Sharing is daring: An experiment on consent, chilling effects and a salient privacy nudge

Hermstrüwer, Yoan ; Dickert, Stephan

Sharing is daring: An experiment on consent, chilling effects and a salient privacy nudge

Yoan Hermstrüwer [a],[*], Stephan Dickert [b]

[a] *Max Planck Institute for Research on Collective Goods, Bonn, Germany*
[b] *Queen Mary University of London, United Kingdom*

ABSTRACT

Privacy law rests on the assumption that government surveillance may increase the general level of conformity and thus generate a chilling effect. In a study that combines elements of a lab and a field experiment, we show that salient and incentivized consent options are sufficient to trigger this behavioral effect. Salient ex ante consent options may lure people into giving up their privacy and increase their compliance with social norms – even when the only immediate risk of sharing information is mere publicity on a Google website. A right to be forgotten (right to deletion), however, seems to reduce neither privacy valuations nor chilling effects. In spite of low deletion costs people tend to stick with a retention default. The study suggests that consent architectures may play out on social conformity rather than on consent choices and privacy valuations. Salient notice and consent options may not merely empower users to make an informed consent decision. Instead, they can trigger the very effects that privacy law intends to curb.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

'. . .the very minute a thought is threatened with publicity it seems to shrink toward mediocrity.'[1]

Sharing is daring. Any piece of personal information shared on Facebook, Google or Netflix might sooner or later be subject to an unwanted gaze (Rosen 2001; Levmore and Nussbaum, 2010). A major problem of privacy law is that we lack compelling empirical evidence of how the widespread willingness to share personal information is related to behavioral conformity. Some scholars have rejected the sense of "creepiness" associated with the free flow and permanent storage of personal information (Posner, 1978, 1981, 2008; Stigler, 1980). Conceptually, however, most scholars agree

that privacy losses might trigger a pressure to comply with normative expectations and that this can be defined as a chilling effect on civil liberties (White and Zimbardo, 1975; Foucault, 1977; Schauer, 1978; Goodin and Jackson, 2007; Richards, 2013; Marthews and Tucker, 2015).

One of the main objectives of privacy law is to forestall chilling effects (Sunstein, 2003, 2009; Zittrain, 2008; Richards, 2013). The U.S. Supreme Court has been reluctant to recognize chilling effects as a sufficient ground to find standing for the plaintiff in a recent Fourth Amendment case (*Clapper v. Amnesty International USA*, 133 S. Ct. 1138 [2013]). The German Constitutional Court, however, acknowledges that the uncertainty and opacity of data processing and storage may cause a "diffuse threat" and an excessive level of conformity, thereby impinging on fundamental rights (BVerfGE 65, 1 [1983]; BVerfGE 125, 260 [2010]; Bull, 2011). The upshot of these statements is that they refer to government surveillance.

---

* Corresponding author.
  E-mail address: hermstruewer@coll.mpg.de (Y. Hermstrüwer).
[1] O.W. Holmes, The Poet at the Breakfast Table, 1872, p. 403.

In this paper, we investigate the behavioral dimension of chilling effects and, more specifically, the factors driving these effects in case of private-sector surveillance. Data processing on digital platforms, such as Facebook or Google, usually results from consent (Benndorf et al., 2015; Hermstrüwer, 2017). Consent is one of the core legitimizing tools for the processing of personal information under European privacy law (Art. 8 § 2 of the Charter of Fundamental Rights of the European Union; Art. 8 of the European Convention on Human Rights; Art. 6 of the EU General Data Protection Regulation [EU-GDPR][2]). Privacy, as guaranteed by the consent principle, reduces the expected cost associated with norm-deviations, because it mitigates the risk of social image harm or of social sanctions (McAdams, 1997; McAdams and Rasmusen, 2007). When people give their consent, they incur a risk of revealing information about their non-compliance with social norms, broadly understood as jointly recognized understandings or expectations about types of behavior that are pro- or prescribed (Elster, 1989; Bicchieri, 2006). This risk of having their social image stained or being sanctioned for norm deviations is exacerbated by the fact the personal information is hardly ever deleted. People who once deviated from a norm incur the risk of being (mis-)judged on the basis of bits and pieces of outdated information (Rosen, 2001; Nissenbaum, 2010; Mitrou, 2010) and having their social image stained for a lifetime (Zittrain, 2008). To allay the fear of social sanctions and dampen chilling effects, privacy scholars and activists have proposed to establish a right to delete personal information: a right to be forgotten (Bannon, 2006; Lessig, 2007; Mayer-Schönberger, 2009; Rosen, 2011, 2012; Cohen, 2012, 2013).[3] Some of these scholars have claimed that chilling effects may best be cushioned in case of automatic deletion (Mayer-Schönberger, 2009; European Data Protection Supervisor, 2011; de Terwangne, 2012).[4] Without really settling the legal dispute, Art. 17(1) of the recently adopted EU-GDPR allows for two default designs of the right to be forgotten: a *retention default* (information deleted on request after a set period) and a *deletion default* (information automatically deleted after a set period).

