

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

06/2018

Harvard Law School
Cambridge, MA 02138

Contributors to this series are John M. Olin Fellows or Terence M. Considine Fellows in Law and Economics at Harvard University.

This paper can be downloaded without charge from:

The Harvard John M. Olin Fellow's Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

Why Are Privacy Preferences Inconsistent?

Dan Svirsky[†]

June 28, 2018

Abstract: People say they value privacy, but they routinely share copious amounts of personal data without much conscious thought. The two most common explanations for this inconsistency are revealed preference (preferences are revealed by actions, not words) and ignorance (people are unaware of their privacy loss). This paper presents an experiment that shows the limits of these explanations. In an online experiment, participants must decide whether to share their Facebook profile data with the survey-taker in exchange for a higher payoff. When directly making a tradeoff between 50 cents and privacy in a *Direct Tradeoff* treatment, 67% of participants appear to value privacy by refusing to share their Facebook data for 50 cents. By contrast, when participants must “click to reveal” to learn whether privacy is free or costs 50 cents in a *Veiled Tradeoff* treatment, only 40% of participants appear to value privacy by refusing to share their Facebook data for 50 cents. Consistent with information avoidance driving these differences in apparent valuations of privacy, 56% of participants did not click to reveal to learn which payment option was associated with privacy. The experiment was repeated 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 Tradeoff* treatment became much less effective. One month later, the results returned to their pre-scandal values, and the treatment was as effective as before. Taken together, these findings show how individuals’ valuations of privacy are sensitive to the elicitation method. 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 suggest reason for skepticism about current federal privacy law, which relies on giving consumers more information about data practices to help them make better choices.

Preliminary Draft

Please do not circulate

[†]Harvard Economics Department, dsvirsky@hbs.edu. I thank Jennifer Arlen, Oren Bar-Gill, Varanya Chaubey, Christine Exley, Louis Kaplow, David Laibson, Florencia Marotta-Wurgler, 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

Over the course of a day, people constantly emit personal data without conscious thought. At the same time, many people express anger over the loss in privacy. The idea of a third-party app on Facebook accessing user data angers many Facebook users, but many have not taken the thirty seconds necessary to limit such apps' access to their data.

Why the inconsistency? The two simplest explanations are revealed preference and ignorance. Economists typically rely on observing people's choices to measure preferences – not what they say. Maybe people *say* they value privacy but really don't. Alternatively, maybe people *do* value privacy, but they don't have a good understanding of how much data they routinely give away, or they don't know how to change their privacy settings. There is evidence for both views (Acquisti et al. (2015)).

In this paper, I use an experiment to illustrate the limits of these explanations. Even people who value privacy – in the economic sense of being willing to pay for it – are simultaneously able to close their eyes and give away their data at low prices. In the experiment, participants complete a survey, but first decide whether to share their Facebook profile with the survey-taker in exchange for a higher payoff.

When participants in a *Direct Tradeoff* Treatment face a choice between a 50 cent bonus and privacy, 67% 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 a \$5.00 offer.

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 Facebook profile, and they can click a button to check which option involves a loss in privacy. I find that most people (56%) 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.

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. The two simplest explanations of privacy inconsistency – revealed preference and ignorance – are insufficient to explain people's behavior in this experiment. 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 differentially. 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

The lodestar of privacy law in the United States is Notice and Choice: firms must inform users about privacy settings and let consumers make decisions for themselves. Privacy scholars have increasingly questioned this approach, pointing to empirical evidence on people's inconsistency in privacy choices. From this inconsistency, scholars conclude that people are susceptible to cognitive biases or bounded rationality when making privacy choices.

However, evidence from psychology and economics suggests a different reason for skepticism of Notice and Choice which would also explain people's inconsistency: information avoidance. Even people who value privacy in one context might also be willing to

maintain strategic ignorance. That is, they may be content to give up their data as long as they do not have to think hard about the loss of privacy.

Given the robust evidence of such behavior in other domains, like altruism and health, and given the Notice and Choice framework, it is important to test whether information avoidance occurs in privacy decisions.

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

¹There 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.

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 & 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 trade privacy for money, or convenience, because this is what they actually prefer, regardless of what they say. If information has some cost, then consumers' decision to avoid privacy information is itself an illustration of revealed preference.

For either explanation – revealed preference or ignorance – more information is better. If it's costless, people will always opt for better information about privacy settings.

