Sprigman, Christopher Jon; Tontrup, Stephan

Working Paper

Privacy Decision-making and the Effects of Privacy Choice Architecture - Experiments Toward the Design of Behaviorally-Aware Privacy Regulation

Law and Economics Research Paper Series Working Paper, No. 23-22

NEW YORK UNIVERSITY

SCHOOL OF LAW

LAW AND ECONOMICS RESEARCH PAPER SERIES WORKING PAPER NO. 23-22

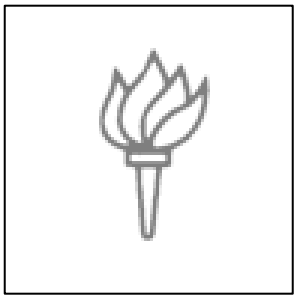

Privacy Decision-making and the Effects of Privacy Choice Architecture - Experiments Toward the Design of Behaviorally-Aware Privacy Regulation

Christopher Jon Sprigman and Stephan Tontrup

March 2024

Privacy Decision-making and the Effects of Privacy Choice Architecture - Experiments Toward the Design of Behaviorally-Aware Privacy Regulation

Christopher Jon Sprigman

Murry and Kathleen Bring Professor at New York University School of Law

Stephan Tontrup[¥]

Law and Business Fellow at New York University School of Law

Correspondence may be directed to: Stephan.Tontrup@nyu.edu

Abstract

The current notice and choice privacy framework fails to empower individuals in effectively making their own privacy choices. In this Article we offer evidence from three novel experiments showing that at the core of this failure is a cognitive error. Notice and choice caters to a heuristic that people employ to make privacy decisions. This heuristic is meant to distinguish between a party's good or bad intent in face-to-face-situations. In the online context, it distorts privacy decision-making and leaves potential disclosers vulnerable to exploitation.

From our experimental evidence exploring the heuristic's effect, we conclude that privacy law must become more behaviorally aware. Specifically, privacy law must be redesigned to intervene in the cognitive mechanisms that keep individuals from making better privacy decisions. A behaviorally-aware privacy regime must centralize, standardize and simplify the framework for making privacy choices.

To achieve these goals, we propose a master privacy template which requires consumers to define their privacy preferences in advance—doing so avoids presenting the consumer with a concrete counterparty, and this, in turn, prevents them from applying the intent heuristic and reduces many other biases that affect privacy decision-making. Our data show that blocking the heuristic

enables consumers to consider relevant privacy cues and be considerate of externalities their privacy decisions cause.

The master privacy template provides a much more effective platform for regulation. Through the master template the regulator can set the standard for automated communication between user clients and website interfaces, a facility which we expect to enhance enforcement and competition about privacy terms.

I. Introduction

Between 2013 and 2015, about 270,000 Facebook users allowed an app called "thisisyourdigitallife" to access and collect data from their Facebook profiles. Developed by a University of Cambridge professor, the app was advertised to users as a personality test based on an analysis of the user's Facebook activity. Individuals accepting the app's terms of service were also (whether they understood it or not) permitting it to collect data about their Facebook friends. This led to a huge multiplier effect: In total, the app collected data from up to 87 million Facebook users—the vast majority of whom had never downloaded "thisisyourdigitallife." Disclosure of this data created substantial privacy risks: analysis of Facebook "likes," for example, can provide information about an individual's personality traits, sexual orientation, political views, mental health status, and substance abuse history. And then there was a separate risk that no user could possibly have anticipated. A company called Cambridge Analytica later used the data to build voter profiles it gave to the Trump campaign, which used them to target political advertisements in the 2016 U.S. presidential election.

Why do so many people hand over so much information to online platforms like Facebook? In this Article we will argue that the answer has a lot to do with the way the current "notice and choice" framework of U.S. privacy law frames privacy decisions.

Notice and choice requires that firms disclose what private information they are collecting and how they will use it, with the expectation that consumers, having read and understood the notice, will choose whether to consent (Strandburg, 2013; Solove, 2013; Obar & Oeldorf-Hirsch, 2020).[1] One foundational problem with this approach is that people often do not read privacy notices,[2] and those who do not read cannot make informed privacy decisions. But while not reading notices suggests that the regulation has little influence on actual privacy decisions, it does not by itself explain why consumers hand over sensitive personal information to unknown online

[1] It would take seventy-six days to read all the privacy policies an average person encounters in a year, with an annual opportunity cost of an estimated $781 billion (McDonald & Cranor 2008).

[2] *See* Pew Research Center, 2019, finding that only 22% of Americans report reading privacy policies "always" (9%) or "often" (13%), while 74% read them "sometimes" (38%) or "never" (36%)). This data may be optimistic. Research on the actual internet browsing behavior shows that only 1 or 2 out of 1000 retail software shoppers access the online end user license agreement (EULA). Marotta-Wurgler & Trossen, 2014). There is little reason to believe that consumer attention to online privacy policies is significantly greater than for EULAs.

actors. Individuals readily reveal information that, under ordinary circumstances, they would share only with their closest family and friends, if at all.[3]

This Article provides evidence that may help better understand how people make privacy decisions and explains why so many seem to disregard personal and social privacy risks. We report the results of three experiments that together begin to illuminate how "notice and choice" impacts the decision-making tools individuals use when making privacy decisions. Notice and choice frames privacy choices as part of a bilateral relationship between the user and the platform or service requesting disclosure. This framing triggers a form of social heuristic decision-making that emerged in and is often successfully employed in real-world face-to-face interaction. The heuristic helps navigate social relationships by estimating the good intent and personal trustworthiness of the partners a person is interacting with. Given its focus we will refer to the heuristic as the "intent heuristic." To identify intent the heuristic looks for a situation in which the partner has multiple choices available to her and must then decide which option to pick from that set of options. If she picks an action more beneficial for the counterpart while more costly to her, compared to the other available options, she reveals good intent. If she chooses an action less advantageous for her counterpart, by comparison to the other available options, she reveals bad intent. If the heuristic suggests that the partner has good intent, then the decision maker trusts her in the interaction. We will see that the intent heuristic is used in many situations of social decision making, in fairness judgments for example, and it also determines positive or negative reciprocity. The prevalence suggests that intent is in general a valid heuristic to navigate social interaction (Schaechtele, et al., 2011; Cushman, et al., 2009).

In the online world however, we will show that the heuristic and its foundation in trustworthiness based on perceived good intent is maladaptive. There are no real personal relationships in this domain: instead people interact online with algorithms and multiple third parties with unobservable choice options. So the heuristic has nothing real to refer to: i.e., no identifiable single intent. For instance, Facebook may (honestly or not) represent that it will not sell private consumer data. However, Facebook may rely on and share data with third party contractors who may also share data with other third parties. The more parties are involved the larger the risks that data is revealed and the less control is exercised by any single party. Each of these parties may breach or exploit lax data sharing provisions and impose externalities. Basing privacy decisions on the perceived intent of the one party the data is directly disclosed to, does not capture this complex online reality.

This example lets us draw a distinction between two alternative notions of trustworthiness: we distinguish between trust placed in good intent and trust placed in institutions. The heuristic leads to trust in a *party's* perceived intent or benevolence, not in the institutional framework

[3] There is a substantial literature critiquing data privacy regulation premised on a one-time exercise of informed individual choice. A portion of this literature goes beyond skepticism of the effectiveness of notice, and extends the critique to various ways in which individualized consent fails to secure anything close to a social optimum with respect to data disclosure and use. *See, e.g.*, Nissenbaum, 2010; Regan, 1995; Cohen, 2013; Richards & Hartzog, 2019; Bietti, 2020. Another portion focuses on technical aspects of the design of particular digital environments, including so-called "dark patterns," that encourage disclosure. *See, e.g.,* Waldman, 2020; Hartzog, 2018. On dark patterns, *see* Narayanan et al., 2020. For a detailed review and categorization of the privacy literature, *see* Viljoen, 2021.

surrounding her, or in the legal regulation that applies to the institution or the technological safeguards linked to the institution.[4]

Trust in the institutional framework by contrast refers to trust as *reliability*. Reliability is a function of the legal and technological environment in which a party operates: for example, the technology and the partners that help the business provide services it may or may not have control over. So a person may not trust Facebook to protect personal data by deliberate choice, but because she believes (correctly or otherwise) that the technology Facebook uses protects her privacy, or the government authorities that regulate Facebook will safeguard it. In this hypothetical consumer's view, Facebook is trustworthy because it operates in a privacy-protective legal framework and technologies are in place that enable Facebook to fulfill these legal requirements. If these conditions obtain, it is the framework that makes Facebook reliable, not individual intent.

In an online environment, individuals contemplating disclosure of private information should rationally look for and be responsive to signals of trustworthiness in the institutional framework within which services or platforms operate. Then individuals may rely on these signals to estimate the risks of disclosure and weigh them against the benefits of using the service. Use of *personal* intent would be less impactful if signals of institutional trustworthiness were considered in this way.

Yet, heuristic decision-making is not an overall assessment of risks and benefits. It is meant to be fast and frugal and economizes by limiting the information processed. Often it focuses on a few cues, sometimes only on a single cue and gives it significant weight, crowding out other valuable information. And indeed our experiments will show not only that many users rely on the intent heuristic in their privacy decisions, but that use of the heuristic also largely blinds users to potentially valuable signals of institutional trustworthiness.

Moreover, we will show that the heuristic also leads people to put little weight in their decision making on information about the social component of privacy—that is, information about on how their personal privacy choices affect the privacy interests of others. In a social network like Facebook, an individual's decision whether to disclose information will often affect others' privacy interests—for example, the privacy of other people whose image is captured in a photograph the user posts on his timeline. The same is true in the other direction—any individual Facebook user's privacy is a product, in part, of other Facebook users' decisions.

Our finding that people base their privacy decisions largely on the intent heuristic explain why many seem to reveal volumes of personal data without being sensitive to structural privacy risks online. We show in our study that when subjects trust in the good intent of the party they are revealing to, they seem to not consider further risks. Online the heuristic is also receptive to manipulation. Facebook for example invests strong efforts in focusing users on their friends as

[4] The heuristic's decision rule demonstrates the difference between the two concepts of trust: the heuristic in order to be used needs to refer to a voluntary decision between choice options, not a decision that is determined by technology or regulation. Only because her choice is unrestricted, the party whose intent is assessed can demonstrate benevolence leading the consumer to conclude she can trust the party.

the parties they disclose data to. If they trust their intent, they reveal information with little regard to the fact that Facebook is analyzing their data and using it to sell targeted advertisements.

We designed a series of three experiments to illuminate the impact of the intent heuristic on how people make privacy decisions in online settings:

In experiment 1, we show evidence that subjects do in fact employ the intent heuristic to make privacy decisions. We designed our experiment as a tag-along to an unrelated study (Tontrup & Sprigman, 2022). When subjects assumed the experiment was over we randomly assigned them into different treatments with varying payment options—some requiring substantial disclosure of confidential financial information, others not.

If we offered a privacy-protective payment method (like an anonymous cash pickup), as opposed to offering only a bank transfer, subjects were more likely to disclose their bank details. In the second treatment, subjects were offered the same payment options but were told that we were *obliged* by university rules to offer them the anonymous payment. In this scenario, subjects were prevented from using the heuristic: because we *must* provide the privacy-protective option, they cannot view our provision of the protection as indicating our good intent, and indeed we found participants to be significantly less likely to disclose their bank data.

In Experiment 2 we investigated whether subjects employing the intent heuristic were more likely to disregard valuable information that could improve the accuracy of their privacy decision. In one treatment, we provided subjects with credible background information, such as the researchers' university affiliations, past work history, and other indicia of trustworthiness while in the second we gave subjects no information at all except the Gmail (i.e., non-university) address of a research assistant. We then varied, for each treatment, whether participants could apply the heuristic.

The results show that when subjects can employ the heuristic, whether they received a positive or a negative signal did not change the rate of disclosure—subjects in the grip of the heuristic did not seem to pay attention to otherwise credible information about trustworthiness. In contrast, when we made it impossible for subjects to use the heuristic subjects paid attention to the signals: they were significantly less likely to disclose their financial data when not given credible information about their counterparty's trustworthiness, and more likely to disclose when they received positive information, as should rationally be expected.

In Experiment 3 we analyze whether the intent heuristic leads subjects to disregard the *collective risks* of their privacy choices. We employed a web application (specifically, the Osint "Social Links" app[5]) that allowed us to elicit all publicly accessible information associated with subjects' Facebook accounts.

In the first of two treatments we offer subjects privacy protection; they can register for our studies while remaining anonymous and without providing Facebook information. In the second treatment, we offered subjects the same privacy protection but prevented them from using the intent heuristic. In the first treatment more than half of the subjects allow us to harvest their

[5] *See* for more detail Section III.C.2.

Facebook data, while when we prevented subjects from using the heuristic, they were significantly less likely to give us access and the subjects who revealed their data have *less Facebook data to reveal*, on average, compared with those who decide not to reveal. By comparison, in the first treatment where subjects can employ the heuristic, we see not only more subjects disclosing their data but also less of a difference in disclosure rate between subjects with a large amount of Facebook data and those who have less. Subjects who have more Facebook data create larger privacy risks for others when they disclose that data. Thus, Experiment 3 demonstrates that blocking subjects' heuristic decision-making can make their choices more sensitive to the privacy risks they impose on others, and that others impose on them, in online settings.

Implications of our Findings for Privacy Law and Policy Reform. Our findings provide guidelines for a behaviorally-aware redesign of privacy law and policy to better take account of how individuals actually make their privacy choice.

Our findings suggest that people are not reckless in imposing privacy externalities on others or indifferent to their own privacy; instead, their behavior is driven by the structure of a privacy law tailored to bilateral relationships and the use of the heuristic it induces.

Where intervention succeeds in blocking the intent heuristic, people are more receptive to credible institutional and technological cues in assessing risk and they are more considerate of and less likely to impose privacy externalities on others.

We suggest that any effective redesign of privacy law that can prevent people from using the heuristic requires that consumers make privacy choices in advance of the disclosure requests. This in effect means that privacy law must take the design of privacy requests out of the hands of data-seekers, and move toward a framework that elicits privacy preferences in advance, and within a standardized menu of options.

A standardized, ex-ante approach has two goals: first, it pervasively disrupts the notice and choice regime's framing of privacy decisions as bilateral transactions—a framing on which the heuristic depends. But it is not only the bilateral framing that is disrupted. Use of the heuristic is also disabled in a very practical sense: If privacy decisions are made in advance for various data uses and or purposes that affect all potential future data seeking parties, there are no actions or choices the heuristic could process to identify the intent of any particular counterparty.

In our re-setting of privacy law, requesters seeking disclosure would query a database or user client containing a person`s privacy choices. The client would present a centrally designed menu of privacy options with simple privacy rules such as "don't sell my data" or "don't store my data." It could use nudges such as opt-in defaults for data sharing and decision trees for identifying and responding to various potential categories of data usage. The design of these consent templates would be standardized and centralized under an agency, such as the FTC, which could mandate compatible interfaces between those who collect data and those who share it. Information would be shared when a data disclosure request and the discloser's preferences match. Any request for data beyond that consent would require additional permissions channeled through the template. In this way, the law could create a new equilibrium in privacy decisions in which companies asking for more disclosure must do so against a default already established by the potential discloser.

In recent years, there has been a trend toward substantive privacy legislation in both the U.S. and Europe. The new laws grant rights to access and delete data, and to mandate secure data processing, categories of more sensitive data that they regulate more strictly. Consent and notice, however, remain cornerstones of both the GDPR and California's new privacy laws. They assure consumers that their privacy cannot be compromised without their consent. But for consent to provide legitimacy, it must be meaningful. We argue that in order to renew the original notice and choice idea of better-informed privacy choices, a behavioral reform is needed that enables consumers to make their privacy decisions ex ante and without facing a particular counterparty, and prevents them from using the intent heuristic.

In Part II we describe our study's theoretical framework. In Part III we report experimental method and results. Part IV discusses the implications of our findings for privacy regulation.

II. The Cognitive Background of Privacy Decision-making

A. Rational Privacy Preferences and Cognitive Errors

Much of the academic research on privacy has, like the origins of U.S. privacy law itself, been based on a rational choice framework. That framework assumes (1) that people act on their stable preferences to make value-maximizing trade-offs between privacy and other concerns (Derlega et al., 1993; Posner, 1981; Rosenfeld, 2000; Stigler, 1980), (2) that decisions whether to disclose private information are made by balancing "the usefulness of privacy with the utility of openness (Petronio, 2000), and (3) that personal information is revealed only when the individual considering disclosure expects a net benefit (White, 2004). This perspective assumes that individuals can be relied on to rationally decide between disclosing or secreting their personal information.