Lawyers and economists have largely ignored the behavioral link between salient consent options, chilling effects and the right to be forgotten in case of private-sector surveillance (see Britz, 2010; Hermstrüwer, 2016). To the best of our knowledge, neither the behavioral impact of salient and incentivized consent options nor the effect of the proposed default designs of the right to be forgotten have been studied in an empirically rigorous way. Our study tackles this lack of empirical evidence in a stylized experimental setting. On the one hand, our study aims at identifying whether salient consent options trigger chilling effects in a standard dictator game (i.e. adherence to the norm of a 50–50 split) and how they affect privacy valuations. On the other hand, we test whether a *deletion default* facilitates deviations from the 50–50 norm and how it affects privacy valuations.

Our experimental results suggest that people may be subject to a chilling effect when facing a salient consent option and consenting to the publication of information about their real identity and their behavior in the dictator game. This effect manifests itself through different fractions of selfish choices. Hence, exploiting social image concerns through salient consent options by itself does not spur

egalitarian choices. By contrast, the design of the right to be forgotten affects neither privacy valuations nor behavior in the dictator game. Finally, in our study we found that most users do not actively claim deletion.

The remainder of this article is organized as follows. In Section 2, we briefly review the relevant literature on behavioral economics of privacy. In Section 3, we present our experimental design. Section 4 provides an overview of our behavioral predictions. In Section 5, we report our experimental results. We conclude with a discussion of the results and their legal implications in Sections 6 and 7.

2. Background in behavioral economics of privacy

While rational choice theorists have claimed that individuals make consistent and utility-maximizing choices based on stable privacy preferences (Stigler, 1980; Posner, 1981), behavioral economists have demonstrated that preference uncertainty is particularly strong in the domain of privacy decision making (Acquisti, 2004; John et al., 2011). Hence, privacy preferences are endogenous to the elicitation methods, the cues and the framing used in privacy-relevant decision contexts (Acquisti and Grossklags, 2005b; Acquisti and Grossklags, 2007; Hui et al., 2007; John et al., 2011; Grimmelmann, 2009). Our experiment dovetails with previous studies in behavioral economics of privacy, but it aims to fill some of the gaps and addresses some of their shortcomings.

First, privacy experiments have often been conducted as survey studies (Ben-Shahar and Chilton, 2016) without being incentivized. Such studies, however, are prone to hypothetical bias: individuals may express privacy preferences that reflect their privacy attitudes, but not their actual privacy choices (Kahneman and Ritov, 1994; Hui et al., 2007; John et al., 2011; Brandimarte et al., 2013; Benndorf and Normann, 2014). In addition, some surveys elicit privacy valuations of non-verifiable personal information. When the truthfulness of information cannot be verified, the dominant strategy for privacy-sensitive individuals is to disclose false information (Jentzsch et al., 2012), which hampers the analysis of privacy valuations. We solve the problem of hypothetical bias by incentivizing consent options. In addition, we use *real* and *verifiable* information including the real names of participants to strengthen the external validity of our results. Following the ethical standards of economic experiments, however, no deception of participants is involved.

