
ARTICLE

WHY DO PEOPLE AVOID INFORMATION ABOUT PRIVACY?

DAN SVIRSKY†

Why do people keep their head in the sand when making data sharing decisions? There is a widespread intuition, supported by copious research, that people are inconsistent in their behavior around internet privacy. Anger about privacy scandals dominates newspaper headlines, but most people don't change their default privacy settings, even when it's easy to do so. New evidence confirms that this inconsistency is real, and that information avoidance helps drive the inconsistency. This raises a new question: how does information avoidance work? This paper presents a new experimental design to start unpacking how information avoidance operates. There are two main results. First, the experiment replicates existing information avoidance experiments: people who value privacy are willing to deal away their data for small money amounts if given a chance to avoid seeing the privacy consequences of their actions. Second, the experiment shows that while people are comfortable avoiding information about privacy in a passive way, they are not comfortable actively hiding it. These results show that people's ability to keep their head in the sand is fragile: it is a preference people are not willing to exercise conspicuously.

INTRODUCTION..... 24

I. PEOPLE'S PRIVACY PREFERENCES ARE INCONSISTENT, AND INFORMATION AVOIDANCE CAN EXPLAIN THIS INCONSISTENCY..... 27

 A. *Privacy Preferences are Inconsistent* 27

 B. *Information Avoidance Can Explain this Inconsistency* 30

† Dan Svirsky, Uber Technologies, Inc. (dsvirsky@uber.com).

II. MULTIPLE MECHANISMS CAN EXPLAIN INFORMATION AVOIDANCE BEHAVIOR.....	32
A. <i>Modeling the Wiggle Room Decision</i>	32
B. <i>Mechanism: Thinking about Privacy Losses is Upsetting</i>	35
C. <i>Mechanism: Signaling</i>	36
D. <i>Mechanism: Psychic Choosing Costs</i>	36
E. <i>Mechanism: All of the Above</i>	37
III. EXPERIMENTAL DESIGN AND EMPIRICAL APPROACH	37
IV. RESULTS.....	42
A. <i>Summary Statistics</i>	43
B. <i>Main Results: Average Treatment Effects</i>	43
CONCLUSION	46

INTRODUCTION

Why do people keep their head in the sand when making data sharing decisions?

There is a widespread intuition, supported by copious research, that people are inconsistent in their behavior around internet privacy.¹ Anger about privacy scandals dominates newspaper headlines, but most people don't change their weak, default privacy settings, even when it's easy to do so.²

New evidence confirms that this inconsistency is real, and that information avoidance helps drive the inconsistency.³ Using an

¹ See Alessandro Acquisti et al., *Privacy and Human Behavior in the Age of Information*, 347 SCI. 509, 510 (2015) (explaining the widespread discrepancies between online privacy attitudes and behaviors); Susan Athey et al., *The Digital Privacy Paradox: Small Money, Small Costs, Small Talk* 17-18 (Nat'l Bureau of Econ. Research, Working Paper No. 23488, 2017) ("Consumers say they care about privacy, but at multiple points in the process end up making choices that are inconsistent with their stated preferences.").

² See, e.g., Kevin Lewis et al., *The Taste for Privacy: An Analysis of College Student Privacy Settings in an Online Social Network*, 14 J. COMPUTER-MEDIATED COMM. 79, 95 (2008) (finding that only one third of college students using Facebook changed their default privacy settings); Ralph Gross & Alessandro Acquisti, *Information Revelation and Privacy in Online Social Networks*, 2005 ACM WORKSHOP ON PRIVACY IN THE ELEC. SOC'Y 71, 78 (2005) ("We can conclude that only a vanishingly small number of users change the (permissive) default privacy preferences.").

³ See Dan Svirsky, *Why Are Privacy Preferences Inconsistent?* 24 (JOHN M. OLIN CTR. FOR LAW, ECON., & BUS. FELLOWS' DISCUSSION PAPER SERIES, HARV. LAW SCH., Discussion Paper No. 81, 2018) ("This paper presents an experiment that adds to the literature documenting inconsistencies in people's privacy preferences.").

experimental design adopted from research on altruism,⁴ this research finds that people are willing to give up nearly an hour's worth of wages to keep their Facebook data private.⁵ At the same time, participants in a treatment group are *also* willing to trade their data for 52 cents if given a chance to avoid seeing the privacy implications of their choice.⁶ Hence, information avoidance behavior can recreate, in a controlled experimental setting, the pattern of behavior commonly seen in field settings where people are inconsistent about privacy.

This raises a new question: why does information avoidance with respect to privacy online happen?

While the experiment gives strong evidence that people avoid (nearly) costless information about privacy, there are multiple ways to understand this behavior. One possibility is that thinking about losing privacy is inherently unpleasant. There are many topics outside of internet privacy that are inherently upsetting to consider, like cockroaches, death, and one's own moral failings.⁷ Avoidance thus reduces the time to consider those unpleasant facets of life. For example, many people eat at restaurants without looking at public health inspection reports on vermin in kitchens. Privacy might be like that.

Another possibility is that even when people know that they *should* care about privacy, they don't really care.⁸ Evidence from altruism experiments, for example, demonstrate that people will give money to a Salvation Army volunteer ringing a bell at a supermarket entrance, but

⁴ See Jason Dana et al., *Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness*, 33 *ECON. THEORY* 67, 70-74 (2007) (describing experimental design of a modified dictator game to test wealth allocation); Zachary Grossman & Joel J. van der Weele, *Self-Image and Willful Ignorance in Social Decisions*, 15 *J. EUR. ECON. ASS'N* 173, 197-206 (2017) (“[analyzing] a Bayesian signaling model of an agent who cares about self-image and has the opportunity to learn the social benefits of a personally costly action”); Lauren Feiler, *Testing Models of Information Avoidance with Binary Choice Dictator Games*, 45 *J. ECON. PSYCHOL.* 253, 256-260 (2014) (extending the moral wiggle room experimental design by manipulating the probabilities of different money payoffs). This paper extends the moral wiggle room experimental design in the privacy space in a similar way to the Grossman & van der Weele paper, which also tests the effects of differing the default amount of information presented, albeit in the social preferences space.