But people's actual privacy choices appear to contradict this account. In a poll conducted by Gallup in 2013, 83% of American internet users surveyed reported either being "very concerned" (48%) or "somewhat concerned" (35%) about privacy (Gallup, Computers and Internet). And yet, a growing body of empirical evidence shows that individuals will voluntarily turn over volumes of personally identifiable information in a range of settings, including when using social media websites (Smith, 1994). Indeed, some evidence suggests that when consumers are asked to pay for privacy protections, few seem to be prepared to do so. In a field experiment, participants were given the choice to buy a DVD from one of two online stores that were identical except that one required the disclosure of substantially more sensitive data than the other (Beresford, 2010). In a first treatment, DVDs were offered at a price one Euro cheaper in the store requesting more personal information; almost all buyers chose the cheaper store (Id. at 3). In a second treatment, prices were identical and participants bought from both shops equally often, suggesting that subjects placed no value on increased privacy. Another study measured users' app choices when a commercial and a free version were offered. Few users were willing to pay for the app, even though most of the participants expected the paid app to be less likely to share their data with advertisers. And only 6% of the users who paid for the app said they paid to improve their privacy (Bamberger, 2020).

One interpretation of the inconsistency between individuals' stated preferences and their disclosures is that they may have unstable or simply weak preferences about privacy. In particular, studies have shown that individuals' preferences often appear to be unstable; that is, affected by the framing of alternatives (Dhar & Simonson, 1992; Kahneman & Tversky, 1974) and by the elicitation method employed (Tversky et al. 1990). Preferences tend to be most unstable when the costs of a choice are abstract and uncertain (Fox & Tversky, 1995; Griffin et al., 2005; Hsee, 1993; Hsee et al. 2003), which may often be the case for privacy decisions, especially when the risks are difficult to assess (Acquisti, 2004). In addition, the literature identifies particular psychological distortions, problems with self-control including hyperbolic discounting (i.e., the tendency to choose smaller, immediate rewards, while incurring larger privacy risks realized later), the tendency toward optimism bias (i.e., the underestimation of the risk that personal privacy disclosures may cause harm), the difficulty of estimating the weight of cumulative privacy risks (Acquisti, 2004), and a "control paradox" whereby people who experience more perceived control over some particular aspect of their privacy sometimes respond by mistakenly generalizing that perception of control and revealing more information than they would otherwise (Brandimarte et al., 2012).

Together, these biases undermine rational choice as an explanatory framework for privacy decisions. They also suggest that a legal approach solely based on notice and choice in its current form was always bound to fail.

However, if cognitive mechanisms are part of the explanation for the variance between expressed privacy preferences and real-world privacy behavior, the picture may not be quite so dire. Cognitive errors can often be mitigated by rearranging the decision-making environment. Yet, for consent to be part of an effective privacy regulation, it must be built on different premises than the current form of notice and choice.

B. Trust and the "Heuristic" Model of Privacy Decision-making

How do people decide whether to disclose personal data online? We think they rely on decision mechanisms they also use in other social domains and therefore look for guidance to the vast economic literature examining fairness and reciprocity in individual decision-making (for example Fehr & Falk, 1999; Bewley, 1999; Fehr et al., 1997). The economic literature suggests that when deciding whether an action was fair or unfair, or whether to reciprocate or not, individuals appear to base their judgment not only on the outcome they received, but on the perceived intent of their interaction partner (Falk et. al., 2003; Dufwenberg & Kirchsteiger, 2004; Rabin, 1993; Charness & Rabin, 2002; Güth & Kirchkamp, 2012. Berg et al., 1995; Roth et al., 1991; Cameron, 1999).

Many studies use the "ultimatum game" for analyzing fairness preferences. In this game, one player, the "proposer," is endowed with a sum of money. The proposer must propose a split of the endowment with a second player, the "responder" and the responder may accept it or reject the suggested split. Both players know in advance that if the responder accepts, the money is split according to the proposal, but if the responder rejects, both players receive nothing. If both proposer and responder are and expect each other to be purely rational, the proposer should make the lowest possible offer and the responder should accept any split that gives him more

than zero. But participants almost never behave in this way; rather, fairness preferences or at least anticipated fairness preferences on the side of the responder lead proposers to make more generous offers, and responders to reject (even at cost to themselves) offers that they perceive as unfair (Güth, 1982; Thaler, 1988; Cameron, 1999).

Perceived intentions play an important role in these fairness decisions. One study limited the offers available to the proposer; she could only choose between two low offers. While such a low offer would likely otherwise be rejected, the responder will often accept the offer if she assumes that the proposer made the fairest offer *that was available to her* (Falk et al., 2003). In other words, identical actions may signal entirely different information about the proposer's intent and the interpretation of that intent depends on the responder's perception of what actions were available to the proposer. For that reason, responses to proposals offering identical payoffs may vary, even though responders' preferences about the payoff itself are stable. [6]

In a series of experiments, Schaechtele et al. (2011) find strikingly that the effect of intention on outcomes was at least as large as the effect of expected payoff. In other words, seemingly good intentions are so important for fairness decisions that they can, and often do, outweigh a comparatively bad payoff expectation. These observations align with a simple decision making strategy: the individual looks for a voluntary choice of action of the party she wants to ascribe good or bad intent to. A voluntary choice requires that more than one action is available to the party, so the party can choose the one or the other. For the individual who ascribes good or bad intent, these actions must have different value, one must benefit her more than the other(s). And the action that benefits the individual more must be more costly for the party performing it than the alternative action that benefits the individual less. If these conditions are met, the individual draws a straightforward conclusion: if the party chooses the more costly action that benefits the individual more, the individual ascribes good intent to the party; if by contrast the party chooses the action that is less costly for her and benefits the individual less, than the individual ascribes bad intent. If used in the privacy setting, the individual would be more like to trust the good-intent party with her data, relative to the party she ascribes bad intent to.

One crucial part of the decision making strategy is what Schaechtele et al. refer to as "constructing the action space". That means the individual must identify the set of actions that she believes (correctly or incorrectly) the counterparty is free to take in the context of their interaction. But it can be difficult to distinguish what actions were truly available to a party: in real word contexts the true action space is often hidden and parties may also deliberately misrepresent their action space.

If individuals use this heuristic to make their privacy decisions, that could account for why people seem to act as if they have inconsistent privacy preferences. A person may express strong preferences for privacy protection and yet be willing to disclose sensitive information to a counterparty she perceives as manifesting positive intent. The heuristic may drive seemingly

[6] The evidence how strongly perceived intentions may affect decision-making independently from expected payoff is mixed. Cushman et al. (2009) found expected payoffs to dominate intentions, whereas Charness and Levine (2007) report that perceived intent dominates and was even more important than actual payoffs.

inconsistent privacy choices in transactions that impose similar objective privacy risks and without the individuals' privacy preferences changing.

In our study, we draw on these insights and analyze whether perceived good or bad intent do in fact drive privacy decisions. Perceived good intent may lead subjects to disclose personal information; bad intent may cause them to abstain from sharing data. Specifically, we gather novel experimental evidence supporting our hypothesis that people facing privacy decisions employ a decision-making strategy that allows them to ascribe good or bad intent to the party requesting disclosure. The subject decides whether the intent of the party requesting information is benign or not in a way that is analogous to the behavior of subjects in the ultimatum game: they compare the terms the counterparty offers for their privacy protection with the terms that the subject believes are available to the counterparty. If the counterparty's proposal is more beneficent than other offers she could have made, then the subject is (1) likely to ascribe good intent to the counterparty, and, (2) more likely to entrust her with her personal information. If, by contrast, the counterparty fails to offer protections that would have been available to her, then the subject is likely to ascribe a bad intent to the counterparty, and will be correspondingly less likely to trust that the counterparty will protect her data in good faith and also less likely to disclose information.

III. Experiments

We designed three experiments to explore our theory that individuals' privacy choices are often based on heuristic decision-making, and, if indeed that is the case, to understand the implications for privacy in the online environment.

A. Experiment #1: The Intent Heuristic and Risk Perception

1. Motivation

In the real world, trust in another' party's good intent may grow out of friendship, past interactions, mutual disclosure of information, or many other factors. But in the online world people often have no (long-term) relations with those they interact with and information about counterparties may be scarce or costly to access. If people interacting online base privacy decisions on personal trust, they must use less individualized and context-rich criteria to decide whether they should trust someone seeking their private data. As explained, we assume that people employ an intent heuristic to make that decision. If we were able to identify this heuristic as a mental strategy people often use to make privacy decisions, then regulators could design interventions that would interfere with this decision-making process—most likely by preventing people from using the heuristic and attempting to induce privacy holders to look for information more relevant to privacy decisions.

Our first experiment aims to provide evidence that individuals indeed employ the intent heuristic when making privacy decisions. For a systematic test, we implement a set of three treatments: In the first, we have constructed a setting in which we expect subjects if they employ the heuristic to conclude that their counterparty is trustworthy. In the second treatment use of the heuristic should suggest the counterparty is not trustworthy and subjects should therefore be less likely to disclose information. Finally, in our third treatment we switch the heuristic off

by restricting the action space. The heuristic exploits whether the decision-maker chooses the more or the less benevolent of her options. However, if the decision maker has only one option, for example because a specific behavior is legally required or prohibited, the heuristic can, in effect, not be applied. Once deprived of the heuristic, we expect subjects to look for cues other than perceived intent to estimate their actual privacy risk and be less likely to disclose their data. The experiment will thus show that people use the heuristic and that it makes them more readily disclose their data. On the policy end, it shows that blocking the heuristic might be an effective yet simple and cheap way to push people toward making more fully-considered privacy decisions in settings.

2. Experimental Design

We used as a platform for Experiment 1 an unrelated online study to which subjects had previously been invited. The unrelated study offered subjects the opportunity to enter into a contract and to complete a real effort task (counting numbers in a table) in return for a payment.[7] After subjects had completed the real effort task, the experiment was, in their perception, over.[8] They entered the experiment's payment screen unaware that the choices they would be making were the object of the study we report here.[9] This element of fieldwork increases the reliability of our findings (List & Harrison, 2004).

We presented subjects with a choice of different payment methods. Two methods required disclosure of confidential information while one fully preserved subjects' anonymity. The first, bank transfer, required that subjects provide their full name and banking information.[10] The second, PayPal, required subjects to reveal their name and PayPal address. When we invited subjects, we made it a condition for participation in all treatments that subjects had a PayPal account (or opened one before participation) associated with an email address that revealed their true name. The third method, anonymous payment, required no disclosure of personal information. Subjects were asked to specify a 5-digit code after the experiment. With this code, but without revealing their name or any other information linked to their identity, subjects could pick up their payment in the office of the University's student government. This payment protocol ensured that experimenters neither learn the names of the subjects nor can connect them with their choices in the experiment.

The bank transfer payment was fast and convenient, but the anonymous payment was designed to impose significant transaction costs. Subjects choosing anonymous payment had to make a trip of fifteen minutes or more to the student government office and then invest an additional five minutes in order to be paid. These transaction costs assure us that the participants

[7] The study demonstrated that people can employ their own loss aversion to improve effort and reach their work goals—an ability we refer to as "behavioral-self-management." *See* Tontrup & Sprigman, 2022.

[8] *Id.*, at 599.

[9] In our data analysis, we control for whether the real effort task has an influence on the subjects' payment choices. For example the amount subjects have earned may influence which payment option they select, or how the way the contractual threshold was set may have an impact on their choices.

[10] Subjects choosing bank transfer were required to indicate their name, the name of their bank, and their account number.

did not randomly choose anonymous payment, but that they chose it in order to preserve their privacy and had a willingness to pay with their time for this protection.[11]

Our experimental treatments vary subjects' choice set of payment options. We randomly assigned participants to one of four treatments. All treatments present participants with two specified payment options to choose between. They also permit participants to opt out of these two fixed options. Subjects who opted out were required to provide an anonymous email address for experimenters to work out an individual payment solution with them (we provided a link to a service that offers anonymous email accounts to users at no charge).

While bank transfer was offered in each treatment, our manipulation varied the second payment option. This second option was either anonymous payment or PayPal, which required subjects to disclose personal information. The four treatments were structured as follows:

Negative Intent. Participants chose between (1) bank transfer and (2) PayPal. Anonymous payment was not offered; subjects had to indicate their full name and then either their bank details or an email address registered with PayPal. As in each treatment, subjects could opt out and reject the two offered payment options and arrange an alternative payment solution with our assistant. Because the counterparty could have chosen to offer an anonymous payment method but did not, the counterfactual action appears to be more benevolent than the action the counterparty did offer. The treatment is therefore designed to suggest that the counterparty is *not* trustworthy and we expect subjects to ascribe a negative intent to their counterparty and to be hesitant to disclose confidential information.[12]

Positive Intent. In the second treatment, subjects were offered (1) bank transfer and (2) anonymous payment and were thus empowered to conceal their identity. Again, subjects could opt out of these two specified options and arrange an alternative solution. Because the experimenter's available action space in this treatment includes the possibility not to offer anonymous payment but this option is in fact offered, the counterfactual action should appear less benevolent to subjects than the offer the experimenters actually made. Therefore this treatment is designed to suggest that the counterparty is trustworthy, and subjects should more willingly disclose private information.

Legal Requirement. In the third treatment, *Legal Requirement*, we offered participants the option of (1) bank transfer and (2) anonymous payment—the same two payment options as in the *Positive Intent* condition—and subjects were also able to opt out of these specified payment methods (as in all conditions). The payment options offered thus should, as in *Positive Intent*, suggest the counterparty is trustworthy. However, we instructed the subjects that the data protection policy of the University of Münster obliges all experimenters to allow participants in

[11] Given that student jobs in Münster pay approximately €8 per hour, the price (framed as an opportunity cost) for protecting their privacy they were willing to pay in the form of transaction was about €2.20.

[12] Even though subjects were offered the option to arrange an alternative payment method with the assistants, it is possible that some subjects did not consider that experimenters could have offered an anonymous payment option. These subjects may have thought that the experimenters did not have options available to better protect their data. In that case they would not have been able to use the heuristic (see LR treatment below). Since we are analyzing treatment effects and the possibility that subjects hold this belief would work against our hypothesis and weaken our findings, it would not affect the validity of our results.

scientific studies to receive payment without revealing their identity.[13] This message constrains the action space of the counterparty; experimenters must select the benevolent action as it is their only choice in compliance with the policy. Consequently, subjects cannot apply the heuristic: they cannot ascribe an intention to their counterparty based on the relative benevolence of the selected behavior if the counterparty has only one course of action available. Thus, the treatment is designed to strip the subjects of their decision making strategy to estimate personal trustworthiness and we expect them to reveal less information relative to subjects in *Positive Intent*.

Expressive Signal. The final treatment provides a robustness check aiming to rule out an alternative explanation for why subjects may be less likely to release information in the *Legal Requirement* treatment. It is possible that by informing subjects of the university's data policy, we are sending a signal that subjects should consider privacy protection to be important and desirable. This might reduce the willingness of subjects in the *Legal Requirement* treatment to disclose personal information.[14] To rule out this possibility, we designed an *Expressive Signal* treatment where subjects could again (as in both *Positive Intent* and *Legal Requirement*) choose between (1) bank transfer and (2) anonymous payment, or arrange an alternative solution. However, subjects in the *Expressive Signal* treatment were given a different message about the university's data policy: they were told that a data policy had been enacted obliging experimenters to offer anonymous payment, but that the regulation did not apply to the study they are participating in "as this study is conducted before the provision comes into effect." Subjects are further told that the experimenters "decided to offer an anonymous payment option even though the current policy does not oblige us to do so."[15]

The treatment should send the same signal about the importance of privacy as *Legal Requirement*—if indeed our treatments are sending any such signal. Even though the provision has not yet come into effect, the policymaker has already expressed its intent and subjects were informed of it. At the same time, and in contrast to *Legal Requirement*, this treatment allows subjects to apply the intent heuristic. As the provision is not yet in effect, it cannot limit the experimenters' current action space, and if they provide anonymous payment this choice can be compared to the less beneficent option of not offering privacy protection. Thus, the *Expressive Signal* treatment enables us to discriminate between whether subjects base their decision on the intent heuristic or on the potential expressive signal. If subjects decide based on the heuristic, they should trust their counterparty and disclose private information more readily compared to subjects in the *Legal Requirement* treatment. On the other hand, if subjects respond to the

[13] Typical lab policies in Germany require that data are anonymized such that names and experimental behavior can never be connected. This applies also for financial data. However, there are different ways to conform to this requirement. In the lab, for example, participants typically are paid in cash. Often the lab assistant who pays the participants is a different person than the assistant who has helped subjects with the study. For online studies final payments can be transferred with the consent of the participants; yet again the connection between payment and name and behavioral choice has to be destroyed.

