment. Participants faced 10 binary decisions where they could sell their Facebook data for $X.XX, with $X.XX ranging from $0.25 to $5.00. Participants were informed that one of their 10 decisions would be randomly selected and enforced. From these decisions, a WTP price is calculated by finding the switching point at which a person is willing to begin selling her data. People with multiple switching points are omitted (18 out of 124 participants had multiple switching points). People who refused to sell data at all prices are categorized as having a WTP of $5.00.

Table 4: Click to Reveal Behavior in Placebo Veiled Tradeoff Treatment

Size of Second Money Bonus	Did Not Click to Reveal	Did Click to Reveal
$0.25	39	29
$0.50	2	1
$1.00	1	2
$2.00	2	8
$3.00	4	8
$4.00	0	9
$5.00	8	51

N = 164

This table shows the number of participants who "clicked to reveal" the second money bonus in the Placebo Veiled Tradeoff Treatment, by the size of the second bonus.

Table 5: Summary Statistics: Before, During, and After Facebook Scandal

	Before Scandal	During Scandal	After Scandal	P-Value
Age	33.54	32.40	33.06	0.12
	(10.67)	(9.711)	(8.356)	
Diet Attempts in Lifetime	2.32	2.26	2.31	0.85
	(1.58)	(1.49)	(1.56)	
Exercise Workouts in a Typical Week	2.48	2.36	2.40	0.55
	(1.33)	(1.34)	(1.44)	
Annual Income	1.33	1.33	1.23	0.20
	(1.11)	(1.15)	(1.12)	
Credit Card Debt	0.64	0.71	0.64	0.98
	(0.99)	(0.95)	(1.01)	

Summary statistics for all participants, broken down by whether the sample was from before, during, or after the Facebook scandal. To calculate p-value for a row, the variable for the survey response was regressed on indicator variables for two of the three treatments. The p-value reported is the p-value for the F-test, or the joint hypothesis that all the coefficients are insignificant.

Table 6: Privacy Decisions Before, During, and After Facebook Scandal

	(1)	(2)	(3)	(4)
Sample	Full Sample	Passed Comprehension Check	Answered Carefully	Excludes 'Didn't Click, Chose 0 cents'
Privacy Setting Hidden	-0.26***	-0.21**	-0.19**	-0.30***
	(0.06)	(0.08)	(0.07)	(0.07)
During Facebook Scandal	-0.03	0.00	-0.00	-0.03
	(0.05)	(0.06)	(0.06)	(0.05)
During Facebook Scandal	0.20*	0.14	0.15	0.21*
* Veiled Tradeoff Treatment	(0.08)	(0.11)	(0.10)	(0.10)
Post-Facebook Scandal	-0.04	-0.01	-0.00	-0.04
	(0.05)	(0.06)	(0.06)	(0.05)
Post-Facebook Scandal	0.09	-0.01	0.02	0.09
* Veiled Tradeoff Treatment	(0.08)	(0.10)	(0.10)	(0.09)
Constant	0.67***	0.68***	0.68***	0.67***
	(0.03)	(0.04)	(0.04)	(0.04)
Observations	755	689	619	734
Adjusted R^2	0.03	0.03	0.02	0.04

Table reports a linear probability regression of a binary variable for whether the participant ended up staying anonymous on binary variables for whether the participant was in the Veiled Tradeoff Treatment, whether the experiment occurred at the height of the Facebook scandal, whether the experiment occurred one month after the Facebook scandal, and interactions between the treatment indicator and date indicators. Each column represents the same regression but with different samples for robustness checks. Block bootstrap standard errors bootstrapped at the treatment level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

6 Appendix A

This section shows the instructions shown to each participant.

All participants start off by seeing the same introductory slide, shown in Figure A1.

Each group is then shown one of three possible instructions slides, shown in Figure A2. Participants then perform the task itself. Finally, after completing the privacy task associated with their treatment, all participants completed a short survey, shown in Figure A5. Participants who agreed to give up their Facebook data would see the "Log In with Facebook" button above the survey; participants who opted to remain anonymous would not see the button.