Second, incentivized experiments have investigated the willingness to accept money (*WTA*) in exchange for their disclosure of personal information (Chellapa and Sin, 2005; Huberman et al., 2006; Wathieu and Friedman, 2007), the willingness to pay money (*WTP*) for privacy protection (Acquisti and Grossklags, 2005a; Tsai et al., 2011) or the discrepancy between *WTA* and *WTP*, i.e. endowment effects with regard to privacy (Hui and Png, 2006; Acquisti et al., 2013). Other privacy studies show that the salience of privacy, for instance through privacy framing, may significantly increase privacy concerns (John et al., 2011) or privacy valuations (Beresford et al., 2012; Preibusch et al., 2013; Benndorf and Normann, 2014; Benndorf et al., 2015). The only possibility to adjust to the perceived risks associated with information disclosure in these experiments is to de- or increase privacy valuations. A closely related study on social signaling investigates how social image concerns affect the reservation price for posting information about choices in a public goods game including a selfie on a public webpage (Holm and Samahita, 2016). This study, again, only allows for an adjustment of privacy valuations ex post. We go a step further and manipulate the sequence of choices: in addition to exploring ex post adjustments of privacy valuations after the socially relevant choice was made, we investigate whether a subtle cue hinting at optional publicity is sufficient to trigger prosocial or norm compliant behavior. Offering participants a consent option before playing a standard

---

[2] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).

[3] The European Court of Justice ruled in favor of a right to be forgotten under the EU Data Protection Directive (*Google Spain, S.L., Google Inc. v AEPD, Mario Costeja González*, Case C-131/12 [2012]).

[4] For instance, by using automated deletion technologies like X-Pire! (Backes et al., 2011) or Vanish (Geambasu et al., 2009).

dictator game, we allow for social image concerns to play out in the very game and not just in privacy valuations. This design feature allows for a more realistic setup: participants have to trade-off the monetary benefits from antisocial behavior against the social image costs associated with their consent.

Third, privacy valuations are often influenced by the perceived immediacy of risk, the strategic nature of uncertainty or the intangible nature of privacy harms (John et al., 2011; Rivenbark, 2012; Schudy and Utikal, 2015). Models of bounded willpower (Loewenstein, 1992; Laibson, 1997; O'Donoghue and Rabin, 2001) suggest that people are likely to hyperbolically discount the (often intangible) costs of privacy losses and opt for the immediate gratification associated with consent (Acquisti, 2004; Acquisti and Grossklags, 2004; Strandburg, 2006; Jolls, 2010). These models are in line with empirical evidence showing that people have cognitive difficulties in anticipating potential (secondary) uses of personal information and might suffer from optimism bias, especially when the risks associated with secondary or aggregated data use are impalpable (Jolls, 2010; Willis, 2014; Schudy and Utikal, 2015) or when revealing information about non-conforming social preferences to strangers (Posner, 1981; John et al., 2011). Driven by a concern for realism, we keep the purposes of consent ambiguous and test the mere impact of optional public scrutiny. We do not directly draw attention to normative expectations. Rather, we manipulate the salience of consent options to trigger a feeling of uncertainty about publicity and social image concerns.

3. Experimental design

In combining elements of a lab and a field experiment, our study is intended to increase external validity without losing control over the experimental setting. On the one hand, we made the decision context as realistic as possible by publishing personal information on a publicly accessible Google website. Publicity was operationalized in a natural way so that participants would easily understand that their real names and choices could be visible to any internet user. On the other hand, we did not provide any explanation as to why we elicited privacy valuations in order to prevent experimenter demand effects.

The experiment was implemented using a 2×2 factorial design (Table 1) and consisted of three stages. In the first stage, we provided participants the opportunity to act selfishly (in a dictator game). In the second stage, we offered to pay participants for publishing information about how they behaved. In the third stage, we gave participants the right to have this information "forgotten". Our experiment involves two types of novel manipulations. On the one hand, we test how salient consent options affect behavior in a standard dictator game (as a tool to measure compliance with social norms), the willingness to give consent to the (probabilistic) publication of information, and privacy valuations. On the other hand, we test whether privacy valuations are affected by default rules of deletion, whether a deletion default dampens chilling effects, and how sticky the retention default is.