[14] We are indebted to Florencia Marotta-Wurgler for suggesting this point.

[15] We debriefed subjects after the session that the policy had already been in place. The subjects in this treatment were not included in the general subject pool.

expressive signal sent by the privacy policy, their willingness to release sensitive content should decrease relative to the *Positive Intent* treatment.

3. Methods

Participants were either current or former students of the University of Münster in Germany (about 30% were recent graduates) who we had recruited for our empirical legal studies subject pool. We sent potential participants an email invitation with a link directing them to the study. We used the open-source LimeSurvey web app as the online platform to conduct the experiments.[16] As the main interest of our study is to learn about how privacy choices are made online (i.e., in social media platforms, cloud computing, and the like), we chose to conduct the experiment online to improve ecological validity. The interaction online is less personal than in the laboratory and better resembles the settings in which online privacy decisions are made. In a laboratory, by contrast, participants would learn who operates the lab and might build trust based on those personal contacts before and during their session.

Our study design offers the ecological validity advantage of field work, as subjects are not aware that they are still participating in an ongoing experiment when they make their payment choices. This design reduces demand effects—i.e., subjects conforming their behavior to what they assume is socially desired.[17] Also, in many lab experiments subjects perform artificial tasks that mimic realistic behavior and they are paid according to a payoff function. In our study, subjects are presented with a realistic privacy choice similar to those they often make outside the context of a research study: they are asked to reveal actual personal information and the consequences of that action can be the same as in other social contexts—when, for example, they purchase goods in an online shop and are asked to decide whether to pay with PayPal or with their credit card. At the same time our design also affords better control over the treatments than is typical in field studies, which often must contend with a host of interfering factors.

4. Hypothesis and Results

The goal of Experiment 1 is to provide evidence that subjects rely on heuristic decision-making when they make privacy choices online. To examine whether that is true, we first compare subjects' behavior in the *Negative Intent* and *Positive Intent* treatments. Recall, in the *Negative Intent* treatment subjects are offered only payment options that require them to reveal their identity and confidential information. If they apply the heuristic, they should perceive the counterfactual—the offering of a privacy-protecting payment option—as comparatively beneficent. On that basis subjects should ascribe to the counterparty a negative intention and treat her as not trustworthy. In the *Positive Intent* treatment, by contrast, subjects offered anonymous payment should perceive the counterfactual—the counterparty might have withheld

[16] *See* http://www.limesurvey.org/. Note that we took precautions to ensure the validity of responses. After participants log in, the link becomes inactive, ensuring that the same participant can complete the study only once. The email informs the participants about the time a typical subject requires to complete the experiment, thereby ensuring that subjects do not discontinue participation because it is taking them longer to finish than expected.

[17] Field experiments are less likely to induce demand effects, as they are designed to unobtrusively assess the effects of realistic treatments on subjects who would ordinarily be exposed to them, often with participants not being aware that their behavior is subject to research.

any privacy-protecting payment option—as less beneficent. Thus, if subjects apply the heuristic, they should ascribe a positive intent and trustworthiness to the counterparty.

This leads us to our first hypothesis[18]: we expect subjects in the *Positive Intent* treatment to choose bank transfer significantly more often and thereby disclose more personal information than subjects in the *Negative Intent* treatment.

Figure 1. Percentage of subjects who reveal confidential information across treatments

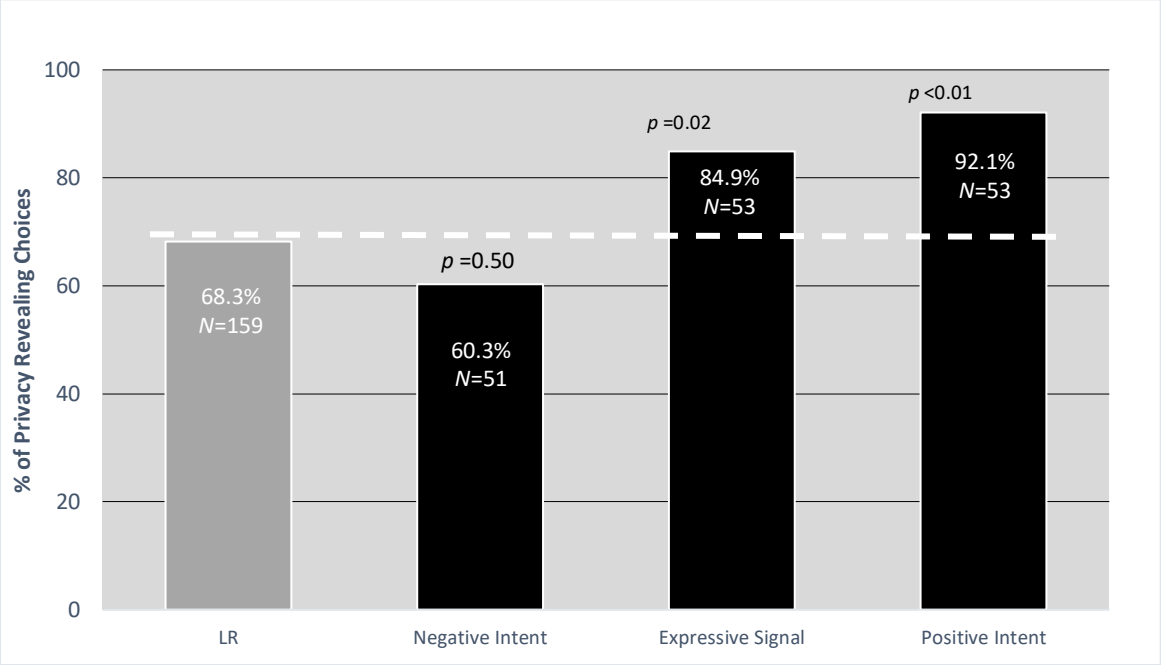

Our results support this hypothesis. In the second column you see the *Negative Intent* treatment, on the right side the *Positive Intent* treatment. The graph shows that 60.3% (32/53) of the subjects in *Negative Intent* choose bank transfer, while 39.7% (21/53) chose to opt out of the offered payment methods (bank transfer and PayPal).[19] By contrast, in the *Positive Intent* treatment 92.1% (47/51) of the subjects chose bank transfer, compared to the 60.3% we observed in *Negative Intent*. The difference between the treatments is strongly significant ($p<0.01$ Fisher Exact).

In a second step we investigate whether we can block subjects from relying on the intent heuristic, and, if so, how that affects their behavior. In our *Legal Requirement* treatment we offer subjects the same payment options as in *Positive Intent*, but we manipulate their perception of the counterparty's action space by instructing them that the law requires the counterparty to offer a privacy-protecting payment option. As our manipulation eliminates the counterfactual, it should preclude subjects from employing the heuristic. Our second hypothesis therefore

[18] Note that all hypothesis were formed prior to the data collection, unless explicatively stated. Hypothesis are almost exclusively treatment comparisons showing that the experiment was designed to test exactly these hypothesis. Note that the N for the LR treatment is much larger than for the other treatments. This has the reason that the study`s treatment we used to run this t, had this size.

[19] Recall that subjects experienced a real cost in opting out of the bank transfer and PayPal options: the fully-anonymous method was costly in terms of time and effort required to collect payment.

predicts that subjects in the *Legal Requirement* treatment should be less willing than subjects in the *Positive Intent* treatment to opt for bank transfer, even though subjects in both treatments are offered the same payment options.

Our results strongly support our second hypothesis: while we have seen that in the *Positive Intent* treatment 92.2% (47/51) of subjects choose bank transfer and only 7.8% wish to conceal their identity and choose to withhold confidential information, disclosure drops significantly in the *Legal Requirement* condition, where only 68.6% (35/51) of the subjects select bank transfer and 31.4% conceal. This treatment difference is clearly significant (p=0.005; Fisher Exact).

The results show, as we had theorized, that many subjects indeed appear to use the intent heuristic to determine their counterparty's intent and that subjects base their privacy decision on this information. When the heuristic suggests subjects should trust their counterparty, as in *Positive Intent*, participants are indeed much more likely to disclose private information compared to when the heuristic suggests not to trust the counterparty, as in the *Negative Intent* treatment. And when we block subjects' use of the heuristic in *Legal Requirement*, subjects become more cautious, even though they are offered the same payment options as in *Positive Intent*.

Finally, we want to confirm our theory using the robustness check provided by the *Expressive Signal* treatment. In this treatment, we aim to reject the alternative explanation that the results in *Legal Requirement* may be driven by a normative message sent by the laboratory policy that a privacy-protecting payment option should always be made available. *Expressive Signal* is identical to *Legal Requirement* except that it instructs subjects that even though the data protection rule was enacted it is not yet in effect. If subjects are responding to the information the law's expressive signal seems to convey, the fortuity of timing should not change that information. But unlike in *Legal Requirement*, subjects can apply the intent heuristic: as the provision does not yet limit the counterparty's action space, subjects should ascribe to the counterparty who voluntarily offers a privacy-protecting payment positive intent and trustworthiness.

Thus, our third hypothesis assumes that subjects in *Expressive Signal* can use the intent heuristic and will disclose more personal information than subjects in *Legal Requirement,* where use of the heuristic is blocked, even though subjects in both treatments should have received the same expressive signal. The results support our hypothesis. As Figure 1 shows, 84.9% (135/159) of the participants choose bank transfer in the *Expressive Signal* treatment, and only 24 subjects or 15.1% opt for the privacy-protecting anonymous payment method. In *Legal Requirement* by contrast we find only 68.6% choosing bank transfer. The treatment difference is significant (p=0.02 (Fisher Exact)). So the data rejects the possibility that subjects in *Legal Requirement* were responding to an expressive signal.[20] As predicted by our theory, the frequency of subjects in *Expressive Signal* choosing bank transfer (84.9%) is statistically indistinguishable from the frequency we observe in *Positive Intent* (92.1% (p=0.27; Fisher Exact), suggesting that subjects in

[20] Indeed, it might be argued that in the *Expressive Signal* treatment we *amplify* whatever expressive signal is sent by the legal mandate. That is because we indicate, through our provision of privacy protections, that we agree with the expressive content of the legal rule even though we are not bound by it. As a result, our *Expressive Signal* treatment is essentially conservative—if the effect on subjects' behavior is driven by the law's expressive content, we would be more likely to have picked it up.

both treatments apply the intent heuristic to ascribe positive intent and trustworthiness to their counterparty.

In sum, the results of Experiment 1 confirm our hypothesis that subjects make online privacy decisions using the intent heuristic. The heuristic is impactful: when its application suggested that the subject should trust the counterparty, subjects' rate of disclosure increased by almost 25%.[21] Furthermore, we can switch off the heuristic which leads to subjects making more cautious privacy decisions.

Note that heuristic decision-making by definition does not process all relevant information. That fact leads us to our next study: Will the heuristic lead subjects to give little weight to relevant information about technological or organizational privacy risks? That is the question we will investigate in Experiment 2.

B. Experiment #2: "Crowding Out" Credible Signals

1. Motivation

Heuristic decision-making is frugal; its decision making rule does not consider some or even most of the information that might be available and possibly relevant to the decision. It may nevertheless be an efficient decision-making strategy (Gigerenzer & Gaissmaier, 2011) if it can exploit a match between cues used and the decision environment. The so-called "recognition heuristic" is an example of a decision-making short-cut that is adaptive at least in some contexts. The recognition heuristic holds that "if one of two objects is recognized and the other is not, then infer that the recognized object has the higher value with respect to the criterion." Recognition can be a valid cue if there is a high correlation between the recognition and a criterion that predicts the outcome well—for example firms with better products often survive longer and are therefore more likely to be recognized by the consumer. In such cases, the correlation of recognition and criterion can make the heuristic ecologically adaptive. On the other hand the recognition heuristic is likely to fail when there is a bad fit between recognition and criterion. Consider an American using the heuristic to decide whether Santa Barbara, California or the Chinese city Chongqing is larger. Chongqing is one of the largest cities in the world, yet the average American is likely not to have heard of it. The heuristic fails because China is so distant from the United States that even large cities may not be recognized. In these circumstances, the connection between recognition and criterion is too loose for the heuristic to work reliably (Czerlinski et al., 1999). While the heuristic itself processes only one informational cue, decision makers may still consider other information the heuristic does not process. However, participants were shown to often not override recognition, even when they thought another cue might be equally or even more valid (Pachur et al. 2008).

In the same way the intent heuristic may also be prone to failure in the online world. The perception that the counterfactual action of the disclosure-seeking counterparty may have been more or less benevolent must be predictive of the potential discloser's true privacy risks. Yet, in

[21] Also, in this experiment subjects were fully informed about the researchers' background; the baseline of disclosures when subjects cannot use the heuristic is therefore comparatively high. Study 2 presents an environment where baseline trust is lower, and we see that the impact of the heuristic gets stronger.

a digital world where the privacy risk is diverse and affected by the choices of many and not just an identifiable counterparty the intent heuristic may be just as loosely connected to true privacy risks as the recognition heuristic is with an American's understanding of the size of a Chinese city.

That alone may not render the use of the intent heuristic problematic, if other relevant information was not crowded out. For example the recognition heuristic when applied to products does not consider whether the technical description of one product suggests more reliability than the product of a recognized brand, or whether the portfolio of one company may suggest more expertise in the product's area (Gigerenzer & Gaissmaier, 2011).[22] If decision makers do not override the heuristic and give those other cues only little weight, valuable information can be lost—in particular, when the heuristic is not adaptive to the decision-making environment. In the same way, the intent heuristic might lead individuals to put little weight on reputational information that would help them to better estimate their true privacy risks.

That possibility is what Experiment 2 investigates: do subjects base their decision primarily on the heuristic and give other valid information that might otherwise lead them to make the opposite decision only little weight?

2. Design and Methods

In Experiment 2 we analyze whether signals of high or low credibility affect subjects' privacy decisions when they apply the intent heuristic. We recruited new subjects via email and invited them to participate in future studies. When invitees entered the registration website, they were instructed how to register and prompted to provide payment information. Invitees could only successfully register if they provided the information. We measure how many invitees provided the information versus how many entered the website but opted out and did not register.

The design of Experiment 2 varies three different payment options for subjects to choose among: bank transfer and PayPal, which were offered in all treatments, and either anonymous payment or mail, which were varied across treatments. We implemented three treatments—*Positive Intent, Negative Intent* and *Legal Requirement*—that are structured similarly to those treatments in Experiment 1. In *Positive Intent* and *Legal Requirement* subjects were offered bank transfer, PayPal, and the anonymous payment method allowing them to use a code to pick up their earnings, as described above.[23] And just as in Experiment 1, subjects in the *Legal Requirement* treatment were informed that the University's data policy required experimenters to enable anonymous participation in the study, while by contrast in *Positive Intent* the experimenters appeared to provide the anonymous payment method voluntarily. Subjects in the *Negative Intent* treatment were offered only payment options that required them to indicate their full name and additional personal information: they could choose between bank transfer or PayPal, or they could indicate their home address to receive payment by mail.

[22] H.A. Simon points out the trade-off between rational decision-making and heuristics that may perform well under uncertainty but that also can lead to failures when most relevant information is not considered. Simon speaks of bounded rationality (Simon, 1977).

[23] *See* Section III.A.3.

Experiment 2 deviates from Experiment 1 by manipulating the availability of other information that subjects could utilize to make a better informed privacy decision. We implement for each of the treatments (*Positive Intent, Negative Intent* and *Legal Requirement*) two informational conditions: *positive signal*, in which we provide credible information about the counterparty's trustworthiness, and *negative signal*, in which we do not provide any information about the counterparty at all, which should send a negative signal about the counterparty's trustworthiness. In *positive signal*, we presented subjects with biographical information about the researchers conducting the experiment, including our ties to academic institutions (NYU and University of Münster), and displayed the emblems of the researchers' universities on top of each screen presented to the subjects. We also provided online links to our academic work and gave subjects the University of Münster email address of our research assistant, informing them that they were free to contact her. By contrast, in *negative signal*, we provided only our RA's *non-university* email address as contact, and gave no information about us (that is, nothing about affiliations, past research, etc.).[24] For a real-world analogue, imagine a website which does not allow the user to trace who is running it and who is responsible for the platform and its content.