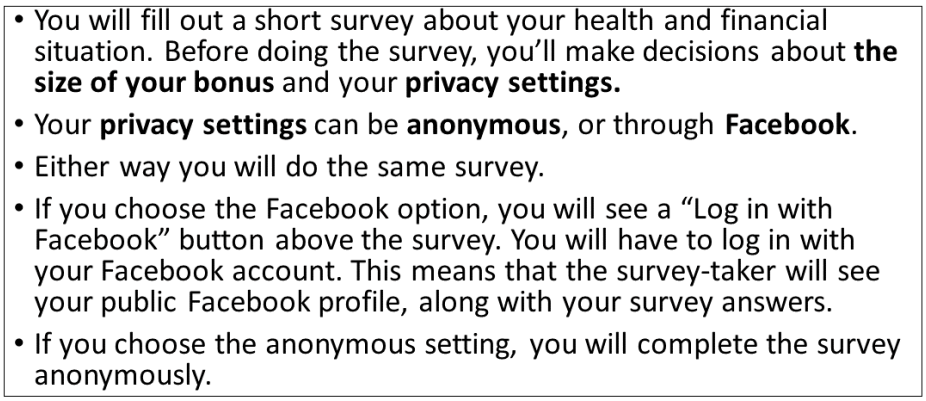

Figure A1: This figure shows the introductory page of the instructions, which was shown to all groups.

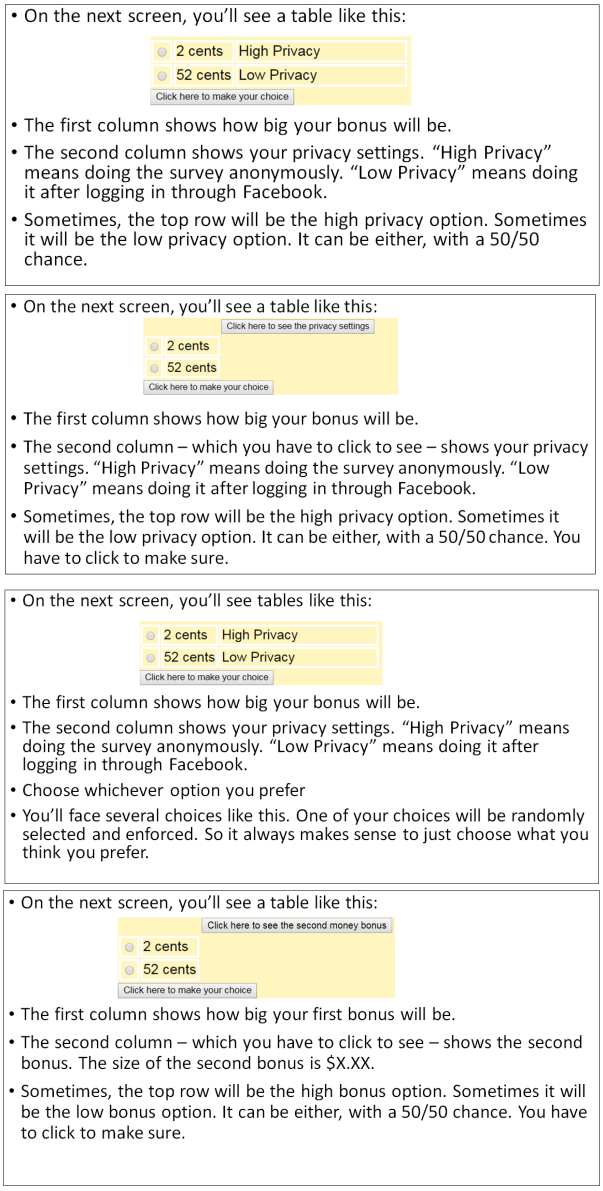

Figure A2: This figure shows the instructions page for each of the four treatments. The Direct Tradeoff Treatment group was shown the instructions in the top panel. The Veiled Tradeoff Treatment group was shown the instructions in the second panel. The Elicitation Treatment group was shown the instructions in the third panel. The Placebo Veiled Tradeoff group was shown the instructions in the bottom panel.

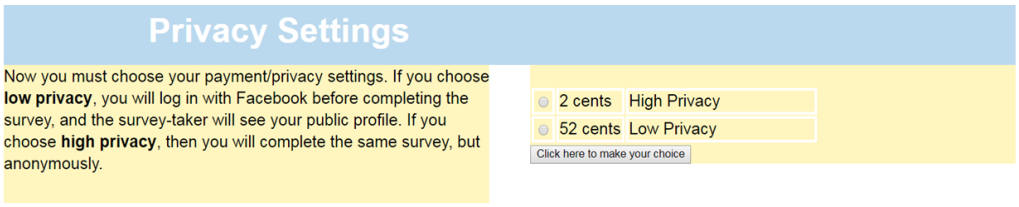

Figure A3: This is a screenshot of the privacy task page completed by participants in the Direct Tradeoff Treatment, after reading instructions and completing practice rounds. Participants in the Elicitation Treatment faced an identical task, but completed it multiple times, with monetary bonuses ranging from \$0.25 to \$5.02. As described in the instructions in the bottom panel of Figure A2, one of these choices would be randomly selected and enforced.