In the first stage, participants took part in a one-shot standard dictator game (Forsythe et al., 1994; Camerer, 2003). At the beginning of this stage, each player was endowed with an initial endowment e_i of 100 tokens. The conversion rate for one token was 0.09 €. Each player then had to decide how much of the endowment e_i they were willing to share with the recipient. The recipient earned the amount $s_j \in [0, 100]$ shared by the dictator. The dictator received the initial endowment minus the amount shared, i.e., $\pi_i = e_i - S_j$. Economists have used the dictator game to show that the standard economic assumption of money-maximizing behavior is not well founded (Engel, 2011) and that behavior in the dictator game cannot only be explained by social preferences, but also

by situational properties (List, 2007). However, it is also a robust result that many dictators keep higher amounts than they share with recipients. We exploit the latter effect and use the game as a technique to induce behavior that deviates from the fairness norm of equal split. Unlike in other neutrally framed experiments, our instructions explicitly referred to the concept of "sharing" in order to verbally illustrate that the dictators' decision may indicate the level of compliance with the fairness norm of sharing.

In the second stage, participants could give their consent to the publication of personal information on a publicly accessible Google website. Participants had to decide whether to agree with the publication of their real name and the amount shared in the role of dictator. In order to prevent inferences regarding the identity of non-consenting participants, valid consent would entail publication only with a probability $p = 0.80$. With the complementary probability $1 - p = 0.20$, valid consent would not be followed by publication. Participants were informed about these probabilities and told that they would be paid for giving their consent regardless of whether their information was published. While privacy choices are usually made on the basis of take-it-or-leave-it options (Jentzsch et al., 2012), we offered participants a continuous valuation space to allow for more granular measures of privacy valuations. Hence, we used an incentive-compatible Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) to elicit participants' reservation prices for privacy – a mechanism with features of a second-price auction. Each participant was asked to state the minimal amount between 0.00 and 9.00 € (in increments of 0.01 €) that they would be willing to accept in exchange for their consent to the publication (*WTA*).[5] The participants' *WTA* was then compared with a randomly determined bid price b, all bid prices b being equally probable in the interval [0.00, 9.00]. If participants' *WTA* was less than or equal the bid price b, their consent was considered to be valid and information would be published with probability $p = 0.80$. In this case, participants received the randomly determined bid price b. Given a random bid b, their payoff from consenting was $b - WTA \geq 0$ if $WTA \leq b$. If $WTA > b$, their consent was considered to be invalid and they obtained a payoff of 0. Participants who wanted to refuse consent could do so by expressing a $WTA > 9.00$ €. Reporting the truthful *WTA* is the optimal choice (Rivenbark, 2012). Before making their decisions in the consent stage, participants were directed to a screenshot of the Google website on which their full name and their decision in the role of dictator would be published in case of valid consent. During the experiment the screenshot did not contain any information about names and amounts shared. Fig. 1 shows a screenshot of the website after the experiment with anonymized names of participants. On the real website, nothing was anonymized.[6]

The first treatment variation consisted in changing the order of the first two stages. In the baseline treatment, participants made their decisions in the dictator game before deciding whether to give their consent (dictator game before consent). Before playing the dictator game, participants knew that another stage would follow without being informed about the specifics. In the chilling effects

---

[5] We only measure the *WTA*, because EU privacy law is based on opt-in rules that allow users to disclose personal in exchange for a reward (cf. Section 1). EU privacy law does not include opt-out rules or "pay-for-privacy provisions".

[6] The website *Kölner Internet-Experimente 2012* can be accessed under the following URL:. The headline states: "Amounts that participants shared with player B". The paragraph below contains a succinct description of the experiment: "On this website you are being presented the results of an experiment. In the experiment, player A was endowed with 100 tokens. Every token was worth real money which was paid to the participants at the end of the experiment. Player A had the opportunity to share any amount of his endowment with player B. The participants listed below have consented to the publication of their name and the amount shared with player B."

Table 1
Treatments.

Variation	Dictator game before consent	Consent before Dictator game
Retention default	Baseline/Non-automatic deletion	Chilling effects/Non-automatic deletion
Deletion default	Baseline/Automatic deletion	Chilling effects/Automatic deletion

Fig. 1. Screenshot of the Google website (anonymized).

treatment, participants had to decide about their consent and state their privacy valuation (*WTA*) before deciding over the split in the dictator game (consent before dictator game). Decisions over the split in the dictator game had thus to be made in light of the decision made in the consent stage. This allows us to test whether prior awareness of potential publicity entails a change in subsequent behavior. It is important to note that participants had to make their decision under risk in both stages. They were informed about the (in-)validity of their consent only after having made their choices in both stages. We use this design feature to test the *behavioral effect* of being offered a risky choice before and after the dictator game. However, our chilling effects treatment does not aim at identifying a *cognitive effect*. We do not and cannot test with our design whether it is awareness, the perception of risk or something else that eventually feeds into behavior.