As in Experiment 1 it was crucial that we conducted Experiment 2 online. The condition in which we provide little or no information is foundational to our design because it sends to subjects a negative signal suggesting that the counterparty's trustworthiness is low. Such a low-information condition is easily implemented online; it is very difficult, however, to reproduce in the laboratory, where institutional information (the university responsible for the study; the researchers who operate the lab, etc.) is a given, and where personal face-to-face interactions tend to build trust.

Like Experiment 1, Experiment 2 has strong field elements. Subjects make decisions in a realistic environment whether or not to disclose actual personal information. And when subjects make their decision, they are likely to perceive the solicitation of their choice of payment method to be an administrative act, and not an experimental task.

3. Hypothesis and Results

We measure first whether subjects being able to use the intent heuristic respond to signals suggesting either low or high trustworthiness of their counterparty. Rationally, one would expect that subjects are less likely to disclose personal information if they receive a negative signal suggesting low credibility and are more likely to disclose if they receive a positive signal. However, like many heuristics (Gigerenzer & Goldstein, 1996) the intent heuristic focuses judgments on a single criterion that is perceived to discriminate effectively—in this case, whether the heuristic ascribes good or bad intent to the counterparty—while all other cues may not be considered. The two informational conditions in Experiment 2, *positive signal* and *negative signal*, test whether both positive and negative signals fail to affect privacy choices, leading us to a first

[24] Since the invitations were sent over the university server, the subjects may assume that we had to meet some minimum privacy standards to be given access. This may increase baseline trust, but it does not affect treatment effects; if anything, a higher baseline works against our hypothesis. Note that we recorded whether subjects provided financial information, but this information was not initially stored; invitees were contacted again and fully informed.

hypothesis: when subjects apply the heuristic, neither negative nor positive signals will substantially affect their privacy choices.

Figure 2. Subjects' Disclosure Choices by Type of Signal Received and Treatment

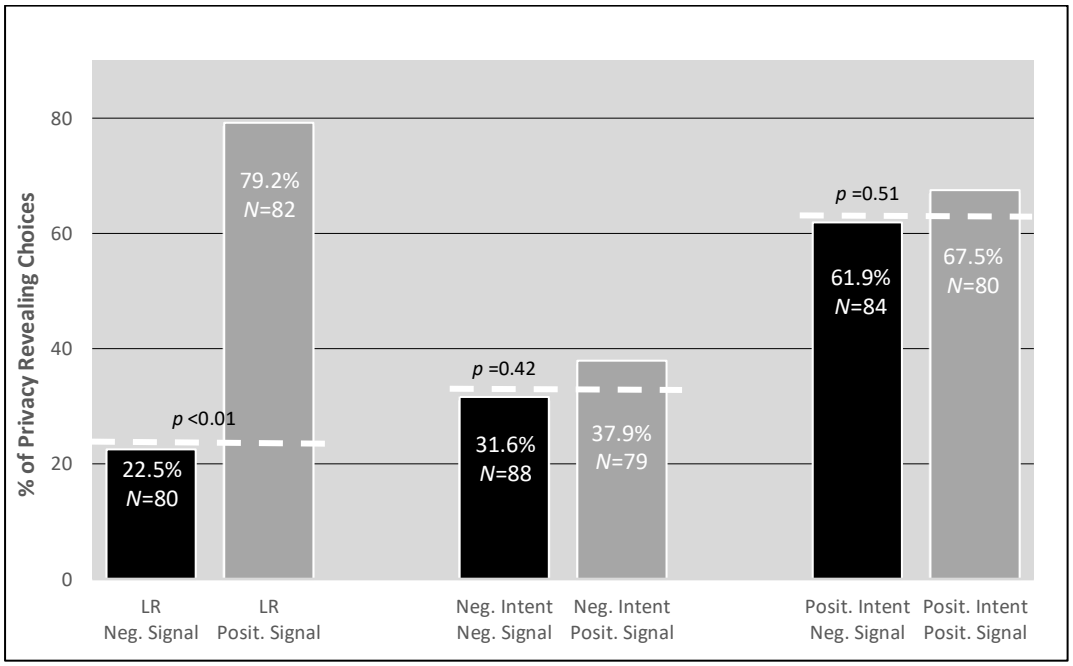

Our results support this hypothesis. In Figure 2 you see for each treatment (Legal Requirement, Negative Intent and Positive Intent) two bars reporting the two informational conditions, negative signal and positive signal. The more similar the bars within each treatment, the smaller the differential impact of the informational signals on subjects' privacy decision. We focus first on Negative Intent and Positive Intent, because in both treatments we expect subjects to apply the heuristic. Percentages reported in Figure 2 refer to the frequency of decisions to disclose personal information—i.e., to choose the payment methods (bank transfer, PayPal) that require disclosure. For our analysis, we employ a Fisher test (2x2 contingency table) for the Negative Intent treatment; that is, we test whether the two informational variables (negative signal and positive signal) and the subjects' privacy choices (i.e., disclosure=bank transfer, PayPal as opposed to non-disclosure=the anonymous payment method) influence one another and whether the informational signals affect subjects' disclosing behavior. The result is insignificant (p=0.45 Fisher Exact), indicating that the informational cues do not change the likelihood of disclosure in the Negative Intent treatment. When we perform the same analysis for the Positive Intent treatment, we get the same result: the relationship is again insignificant (p=0.24 Fisher Exact), which means that informational signals do not appear to affect privacy choices in Positive Intent either.

Note, we do not conclude from these data that the intent heuristic is non-compensatory, i.e., that it causes people to ignore and not process other information. We can observe that the informational signals we provide do not change privacy choices when the heuristic can be used,

while they change privacy choices substantially, when subjects cannot use the heuristic. We interpret this result to mean that decision makers seldom override the heuristic. We do not imply that no other potential information could have moved the subjects` privacy decisions, but apparently it would have to be strong cues. These results support our first hypothesis: subjects applying the heuristic seem to give negative as well as positive informational signals little weight in their decision-making. That negative informational signals did not affect subjects' privacy decisions in the context of the Positive Intent treatment is perhaps most concerning, for there the heuristic suggests disclosure and subjects are ignoring signals at odds with the heuristic's conclusion that the counterparty is trustworthy. But the opposite case, where subjects in the Negative Intent treatment fail to heed the positive informational signals and remain reluctant to disclose, also produces costs. In these instances, the heuristic can strip individuals of opportunities for beneficial disclosures and/or lead to unnecessary transaction costs, as we see in our study, where subjects who did not base their decision on the positive signals invest time and effort to pick up their earnings in person.

We next investigate whether subjects will heed valid positive or negative informational signals when the intent heuristic is unavailable to them. In the Legal Requirement treatment, we prevent subjects from using the heuristic, as we have seen above. If the heuristic is blocked the subjects need to search for and rely on other information in order to make their privacy decisions. This may push our participants to process and respond to the credibility cues they were given. In our second hypothesis, we posit that if subjects cannot rely on the heuristic, their privacy choices will be affected by the type of signal—positive or negative—the subjects receive.

Our results support this second hypothesis. For subjects in the *Legal Requirement* treatment we observe a strong impact of both the negative and positive signals on privacy choices; in Figure 2 we can see that 79.2% of the subjects in the *Legal Requirement* treatment who receive a positive signal disclose private data, but only 22.5% who are presented with a negative signal disclose. The test for statistical relationship is highly significant: $p < 0.01$ (Fisher Exact)—indicating that the informational signals clearly had an effect on the privacy behavior of the subjects.

This is an important finding: by precluding use of the heuristic, our *Legal Requirement* intervention was able to open up the decision-making process, extracting it from the intent heuristic and making effective important informational cues, which appear to not affect privacy choices when subjects' decision-making is based on the heuristic. As a consequence, blocking use of the heuristic appears, at least in settings where otherwise credible informational signals are available, to lead individuals to better estimate their privacy risks.

C. Experiment #3: The Heuristic and the Public Good Characteristics of Individual Privacy Decision

1. Motivation

In Experiments 1 and 2 subjects could disclose strictly-defined types of information: their name and bank information, their PayPal address, their home address. They had full awareness and control over what information was disclosed if they decided to give us their data. In contrast, the risks of privacy decisions in social media and other social networks are inherently less predictable: *data is typically not in the control of the discloser alone*, but can be accessed, altered, or even

created in the first place by other users of the network, and even beyond by the wider audiences with which the data is shared, often without the discloser's knowledge.

A concise way to say this is that in social networks, data (and privacy) are a product of user *interactions*. In these settings, privacy is not an individual choice. It is a *collective good*, or, as some have put it, a *public good* (Fairfield & Engel, 2015; Barocas & Nissenbaum, 2014). In settings where people are connected in networks, an individual's decision to protect her privacy can lead to a "positive externality" whereby others' privacy is protected as well. And conversely, one individual's decision to sacrifice her own privacy can give rise to a *public harm*—i.e., it can result in others suffering "negative externalities" in the form of damage to their privacy. These externalities strip the individual of the power to protect her privacy alone. Privacy protection is a group effort.[25]

We capture this dynamic in Experiment 3. In this experiment, the Facebook information subjects can reveal to us is created and altered by a potentially unlimited number of other users, who may link their own content to the subjects' accounts, make public formerly private interactions with the subject, or share subjects' Facebook postings with other audiences who may then add content themselves. As a consequence, from the perspective of any individual subject, the scope of the data that subjects disclose if they allow us access to their account is fundamentally uncertain.

Whether users realize it or not, this is a typical scenario for making privacy decisions in the context of social media and network interactions. An individual may decide to give a company access to their Facebook data in exchange for being allowed to use a web service. Or they may log into a service using their Facebook credentials for convenience, and as a result the company and Facebook will own her user data and likely will share it (at least) indirectly via targeted advertising with other companies or services. Because the data disclosed may include content

---

[25] Gmail is another illustration. In exchange for free email service, Gmail users agree to have their emails scanned (the searches are used by Google to target ads). But it isn't only Gmail users' emails that are searched—Google also searches email sent to Gmail users by others, even if the sender is not using Gmail. As a consequence, the privacy behaviors of an individual Gmail user affect not just that individual user's privacy, but also that of the people who correspond with her.

The Gmail example illustrates that privacy harms are often *nonexcludable*—that is, the harms from an individual's decision to sacrifice her privacy cannot be confined to the individual making the decision. And the same is true of decisions to protect privacy—the benefits are nonexcludable because they flow not only to the person making the privacy decision, but to others whose privacy is preserved as a result of that decision. To avoid the negative externalities, an individual who doesn't use Gmail must stop corresponding with Gmail users. This possibility becomes more costly as the number of Gmail users grows. Unless many Gmail users decide to leave Gmail, the individual who doesn't use Gmail cannot effectively protect her privacy through her own decision-making.

Similarly, users have only partial control over their privacy on Facebook. No individual Facebook user is able to understand fully, at the time they consent to Facebook analyzing their data, what data will be created and associated with them on Facebook. The data that is produced and associated with their profile depends on the actions of many who post their comments and likes on the individual's page, or who otherwise link content to the account of another Facebook user—for example, by tagging that user in a photograph. Once the individual is part of the network she basically loses the control over her privacy, and the more she interacts and the larger the group with which she interacts, the less control she retains over what is disclosed.

Note, however, that as Helen Nissenbaum's "contextual integrity" theory teaches us, privacy is as much about appropriate sharing of information as it is about not sharing. That is, whether a particular data flow is appropriate depends on context and a careful assessment of costs and benefits, both private and social. *See* Nissenbaum, 2010.

generated by the people the discloser has been interacting with, disclosure *imposes a privacy risk on these other users of the network*. And the risk flows in the other direction as well: data produced or modified by other users on the network but linked to the discloser's account—e.g., a Facebook comment on the discloser's post made by a Facebook friend—*presents a privacy risk to the discloser*. Thus, individual privacy decisions in a networked environment are always intertwined with the privacy decisions of others, and likely, for that reason, to create externalities.

The story of "thisisyourdigitallife," which we discussed in the Introduction to this Article, is an apt illustration of the social dimensions of individual privacy decisions in the context of online networks and platforms (Tufekci, 2018). The privacy implications of "thisisyourdigitallife" may be unusual in scale, but the social privacy risk is, in an important way, entirely typical: the app didn't steal data; rather, people consented (at least notionally) to disclosure.[26] Lax Facebook rules attribute all content associated with a posting to the user who made that posting. And that means that the authors of those reactions need not consent when their information is disclosed alongside with the original posting. Under these rules, Facebook permitted the "thisisyourdigitallife" app to access and collect Facebook data not only from the 270,000 people who downloaded it, but also from up to 87 million of the downloaders' Facebook friends—people who had never given their consent or downloaded "thisisyourdigitallife" (Chang, 2018).

Experiment 3 investigates this social dimension of privacy decision-making. We hypothesize that the heuristic will lead subjects to disregard the uncertainties of their own privacy risks in networks and in particular the risks they may impose on others, as it reduces the disclosure decision to a judgment of the trustworthiness of the person or service they are directly interacting with. We further hypothesize that when we block subjects from using the intent heuristic, they will be better able to consider the social dimension of their privacy decision, and, further, to better consider whether they are willing to limit their own disclosures to reduce the externalities they impose on others and also the externalities the network may impose on them. Thus they may begin to contribute to preserving privacy as a public good.

2. Design and Methods

To implement Experiment 3 we partnered with the company Social Links, which has developed a web application that allowed us to elicit all publicly accessible data linked to subjects' Facebook accounts.[27] The application is registered with Facebook and complies with Facebook's current data protection rules. The information we are able to collect using Social Links includes all of an individual's Facebook postings, comments and likes that subjects have published and to which subjects have not applied a privacy setting that restricts access. We are also able to collect content that other users have posted about the subjects; for example, photos on which the subjects are tagged and that are not subject to a restrictive privacy setting. To collect the data the application accessed the Facebook profile of participants who consented to disclosure; once the information was received the app was disconnected by Facebook. This consent applies also

[26] The app's user agreement stated specifically that the app might gather information on the user's Facebook friends (Etter & Frier, 2018). And yet it is likely that virtually none of the people who used the app actually read it. *See* Bakos et al., 2014.

[27] *See* https://sociallinks.io/.

to the comments and likes of others, if they are linked to the subjects' account (we explain this in detail below).

Since not all potential participants have an active Facebook account, we first prompted subjects to indicate which social media platforms they use and in particular whether they use Facebook. This information allowed us to consider only subjects for our experimental sample who could give us access to their Facebook accounts if they wanted (22.5% of the participants did not have a Facebook account).

We then asked subjects to indicate whether they would allow our web application to collect all publicly available Facebook data associated with their Facebook profile. We incentivized their decision, offering subjects a payment of €4.00 (~$4.50 at the time) if they agreed to disclose. We informed subjects that the information we would collect might (if made public) include the groups they have joined on Facebook, the interests they indicated, the discipline they study, and the Facebook likes and comments that they have published, as well as content that other users have posted and linked to their profile. We instructed the participants that we would use the information to better understand which of our future (social media) studies might be appropriate both for them and for their Facebook friends. We also instructed subjects that we would reach out to their Facebook friends and invite them to register or our experiments—without informing their friends that their participation in our study was how we learned about and were able to contact them.

Since the information is publicly available, one might ask whether collecting the data that individuals have already made public exposes them to an additional privacy risk. The risk, however, relates to the personal information we could derive from the totality of the data w get access to: we are scraping all public Face book data, enough to derive very personal information that users would likely want to keep private. So rather than only learning what our participants posted directly, we could infer their political beliefs, sexual orientation, character traits, and so on.[28]

Subjects made two decisions. First, they decided whether to permit us access to their account. We then asked them on a separate screen to provide us with their list of Facebook friends (to the extent it was publicly available on their Facebook profile). We incentivized this decision by informing participants that if they disclose only partial information (either their friends list or their Facebook data, but not both), we may reduce compensation accordingly.