⁵ See Svirsky, *supra* note 3, at 14 (“[F]or these participants, sharing their Facebook profile entails a privacy cost equal to roughly one hour of labor.”).

⁶ See *id.* (finding that nearly a third of participants chose to share their Facebook profile for 50 cents).

⁷ See Russell Golman et al., *Information Avoidance*, 55 *J. ECON. LIT.* 96, 106-07 (2017) (explaining the use of information avoidance as a defense against disappointment).

⁸ Cf. Christine Exley, *Excusing Selfishness in Charitable Giving: The Role of Risk*, 83 *REV. ECON. STUDIES* 587 (2016) (demonstrating how participants use risk as an excuse to avoid donating money); Dana et al., *supra* note 4 (showing how people exploit wiggle room to avoid behaving altruistically).

people will also avoid that entrance if there are multiple ways to enter the store.⁹ Perhaps in both this domain and privacy, people simply want to *seem* like the type of person who values an important social good (altruism, data security, privacy).¹⁰

Another alternative is that making a tradeoff between money and privacy is difficult, and people are happy to avoid undergoing this psychic cost.¹¹ If someone is offered a cup of coffee for \$0.25, or for \$5.00, she can tell that the first price is somewhat low and the second price somewhat high. The same might not be true for sharing data.

Yet another alternative is that all these explanations hold, to different degrees and with different interaction effects, depending on the person and the context. Perhaps someone wants to *seem* like she cares about privacy, doesn't like thinking about it, and has no real idea what a fair price for data is. All these mechanisms can push her to avoid information. For one person, the first mechanism might dominate. For the same person, the third mechanism might dominate for certain types of data.

This paper presents a new experimental design to start exploring these questions in two steps. First, it replicates the initial two-group experiment on information avoidance in privacy, and second, it adds a third group which has an active choice about hiding information. In the design, participants make decisions about the privacy settings and potential money bonuses for a survey they must complete. They can either complete the survey anonymously or after sharing their public Facebook profile with the survey-taker. Different money bonuses can attach to different privacy settings.

⁹ See James Andreoni et al., *Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving*, 125 J. POL. ECON. 625, 628 (2017) ("When avoidance was easy because only one door had a solicitor, nearly one-third of those intending to pass through the occupied door instead chose to use an unoccupied entrance."). See also Edward Lazear et al., *Sorting in Experiments with Application to Social Preferences*, 4 AM. ECON. J. APPLIED ECON. 136, 136 (2012) ("[A]llowing subjects to avoid environments in which sharing is possible significantly reduces sharing."); Stefano DellaVigna et al., *Testing for Altruism and Social Pressure in Charitable Giving*, 127 Q.J. ECON. 1, 1 (2012) (finding that individuals who knew when fundraiser solicitations would arrive at their homes were more likely to avoid the giving scenario).

¹⁰ See Christine Exley & Judd Kessler, *Motivated Errors* 1-2 (Harv. Bus. Sch., Working Paper, No. 18-017, 2017) (finding that individuals motivated to act in their own self-interest display behavioral biases yet act more rational when these self-serving motivations are removed).

¹¹ See Cass Sunstein, *Choosing Not to Choose*, 64 DUKE L.J. 1, 1 (2014) ("In part because of limitations of [cognitive resources,] and in part because of awareness of their own lack of information and potential biases, people sometimes want other people to choose for them.").

There are three experimental groups: a *direct* tradeoff group, a *veiled* tradeoff group, and a *choice* tradeoff group. Importantly, there is no real difference in the choices the three groups make. In all three cases, participants decide whether to share their Facebook data for 52 cents, a decision participants can be made aware of if presented the option to view privacy settings. For the *direct* tradeoff group, privacy settings on data sharing are hidden by default; for the *veiled* tradeoff group, privacy settings are visible by default but can be actively hidden; and for the *choice* tradeoff group, privacy settings are visible by default but can be hidden or randomized.

There are three main results. First, the findings replicate the original information avoidance experiment: the *direct* tradeoff group chose privacy 70% of the time, while the *veiled* tradeoff group chose privacy 40% of the time. Second, the findings for the *choice* tradeoff group are directly in between the *direct* and *veiled* groups: participants choose privacy 56% of the time. Third, I find that participants in the *choice* tradeoff group very rarely made the active choice to hide information: they clicked the button to hide privacy settings only 9% of the time, whereas the *veiled* tradeoff group accepted the default of keeping privacy settings hidden 44% of the time.

Taken together, these results shed light on how information avoidance works. People are comfortable avoiding information about privacy, but they are not comfortable actively hiding it. Strangely, even the option of actively hiding makes people less likely to choose privacy.

Section I of the paper discusses current privacy law in the United States as well as the literature on privacy inconsistency and what causes it. Section II unpacks different mechanisms that can drive information avoidance. Section III details the experimental design. Section IV presents the results of the experiment. Section V concludes.