The second treatment variation consisted in changing the default of the right to be forgotten. Participants were informed about the default prior to making their choice in the consent stage. The actual deletion was implemented in the third stage outside the lab. In the non-automatic deletion treatment, participants were told that they would have the opportunity to have their personal information deleted by addressing an informal deletion request to the administrator of the Google website. They were informed that their personal information would remain public for at least four weeks starting from the date of publication and that they would be able to request deletion via Email afterwards. It was made clear that without actively claiming deletion their information would remain on the Google website for an indeterminate period. In the automatic deletion treatment, participants were informed that the information published on the website would automatically be deleted from the website after a period of four weeks. All participants were informed about the exact date of publication. Finally, we provided all participants with a concise summary of the privacy policy of the Cologne Laboratory of Economic Research and an extract of the Google privacy policy in effect at that time.

To observe the sharing decisions and elicit the privacy valuations of each participant, we implemented the strategy vector method (Selten, 1967; Brandts and Charness, 2011). Unlike under a sequential decision protocol, participants had to provide a strategy profile

for the first two stages prior to being randomly assigned the role of "dictator" or the role of "recipient". Accordingly, participants had to make their decisions in the role of "dictator" in both the dictator game stage and the consent stage before the random assignment of roles.

At the end of the experiment, we conducted a survey measuring participants' privacy attitudes and personality traits, including the Westin Privacy Index (Kumaraguru and Cranor, 2005). The objective of these surveys was to determine whether privacy concerns are actually related to privacy valuations and compliance with social norms.

The experiment was conducted at the Cologne Laboratory of Economic Research at the University of Cologne. We used the experimental software z-Tree (Fischbacher, 2007) and invited participants from a subject pool of approximately 4500 individuals using ORSEE (Greiner, 2004). Recruited participants were randomly assigned to one of four treatments (between-subjects design). We ran four sessions with a total of $n = 122$ participants[7] (67 of them female). When arriving at the lab and before every stage of the experiment, participants received paper instructions that were also read aloud. Each session lasted approximately one and a half hours. To guarantee the confidentiality of our experimental procedure and protect the reputation of the lab, participants could not disclose their names during the experiment. Instead, they were asked to assign themselves a personal identifier number (PIN) at the beginning of the experiment. After the experiment, each participant was individually invited to fetch their respective payoff. To avoid the feeling of surveillance and social pressure within the lab, during the payment procedure neither participants nor the experimenter monitoring the sessions could learn about the decisions that (other) participants had made. Participants who validly consented had to present their PIN and an ID card in order to obtain their payoff. All participants received a show-up fee of 4.00 €. Across all treatments, participants earned 13.80 € on average. Those who gave their consent (*WTA* ≤ 9.00 €) earned 14.57 € on average,[8] while those who refused their consent (*WTA* > 9.00 €) earned 9.64 € on average.

4. Hypotheses

If dictators hold textbook preferences, they should not share anything in the dictator game. This prediction, however, is not in line with empirical findings on dictator games (Engel, 2011). Previous findings indicate that an equal split of the endowment can be considered as reflecting a normative value in the dictator game (Krupka and Weber, 2013) for the very reason that it is customary (Levitt and List, 2007; Young, 2008; Andreoni and Bernheim, 2009). However, experiments have also shown that adherence to this 50–50-norm is reduced by double-blind protocols (Hoffman et al., 1994; Hoffman et al., 1996; Cherry et al., 2002; Levitt and List, 2007). By contrast, subtle reductions of social distance may be sufficient to trigger altruistic behavior in the dictator game. This holds for giving recipients the opportunity to visually identify dic-

---

[7] 28 in the baseline/non-automatic treatment, 32 in the baseline/automatic treatment, 30 in the chilling effects/non-automatic treatment, and 32 in the chilling effects/automatic treatment.