We classify the information we collect into two categories. In the first are subjects' own activities: their own postings, comments, photos and likes, the groups they join, pages from organizations or notable people they liked, event pages they visited, etc. The second category is information produced as a result of the subject's connections to other Facebook users: that is, data from all activities other users have linked to a subject's profile like posts, likes or photos in

[28] See, for a demonstration, "Magic Sauce," a project developed at the University of Cambridge that uses Facebook data to predict personality traits. See Apply Magic Sauce, The Psychometrics Centre (cam.ac.uk), https://applymagicsauce.com/demo.

which the subject is tagged, or the fact that the third person has become Facebook "friends" with the subject.

Both categories of information—activities and connections—pose social privacy risks. First, subjects have incomplete control over or knowledge of what data they disclose to us. To be aware in full of what posts they made in public or other users (later) linked publicly to their Facebook profile is practically impossible; those postings may be public or private, depending on the privacy settings of yet again multiple other users, and the subjects may not even know of the content being linked to their profile. The risk that information that subjects would have preferred to keep private is revealed obviously rises along with the number of contacts and interactions a subject has. It is therefore difficult for any subject to have a complete picture of the true social risks of disclosure.

Subjects face similar difficulties in assessing the privacy risk their own decision to disclose imposes on others. Subjects who opt to disclose are authorizing us to collect other users' postings linked to the subject's Facebook profile. This is consistent with Facebook's own rules.[29] The more content subjects disclose the larger this externality becomes: the more data collectors learn about users in general, the better-situated they are to draw conclusions about the behavior of any particular individual. For example, Facebook has acquired a patent for inferring the creditworthiness of particular individuals. Each financial decision a user makes online and for which she allows Facebook to collect data is fed into an algorithm that estimates creditworthiness and that may ultimately decide whether a friend in the same social network—i.e., someone who shares similar characteristics with the user—is extended a loan (Fairfield & Engel, 2015). The data has established a digital redline that the individual seeking the loan is not aware and cannot influence with her personal behavior. These are the characteristic risks of a public good: the privacy decisions of one user affect not only her own privacy but potentially the privacy of the whole community. And then, breaking down the social privacy risk to our study, the point illustrated earlier applies: the more we learn about an additional user, the more we would be able to infer beyond what is explicitly posted on Facebook.

We compare the two treatments *Positive Intent* and *Legal Requirement*, structured similarly to those treatments as they are administered in Experiments 1 and 2. In Experiment 3's version of both treatments, we inform subjects that if they wish to register to participate in future studies, they must choose how they want to collect earnings from those studies. We offer them anonymous payment (using a code as in Experiments 1 and 2) alongside three payment methods that require them first to indicate their full name and then confidential information: bank transfer, PayPal, and payment via mail to their home address.[30] Additionally, we inform subjects in both treatments that they can also register to participate in future studies without disclosing the requested Facebook data.

[29] *See* https://www.facebook.com/privacy/policy/?subpage=1.subpage.2-FriendsFollowersAndOther

[30] At each stage of the experiment—i.e., when we (1) ask subjects to list the social media platforms they use, when we (2) prompt them to provide access to their public Facebook data and when they are requested to (3) disclose their friends list—subjects are offered the opportunity to participate in our studies while remaining fully anonymous.

The *Legal Requirement* treatment differs from *Positive Intent* in that the subjects are told, as in Experiments 1 and 2, that the University's rules require experimenters to enable anonymous participation.

We compare across the two treatments how many of the subjects (1) disclose their Facebook data and (2) disclose their friends list. For subjects who choose to disclose data, we compare the amount of data participants reveal.

Experiment 3 is a partial field study: while subjects are aware that their privacy choices are subject to our research, they perform an everyday decision they are accustomed to and that decision could affect their real privacy, as they disclose all their publicly available Facebook data to us.[31] This is the same decision they would make in the real world if, for example, an app on Facebook asks for access to their Facebook data in exchange for allowing them to use the app.

3. Hypothesis and Results

We assume that subjects in the *Positive Intent* treatment will employ the intent heuristic and frame their privacy decision as a bilateral interaction with the experimenters. Basing their decision on the heuristic, we expect them to trust the experimenters and to disregard the privacy risks that the decision's social dimension may impose on them as well as on others.

By contrast, we expect that subjects in the *Legal Requirement* treatment, who cannot apply the intent heuristic, will be more likely to consider both the public good nature of their privacy decision and the privacy risk they would take on and impose on others. We further expect, as a consequence, that these subjects will be more hesitant to disclose their Facebook profile information and friends list.

This leads to our first hypothesis for Experiment 3**:** Subjects in *Legal Requirement* are less likely to give us access to their public Facebook data relative to subjects in the *Positive Intent* treatment.

The results shown in Figure 3 support our hypothesis. In the Positive Intent treatment where subjects can ascribe good intent and trustworthiness to a counterparty, subjects are significantly more likely to allow us access to their Facebook data relative to subjects in the Legal Requirement treatment whose use of the heuristic is blocked ($p < 0.01$ Fisher Exact). We conclude that subjects who employ the heuristic seem to be less considerate of the social risks of disclosure. In contrast, when we block the intent heuristic, subjects appear to change their decision-making strategy: They are more likely to take the social risks of their privacy decision into account.

[31] We did not provide subjects with a sample report from Social Links that would demonstrate the capability of the software, as this would have given subjects a state of precise information that is not typical for privacy choices—in most cases data protection protocols remain an abstract description and do not demonstrate vividly the scope of data that is disclosed and how it is going to be used. This uncertainty and vagueness are part of the experiment's ecological validity. After the experiment was completed, all Facebook data was anonymized and will be stored only as evidence for the sincerity of the study.

Figure 3: Subjects Disclosing Facebook Data by Treatments

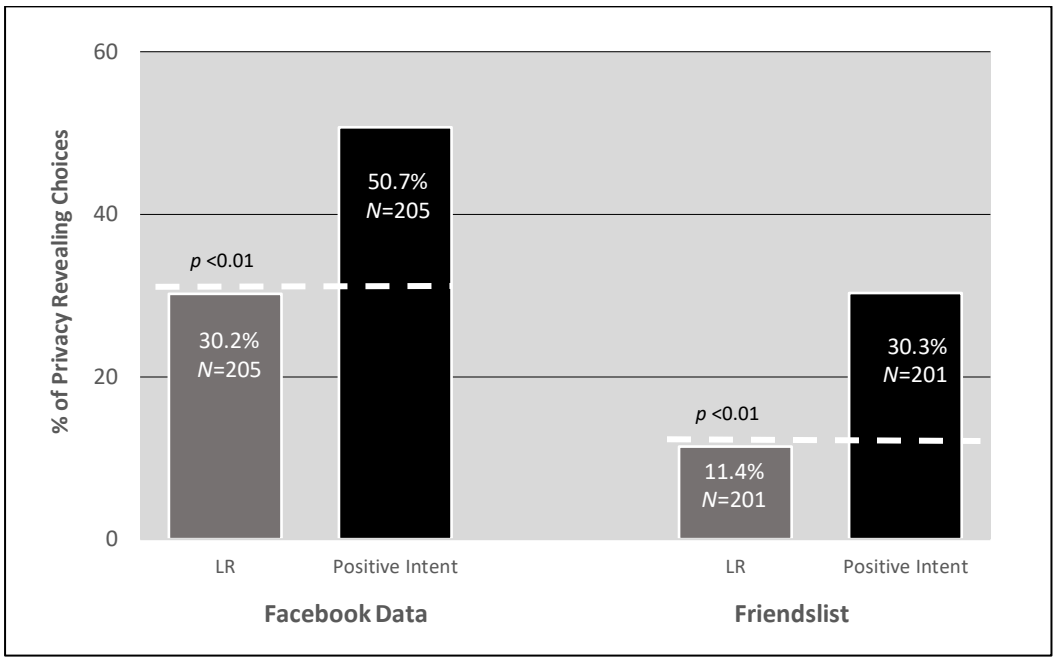

To better understand whether subjects indeed consider the social risks of their privacy choices once they are blocked from using the intent heuristic, we analyze the relative privacy risk of the content that subjects reveal to us. Based on the Social Links reports, we are able to calculate "privacy exposure" scores for each individual subject—i.e., a measure of how much Facebook data an individual subject is disclosing to us that allows us to compare how the privacy risk posed by disclosure varies across treatments.

We calculate three privacy exposure scores for each subject: (1) *own activity*, (2) *number of connections* and (3) *number of friends*. Calculating these scores gives us a way to assess whether subjects in *Legal Requirement* are indeed considering the social risks of privacy more: If they are, then subjects in *Legal Requirement* with lower privacy exposure scores should be more likely to disclose Facebook data to us relative to subjects in that same treatment with higher privacy exposure scores, while we expect subjects in *Positive Intent* to be relatively insensitive to their privacy exposure scores when making their disclosure decision. This is the second hypothesis for Experiment 3.

Figure 4: Privacy Exposure by Choices and Treatments

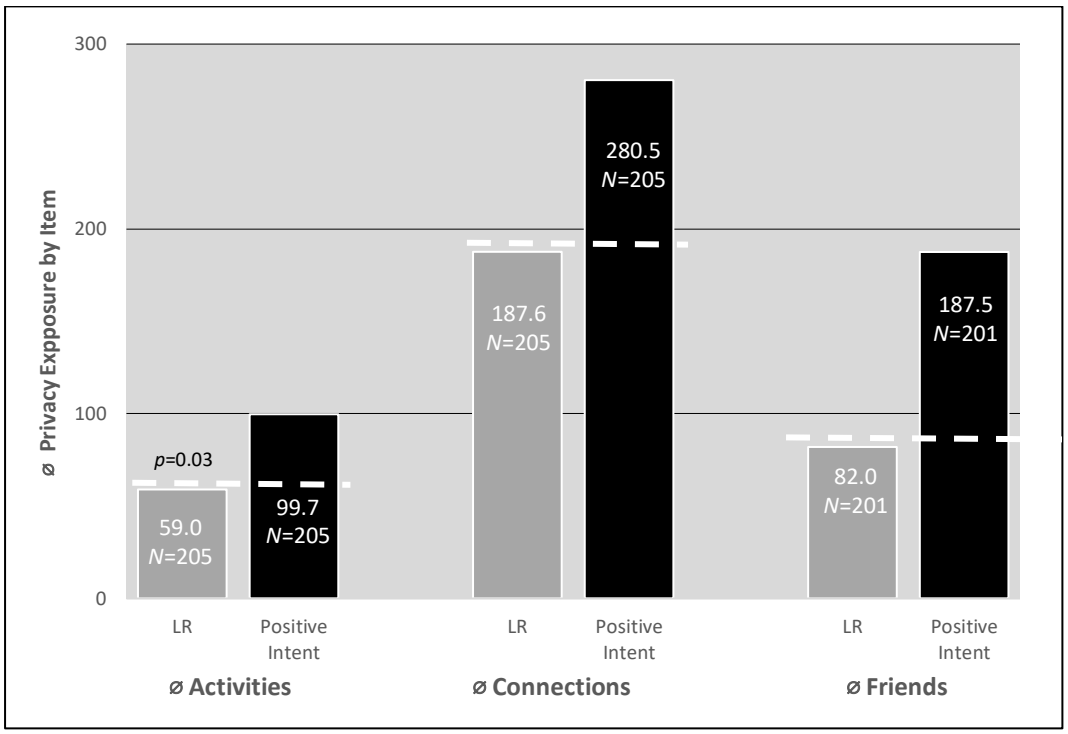

The data reported in Figure 4 support our hypothesis. On the left side of the figure, we see the participant's own activities: i.e., the subject's own postings, likes etc.

In the middle of the x-axis we see the "connections" score, which refers to comments or likes other Facebook users have linked to the participants' postings or photos on which users tagged the participant. Of course, when we aggregate this data, we cannot determine the relative weight of the privacy concern a particular posting or photo carries for a subject; that is private knowledge. We therefore assign all postings, comments and likes the same value and aggregate all values into separate scores for the subjects' own activities and their connections.

The results show that subjects who permit us to assess their data in the *Legal Requirement* treatment disclose significantly less data on both dimensions. First, subjects in *Legal Requirement* revealed fewer own-activities recorded on Facebook than subjects in *Positive Intent* (p=0.026 t-test), and also disclosed fewer connections to other Facebook users than subjects in *Positive Intent* (p=0.036 t-test).

Our data also shows directly that subjects in Legal Requirement were more considerate of the severity of the privacy risk they would be creating by disclosure than subjects using the heuristic in Positive Intent. We find not only that (1) fewer subjects in Legal Requirement disclose data at all, but also (2) those who decide to disclose, have less data to divulge, on average as indicated by their scores for both own activities and connections, than subjects in Positive Intent who decide to give us access. This suggests that in Legal Requirement subjects were considering how much data they would reveal and thus how much of a social privacy risk they would impose. And those subjects who use Facebook often and would have revealed a large quantity of data in their accounts were more likely not to disclose than in Positive Intent.

There is one remaining uncertainty. When we look at the two scores, we see about the same reduction in comparison to *Positive Intent* on both dimensions (own activities and connections). So, potentially subjects did not consider the externalities and social risks of their privacy decision but considered only their own activities and personal risks in their disclosure decision—not wanting to reveal the postings they had made and could remember.

Our data also shows directly that subjects in Legal Requirement were more considerate of the severity of the privacy risk they would be creating by disclosure than subjects using the heuristic in Positive Intent. We find not only that (1) fewer subjects in Legal Requirement disclose data at all, but also (2) those who decide to disclose, have less data to divulge, on average as indicated by their scores for both own activities and connections, than subjects in Positive Intent who decide to give us access. This suggests that in Legal Requirement subjects were considering how much data they would reveal and thus how much of a social privacy risk they would impose. And those subjects who use Facebook often and would have revealed a large quantity of data in their accounts were more likely not to disclose than in Positive Intent.

Importantly, as we only collect public data, subjects who give us access to their friends list do not disclose any personal information they have not already revealed: their friends list is posted in one spot on their personal Facebook page for anyone easily to access along with the number and identity of their friends. So the upshot is by giving us access, they primarily cause a social privacy risk, and not a personal one. As is characteristic for a public good, their personal risk is not primarily driven by their own actions, but by the choices of others who may also reveal their friends list to us. Subjects have to trade off taking the money for personal benefit versus contributing to the public good; here, the privacy of the collective. Of course, plausible alternative explanations exist that answer why subjects may refuse to disclose the list of their friends. On one hand, subjects may perceive it as inappropriate to involve their friends in the study without having their consent. On the other hand, subjects may also reveal the list because they want to support research.[32] However, these alternative motivations should, if they exist, be present in both the *Positive Intent* and *Legal Requirement* treatments. It is key to our research question that the overall impact of these alternative explanations may change with the treatment—for example subjects may be more willing to involve their friends even without having their consent in *Positive Intent* because they trust the counterparty. This is exactly the treatment effect we expect to measure.

The first hypothesis, as in Experiments 1 and 2, suggests that subjects in *Legal Requirement* will be less likely to disclose their friends list than the participants in *Positive Intent*. The data support this hypothesis (see Figure 3). As expected, we find that the rate of disclosure is significantly lower in *Legal Requirement* with 11.4% compared to *Positive Intent* with 30.3% ($p < 0.01$ t-test), suggesting that subjects in *Positive Intent* rely on the heuristic while the use of the intent heuristic in *Legal Requirement* is blocked, making subjects more cautious. Generally, in both treatments fewer subjects disclose their friends list, compared to the first decision whether or not to reveal their Facebook data. The lower rates of disclosures may suggest that subjects are

[32] Subjects may also think that the study is interesting and they want their friends to have the same experience or be able to earn the 4€.

reluctant to involve their friends in the study. However, this hesitation appears to be equally present in both treatments (comparing decisions 1 (activities and connections) vs. 2 (friends list) in *Positive Intent* -20.4%; $p < 0.01$ (Fisher) and *Legal Requirement* -18.8%; $p < 0.01$ (Fisher).

Finally, we want to confirm whether subjects indeed considered in their decision the risk their behavior would impose on others. The risk corresponds with the amount of data they would disclose—the longer their list of friends that would be contacted and potentially participate in the study, the larger the data pool would grow. We expect that subjects in the *Legal Requirement* treatment who decide to pass on their friends list will have a relatively shorter list of friends to reveal than participants in the *Positive Intent* treatment.

The results reported in Figure 4 support this hypothesis: the 23 participants in *Legal Requirement* who disclose their friends list reveal an average of 82.03 friends, while, by contrast, the 61 participants in *Positive Intent* who permit us access to their list, disclose with an average of 187.62, which is significantly more contacts ($p = 0.01$ t-test).