I. PEOPLE'S PRIVACY PREFERENCES ARE INCONSISTENT, AND INFORMATION AVOIDANCE CAN EXPLAIN THIS INCONSISTENCY

A. *Privacy Preferences are Inconsistent*

Extensive experiments and surveys document that people say they value privacy but also give up their data for small amounts of money or convenience.¹² For example, people claim to care greatly about protecting

¹² Athey et al., *supra* note 1, at 2 (“Whereas people say they care about privacy, they are willing to relinquish private data quite easily when incentivized to do so.”); Leslie John et al., *Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information*, 37 J. CONS. RES. 858, 858 (2011) (“[D]isclosure of private information is responsive to environmental

their data, yet are much less likely to choose a privacy-preserving option if it is listed second on a menu instead of first.¹³ Even privacy disclosures that are strikingly clear and scary have limited impact on how much data people give away.¹⁴

These empirical findings have legal importance because privacy law in the United States relies on a Notice and Choice framework.¹⁵ Firms in the United States can legally harvest data from consumers so long as consumers receive proper notice and agree to the exchange. This approach was first outlined in a 1973 report by the U.S. Department of Health, Education and Welfare.¹⁶ The reliance on notice and voluntary consent was a departure from how privacy law originally developed. Before the rise in internet commerce and telecommunications, privacy concerns in transactions between non-state actors were governed by tort law.¹⁷ As internet transactions have come to dominate private data, contract law principles have come to increasingly govern privacy law.¹⁸ Since privacy is governed by consumer choice, the well-documented fickleness in how consumers make privacy decisions has policy importance.

There are exceptions to the Notice and Choice framework. Banks send annual privacy notices because of the Gramm-Leach-Bliley Act.¹⁹ Doctors require patients to sign an extra form because of the Health Insurance Portability and Accountability Act.²⁰ Websites ask users if they are older

cues that bear little connection . . . to objective [privacy] hazards.”). *See, e.g.*, Alessandro Acquisti et al., *What is Privacy Worth*, 42 J.L. STUD. 249, 268-69 (2013) (discussing how individuals make inconsistent decisions in privacy contexts in part because of default privacy settings). *Cf.* Adam Chilton & Omri Ben-Shahar, *Simplification of Privacy Disclosures: An Experimental Test* 566 (Coase-Sandor Working Paper Series in Law and Economics, No. 737, 2016) (describing the failure of simplified privacy disclosures to effect meaningful change in participants’ behavior in disclosing private information).

¹³ Athey et al., *supra* note 1, at 12 (“[W]hen wallets that would maximize privacy from the public are not listed first, students are 13% less likely to select them . . .”).

¹⁴ Chilton & Ben-Shahar, *supra* note 12, at 541 (“[B]est-practice simplification techniques have little to no effect on respondents’ comprehension of the disclosure, willingness to share personal information, and expectations about their rights.”).

¹⁵ *See generally* Chilton & Ben-Shahar, *supra* note 12, at 572-73 (discussing the emphasis in American privacy law on giving proper notice to consumers).

¹⁶ Records Computers and the Rights of Citizens, U.S. DEP’T OF HEALTH, EDUCATION AND WELFARE, SUMMARY AND RECOMMENDATIONS, xxx-xxxii (1973).

¹⁷ *See, e.g.*, Richard A. Posner, *The Right of Privacy*, 12 GA. L. REV. 393, 410 (1978) (discussing the general features of tort-based commercial privacy law); William L. Prosser, *Privacy*, 48 CAL. L. REV. 383, 389 (1960) (highlighting the four types of privacy torts). *Cf.* Samuel D. Warren & Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193, 195 (1890) (articulating the need for common law to grow to cover an individual’s right ‘to be let alone’ and provide a remedy for invasions of privacy by the press”).

¹⁸ There is more stringent regulation of certain consumers and certain industries.

¹⁹ 15 U.S.C. §§ 6801(b), 6805(b)(2) (2000).

²⁰ 42 U.S.C. § 1320d-2(d)(2) (2000).

than 13 -- not 18, not 12, not 16 -- because of the Children's Online Privacy Protection Act.²¹ Outside the United States, there is even more stringent regulation. 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.²² Hence, more muscular regulation does exist, and the political will for it is increasing. But in the United States, such regulation is the exception.

The standard explanations for the inconsistency in measures of how people value privacy are bounded rationality and revealed preference.²³

Under bounded rationality, people are unaware of how much data they are emitting or they struggle to value privacy. The latter may be because privacy is abstract, or because privacy costs are inchoate and uncertain, both in scope and timing.²⁴ Either way, people 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.²⁵ 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.²⁶

Under the revealed preference explanation, people give up privacy simply because this maximizes their utility.²⁷ 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 a consumer's decision to avoid privacy information is itself an illustration of revealed preference.

²¹ 15 U.S.C. § 6502(b)(1)(D) (2000).

²² Regulation 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Advancement of Such Data, and repealing Directive 95/46/EC, 2016 O.J. L 119/1, Art. 83 § 5.

²³ See Svirsky, *supra* note 3, at 2 (noting that scholars point to bounded rationality or cognitive bias to explain inconsistency in privacy choices).

²⁴ See Acquisti, *supra* note 12, at 251-52 (stating that privacy violation costs are amorphous and difficult to assess even when quantifiable).

²⁵ See *id.*, at 267 (explaining that data from one experimental design shows subjects were five times more likely to choose privacy when the trade-off to not doing so was framed as an opportunity to add to an initially gifted amount of money as opposed to retaining the entirety of the initially gifted amount).

²⁶ See *id.*, at 252 (showing empirically that endowment effects and order preferences affect privacy valuations).

²⁷ See Athey, *supra* note 1, at 4 ("The second policy our results document is that there is a disconnect between stated privacy preferences and revealed preference, but that revealed preference is actually closest to the normative preference.").

For either explanation – revealed preference or ignorance – more information is better. If its costless, people will always opt for better information about privacy settings. More recent research suggests a third explanation.

B. Information Avoidance Can Explain this Inconsistency

Recent research demonstrates that information avoidance can explain inconsistency in people’s privacy decisions.²⁸

There is a robust literature from psychology and economics on information avoidance.²⁹ While economists typically model information as an intermediate good³⁰ – 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.³¹ This leads to a recognition that more information is not always better.