2.3 Fickle privacy preferences can be explained by bounded rationality or by information avoidance

There is a robust literature from psychology and economics on information avoidance and self-serving biases (Golman et al. (2017)). This literature runs counter to classic economic models of information. Economists originally modeled information as an intermediate good (Posner (1978), Stigler (1961)). Information is valuable – i.e., is worth hiding or discovering – because it helps us achieve ends. But scholars in psychology and economics increasingly recognize that people sometimes behave as if information has emotional valence (Oster et al. (2013)), especially in situations where there is a gap between what a person wants to do and what she feels she ought to do. 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)). These findings are in line with models of behavior where actors are uninformed, but have some control over what they learn and make optimal choices over their decisions and the information they see (Grossman & van der Weele (2016), Rabin (1995); *c.f.* Kaplow (1990)).

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 estate planning, the majority of Americans do not have wills (Jones (2016)). In sum, people avoid information that upsets them, even if in theory it should help them make a more optimal decision.

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, it is worth asking how consumers respond to privacy disclosures.

3 Experimental Design

The core of the experiment looks at people’s privacy choices in one of three treatments: an Elicitation Treatment, a Direct Tradeoff Treatment, and a Veiled Tradeoff Treatment. This section first discusses the overall timeline of the experiment, then describes the three treatments in detail.

621 participants were recruited on Amazon Mechanical Turk to take a short survey about health and financial status.² 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.³ The experiment was conducted on May 30, 2018.

After recruitment, the timeline of the experiment consists of three stages: instructions, privacy settings, and a survey.⁴ 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.

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.⁵ 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 Cambridge Analytica scandal.

²Research 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).

³Median hourly wages for workers was \$13.56 (based on a median payment of \$1.12 for a median completion of 4 minutes 56 seconds).

⁴Appendix A presents the experimental instructions in detail.

⁵Even 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.




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.

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

• On the next screen, you'll see a table like this:

<input type="radio"/> 2 cents	High Privacy
<input type="radio"/> 52 cents	Low Privacy

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Sometimes, the top row will be the high privacy option. Sometimes it will be the low privacy option. It can be either, with a 50/50 chance.

• On the next screen, you'll see a table like this:

<input type="radio"/> 2 cents	Click here to see the privacy settings
<input type="radio"/> 52 cents	

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column – which you have to click to see – shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Sometimes, the top row will be the high privacy option. Sometimes it will be the low privacy option. It can be either, with a 50/50 chance. You have to click to make sure.

• On the next screen, you'll see tables like this:

<input type="radio"/> 2 cents	High Privacy
<input type="radio"/> 52 cents	Low Privacy

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Choose whichever option you prefer
- You'll face several choices like this. One of your choices will be randomly selected and enforced. So it always makes sense to just choose what you think you prefer.

Figure 2: This figure shows the instructions page for each of the three 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 middle panel. The Elicitation Treatment group was shown the instructions in the bottom panel.

In the Elicitation Treatment, participants faced a multiple price list where they decided whether to share their Facebook data at different price levels ranging from \$0.50 to \$5.00, knowing that one of their choices would be enforced. This is a standard technique in applied microeconomics to elicit a willingness-to-pay (WTP), 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 in equal increments from \$0.52 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.

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.

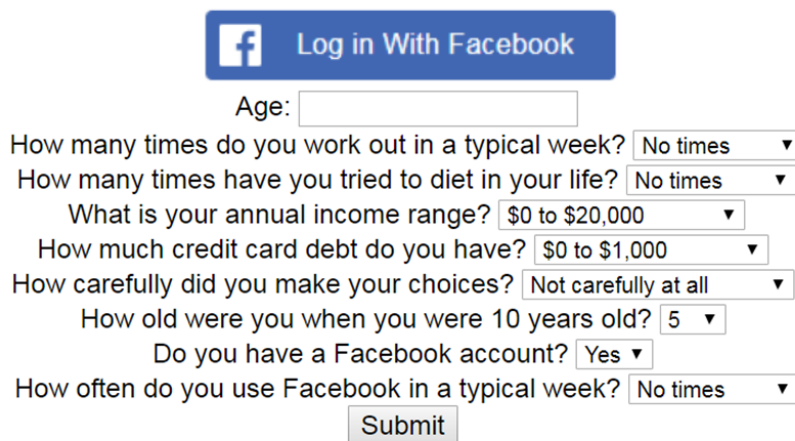
In the Veiled Tradeoff Treatment, participants faced the same decision as in the Direct Tradeoff Treatment, but the privacy setting was initially hidden. Participants had to click to reveal the column describing the privacy settings, and there was a 50% chance that the higher money bonus would mean losing their anonymity.⁶

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.⁷

⁶Note 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.