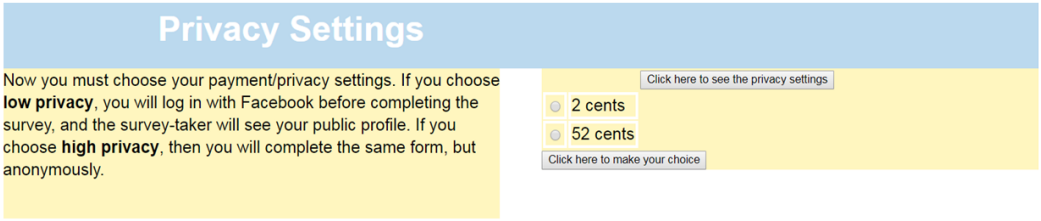

Figure A4: This is a screenshot of the privacy task page completed by participants in the Veiled Tradeoff Treatment, after reading instructions and completing practice rounds.

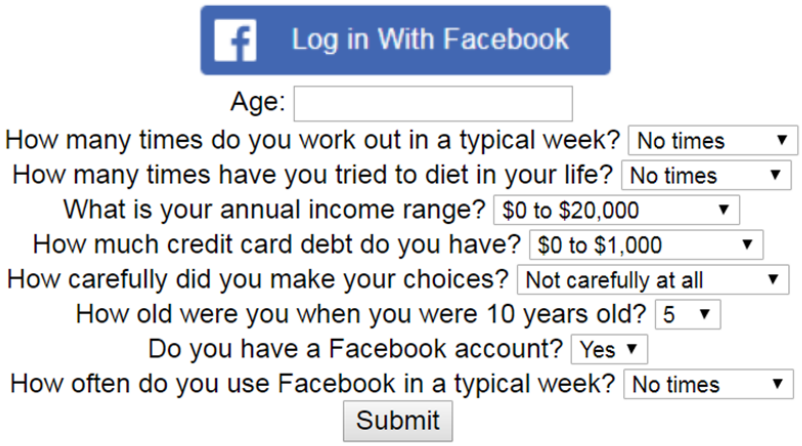

Figure A5: After making their privacy choices, all participants completed the survey above. Those participants who opted for the anonymous survey were not shown the Facebook login button. Those that opted for the low privacy setting saw the login button, as in the picture above.

7 Appendix B: Model of Clicking Behavior in Veiled Tradeoff Treatment

Consider the decision of whether a participant should click to reveal the privacy column. Let her utility from privacy be v, her utility from 50 cents be $u(50)$, and her utilitity from clicking be c.

If she does not click to reveal the privacy column, the participant will choose the 50 cent option, face a 50% chance of losing her privacy, and incur no clicking costs:

$$U_{\text{no click}} = u(50) + 0.5 \cdot v$$

Indeed, if $v < u(50)$ (she values 50 cents more than privacy), there is no need to click to reveal. If she clicks and faces a tradeoff between money and privacy, she will choose the money anyway, so if $v < u(50)$, **she will never click to reveal**, since this will allow her to avoid clicking costs.

Suppose instead that $v > u(50)$, so the participant may want to click to reveal. The participant will click to reveal if the expected utility of clicking is larger than the expected utility of not clicking.

If she clicks to reveal, she will choose privacy with certainty, face a 50% chance of not getting the 50 cents, and incur clicking costs:

$$U_{click} = 0.5 \cdot u(50) + v + c$$

She will click if U_{click} is greater than the utility of not clicking:

$$0.5 \cdot u(50) + v + c > u(50) + 0.5 \cdot v$$
$$v - u(50) > 2 \cdot c$$

This demonstrates the second conclusion. If clicking costs are zero, a participant who values privacy more than 50 cents **will always click to reveal**. In general, she will click to reveal so long as the difference between her privacy valuation and the utility from 50 cents is larger than two times the clicking costs.

Given this decision rule, and given the distribution of WTP prices in the Elicitation Treatment, how big would clicking costs have to be to explain the treatment effect? If $c = 0$, then everyone with a WTP for privacy above 50 cents should click to reveal, leading to a click rate of 65%. If clicking costs are \$1.00, then anyone with a WTP for privacy of \$2.50 and above would click to reveal, leading to a click rate of roughly 50% – still higher than the observed click rate of 42%. Clicking costs would have to be nearly \$2.00 to end up with a click rate observed in the main experiment.