This pattern of information-avoiding behavior is important across information-sharing domains. People will give money to a non-profit when a fundraiser goes door to door; many of the same people will find an excuse not to answer the door if they are warned ahead of time that a fundraiser is coming.³² In the health sector, one study found that 27% of intravenous drug users at risk of HIV who got tested for the disease did not return to the clinic to see their results,³³ even though knowing ones HIV positive status can lengthen ones life. People avoid information that upsets them, even if in theory it should help them make a more optimal decision.

Indeed, such behavior appears to be at play in privacy choices as well. Svirsky (2018) demonstrates that even people who value keeping data private at willingness-to-pay (“WTP”) prices of several dollars are willing to give up their data at nominal prices if they can avoid immediately seeing

²⁸ See Svirsky, *supra* note 3, at 24 (concluding that information avoidance may drive privacy decisions).

²⁹ See e.g., Golman, *supra* note 7, at 110 (summarizing the literature in economics and psychology related to regret aversion and optimism maintenance).

³⁰ See generally Posner, *supra* note 17 (analyzing the economics of information from an individual perspective to improve privacy analysis); George Stigler, *The Economics of Information*, 69 J. POL. ECON. 213 (1961) (modeling the ascertainment of market price in order to improve economic organization techniques).

³¹ See Emily Oster et al., *Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease*, 103 AM. ECON. R. 804, 806 (2013) (analyzing the impact of an individual’s expectations in determining whether to undergo genetic testing).

³² See DellaVigna et al., *supra* note 9, at 3 (finding that individuals who knew the exact time at which fundraising solicitors would arrive at their homes were more likely to not open the door to the solicitors).

³³ Patrick Sullivan et al., *Failure to Return for HIV Test Results Among Persons at High Risk for HIV Infection*, 35 J. ACQUIRED IMMUNE DEFICIENCY SYNDROMES 511, 515 (2004).

the result of their decision.³⁴ In the experiment, participants completed a survey after first deciding whether to do the survey anonymously (“high privacy”) or after giving their public Facebook profile data to the survey-taker (“low privacy”) for a bonus.³⁵ A control group chose between {0 cents, high privacy} and {50 cents, low privacy}.³⁶ Roughly two thirds of participants opted for “high privacy”, and in follow-up treatments, most participants refused to opt for “low privacy” until offered at least \$2.50.³⁷

In a treatment group, participants faced a choice between {0 cents, privacy option A} and {50 cents, privacy option B}.³⁸ They knew that the two privacy options were randomized so that “privacy option A” could be “high” or “low” privacy with a 50% chance, and vice versa.³⁹ Importantly, participants in the treatment group could click to reveal the privacy options *before choosing*, at no monetary cost.⁴⁰ If they click a button, they know that they will then either see {0 cents, high privacy} and {50 cents, low privacy} as in the control group, or they will see a more obvious choice between {0 cents, low privacy} and {50 cents, high privacy}.⁴¹

The key finding was that hiding the potential privacy settings behind a veil – even when removing the veil is costless – causes a drop in people’s willingness to keep their data private.⁴² The percentage of people who refused 50 cents to stay anonymous dropped from 67% in the control group to 40% in the treatment group.⁴³

Importantly, this treatment effect that occurs for decisions between a money bonus or privacy does not replicate for decisions between two privacy settings, both associated with money bonuses (with the second money bonus drawn from the distribution of people’s willingness-to-pay prices for privacy). When a second money bonus is hidden by a costless veil, participants do not evince the same willingness to engage in information avoidance.

While the treatment effect is large and robust – it was documented across four experimental rounds across several months in a sample size of over 1000 subjects⁴⁴ – it leaves open important questions of what specific mechanism drives information avoidance behavior.

³⁴ Svirsky, *supra* note 3, at 24.

³⁵ *Id.*, at 6.

³⁶ *Id.*, at 9.

³⁷ *Id.*, at 13.

³⁸ *Id.*, at 9.

³⁹ *Id.*

⁴⁰ *Id.*

⁴¹ *Id.*

⁴² *Id.*, at 21.

⁴³ *Id.*, at 15.

⁴⁴ *Id.*, at 11.

II. MULTIPLE MECHANISMS CAN EXPLAIN INFORMATION AVOIDANCE BEHAVIOR

People engage in information avoidance when making privacy decisions.⁴⁵ That is, they avoid looking at low-cost information about how their data will be shared, even when they value keeping their data private. But why?

This section begins by modeling how a participant makes choices in the information avoidance experiment before then discussing competing mechanisms to explain the treatment effect and how the model can be extended to incorporate these mechanisms.

A. Modeling the Wiggle Room Decision

An agent is making a tradeoff between a payoff and an uncertain cost. For example, she might be deciding whether to download the Uber app, knowing that her data might be sold or her location tracked. Suppose the app brings some value v and a privacy cost c which occurs with probability π . Then, in a standard expected utility model, her utility is as follows:

$$u(\pi) = v - \pi c$$

Now suppose that there is a psychic cost to potentially losing privacy. The thought of something upsetting is itself upsetting. Let the function $\psi()$ map π onto disutility, with

1. $\psi(\pi) > 0 \forall \pi \in [0,1]$
2. $\psi'(\pi) > 0 \forall \pi \in [0,1]$

The first condition says that the possibility of something upsetting is in itself upsetting. The second condition says that the agent gets more upset as the upsetting possibility becomes more likely -- a 100% chance of getting an electric shock upsets the agent more than a 10% chance. Throughout, I will assume that $\psi(0) = 0$. If $\psi(\pi) = 0 \forall \pi \in [0,1]$, we are in the case of classical preferences, where information is only valuable for instrumental reasons but has no valence in and of itself.