These results show that our mild intervention—blocking the use of the intent heuristic—is effective in making subjects more considerate of social privacy risks that do not directly threaten the subject's personal privacy interest. Subjects in *Legal Requirement* appear to consider their privacy decision as a contribution to privacy as a public good and more often refuse to take the personal benefit of disclosure relative to subjects in *Positive Intent*, who are significantly less likely to consider social privacy risks in their decision. [33]

While the Legal Requirement treatment suggests a strategy for making privacy decisions more attentive to social risks, the treatment also conveys another important message for privacy regulation: to the extent that individuals make privacy decisions that seem to disregard the externalities that their choices impose on others, for many subjects that does not necessarily suggest that they are unconcerned about the welfare of others. Rather, the cognitive tools that individuals employ to make privacy decisions have a tremendous influence on how their perception of privacy protection is framed. As long as the heuristic indicates that the requesting party can be trusted, social privacy risks are likely to be less salient: a trustworthy counterparty is expected to protect also the privacy of those whose data the subject is revealing.

For these reasons, our data is helpful for building a regulatory approach that empowers individuals to make more socially-conscious privacy choices. We will discuss the wider implications of our results for privacy policy in the next section.

* * *

Let us summarize the results of our three experiments. Our data suggest the following:

Experiment 1:

(1) Subjects tend to make privacy choices by relying on an intent heuristic that suggests whether to treat their counterparty as trustworthy or not.

[33] Our theory also suggests that if regulation can disrupt the framework that induces individuals to use the heuristic, then individuals may move toward decisions that are more considerate of privacy externalities and which contribute to the public good of privacy. *See* Fairfield and Engel, 2015.

(2) We can stop subjects from employing the heuristic by blocking the information the heuristic processes: that is, information about whether the counterparty offers them privacy protection voluntarily.

Experiment 2:

(1) Subjects using the heuristic tend not to give little weight to otherwise credible information—whether positive or negative signals—regarding the safety of a potential disclosure. The heuristic appears to crowd out consideration of such information.
(2) When we block the heuristic, subjects become more likely to heed otherwise credible information .

Experiment 3:

(1) The heuristic makes subjects perceive privacy decisions as bilateral trustworthiness assessments, and therefore they tend not to consider the social/public goods aspects of their privacy decisions.
(2) Subjects blocked from using the heuristic are more likely to forgo personal benefits from disclosure and contribute to preserving privacy as a public good in a networked setting.

IV. Discussion of Results and Implications for Legal Reform

A. On the Validity of our Results

We first discuss factors that relate to both the external and internal validity of our results.

1. External Validity.

If we hope to discover ways to empower individuals to more effectively understand privacy risks, it is important that we can generalize our findings to behavior outside of the laboratory. Therefore we have used a partial field design that implements several elements of field work.[34]

First, our subjects are not aware that a study is under way when we begin to observe their behavior. When they decide how they want to be paid, they assume either that the experiment is over (Experiment 1) or that the experiment has not yet begun (Experiments 2 and 3) and that their choices are not the subject of a study. This reduces demand effects—i.e., the phenomena that participants may try to help experimenters reach the assumed scientific goal or attempt to conform their behavior to what they perceive is socially desirable.

Second, our subjects are not assigned experimental roles they are not accustomed to. We do not, for example, put students in the role of managers or entrepreneurs. Instead, we ask subjects who hold financial accounts and are familiar with using them to consider revealing their account information, and we ask subjects who use Facebook to consider disclosing their data and friends-list. Thus subjects are presented with decisions that they have already confronted in their

[34] Harrison and List (2004) develop criteria for field experiments referring to the elements our experiments meet: that subjects are not aware of being studied, that the task subjects are presented with is real and consequential ("nature of commodity and stakes") and that the task is theirs to solve in the real world ("match of task and subject pool").

ordinary experience before participating in our study. The direct similarity of the experimental and real-world tasks supports the external validity of the studies.

Third, participants make decisions that directly affect their situation in the real world. Instead of completing an abstract laboratory task, they disclose a relevant part of their private data—in the first two experiments their financial contact information and in the third experiment their Facebook data.

The external validity of our findings might nevertheless be limited. First, our subjects are current and recently-graduated university students. While not a representative sample, our academically trained subjects have experience in deciding whether to share their data. It is plausible, that our study population is thus more aware of privacy risks and better prepared to protect themselves and others relative to individuals with less experience. That our comparatively sophisticated subjects nonetheless appear to be making privacy decisions using the intent heuristic may therefore rather strengthen than weaken the external validity of our results and reinforce the importance of a regulatory response. Because we are academic researchers, participants may trust us more relative to a for-profit firm which might benefit from analyzing and selling their data. That baseline trust may exaggerate subjects' tendency to make privacy choices by relying on a heuristic based on a single cue regarding trustworthiness while ignoring other information. However, we observe and are only interested in *relative* treatment effects. That is, we compare the change we observe in our subjects' choices when they can rely on the intent heuristic versus when we block their use of the heuristic. Even if subjects may be more likely to disclose their data to us than to a for-profit firm, this higher level of trust would affect both treatments equally. Therefore, it cannot explain that subjects more readily disclose their data in the *Positive Intent* treatment where they can rely on the heuristic than in the *Legal Requirement* where they cannot use the heuristic.

Another potential limitation is whether we can extrapolate our results to other jurisdictions like the US that have a different cultural and privacy law background.[35] People in different countries may value their privacy differently. However, while law and culture differ, users around the world are exposed to and interact with the same big tech players (Google, Facebook and many others), they visit the same websites, and with global media coverage, they learn about major data breaches like the Facebook scandal. Forbes reported that over 90% of Americans say they are generally aware that their personal information is being collected and stored online for commercial use.

There are few empirical studies that directly compare American and German users' assessment of their online privacy. One study (Prince & Wallsten, 2022) analyzes consumers'

[35] The legal landscape in the US is generally much less uniform than in Germany. While the EU has overarching legislation with the GDPR, US legislation and protection standards vary widely from state to state, with some federal protections such as Section 5 of the Federal Trade Commission Act.

However, when we compare California's CPRA and CCPA to the protections of the GDPR, the similarities with German/EU regulation are significant. And the CCPA has been the model for privacy laws recently passed in Colorado, Connecticut, Utah and Virginia (although not all of these laws, particularly Virginia's, offer the same suite of consumer protections). CCPA is likely to be influential in shaping additional states' privacy laws in the years to come. It is also likely to shape the conduct of businesses across the country who provide goods and services to California residents.

WTA for the protection of a number of different privacy aspects. The data show that Germans seem to value their privacy more than Americans, demanding on average more to give Facebook the right to read their texts or share their contact information or network data (i.e., connections between Facebook users as in our Experiment 3). For financial information the effect is larger. However, the difference between the US and Germany shrinks when controlling for factors such as age and online experience; for the age cohort we analyzed in our study (mostly students and younger professionals), the difference compared to a similarly-aged American subject group should be small.

A second study on privacy concerns finds opposite trends: It compares how American and German consumers perceive the sensitivity of different types of personal data (Schomakers, 2019). In this study, Americans seem to perceive their social media profiles as more sensitive than Germans; this difference was even greater for financial data, especially for bank account information.

Different privacy valuations and concerns may also not change our results and our policy recommendations, because these baselines may not affect how people make their privacy decisions, which is what we analyze in our study. Our main question is whether participants use the heuristic to estimate trustworthiness in an online privacy context. Many studies (for example McCabe et al. 2003; Charness et al 2007; Cushman 2009) that analyze intent-focused fairness judgments and reciprocity were conducted in the US and the results were largely identical to similar experiments conducted in Germany. Thus, there is no particular reason, either theoretical or empirical, to expect that German participants would use an intent-based heuristic in the domain of privacy decisions, while Americans would not.

Second, we show a treatment effect: i.e., the difference between the willingness to disclose data when our subjects can use the heuristic versus when we prevent them from using the heuristic. Even if the German participants' baseline valuation of their privacy is somewhat higher than that of their US counterparts (or vice versa), this does not mean that the relative effect of preventing them from using the heuristic must be different. That is, the heuristic-blocking intervention may still be just as effective at getting users to consider the externalities of their privacy choices and to consider institutional signals of trustworthiness. Thus, while there is no data, there are good reasons to expect that our main findings may similarly apply to the US, as they do to Germany, in spite of differences in the legal and social landscape.

Another limitation may arise from potential subject misconceptions. Even though we instruct subjects clearly, they may not fully understand (or pay attention to) what content they are revealing. With respect to the Facebook information, for example, a subject may not recognize that the data revealed includes all comments or photographs that others have associated with the subject's profile for the full time-span that the subject has operated a Facebook account. It is possible that had they fully considered these threats, subjects may not have chosen to disclose their data. This failure to fully consider risks, however, is not a limitation

to our experiment but rather an element of its ecological validity as these misconceptions are typical for privacy decisions on media platforms and services.[36]

To test the intent heuristic's impact on privacy choices, we implemented the simplest possible case where the counterparty has only a binary choice: that is, where she can only either offer privacy protection or not. This might seem artificial, as in reality businesses will typically be able to choose between more actions. However, for our experimental intervention this makes no difference. To block the heuristic all action space must be eliminated whether that includes just one alternative action or (many) more. If the law prescribes a particular act, the heuristic is no longer applicable no matter how large the counterparty's action space might have been before.

2. Internal Validity.

With respect to our study's internal validity, we account in our experimental design for the possibility that the legal policy requiring the provision of anonymous payment, which we informed subjects about in *Legal Requirement*, may not only have prevented subjects from using the heuristic but it may also have suggested to subjects they should be cautious in disclosing their data. We addressed this potential confound in Experiment 1 in the *Expressive Signal* treatment. Our results did not reveal signs of an expressive effect of the enacted legal provision; subjects were significantly more likely in *Expressive Signal* than they were in *Legal Requirement* to permit us access to their data.[37] Since the results in *Expressive Signal* were clearly significant, we did not repeat this treatment in Experiments 2 and 3.

It is also possible that some participants in our study had little data to disclose, and therefore had little at stake in their privacy decisions. For example, some subjects might be willing to disclose their data because they use Facebook only rarely. Similarly, with respect to the choice of payment methods in our experiments, some subjects may have no bank account or PayPal address and may therefore choose to pick up their earnings anonymously in cash even if they would have preferred to receive the money more conveniently. However, we assign our subjects randomly to the treatments we test and therefore should see in each treatment approximately the same number of subjects who have no bank account or PayPal address, or subjects who seldom use Facebook. They should thus not influence the results of the treatment comparisons.

An additional motivation for disclosure might be that parties want to reciprocate for the perceived effort to protect their privacy, and so disclose their data when offered anonymous payment. However, we believe that our experimental design makes reciprocity an unlikely motivation. Our instructions explicitly state that participants are not doing the researchers a favor by not choosing to be paid anonymously: we must provide the logistics for the anonymous payment regardless of whether participants choose it or not. It also seems that reciprocity as a motivation depends on trust: When you use a service and that service offers to protect your personal information, then you may trust that service and disclose, but you will not disclose your

[36] Herbert Simon coined the term ecological validity, which means that heuristics are not per se rational or irrational. Instead, whether they lead to good outcomes depends on their match to the environment in which the heuristic is used. *See also* Hertwig et. al, 2022.

[37] *See* the discussion in section III.B.4. about this point.

data out of reciprocity unless you also trust in the service to try in good faith to actually protect your privacy. Reciprocity may be an additional motivation, but it presupposes that people trust.

However, while our data show that trust in a party's good intent is driving privacy choices, we do not suggest that the heuristic may mute all doubts in the participants and that they are fully confident that the researchers will protect their privacy. We make an empirical claim: We measure that participants are more likely to disclose their data when they can use the heuristic compared to when they cannot. They may still have doubts but those doubts affect disclosures less than the heuristic does.

B. How Should the Law Intervene?

1. Preventing Use of the Intent Heuristic

Our findings suggest that privacy law should become more behaviorally aware: that is, it should be based in an empirical understanding of the process how people actually make privacy decisions. Notwithstanding recent developments in privacy laws like California's CCPA and the GDPR that include more substantive regulation providing consumers with rights to access or delete their data, and requiring businesses to put safeguards in place to protect personal data, consent remains a major cornerstone of privacy laws. As a consequence, preventing people from employing the heuristic will be important for the success of new regulation that complements notice and choice with explicit privacy rights for consumers (Solove, 2013).

We have explained and confirmed in our experiments that the structure of the heuristic presents policy-makers with a strategy to blunt its effect. The heuristic prompts its users to envision the action space of the party whose trustworthiness is evaluated and then to compare the action she has taken with the action she could have chosen instead. The party is perceived as trustworthy when she has chosen the relatively more benevolent action.[38] So a direct approach for regulation is to eliminate that action space; if the party has no choice, then potential disclosers cannot infer either a positive or negative intent from the counterfactual. In this vein—and as modeled in our *Legal Requirement* treatment—privacy regulation may oblige social media platforms and other companies to reveal when they are forced by law to offer privacy protections to make consumers understand that their behavior is not an act of beneficence.

2. A Centralized Master Template for Privacy Preferences

But how is the disclosure mandate best implemented? One of the major problems of notice and choice is that barely anyone reads the notices. Additionally, websites have learned to effectively use deceptive strategies to manipulate consumers and sanctioning such strategies is difficult as each website may use its own. So leaving disclosure in the hands of those who have all the incentive to make it fail is unlikely to succeed. A second point is that private self-regulatory projects like the "Do Not Track" header or the Platform for Privacy Preferences ("P3P") and other privacy seal programs (TRUSTArc or BBBone) failed to gain much traction due to the lack of legislation that would require companies to legally comply with the schemes (Kamara &

[38] As we pointed out in the external validity section, we have tested the simplest case in our study, where the counterparty can only offer privacy protection or not. However, for our policy question this makes no difference: Whenever the law prescribes a particular action, the heuristic is no longer applicable no matter how large the party's action space might otherwise be.

Kosta, 2006). Without legal enforcement and lack of funding the technology those entities had developed was quickly outdated.[39]

We therefore propose the creation of a centralized "master" privacy template to standardize privacy notices (Cranor, 2012).[40] Responsibility for the design of the master template should be located in an agency such as the Federal Trade Commission, currently the chief federal agency on privacy policy and enforcement.[41] The protocols by which consumers register their privacy preferences and the interface for companies seeking to collect personal data could be hosted by the FTC on a particular website—perhaps "privacy.gov." Consumers would make their choices within their user client; interfaces and protocols should be designed to work together such that devices can automatically communicate with websites or social media platforms and transmit the users' privacy choices regarding what data may or may not be collected, as well as opt outs from tracking and personal advertising.

A master privacy template could define and present a menu of the categories of data that may be collected, as well as the purposes for which the collected data may be used. Defining this menu of choices would require the regulator to break down categories of data and uses into a manageable menu and would be the subject of an ongoing effort to survey and classify to account for data innovations. The user would select the categories of data she is prepared to disclose, and the purposes for which she is willing to have those data used.[42]

Clear and simple choice options can make it much easier and less costly for customers to understand and communicate their privacy preferences. Examples for such rules are the "do not track" headers or GPC's "do not sell," or a "do not share with contractors" rule, or "do not use

[39] Global Privacy Controls is a new self-regulatory candidate that sends a "do not sell" header and that has received legal recognition: French retailer Sephora became the first company to be fined under the California Consumer Privacy Act (CCPA) for failing to honor users' GPC's and sell their data (*see* Kress et al. 2022). But users also need to adopt the GPC and while Firefox has GPC built in, the larger browsers Chrome, Safari and Edge do not. And GPC supporters will need to continue their private funding in the future to maintain the technology. So it is yet unclear whether GPC's can become a standard.

[40] For an early vision of something similar in the context of marketing communications sent to consumers, *see* Goldman, (2006) proposing mechanisms to allow consumers to disclose their marketing preferences and thereby to reduce "marketer-consumer matchmaking costs".

[41] *See* Federal Trade Commission, *Protecting Consumer Privacy and Security*, https://www.ftc.gov/news-events/topics/protecting-consumer-privacy-security.

[42] One may fear that the template could by revealing an individual's privacy preferences convey information allowing companies to exploit vulnerable customers. For example, the literature shows a clear correlation between restrictiveness of privacy preferences and advanced education. Or preferences could be strict in certain respects, such as gender or sexual orientation, while being loose in others compared to the large dataset of preference sets that companies might collect over time.