⁷Contact the author for the ZIP file: dsvirsky@hbs.edu .



The image shows a survey form with a blue button at the top that says "Log in With Facebook" next to the Facebook logo. Below the button is an "Age:" label followed by an empty text input field. The survey contains several questions with dropdown menus:

- How many times do you work out in a typical week? (No times)
- How many times have you tried to diet in your life? (No times)
- What is your annual income range? (\$0 to \$20,000)
- How much credit card debt do you have? (\$0 to \$1,000)
- How carefully did you make your choices? (Not carefully at all)
- How old were you when you were 10 years old? (5)
- Do you have a Facebook account? (Yes)
- How often do you use Facebook in a typical week? (No times)

At the bottom of the form is a "Submit" button.

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.

3.1 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 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 May 30, 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.




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.

4 Results

Table 1 presents summary statistics on the survey answers, as well as a balance check. Nearly all – 94% – 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 average participant accessed Facebook roughly three 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. Attrition was quite low. In the Direct Tradeoff, Veiled Tradeoff, and Elicitation treatments, attrition (defined as people who read the instructions but quit before the survey round) was 5%, 2%, and 3% respectively.

	Elicitation Treatment	Direct Tradeoff Treatment	Veiled Tradeoff Treatment	P-Value
Age	33.72 (10.60)	31.84 (8.868)	35.12 (12.52)	0.11
Diet Attempts	2.310 (1.589)	2.286 (1.553)	2.390 (1.616)	0.90
Workouts Per Week	2.437 (1.315)	2.520 (1.270)	2.524 (1.434)	0.82
Annual Income (0-4)	28.26 (24.01)	24.08 (23.15)	27.56 (23.44)	0.34
Credit Card Debt (0-4)	1.728 (3.844)	1.296 (2.492)	1.768 (3.863)	0.56
Has Facebook Account	0.92 (0.24)	0.95 (0.22)	0.95 (0.22)	0.41
Weekly Facebook Use	2.89 (1.45)	2.82 (1.33)	2.91 (1.51)	0.56
N	111	122	133	

Table 1: Summary statistics for the three treatment groups. 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.

4.1 Results: Elicitation Treatment

The results of the Elicitation Treatment show that many people are willing to pay non-trivial amounts of money to remain anonymous. Table 2 shows people's WTP for staying anonymous in the Elicitation Task. 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 – 42.3% – 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 for these participants, sharing their Facebook profile entails a privacy cost equal to roughly one hour of labor. Nonetheless, the second most-common WTP price was at the lower end, with 30.6% choosing to sell their Facebook profile at 50 cents. The remaining 27% evinced a WTP between 50 cents and \$5.00. The majority of participants required at least \$2.50 before they would share their Facebook data.

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.

Willingness to Pay to Remain Anonymous	Column %	Cumulative %
\$0.50	30.6	30.6
\$1.00	1.8	32.4
\$1.50	0.9	33.3
\$2.00	4.5	37.8
\$2.50	5.4	43.2
\$3.00	2.7	45.9
\$3.50	1.8	47.7
\$4.00	4.5	52.2
\$4.50	1.8	54.0
\$5.00	3.6	57.6
More Than \$5.00	42.3	100.0

N = 111

Table 2: Participants in the Elicitation Treatment faced a multiple price list to elicit their WTP price for remaining anonymous (i.e., not sharing their Facebook profile). This table presents the breakdown of these WTP prices.

4.2 Results: Direct Tradeoff Treatment vs Veiled Tradeoff Treatment

I find a treatment effect from putting a costless veil on privacy settings. 67% of people in the Direct Tradeoff Treatment refuse to sell their Facebook data for 50 cents.⁸ 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 slim majority in the Veiled Tradeoff Treatment (56%) 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 3 reports various regressions where the unit of observation is an individual, the dependent variable is whether the participant chose Low Privacy, and the independent variable is an indicator variable for being in the Veiled Tradeoff Treatment.

⁸This is in line with the results from the Elicitation Treatment group, 69.4% of whom rejected an offer of 50 cents to share their Facebook profile, a slightly higher but statistically insignificant difference.