The person's utility function now incorporates psychic costs:

$$u(\pi) = v - \pi c - \psi(\pi)$$

⁴⁵ *Id.*, at 24.

In cases where a person's information is fixed – she has a belief about π but can do nothing to change this belief – psychic costs are akin to increasing the cost of a harm, albeit in a potentially non-linear way. The comparative statics are straightforward: more psychic costs means an individual is more likely to avoid an action. Where psychic costs will generate more interesting departures from standard models is in decisions over how much information to collect before making a decision.

How does a participant make choices in the experimental design described above? Consider what happens when a participant gets what is commonly described as “wobble room” – or the chance to make a choice where they give up their data without directly seeing that they are giving up their data. That is, they can choose a higher monetary payoff while still telling themselves that they might be keeping their data private.

In the control group, the participant makes a direct tradeoff between money v and the privacy cost c and the psychic cost of losing privacy with near-certainty, $\psi(\pi_H)$, where π_H is close to 1. She chooses to keep her privacy if the monetary payoff v is lower than the cost (psychic or otherwise) of losing privacy:

$$u(\pi) = v - \pi_H c - \psi(\pi_H) < 0$$

In the *veiled* tradeoff treatment, the participant first has to make a choice about whether to lift a veil, or whether to remain ignorant and take a higher payoff.

Suppose the privacy options are randomized, so that if she lifts the veil, then with probability p she will discover she is in the baseline condition (more money means less privacy), and with probability $1 - p$ she will discover she is in an easy situation where she can get more money *and* keep her privacy. If she remains ignorant, her payoff is:

$$v - pc - \psi(p)$$

In words, she gets the value v with certainty, but undergoes a privacy cost c with probability p and has a psychic cost $\psi(p)$. Consider a participant who would have opted for privacy in the control treatment, so the value of privacy is higher than the monetary payoff v . If she lifts the veil, then her expected payoff is:

$$p(0) + (1 - p)v = (1 - p)v$$

That is, with probability p she will face a tradeoff between privacy and money and will keep her privacy as before, yielding payoff 0; the rest of the time she will get a payoff without any privacy costs (psychic or instrumental).

In a classical preferences world – one where people value information solely because it helps them make better choices, and where information has no attendant psychic costs – $\psi(\pi) = 0 \forall \pi \in [0,1]$, and any agent who chose privacy over money in the baseline treatment will always choose to lift the veil. Why? If, in baseline, she chose privacy over money, that means

$$\begin{aligned} v - \pi_H c - \psi(\pi_H) &< 0 \\ v - \pi_H c &< 0 \\ v &< \pi_H c \\ v &< \pi_H c < 1 \cdot c \\ v &< c \end{aligned}$$

In the *veiled* tradeoff treatment, she lifts the veil if

$$\begin{aligned} (1 - p)v + p(0) &> v - pc - \psi(p) \\ v - pv &> v - pc \\ -pv &> -pc \\ v &< c \end{aligned}$$

The last line is true by assumption. Putting the conclusion into plain language: if clicking to reveal the privacy settings is costless, then anyone who values privacy more than the money bonus would make it their business to *see* which privacy options they were agreeing to. The monetary bonus is simply not worth the risk of giving up data.

In sum, the model demonstrates that the experimental result in Svirsky (2018) cannot be obtained from classic preferences, so long as the cost of clicking to reveal is minimal. The following subsections turn to different mechanisms that *can* explain the wiggle room result.

B. Mechanism: Thinking about Privacy Losses is Upsetting

One explanation for the experimental results is from a model of anxiety in which people are upset by probabilistic harms.⁴⁶ Importantly, the magnitude of the psychic harm need not be a linear function of the probability of the harm. Unlike in classic expected utility theory, when a 100% chance of something good is exactly twice as nice as a 50% chance of the same reward, psychic costs can have different shapes.

If someone has convex psychic costs – e.g., a 1%, or 2%, or 50% chance of harm are all treated like a 0% chance – then the wiggle room result can be obtained.

Again, assume the agent chooses privacy over money in the baseline treatment. That means:

$$v - \pi_H c - \psi(\pi_H) < 0$$

In the information avoidance treatment, she remains ignorant if:

$$\begin{aligned} (1 - p)v + p(0) &> v - pc - \psi(p) \\ -pv &> -pc - \psi(p) \\ pc - pv + \psi(p) &< 0 \\ p(c - v) + \psi(p) &< 0 \\ p(v - c) - \psi(p) &> 0 \end{aligned}$$

Unlike before, an agent with psychic costs might opt for privacy in the *direct* tradeoff treatment, but choose to remain ignorant in the *veiled* tradeoff treatment.

For this to happen, we need two conditions: $v > c$ and a functional form for $\psi(\cdot)$ which is convex. This means that in the baseline case, psychic costs are what is driving the agent to opt for privacy. At the same time, her psychic costs are relatively low when losing privacy is uncertain: a 0.01% chance of losing privacy, or a 1%, or 10%, or 50% chance – all these feel distant from a 100% chance. Consider the classic Star Wars quote when an anxious C-3PO warns the heroic Han Solo about the odds of successfully navigating an asteroid field.⁴⁷ Solo shouts back: “never tell me the odds.”⁴⁸ Here, Solo is like an agent with convex preferences over probabilistic harms. Whether

⁴⁶ See Botond Koszegi, *Health Anxiety and Patient Behavior*, 22 J. HEALTH ECON. 1073, 1074 (2003) (describing a model in which a patient’s utility function is defined by her expectations about her future physical outcomes).

⁴⁷ STAR WARS EPISODE V: THE EMPIRE STRIKES BACK (Lucasfilm Ltd. 1980).