However, the template must not disclose preferences to companies at all, or only selectively. The template would require companies to channel their requests through the platform; the template can compare the request with the consumer's preferences; if there is a match, the consumer gets access and the company gets the requested information, but without ever learning the true set of personal privacy preferences.

And in further defense of the template we have to say that inferences would be drawn much easier if the full behavioral data could be used if as mostly the case under the current regime consent is given in a blank form.

data for targeted advertising," "do not store," or "do not use identifiers."[43] Privacy choices could also be organized around services that consumers use.

To make navigating the template easier, rules on the consumer side can be simpler than they need to be on the business end, and must not be concerned with filling gaps that might otherwise be exploited by abusive business strategies. Because the template organizes the communication between consumers (who define their privacy preferences) and businesses, it is up to the data-seeking business to specify its intended use and data request and up to the template to decide whether that use is consistent with the consumer's previously-expressed privacy preferences.

Once those choices are made, any party seeking disclosure would have to communicate according to standardized protocols with the client containing the individual's previously-expressed preferences. The website can give access *only* and the data is disclosed *only* if there is a match—that is, the business does not seek disclosure of data the consumer is unwilling to reveal, or asks to put that data to uses she has not approved.

When a website asks for additional data access in order to grant use of service, the template can support users with making the risk estimation. An example how risk estimations can be structured are decision trees that handle the empirical support for the decision rule in the background (Marewski & Gigerenzer, 2012). The decision tree might contain questions like: is this service located in a jurisdiction with privacy laws? does the shop have a physical branch? The template might also list trustworthy services and suggest other services that appear to fulfill the same purpose and are consistent with the individual's privacy preferences.

A master template following this design would prevent use of the intent heuristic in several ways. First, the master template is the vehicle to clarify which privacy protections consumers are entitled to (now and in the future). For example, the master template could inform California residents of the entitlement the CCPA gives them to opt out of web tracking. This clarification would make salient to consumers that the option to require tracking is not within data requesters' action space.

Second, the master template changes the framing of notice and choice. A single business asking for disclosure creates the framing of a bilateral interaction between the requester and the consumer, which makes the heuristic appear applicable even though it is maladaptive in digital environments. In contrast, when consumers determine a baseline for all potential data requesters alike, grounding privacy choices on assumed intent and *individual* trustworthiness is no longer a decision-making concept that can be applied. Once privacy decisions are re-located to a master privacy template and made ex ante, websites and platforms will have far fewer opportunities to use deceptive techniques, such as so-called "dark patterns," to lure consumers into consenting to vast data collection and use. The template is not supposed to prevent consumers from trusting a service; rather, we want them, instead of relying on perceived intent, to focus on a more reliable form of trust— institutional trust in legal regulation and the technological solutions that are deployed to comply with regulation.

[43] Also interesting is the hhtps://cusomtercommons.com project that aims to improve customer relationship management.

3. Privacy Self-Management and Privacy Nudges

We expect internet users to be more likely to use the template for protecting their privacy compared to the piecemeal choices they have to take under the current notice and choice regime. First, the template reduces personal costs of exercising privacy choices. And compliance with the template's signals and terminology would be legally enforced. This should improve the perceived efficacy of managing one's privacy. Privacy management may more likely appear to be a realistic possibility than now and consumers should be more motivated to engage, but this would be subject to empirical test.

Notwithstanding, privacy choices will always be subject to misperceptions, such as an underestimation of long-term privacy risks. However, the template could help consumers reduce some of the most significant behavioral threats to privacy self-management that we have identified above. The template can help mitigate hyperbolic discounting by enabling consumers to make decisions ex ante and for many websites at once, such that immediate benefits are less salient relative to privacy costs. For particular data requests, the template can explain the privacy risks better to mitigate the difficulty of estimating the weight of cumulative privacy risks. The ex ante decision also could help mitigate the illusion of control bias, by providing simple rules that focus on downstream risks. It should reduce frustration with privacy decisions and avoidance of privacy decision-making, by increasing the efficacy of privacy management. The reduction in cognitive load should also reduce the attractiveness of using the heuristic.

Our centralized approach would also enable the regulator to use framing and nudges to countervail biases and improve consumers' privacy choices. Since currently businesses can design their own notices, unsurprisingly they present consumers with an opt-out framework: if consumers do not want their data being traced and processed, they have to opt out of the many ways to collect and use data. We know from many studies that defaults influence behavior (Mertens et al. 2022); for example, in the willingness to donate organs (Johnson & Goldstein, 2004), enrollment in retirement savings plans (Thaler & Benartzi, 2004), in health care decisions (Narula et al., 2014) or decisions about environmental protection (Pichert & Katsikopoulos, 2008). The regulator could leverage this effect in the master template design and present privacy choices as opt ins to consumers.

We already see this power of defaults at work: In 2021, Apple changed the requirements for apps in their store; apps must now upon installation prompt users to state whether they want to *opt in* to data tracking. The effect of the switched default—from data tracking unless the user opts out, to a ban on data tracking unless the user opts in—was striking. Before the default flip, only 25% of users indicated they wished not to be tracked.[44] After the default flip, 79% of users stayed with the new default and have banned tracking.[45] It will be interesting to observe what equilibrium preferences may settle into overtime.

[44] *See* https://www.nytimes.com/wirecutter/blog/apple-privacy-labels-tracking/

[45] In the beginning 96% stayed with the default. *See* https://www.flurry.com/blog/ios-14-5-opt-in-rate-att-restricted-app-tracking-transparency-worldwide-us-daily-latest-update/. There is also a large variance between apps suggesting that competition over privacy terms is growing.

Defaults can also help diminish the impact of the intent heuristic.[46] An opt-in default suggests to users that initial entitlements lie with them, not the business requesting disclosure. If a website offers to collect only limited categories of information for narrow and seemingly legitimate purposes, users who start with an opt-in default may not view this as a benevolent practice, — because if the initial entitlement lies with the consumer, the requester may be perceived to enjoy a more constrained action space from the start (Depoorter & Tontrup, 1999). The Apple example also shows how crucial the centralized control over the app design is. Websites are testing every day to see which prompt may provide them with the highest opt in rates.[47]

4. Compliance and Enforcement.

The notice and choice framework of privacy regulation was essentially libertarian, and, as such, was supposed to be largely self-enforcing. In contrast, our suggested regulatory reforms of notice and consent, focus on the establishment of a centralized regulation and entail a greater focus on enforcement to make individual privacy choices effective. That effort will focus on several related goals. First, to ensure that the template and its facilitating technologies are not subject to circumvention or abuse, the requester would be legally barred from soliciting disclosure of data unless it has obtained consent through the standardized template that the potential discloser has already used to establish his or her privacy preferences. Standardization of protocols and clients which enable automation also allow for much easier random tests of compliance with the requirements of the master template. It would also be easier to control whether websites respect the limits of the consent the consumer has given.

Moreover, standardized, machine-readable privacy terms not only increase transparency, they should also enable the consumer to compare policies across websites and platforms. Websites whose privacy policy falls short of the user's preferences laid down in the template can be automatically rejected. This could lead to more competition around privacy terms and voluntary compliance with consumers' preferences (see in more detail in the appendix). To get traffic and benefit from disclosure, businesses would be incentivized to align their requests with what most consumers are willing to provide. Failing to do so will put a business in the position of regularly requesting that users flip their already-established (and therefore likely at least somewhat sticky) defaults. That is going to put parties who are making many such requests at a disadvantage vis-à-vis firms that are not.

The template can also add teeth to the FTC's Section 5 authority. Currently FTC's enforcement power to combat dark patterns is limited, even though those strategies very effectively induce consent.[48] Dark patterns do not fit into the FTC's deception and unfairness doctrine. A particular web design that effectively emphasizes some privacy options over others is not deceptive in the law's sense and a practice would typically only be considered unfair, if it caused economic harm to consumers, while the harm may just be inconvenience, leading to disclosure. Also it will often be reasonably avoidable, for instance consumers can ignore nagging

[46] Notice that the CCPA does more than simply framing privacy choices: for example if the consumer exercises her right to get data deleted the business cannot stop providing services to her.

[47] https://www.appsflyer.com/blog/tips-strategy/apps-boost-att-opt-in/

[48] Even thou the FTC has settled in some cases like Age of Learning and AMG Capital Management the current doctrine would need significant terminological development to reliably combat dark patterns.

prompts. And even if applicable, Section 5 could only result in injunctions after the fact, if violations were found. The master template by contrast prevents that strategic design of privacy policies and consent forms can be used. So companies that try reach the same outcomes they otherwise achieve using dark patterns, would have to outright disregard the explicit limits of the consent given. This is easily captured and enjoined under Section 5. The template can thus facilitate enforcement by making the application of existing tools easier.

5. Weighing Costs and Benefits of Data Disclosure.

One characteristic shared by both notice and choice and our approach is that neither suggests an optimal level for either individual privacy protection or data sharing (Simeon, 2007). As the GDPR aims to do, privacy regulation must seek to reduce the risks of disclosure, and also consider the benefits of data sharing. For example, the disclosure of a person's finances and creditworthiness can cause substantial individual harm, but can also produce benefits: disclosure is necessary to create reliable credit ratings and scores, and these services are themselves public goods that foster social trust and cooperation. Collection and sharing of health data (Bentzen, 2021) also obviously produces both harms and benefits, as does collecting and processing virtually any form of confidential information. The costs and benefits of disclosure (that is, on the total social benefits of disclosure versus the risks of disclosure) will vary substantially across domains. One advantage of our approach is that rules intended to regulate the level of disclosure in particular domains can be integrated into the centralized privacy template. By contrast when consumers are employing the heuristic on a website-by-website basis the provision of legally-mandated substantive protections creates particular risks that businesses misrepresent consumer rights as voluntary protections. .

VI. Conclusion.

For many reasons the current notice and choice privacy framework fails to empower individuals in effectively making their own privacy choices. We show in our Article that at the core of this failure is a cognitive error. Notice and choice caters to a heuristic that people employ to make privacy decisions. This heuristic is meant to judge trustworthiness in face-to-face-situations. In the online context, it distorts privacy decision-making and leaves potential disclosers vulnerable to exploitation.

From our experimental evidence exploring the heuristic's effect, we conclude that privacy law must become more behaviorally aware. Specifically, the consent requirement should be redesigned to intervene in the cognitive mechanisms that keep individuals from making better privacy decisions. A behaviorally-aware privacy regime of notice and consent should as we suggest, standardize and simplify the framework for making effective privacy choices. To achieve these goals, we propose a master privacy template which requires consumers to define their privacy preferences in advance—doing so avoids presenting the consumer with a concrete counterparty, and this, in turn, prevents them from applying the intent heuristic and reduces many other biases that affect privacy decision-making. Our data show that blocking the heuristic enables consumers to consider relevant privacy cues and be considerate of externalities their privacy decisions cause.

A master template can provide a more effective platform to make consumer consent feasible, less subject to manipulation, and a more meaningful part of privacy regulation. .

Appendix

In this Appendix we first present several additional arguments detailing why the law should intervene to mute the effect of the intent heuristic in privacy decision-making. We then describe several additional benefits of re-framing privacy law around standardized, ex ante decision-making employing the mechanism of a Master Privacy Template.

A. Should the Law Intervene?

1. Maladaptation Distorts Privacy Decision-making

Our results suggest that any effective privacy regulation must be behaviorally-informed and must respond to the cognitive processes that influence individuals' real-world privacy decision-making. We have seen that the intent heuristic is maladaptive in a digital environment where people interact with complex interfaces and algorithms rather than other individuals. In that environment, a heuristic that tries to infer the intention, goodwill and trustworthiness of a counterparty is a flawed concept. Privacy decisions should be based, for the main part, on reputational and technological cues regarding the safety of a particular online seller or platform.[49]

Our experiments show that the intent heuristic leads individuals to disregard even strong negative signals about credibility (Experiments 2 & 3). In particular, because of its bilateral focus, the heuristic is not suited for domains where privacy choices have strong externalities, as is typical for social media platforms. Here the heuristic appears to blind individuals for the social privacy risks of their privacy choices. In turn, once the heuristic is blocked, individuals' decisions become much more sensitive to the consequences their choices can have for others (Experiment 3).

These findings have an important knock-on effect: If parties do not respond to otherwise valid signals of credibility that the heuristic does not process (as in Experiment 2), then notice and choice does not incentivize social media platforms and data businesses to invest in implementing and providing safety technologies such as end-to-end-encryption or two-factor authentication. The result is that the heuristic affects not just individuals' privacy decision-making, but also overall demand in the market for privacy protection.

2. The Heuristic undermines Development of the Privacy Law

Compounding these concerns, our study suggests that exploitation of the intent heuristic may be relatively easy to accomplish. In our *Positive Intent* treatment, we make no overt representation regarding privacy; we merely offer an anonymous payment method, yet that appears to induce substantial trust. While our trust signals are sent in good faith—we do not

[49] For example, whether stored personal data is encrypted, what the business model of the service provider is, what the provider has used the information for in the past, whether data breaches came to light and so forth.

misuse subjects' confidential bank information or the Facebook data they revealed to us—they are both cheap and easily manipulated. And this is true in the real world as well.

For example, Facebook claims in its privacy policy that it does not sell personal data. What is meant to sound like an affirmative statement of privacy protection obfuscates that Facebook sells identifiable data indirectly: It allows advertisers to exactly choose a target group by race, gender, age, education, partnership, hobbies and location: for example, Hispanic, woman, 25-35 years, college degree, single, surfer from LA. The advertisement is shown only to the target group; if a user clicks on the ad, she is immediately identified as fitting all those categories. The websites that Facebook ads direct people to often have their own trackers, which can uniquely identify a person by IP address and device ID. So the law sets up barriers to selling data, and yet Facebook's privacy policy suggests that Facebook is providing this protection as a result of its choice. This false signal is likely to invoke a trust judgment in users employing the heuristic.

Another example is that companies that fall under Californian jurisdiction may claim that because the business cares about its customers' privacy, the customer may ask at any time to have all personal data previously collected deleted. The seemingly beneficent offer may induce a person employing the heuristic to conclude that the business is trustworthy, as it could apparently also have refrained from making this offer. But this is a framing that relies on a false signal about the action space: in fact, asking for deletion is a right granted by the CCPA[50] This scenario will become increasingly important for the effectiveness of new privacy laws like the CCPA. In response to the failures of notice and choice, privacy regulation continues to move toward granting explicit privacy rights.[51] In the case of the CCPA consumers are guaranteed the right to have their data deleted, and to opt out of their data being sold or shared. Consumers are also given a set of rights to limit the use of their data for purposes they indicate.[52] In Europe the General Data Protection Regulation (GDPR)[53] has also created new privacy rights in Chapter 3 (Art. 12-23) including rights to object against processing personal data for marketing purposes, the right to access stored personal data, and the right to have data erased. The law also mandates technical and organizational measures like using two-factor authentication and end-to-end encryption, as well as regulations governing when personal data is processed and limiting access to that data within organizations that collect it. However, most of these protections are legal rights, so consumers may or may not exercise them. And given individuals' use of the heuristic, the result may be that businesses frame legal mandates as voluntarily provided—i.e., as acts of beneficence. Unless effective disclosure is made that the law is responsible for the protections, the net result may be, rather than more cautious and socially considerate privacy choices, that consumers do not exercise their protections. They may do so in response to the input they get from the heuristic that tells them to trust the requesting party whose beneficence they perceive—incorrectly—to be the source of the protection they are "offered". Thereby, our study allows us to see that the CCPA's rule that contracts may not require users to waive their rights is likely

[50] https://oag.ca.gov/privacy/ccpa#sectiond.

[51] This development is evidenced by the GDPR in Europe or the California Consumer Privacy Act.

[52] https://oag.ca.gov/privacy/ccpa

[53] https://gdpr.eu/what-is-gdpr/

insufficient to ensure that consumers are not manipulated away from effectively exercising their granted rights.

In sum, a company that understands how consumers make privacy decisions can frame privacy protections as voluntary practices even though consumers have a right to that protection (Klusowski, 1978).

3. Intent Heuristic Reinforces other Biases

Privacy decisions are influenced by a number of psychological biases. The intent heuristic likely exacerbates many of these biases. Here are a few examples.

Disclosure is often chosen for immediate benefits, such as access to services, or simply for greater convenience. Privacy choices are therefore often cited as an example of time discounting. The heuristic reinforces time discounting, both because its focus on personal trust draws attention away from the future consequences of disclosure, and because it encourages the implicit assumption that if the counterparty has good intent, she will protect the data in the future.