Figure 5: This figure shows the proportion of participants opting to remain anonymous instead of sharing their Facebook profile for 50 cents, for the Direct Tradeoff Treatment ($N = 122$) and the Veiled Tradeoff Treatment ($N = 133$). 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.

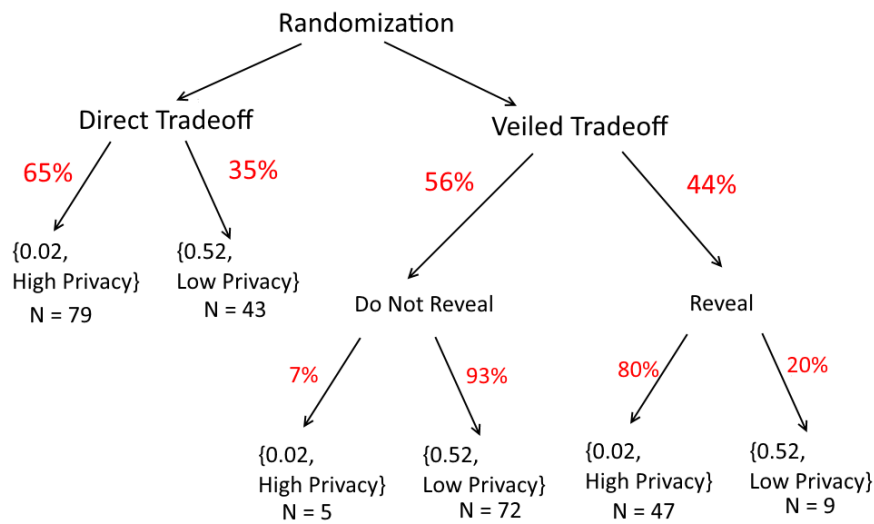


Figure 6: This figure shows the decisions made by participants across the Direct Tradeoff Treatment and Veiled Tradeoff Treatment. This figure excludes participants who were randomized into a degenerate choice between more money *and* high privacy vs less money and low privacy. In the Direct Tradeoff 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.

The treatment effect is robust to using different samples. For example, the survey includes questions to test comprehension and attention. 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 90% 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 80% of participants said they answered the questions “very carefully”, 15% said “a little carefully” and 4.5% said not at all carefully. 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 under 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 3 reports the results of these robustness checks. Columns 2 - 5 report the result of the main regression, described above, but using different models. Column 2 includes controls for survey answers, while columns 3 - 5 exclude participants based on comprehension, attention, and confusion (defined as opting for less money without clicking to reveal the privacy setting). The main results hold throughout.

Table 3: Privacy Decisions In Direct Tradeoff Group and Veiled Tradeoff Group, with Robustness Checks

	(1)	(2)	(3)	(4)	(5)
	Full Sample	Full Sample	Passed Comprehension Check	Answered Carefully	Excludes 'Didn't Click, Chose 0 cents'
Veiled Tradeoff Treatment	-0.26*** (0.06)	-0.25*** (0.07)	-0.21** (0.07)	-0.24*** (0.07)	-0.30*** (0.07)
Constant	0.67*** (0.04)	0.34** (0.13)	0.68*** (0.04)	0.68*** (0.04)	0.67*** (0.04)
Controls for Survey Answers	No	Yes	No	No	No
Observations	255	255	232	245	250
Adjusted R^2	0.06	0.07	0.03	0.04	0.07

Table reports a linear probability regression of (Chose Anonymity over 50 cents) on indicator variable for Veiled Tradeoff Treatment. Omitted group is the Direct Tradeoff Treatment. Block bootstrap standard errors, clustered at the treatment level, in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

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

Roughly a month after the first 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.

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. Further, it is not clear whether the scandal should make people distrust Facebook more, or instead have more faith that because of the scandal, Facebook has to be more careful with its third party applications. Nonetheless, the data is a unique opportunity to see how privacy preferences change during major privacy scandals, and whether any changes persist.

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.




Figure 7: This figure shows the proportion of people in each treatment who choose to keep their Facebook profile private 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, using block bootstrap standard errors at the treatment level.

Table 4 presents regression results and robustness checks. The regression specification is as follows, letting p be the proportion of people choosing privacy over 50 cents, T be an indicator variable for 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.

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

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.