⁴⁸ *Id.*

the probability of crashing is 0.01 or 0.99, he needs to ignore the danger and treat all probabilities as if they are zero. He wants to remain ignorant.

C. Mechanism: Signaling

Another explanation for the experimental result is signaling: people care about privacy, but they also care about *seeming* like they care about privacy.⁴⁹

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 sacrificing privacy. In the *direct* tradeoff group, taking the monetary payoff and rejecting privacy comes with a signaling cost of showing (either to herself or an observer) that the participant does not value privacy. In the *veiled* tradeoff group, meanwhile, taking the monetary payoff without looking at the privacy choices carries no such signaling cost.

In the model, this would mean that the psychic cost of losing privacy depends on how observable her choice is. The monetary value v is the same across groups, but in the *veiled* tradeoff, if the cost of knowingly losing privacy is c , then the cost of losing privacy without being aware of doing so is $c_o < c$. Hence, some people who would choose to keep their privacy in the *direct* tradeoff treatment ($c > v$) would take the money and not click to reveal the privacy settings in the *veiled* tradeoff treatment ($c > v > c_o$). If she gives up privacy without a (costless) veil, the psychic cost is imposed. If there is a veil, then the psychic cost is lower.

D. Mechanism: Psychic Choosing Costs

Some scholars posit that the act of making a choice imposes costs.⁵⁰ The *direct* tradeoff group faces a direct choice between money and privacy, which may be difficult if privacy costs are inchoate or hard to measure. The *veiled* tradeoff group, meanwhile, can opt out of a difficult choice by refusing to consider it. The veil, then, creates a treatment effect by letting people avoid choosing costs.

In the model, this works like signaling in reverse. People in the *direct* tradeoff group face a psychic cost of losing privacy (due to the difficulty of making the choice). If they give away their privacy, they lose cost c , but the

⁴⁹ Zachary Grossman & Joel J. van der Weele, *Self-Image and Willful Ignorance in Social Decisions*, 15 J. EUR. ECON. ASS'N 173, 176 (2017) (concluding that endogenous signaling is one driver of behavior in social situations).

⁵⁰ Cass Sunstein, *Choosing Not to Choose*, 64 DUKE L.J. 1, 40 (2014) (noting that active choice imposes a large burden on the chooser, unlike passive acceptance of a default).

act of choosing imposes a psychic calculation cost c_{choose} . People in the *veiled* tradeoff group, meanwhile, face no such cost unless they click to reveal the privacy settings. The result is that there exist participants who opt to remain anonymous in a control group setting but refuse to unveil – and then give up their privacy – in a treatment group setting. That is, $c > v$, so they choose privacy in the *direct* tradeoff treatment, but c_{choose} is large enough that it is not worth clicking to reveal the privacy settings and choosing to remain anonymous.

E. Mechanism: All of the Above

None of the explanations above are mutually exclusive. They may also be operative, to different degrees in different people. They may interact, so that cases with high choosing costs are also ones where signaling is more powerful. The interactions themselves may differ across people. Hence, the treatment effect might occur for one participant because she finds it unsavory to think about privacy losses, she has a hard time choosing, *and* she really only cares about *seeming* like she cares about privacy. For another participant, the treatment effect might hold because she actually does care about privacy but hates thinking about it, so she does not click to reveal in the *veiled* tradeoff treatment. A different participant might actually want to seem like she is *not* worried about privacy, but also has a hard time making tradeoffs between privacy and money, so the mechanisms would work in opposite directions.

In short, while the existence of a treatment effect from the wiggle room experimental design has been demonstrated for privacy decisions, many mechanisms might be at play. The remainder of this paper turns to exploring these mechanisms with an additional experimental treatment.

III. EXPERIMENTAL DESIGN AND EMPIRICAL APPROACH

304 participants were recruited on Amazon Mechanical Turk to take a short survey about health and financial status. 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.⁵¹ The experiment was conducted on January 7, 2019. The sample of participants was limited to those in the United States.

Research increasingly suggests that, for the purpose of social science experiments, Mechanical Turk users are a reliable sample. One might be

⁵¹ The median wage in the study was \$15.33 (based on a median payment of \$1.02 for a median completion of 3:59 seconds).

concerned about how findings in this population translate to others. Because the setting is Mechanical Turk, one can assume that the sample is quite computer literate and also is comfortable completing short (mundane) tasks for a low wage. However, research suggests that these external validity issues are not of first-order importance. Irvine (2018) replicates three experiments using in-person labs, national online platforms, and Mechanical Turk, and finds that the results are constant across samples.⁵² Nonetheless, as with any experiment, the sample of participants is important to keep in mind when interpreting results.

After recruitment, the timeline of the experiment consists of three stages: instructions, privacy settings, and a survey.⁵³ First, participants were shown an initial introductory screen giving 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 participants 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.

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 persons willingness to answer an intrusive question, a participant in this experiment cannot pretend to give up privacy without actually giving anything up.⁵⁴ Second, the decision is a realistic one. The “Log in with Facebook” button is a

⁵² See Krin Irvine et al., *Law and Psychology Grows Up, Goes Online, and Replicates*, 15 J. EMP. LEG. STUD. 320, 343-44 (2018) (demonstrating the key difference that Mechanical Turk users were significantly more attentive than other samples).

⁵³ For detailed study instructions, please email the author at dsvirsky@uber.com.

⁵⁴ Even if participants have a fake account they can use -- Facebook works hard to limit such behavior, but is not always successful -- 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.