[8] This includes participants whose *WTA* was smaller (valid consent) and larger (invalid consent) than the random bid price *b*.

tators (Frey and Bohnet, 1999), learning the names of the respective counterparts (Charness and Gneezy, 2008) or showing a pair of eyes (Haley and Fessler, 2005). However, these findings have been subject to important refinements. Recent evidence suggests that watching eyes only increase the probability of giving something, but not the mean amount shared (Nettle et al., 2013; Fathi et al., 2014). In line with these refinements, we take a somewhat more conservative approach and relate our predictions to the probability of giving something in the dictator game.

While some behavioral fairness theories assume that people experience disutility from inequitable outcomes in these settings (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000), they cannot explain why publicity or social proximity might decrease the utility from selfish choices. Therefore, we derive our hypotheses from theories that take account of social image concerns (Bénabou and Tirole, 2006, 2010). Accordingly, people have a desire to be perceived as fair by others, and they partly want to reassure themselves that they are not only steered by monetary incentives (Bernheim, 1994; Andreoni and Bernheim, 2009). Increasing publicity or the size of an audience thus triggers 'audience effects' and adherence to social norms. Related research on social norms indicates that publicity enhances signaling of cooperative or socially virtuous behavior, thereby inducing compliance with social norms (Sugden, 1986; Posner, 2000; Young, 2008; Krupka and Weber, 2013; Marthews and Tucker, 2015; Bursztyn and Jensen, 2015). These behavioral patterns may arise through an expectation of future social reward (Bateson et al., 2013) or by fear of being sanctioned for norm deviations (Oda et al., 2015). Other motivational channels discussed in the theoretical literature, such as identity (Akerlof and Kranton, 2000, 2005), normativity (Bicchieri, 2006) or guilt aversion (Dufwenberg et al., 2011) might provide alternative explanations. While our chilling effects treatment is not intended to disentangle these motivational forces, we exclude the possibility of immediate social or monetary sanctions in our experiment. Moreover, our experiment is a one-shot game and therefore excludes both the evolution of social norms and the emergence of social identities. Hence, we expect behavior to be mainly driven by social image concerns and audience effects.

Building on these findings, we expect that a gentle nudge to the mere option of publicity is sufficient to trigger social image concerns before the dictator game. The timing of our nudge is particularly important, since it determines the cognitive availability of information about publicity. In line with self-awareness theories (Duval and Wicklund, 1972), subtle cues ex ante make social image considerations or social norms cognitively accessible right before the choice is made (Rigdon et al., 2009). Experimental evidence on tax compliance, for instance, shows that signing an honesty statement before rather than after filling in a tax form significantly increases honesty (Shu et al., 2012). Therefore, we predict:

H1. Dictators are more likely to give something in the chilling effects treatment than in the baseline treatment.

Concerning privacy valuations, we build on evidence showing that "undesirable traits" may lead to an increase of the reservation price for personal information (Huberman et al., 2005). In a similar vein, valuations for information about behavior in a public goods game have been shown to be negatively correlated with contributions (Holm and Samahita, 2016). Costly concealment of personal information may be considered as an insurance against the expected costs of social misconduct or damage to the social image (Posner, 1978; Posner, 1981). Since we expect less extreme deviations from the 50–50 norm in the chilling effects treatment (H1), the average threat value of consent should be lower in this treatment. We therefore hypothesize:

H2. Dictators have lower privacy valuations in the chilling effects treatment than in the baseline treatment.

Consent and privacy valuations are endogenous variables. Therefore, we cannot identify a causal effect of consent on behavior in the dictator game. However, building on previous findings on social image concerns (Bénabou and Tirole, 2006, 2010; Holm and Samahita, 2016), we assume that consent is positively correlated with prosocial behavior. We predict:

H3. Consenting dictators share higher amounts than non-consenting dictators across treatments.

We consider the deletion default to be a tool enhancing control over the *lifetime* of information. Psychological research on the so-called 'paradox of control' suggests that perceived control may trigger risky behavior (Slovic, 2000). In line with these findings, an increase of perceived control over the *release* of personal information has been shown to mitigate the effect of reactance and enhance people's willingness to disclose personal information (Tucker 2011; Brandimarte et al., 2013). Hence, automatic deletion should reduce the perceived risk of being perceived as a norm violator. Thus, we predict:

H4. Privacy valuations are lower in case of automatic deletion (deletion default) than non-automatic deletion (retention default).[9]

Dovetailing with the literature on status quo bias (Samuelson and Zeckhauser, 1988; Kahneman et al., 1991; Johnson et al., 2002; McKenzie et al., 2006), we expect most participants to refrain from actively making deletion requests. Hence, most participants should stick with the retention default.