Table 4: Privacy Decisions Before, During, and After Facebook Scandal

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

Table reports a linear probability regression of (Chose Anonymity over 50 cents) on indicator variable for the Veiled Tradeoff Treatment, a time fixed effect for the experiment run at the peak of the scandal, a time fixed effect for the experiment run after the scandal, and interactions between the time fixed effects and the treatment. Block bootstrap standard errors, clustered at the treatment level, in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In sum, during the Facebook scandal the experimental treatment became significantly less effective, but this was driven by people in the treatment group opting for privacy more often, rather than by a change in preferences in the Direct Tradeoff Treatment.

5 Discussion

This paper presents an experiment that adds to the literature documenting inconsistencies in people's privacy preferences. The two most common explanations for this inconsistency are revealed preference and ignorance. The findings of this experiment show the limits of these explanations. Instead, the findings here suggest that information avoidance may be an important mechanism driving people's privacy decisions. When directly making a trade-off between 50 cents and anonymity, most people opt for anonymity. But when the privacy settings are occluded, even with a costless veil, people take the money and trade away their data.

These results 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. This result was short-lived.

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 accessible, 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 Preferences 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 Literature* **55**(1), 96–135.
- Grossman, Z. & van der Weele, J. (2016), ‘Self-Image and Strategic Ignorance in Social Decisions \$’, *Journal of the European Economic Association* (February), 1–48.
- 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.
- Jones, J. M. (2016), ‘Majority in U.S. Do Not Have a Will’.
URL: <http://news.gallup.com/poll/191651/majority-not.aspx> **CONTACT:**

- Kaplow, L. (1990), 'Optimal deterrence, uninformed individuals, and acquiring information about whether acts are subject to sanctions', *Journal of Law, Economics, & Organization* **6**(1), 93–128.
- 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.
- 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).
- Rabin, M. (1995), 'Moral Preferences, Moral Constraints, and Self-Serving Biases', *Working Paper* p. 44.
- 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.
- Warren, S. D. & Brandeis, L. D. (1890), 'The Right to Privacy', *Harvard Law Review* **4**(5), 193–220.

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. Screenshots of the task for the Direct Tradeoff Treatment and Veiled Tradeoff Treatment are in Figure A3 and Figure A4. The task screen for participants in the Elicitation Treatment was identical to the task seen by participants in the Direct Tradeoff task, but with monetary bonuses ranging from 2 cents to \$5.02.

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.

- You will fill out a short survey about your health and financial situation. Before doing the survey, you’ll make decisions about **the size of your bonus** and your **privacy settings**.
- Your **privacy settings** can be **anonymous**, or through **Facebook**.
- Either way you will do the same survey.
- If you choose the Facebook option, you will see a “Log in with Facebook” button above the survey. You will have to log in with your Facebook account. This means that the survey-taker will see your public Facebook profile, along with your survey answers.
- If you choose the anonymous setting, you will complete the survey anonymously.

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

• On the next screen, you'll see a table like this:

<input type="radio"/> 2 cents	High Privacy
<input type="radio"/> 52 cents	Low Privacy

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Sometimes, the top row will be the high privacy option. Sometimes it will be the low privacy option. It can be either, with a 50/50 chance.

• On the next screen, you'll see a table like this:

<input type="radio"/> 2 cents	Click here to see the privacy settings
<input type="radio"/> 52 cents	

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column – which you have to click to see – shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Sometimes, the top row will be the high privacy option. Sometimes it will be the low privacy option. It can be either, with a 50/50 chance. You have to click to make sure.

• On the next screen, you'll see tables like this:

<input type="radio"/> 2 cents	High Privacy
<input type="radio"/> 52 cents	Low Privacy

[Click here to make your choice](#)

- The first column shows how big your bonus will be.
- The second column shows your privacy settings. "High Privacy" means doing the survey anonymously. "Low Privacy" means doing it after logging in through Facebook.
- Choose whichever option you prefer
- You'll face several choices like this. One of your choices will be randomly selected and enforced. So it always makes sense to just choose what you think you prefer.

Figure A2: This figure shows the instructions page for each of the three 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 middle panel. The Elicitation Treatment 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 \$0.02 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.


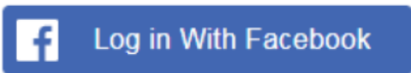


Figure A4: This is a screenshot of the privacy task page completed by participants in the Veiled Tradeoff Treatment, after reading instructions and completing practice rounds.



Age:

How many times do you work out in a typical week?

How many times have you tried to diet in your life?

What is your annual income range?

How much credit card debt do you have?

How carefully did you make your choices?

How old were you when you were 10 years old?

Do you have a Facebook account?

How often do you use Facebook in a typical week?

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.