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.⁵⁵

Each person was randomized into one of three treatments during the privacy settings stage: the *direct* tradeoff treatment, the *veiled* tradeoff treatment, and the *choice* tradeoff treatment. The exact format of the privacy choice made in each of the treatments can be seen in Figure 1 (*direct* tradeoff), Figure 2 (*veiled* tradeoff), and Figure 3 (*choice* tradeoff).⁵⁶

<input type="radio"/>	2 cents	High Privacy
<input type="radio"/>	52 cents	Low Privacy
Click here to make your choice		
<input type="radio"/>	2 cents	High Privacy
<input type="radio"/>	52 cents	Low Privacy
Click here to make your choice		

Figure 1: In the *direct* tradeoff group, participants are aware that they are choosing between privacy and money, as both settings are visible by default.

⁵⁵ The privacy measure is less ecologically valid in the sense that it is about sharing data with a researcher, rather than a corporation or a government. It could be that people are more comfortable sharing data with an academic researcher than with Facebook or a police department. The opposite could also be true. In any case, this would cause all three experimental groups to change how they value privacy, but not impact them differentially. One interesting note is that in the Cambridge Analytica scandal, the malicious actors who harvested data posed as academic researchers.

⁵⁶ One contribution of this paper is replication. The first two groups – the *direct* and *veiled* tradeoff groups – face a decision identical to that in Svirsky, *supra* note 3. That paper finds a treatment effect of information avoidance. This paper expands on that paper by adding a third treatment group and is also an opportunity to replicate and retest the initial findings, which is vital for the health of scholarly disciplines that rely on sound experimental findings. *See, e.g.,* Irvine et al., *supra* note 52.

Click here to see the privacy settings

2 cents

52 cents

Click here to make your choice

Figure 2: In the *veiled* tradeoff group, participants choose between privacy and money. The privacy setting is hidden by default but can be revealed instantly and costlessly.

Click here if you'd rather hide the privacy settings and have them randomized

\$0.02 High Privacy

\$0.52 Low Privacy

Click here to make your choice

Figure 3: In the *choice* tradeoff group, participants choose between privacy and money. The privacy setting is visible by default but can be hidden (and randomized) instantly and costlessly.

In the *direct* tradeoff treatment, participants only made one decision: a direct choice between a 2-cent bonus and privacy option A or a 52-cent 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.⁵⁷

In the *choice* tradeoff treatment, participants faced the same layout as in the *direct* tradeoff treatment, but they had the option of clicking a button to hide (and randomize) the privacy settings. Upon clicking, the privacy

⁵⁷ 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.

settings were hidden, and participants faced the same layout (and choice set) as the participants in the *veiled* tradeoff treatment.

In sum, all participants faced the same choice, but depending on random assignment, they faced a different default layout. Some saw everything – money and privacy options – and had to choose directly. Some started off by seeing everything but through active choice could hide the privacy settings. Some started off with privacy settings hidden, and through active choice, could have revealed these settings. Hence, if clicking is costless, there should be no difference between the three groups.

After completing the privacy stage, all participants completed a nine-question survey, shown in Figure 4. Five questions covered demographics, health, and financial topics. These questions asked about the persons 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 screenshot shows a survey interface. At the top, there is a blue button with the Facebook logo and the text "Log in With Facebook". Below this is a text input field for "Age:". 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 4: Screenshot of the survey that each participant completed. The “Log in With Facebook” button only appears if the participant opted to share her Facebook info. If she instead opted for anonymity, the button would not be included.

The demographic questions were selected somewhat arbitrarily, since they were not the focus of the experiment. The goal was to find questions that were somewhat intrusive (that implicate some privacy concerns) without being offensive. The comprehension and Facebook questions help to interpret any results. If a participant does not have a Facebook account, it is hard to interpret her privacy choices. Similarly, if the treatment effect is driven by people who fail comprehension questions, or who use Facebook rarely, then this is informative in understanding what drove any treatment effect.

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.⁵⁸

Even though the choice is essentially 50 cents vs privacy, it is more accurate to note that this is a 50-cent *bonus*. The participants are foregoing 50 cents, not actually giving away any of their pre-experimental wealth. There is extensive behavioral economics literature noting the distinction between losses and gains.⁵⁹ This point is broadly important but has little relevance here. Participants would be more likely to opt for money over privacy if the monetary change were a loss rather than a gain, but this would affect all three experimental groups equally.

IV. RESULTS

The results are organized as follows: Section A gives summary statistics and balance checks, while Section B shows the primary findings – the average treatment effects as compared to how often each group chose to remain anonymous, as well as how often the *veiled* and *choice* tradeoff groups clicked to hide or reveal the privacy settings.

⁵⁸ All code is available on request from the author and includes survey instructions, experimental module coding, and the raw data. Contact the author for the ZIP file: dsvirsky@hbs.edu.

⁵⁹ See, e.g., Botond Koszegi & Matthew Rabin, *A Model of Reference-Dependent Preferences*, 121 Q. J. OF ECON 1133, 1134 (2006) (expounding on prospect theory as applied to consumer behavior).

A. Summary Statistics

There were no systematic demographic differences between the treatment groups, as expected given the random assignment. Of note, 90% of participants reported having a Facebook account, and the median participant used Facebook three times per week.

	<i>Direct</i> Tradeoff (N = 108)	<i>Veiled</i> Tradeoff (N = 109)	<i>Choice</i> Tradeoff (N = 87)	P- Value
Age (years)	34.15 (10.24)	33.41 ^I (9.112)	34.82 (10.40)	0.61
Diet Attempts in Lifetime (0 – 4)	2.102 (1.646)	1.954 (1.512)	2.517 (1.547)	0.04
Exercise Workouts in a Typical Week (0 – 4)	2.213 (1.454)	2.358 (1.385)	2.276 (1.476)	0.76
Annual Income (0 – 4)	1.519 (1.196)	1.385 (1.053)	1.494 (1.160)	0.66
Credit Card Debt (0 – 4)	0.815 (1.051)	1.018 (1.097)	0.839 (1.066)	0.32
Has Facebook (0,1)	0.917 (0.278)	0.890 (0.314)	0.874 (0.334)	0.61
Weekly Facebook Use (0 – 4)	2.491 ^I (1.568)	2.541 ^I (1.549)	2.793 (1.526)	0.36

Table 1: Summary statistics. Standard deviation reported in parenthesis. With the exception of age, which is reported in years, each variable is categorical. Hence, an answer of 1 for credit card debt corresponds to a range of \$1000 to \$2000 in debt.