Many users have difficulty assessing the weight of cumulative privacy risks, and therefore often fail to take even simple steps to protect their privacy (Svirsky, 2021). The heuristic makes them even less receptive to information relevant to protecting against their own privacy risks and the risks faced by others.

The heuristic's focus on the intent of the data seeking party also likely reinforces the bias that individuals often feel their privacy is protected when they have control over who they share personal information with, while they may have little control over how that information is used once it is shared (Brandimarte, 2015). If the counterparty is perceived as trustworthy to protect the data, concerns about what happens to the data after disclosure may be supplanted, as suggested by Experiment 3.

B. Additional Benefits of the Master Privacy Template

Finally, there are two additional benefits of the shift to a standardized, ex ante privacy regime—increased competition around privacy, and a potential salutary shift in privacy norms—that are worth brief discussion.

Competition. Under the notice and choice regime competition over privacy terms is barely functioning, in part because notices are rarely read and because it is often easier for businesses to manipulate customers into consent than to try to accommodate their preferences. The heuristic also plays a part in this. The heuristic is easy to exploit and once a data-requesting business is classed as trustworthy, it may barely matter what type of data the business is asking for. Businesses might even have to compete over how to manipulate consumers effectively to provoke disclosure rather than try to meet consumers' privacy preferences as they exist in the absence of the heuristic. Obviously, competition always needs a framework to operate within. Notice and choice left creation and maintenance of this framework to self-governance and perhaps to industry standards of fair behavior (Simeon, 2007).

The centralized template requires businesses to submit transparent requests for data collection and data use. Just as consumers have to fill out a master template, businesses would have to do the same. With data requests standardized and implementation of privacy policies monitored through user clients, businesses will compete to better meet consumers' preferences by focusing or scaling down their data requests. We have seen above that our approach should also improve enforcement, so it is less of an option for businesses to simply not comply with the data terms offered to customers and collect more data than agreed upon.[54]

Competition should also be facilitated by an informational aspect of our approach: Connecting privacy clients and browsers enables consumers to focus on a market screened by their personal privacy preferences. They may learn, for example, that a particular mobile phone carrier they used in the past does not accommodate their privacy preferences, while another one they may not have considered before does (Ramachandran & Balasubramanian, 2020).

Social norms and the collective good. We have also seen that the legal regime we propose should help produce more carefully considered privacy decisions, which will in turn reduce the negative externalities of privacy decision-making. Without the heuristic blinding them, individuals become more sensitive to social privacy risks and more likely to refrain from disclosing data associated with those risks. But in addition, the regulation likely causes an indirect effect on social privacy behavior. If, in the absence of the heuristic, more individuals are willing to consider the externalities of their privacy choices, then this may push others to do the same. Most individuals are willing to contribute to collective goods like privacy protection or environmental protection. However, they condition their cooperation and contribute as long as they expect others to do the same (Fairchild et al., 2015; Leni, 2001). Thus, increasing the number of people who are considerate of privacy externalities may lay the foundations for reciprocity and a greater degree of conditional cooperation—making others more likely to contribute to the collective good of privacy as well. Blocking the heuristic may thereby shift the equilibrium of privacy behavior and establish new privacy norms in communities (McAdams, 2000).

References

1. Acquisti, A. (2004). Privacy in Electronic Commerce and the Economics of Immediate Gratification. In Proceedings of the 5th ACM Conference on Electronic Commerce (pp. 21–29).

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

3. Bamberger, K. A., et al. (2020). Can you pay for Privacy? Consumer Expectations and the Behavior of Free and Paid Apps. Berkeley Technology Law Journal, 35, 327.

4. Barocas, S., & Nissenbaum, H. (2014). Big Data's End Run Around Anonymity and Consent. In J. Lane, V. Stodden, S. Bender, & H. Nissenbaum (Eds.), Privacy, Big Data, and the Public Good: Frameworks for Engagement (pp. 44-75).

[54] Most apps collect as much data as they can to survive in the market https://www.digitalinformationworld.com/2022/10/android-apps-collect-way-more-data-than.html.

5. Bentzen, H. B., et al. (2021). Remove obstacles to sharing health data with researchers outside of the European Union. Nature Medicine, 27, 1329.

6. Beresford, A. R., Kübler, D., & Preibusch, S. (2010). Unwillingness to Pay for Privacy: A Field Experiment. https://www.ssoar.info/ssoar/bitstream/handle/ document/23832/ssoar-2010-beresford_et_al-unwillingness_to_pay_for _privacy.pdf?sequence=1

7. Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, Reciprocity, and Social-History. Games and Economic Behavior, 10, 122.

8. Bewley, Truman F. Why wages don't fall during a recession. Cambridge, MA: Harvard University Press, 1999.

9. Bietti, E. (2020). Consent as a Free Pass: Platform Power and the Limits of the Informational Turn. Pace Law Review, 40, 307.

10. Brandimarte, L., et al. (2012). Misplaced Confidences: Privacy and the Control Paradox. Social Psychology & Personality Science, 4, 340.

11. Cameron, L. (1999). Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia. Economic Inquiry, 37, 47.

12. Chang, A. (2018). The Facebook and Cambridge Analytica Scandal, Explained with a Simple Diagram. https://www.vox.com/policy-and-politics/2018/3/23/17151916/facebook-cambridge-analytica-trump-diagram

13. Charness, G., & Levine, D. I. (2007). Intention and Stochastic Outcomes: An Experimental Study. The Economic Journal, 117, 1051.

14. Charness, G., & Rabin, M. (2002). Understanding Social Preferences with Simple Tests. Quarterly Journal of Economics, 117, 817.

15. Cohen, J. E. (2013). What Privacy Is For. Harvard Law Review, 126, 1904.

16. Cranor, L. F. (2012). Necessary But Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice. Journal of Telecommunications and High Technology Law, 10, 273.

17. Cushman, F., Dreber, A., Wang, Y., & Costa, J. (2009). Accidental Outcomes Guide Punishment in a "Trembling-Hand" Game. PLOS One, 4, Article e6699.

18. Czerlinski, J., Goldstein, D., & Gigerenzer, G. (1999). How Good Are Simple Heuristics? In G. Gigerenzer, P. M. Todd, & The ABC Research Group, Simple heuristics that make us smart (pp. 97–118). Oxford University Press.

19. Depoorter, B., & Tontrup, S. (2012). How Law Frames Moral Intuitions: The Expressive Effect of Specific Performance. Arizona Law Review, 54, 673.

20. Derlega, V., Metts, S., Petronio, S., & Margulis, S. (1993). Self-Disclosure. Newbury Park, CA: Sage Publications.

21. Dhar, R., & Simonson, I. (1992). The Effect of the Focus of Comparison on Consumer Preferences. Journal of Marketing Research, 29, 430.

22. Dufwenberg, M., & Kirchsteiger, G. (2004). A Theory of Sequential Reciprocity. Games and Economic Behavior, 47, 268.

23. Etter, L., & Frier, S. (2018). Facebook App Developer Kogan Defends His Actions with User Data. Bloomberg. https://www.bloomberg.com/news/articles/2018-03-21/facebook-app-developer-kogan-defends-his-actions-with-user-data

24. Fairfield, J. A. T., & Engel, C. (2015). Privacy as a Public Good. Duke Law Journal, 65, 385.

25. Falk, A., Fehr, E., & Fischbacher, U. (2003). On the Nature of Fair Behavior. Economic Inquiry, 41, 20.

26. Federal Trade Commission. (1998). Privacy Online: A Report to Congress (Report No. Priv-23a). https://www.ftc.gov/sites/default/files/documents/reports/privacy-online-report-congress/priv-23a.pdf

27. Federal Trade Commission. (2015). Internet of Things: Privacy & Security in a Connected World. https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things-privacy/150127iotrpt.pdf

28. Fehr, E., & Falk, A. (1999). Wage Rigidity in a Competitive Incomplete Contract Market. Journal of Political Economy, 107, 106.

29. Fehr, E., Gächter, S., & Kirchsteiger, G. (1997). Reciprocity as a Contract Enforcement Device: Experimental Evidence. Econometrica, 65, 833.

30. Fischbacher, U., Gächter, S. & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. Economics Letters, 71, 397.

31. Fox, C. R., & Tversky, A. (1995). Ambiguity Aversion and Comparative Ignorance. Quarterly Journal of Economics, 110, 585.

32. Gallup. (2023) Computers and the Internet. http://news.gallup.com/poll/1591/Computers-Internet.aspx

33. Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic Decision Making. Annual Review of Psychology, 62, 451.

34. Gigerenzer, G., & Goldstein, D. (1996). Reasoning the fast and frugal way: Models of bounded rationality. Psychological Review, 103, 650.

35. Goldman, E. (2006). A Coasean Analysis of Marketing. Wisconsin Law Review, 2006, 1151.

36. Griffin, D., Liu, W., & Khan, U. (2005). A New Look at Constructed Choice Processes. Marketing Letters, 16, 321.

37. Güth, W., & Kirchkamp, O. (2012). Will You Accept Without Knowing What? The Yes-No Game in the Newspaper and in the Lab. Experimental Economics, 15, 656.

38. Güth, W., Schmittberger. R., and Schwarze, B. (1982). An Experimental Analysis of Ultimatum Bargaining. Journal of Economic Behavior and Organization. 3 (4): 367–388.

39. Harrison, G. W., & List, J. A. (2004). Title of the Article. Journal of Economic Literature, 42, 1009.

40. Hartzog, W. (2018). Privacy's Blueprint: The Battle to Control the Design of New Technologies. Harvard University Press.

41. Hertwig, R., Leuker, C., Pachur, T., Spiliopoulos, S., & Pleska, T., J. (2022). Studies in Ecological Rationality. Topics in Cognitive Science, 14, 467.

42. Hsee, C. K., et al. (1999). Preference Reversals Between Joint and Separate Evaluations of Options: A Review and Theoretical Analysis. Psychological Bulletin, 125, 576.

43. Hsee, C. K., et al. (2003). Lay Rationalism and Inconsistency Between Predicted Experience and Decision. Journal of Behavioral Decision Making, 16, 257.

44. Johnson, E. J., & Goldstein, D. G. (2004). Defaults and Donation Decisions. Transplantation, 78, 1713.

45. Kamara, I., & Kosta, E. (2016). Do Not Track initiatives: regaining the lost user control. International Data Privacy Law, 6, 276.

46. Klusowski, J., et al. (2021). Does Choice Cause an Illusion of Control? Psychological Science, 32, 159.

47. Kress, L. E., Stine, M. M., & Trifon, T. L. (2022). The Sephora Settlement Is Just the Start. Locke Lord's Privacy & Cybersecurity Newsletter. https://www.lockelord.com/newsandevents/publications/2022/12/ccpa-enforcement-sephora-settlement

48. Marewski, J., & Gigerenzer, G. (2012). Heuristic decision making in medicine. Dialogues in Clinical Neuroscience, 14, 77.

49. McAdams, R. (2000). An Attitudinal Theory of Expressive Law. Oregon Law Review, 79, 339.

50. McCabe, K., Rigdon, M., & Smith, V. (2003). Positive reciprocity and intentions in trust games. Journal of Economic Behavior and Organization, 52, 267–275.

51. McDonald, A. M., & Cranor, L. F. (2008). The Cost of Reading Privacy Policies. Journal of Law and Policy for the Information Society, 4, 543.

52. Stephanie Mertens, Mario Herberz, Ulf J.J. Hahnel & Tobias Brosch (2022). The Effectiveness of Nudging: A Meta-Analysis of Choice Architecture Interventions Across Behavioral Domains, 119 Proceedings Nat`l Acad. Sci. 1, 4.

53. Narayanan, A., Mathur, A., Chetty, M., & Kshirsagar, M. (2020). Dark Patterns: Past, Present and Future. ACM Queue, 18, 68.

54. Narula, T., Ramprasad, C., Ruggs, E. N. & Hebl, M. R. (2014). Increasing colonoscopies? A psychological perspective on opting in versus opting out. Health Psychology, 33, 1426.

55. Nissenbaum, H. (2010). Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford University Press.

56. Obar, J. A., & Oeldorf-Hirsch, A. (2020). The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services. Information, Communication & Society, 23, 128.

57. Petronio, S. (2000). The Boundaries of Privacy: Praxis of Everyday Life. In Balancing the Secrets of Private Disclosures, Lawrence Erlbaum Associates Publishers. (pp. 37-49).

58. Pew Research Center. (2019). Americans And Privacy: Concerned, Confused and Feeling a Lack of Control Over Their Personal Information. https://www.pewresearch. org/ internet/2019/11/15/americans-attitudes-and-experiences-with-privacy-policies-and-laws/

59. Pichert, D., & Katsikopoulos, K. V. (2008). Green defaults: Information presentation and pro-environmental behaviour. Journal of Environmental Psychology, 28, 63.

60. Posner, R. A. (1981). The Economics of Privacy. American Economic Review, 71, 405.

61. Prince, J., & Wallsten, S. (2022). How Much is Privacy Worth Around the World and Across Platforms? Journal of Economics & Management Strategy, 31, 841.

62. Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics. American Economic Review, 83, 1281.

63. Ramachandran, S., & Balasubramanian, S. (2020). Examining the Moderating Role of Brand Loyalty among Consumers of Technology Products. Sustainability, 12, 9967.

64. Regan, P. M. (1995). Legislating Privacy: Technology, Social Values, and Public Policy. The University of North Carolina Press.

65. Richards, N., & Hartzog, W. (2019). The Pathologies of Digital Consent. Washington University Law Review, 96, 1461.

66. Rosenfeld, L. B. (2000). Overview of the Ways Privacy, Secrecy and Disclosure are Balanced in Today's Society. In Balancing the Secrets of Private Disclosures. The University of North Carolina Press.

67. Roth, A., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. American Economic Review, 81, 1068.

68. Schaechtele, S., Gerstenberg, T., & Lagnado, D. (2011). Beyond Outcomes: The Influence of Intentions and Deception. In Proceedings of the 33rd Annual Conference of the Cognitive Science Society (pp. 1860–1865).

69. Schomakers, E.-M., Lidynia, C., Müllmann, D., & Ziefle, M. (2019). Internet users' perceptions of information sensitivity. International Journal of Information Management, 46, 142.

70. Simeon, R., et al. (2007). Private credit in 129 countries. Journal of Financial Economics, 84, 299.

71. Simon, H. A. (1977). The logic of heuristic decision-making. In Models of Discovery and Other Topics in the Methods of Science, Boston Studies in the Philosophy of Science, 54, 154.

72. Smith, J. H. (1994). Managing Privacy: Information Technology and Corporate America. UNC Press Books.

73. Solove, D. J. (2013). Privacy Self-Management and the Consent Dilemma. Harvard Law Review, 126, 1880.

74. Stigler, G. J. (1980). An Introduction to Privacy in Economics and Politics. Journal of Law Studies, 9, 623.

75. Strandburg, K. J. (2013). Free Fall: The Online Market's Consumer Preference Disconnect. University of Chicago Legal Forum, 95, 143.

76. Svirsky, Dan (2021). Why Do People Avoid Information About Privacy?. Journal of Law & Innovation, 2, 23.

77. Thaler, R. (1988). Anomalies: The Ultimatum Game. Journal of Economic Perspectives, 2, 195.

78. Thaler, R. H., & Benartzi, S. (2004). Save More TomorrowTM: Using Behavioral Economics to Increase Employee Saving. Journal of Political Economy, 112, 164.

79. Tontrup, S., & Sprigman, C. J. (2022). Self-Nudging Contracts and the Positive Effects of Autonomy—Analyzing the Prospect of Behavioral Self-Management. Journal of Empirical Legal Studies, 19, 594.

80. Tufekci, Z. (2018). Facebook's Surveillance Machine. The New York Times. https://www.nytimes.com/2018/03/19/opinion/facebook-cambridge-analytica.html

81. Tversky, A., & Kahneman, D. (1974). Judgment Under Uncertainty: Heuristics and Biases. Science, 185, 1124.

82. Tversky, A., Slovic, P., & Kahneman, D. (1990). The Causes of Preference Reversal. American Economic Review, 80, 204.

83. Viljoen, S. (2021). A Relational Theory of Data Governance. Yale Law Journal, 131, 573.

84. Waldman, A. E. (2020). Privacy Law's False Promise. Washington University Law Review, 97, 773.

82. White, T. B. (2004). Consumer Disclosure and Disclosure Avoidance: A Motivational Framework. Journal of Consumer Psychology, 14, 41.