B. Main Results: Average Treatment Effects

Do either of the two treatments lead people to choose privacy more often? When given the option to hide or reveal information, do participants do so?

In all three treatments, participants faced the same choice: participants were offered 52 cents to share their Facebook data, or 2 cents to preserve their anonymity. The only difference was in the default information presented. Nonetheless, I find a significant impact on people's

willingness to sell their data. Roughly 70% of people in the *direct* tradeoff group opted to remain anonymous. In the *veiled* tradeoff group, only 40% remained anonymous. These numbers are almost identical to those found in Svirsky (2018).⁶⁰ Meanwhile, however, participants in the *choice* tradeoff group – who saw both money and privacy settings but had the choice to hide the privacy settings – were halfway between the other groups. Roughly 56% of participants in the *choice* tradeoff group opted to remain anonymous: less than in the *direct* tradeoff group, but more than in the *veiled* tradeoff group. Figure 5 presents the results graphically. Table 1 shows the results of regression models where {ended up staying private} is the binary dependent variable, and there are indicator variables for the *veiled* and *choice* tradeoff groups. Each column presents a different sample, each one representing a robustness check.




Figure 5: This figure shows the proportion of participants who ended up remaining anonymous for 2 cents instead of sharing their Facebook profile for 52 cents, for the *direct* tradeoff group (N = 108), the *veiled*

⁶⁰ See Svirsky, *supra* note 3, at 1 (finding that online survey participants had to make the same choice whether to share their Facebook profile data with the survey taker in exchange for a higher payoff).

tradeoff group (N = 109), and the *choice* tradeoff group (N = 87). These results exclude all participants who, by randomization, faced a degenerate tradeoff of 52 cents and high privacy vs 2 cents and low privacy. Therefore, for the *veiled* tradeoff group, 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.

	Entire Sample	Excludes People Who Fail Comprehension Check	Excludes People Who Did Not Answer Carefully	Excludes People w/o a Facebook account
<i>Veiled</i> Tradeoff	-0.30*** (0.06)	-0.29*** (0.07)	-0.29*** (0.07)	-0.32*** (0.07)
<i>Choice</i> Tradeoff	-0.13* (0.07)	-0.13* (0.07)	-0.13* (0.07)	-0.13* (0.07)
Constant	0.69*** (0.05)	0.71*** (0.05)	0.70*** (0.05)	0.69*** (0.05)
N	304	264	294	272
Adjusted R ²	0.06	0.05	0.06	0.07

Table 2: Regression models of average treatment effect. The dependent variable is a binary variable for whether the person ended up remaining anonymous. There are indicator variables for the *veiled* and *choice* tradeoff groups, so the constant represents the proportion of participants in the *direct* tradeoff group who ended up remaining anonymous. Each column uses a different subset of the sample in order to provide robustness checks. *** $p < 0.001$, * < 0.10 .

Participants in the *choice* tradeoff group by and large did *not* hide information about privacy. In the *choice* tradeoff group, only 9% of participants made an active choice to hide (and randomize) the privacy settings before making a choice. In the *veiled* tradeoff group, where privacy settings were hidden by default (but could be revealed), 45% of participants made a choice without seeing the privacy information. This difference in proportions is statistically significant (Fisher's Exact $p < 0.001$).

CONCLUSION

This paper explores why people avoid information about privacy when making data sharing decisions. Existing work demonstrates that even when privacy settings are easy to read – even just two words long – people who otherwise would pay several dollars to remain anonymous are happy to avoid looking at the settings and take a 50-cent bonus. This paper solidifies this behavior. It finds that while people are happy to avoid information that is already hidden, they are not likely to actively hide information that is in front of them to begin with. At the same time, the option of hiding information makes people marginally more likely to sell their data, even if they do not choose to hide the privacy settings.

The results give more support to certain mechanisms of information avoidance than others. Theories that rely on signaling are consistent with the data presented here. If people care more about *seeming* like they value privacy, then they might take the money if given plausible deniability (as in the *veiled* tradeoff group), but not go so far as to actively hide information (as in the *choice* tradeoff group), as such a choice would signal a willingness to care little about privacy.

Theories that posit that people simply prefer not to think about privacy, or prefer not to choose, are less consistent with the data. If these mechanisms explain information avoidance, then people would opt to simplify their choice if given the option.

The results also suggest that current U.S. privacy law – centered around giving consumers better information – may be difficult to achieve in practice. There is considerable scholarship and policy experimentation around giving people simpler, more effective disclosures.⁶¹ Simpler disclosures is likely a good thing: if it gives people more information at lower costs, this should improve welfare. At the same time, if many people choose to avoid information about privacy, then better disclosures will not be as effective as a classical economics model would suggest. If societies want people to end up with more privacy, it will be difficult to do so by relying on individuals to seek out the information they need and choose accordingly.

⁶¹ See, e.g., Chilton & Ben-Shahar, *supra* note 12, at 1 (describing the failure of simplified privacy disclosures to effect meaningful change in participants' behavior in disclosing private information); Corey Ciocchetti, *The Future of Privacy Policies: A Privacy Nutrition Label Filled with Fair Information Practices*, 26 J. MARSHALL J. COMP. & INFO. L. 1 (2008-2009) (discussing standardization of labels to force all e-commerce homepages to conspicuously post their privacy practices).