5. Results

5.1. Behavior in the dictator game

We first examined the impact of the chilling effects treatment on participants' decisions in the dictator game (in comparison with the baseline treatment). A Wilcoxon-Mann-Whitney test revealed that participants shared larger amounts in the chilling effects treatment (consent before dictator game; $n = 62$, $M = 25.48$, $SD = 19.21$) than in the baseline condition (dictator game before consent; $n = 60$, $M = 19.61$, $SD = 20.98$, $z = 1.74$, $p = .08$).[10] Moreover, a two-part regression model[11] showed that almost twice as many participants ($n = 24$, 38.7%) shared nothing of their endowment in the baseline condition as compared to the chilling effects treatment ($n = 13$, 21.7%) (Table 2, column 1). This result supports H1. However, the chilling effects treatment did not have a significant impact on the amounts shared in the dictator game once participants who kept the entire endowment for themselves are excluded (Table 2, column 2). Therefore, the chilling effect seems to be driven mainly by the different proportions of participants who do not share anything. This difference in participants' behavior in the dictator game is also visible in Fig. 2, which shows the distributions of amounts shared in the baseline and the chilling effects treatment. Eyeballing the graph, the difference at 0 is evident, whereas the distribution looks rather similar for amounts above 0. This difference

---

[9] The behavioral predictions are less clear for the chilling effects treatment. In this case, automatic deletion could also facilitate self-regarding behavior, thereby increasing privacy valuations.

[10] All the results presented in this section are those of a two-sided Wilcoxon-Mann-Whitney test, unless reported otherwise.

[11] First part: logistic regression to predict a dummy variable for giving in the dictator game. Second part: truncated regression to predict giving after excluding those who kept the entire amount for themselves.

Table 2
Treatment effects.

	(1)	(2)	(3)	(4)
	Give something (yes = 1/no = 0)	Amount shared	Give something (yes = 1/no = 0)	Amount shared
Chilling effect (1 = consent before DG)	0.826* (2.02)	0.650 (0.15)	−1.070 (−1.11)	−4.168 (−0.29)
Consent (1 = gave consent)			−0.059 (−0.08)	7.991 (0.80)
Chilling effect × Consent			2.588* (2.37)	4.800 (0.32)
Constant	0.460+ (1.76)	30.10** (8.89)	0.511 (0.70)	23.22* (2.44)
Sigma Constant		18.09** (9.74)		17.85** (9.82)
N	122	85	122	85

Note: *t*- and *z*-values in parentheses. Columns 2 and 4 report the results of a truncated regression which excludes participants who kept the entire endowment for themselves.
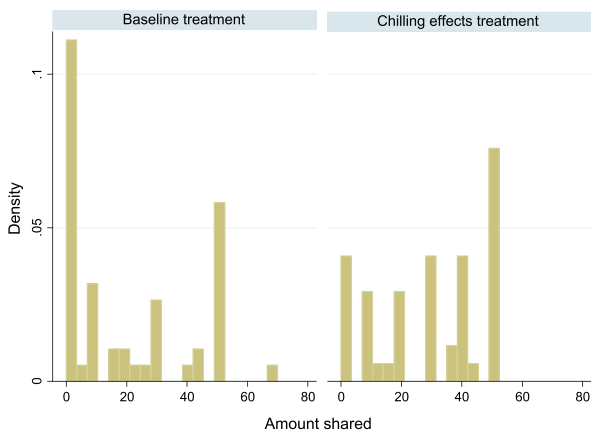
+ *p* < .10.
* *p* < .05.
** *p* < .01.

Fig. 2. Distribution of amounts shared in the baseline and chilling effects treatment.

also implies a shift of the mode from 0 in the baseline treatment to 50 in the chilling effects treatment. An Epps-Singleton test for equality of variances was not significant ($p = .13$). Hence, the distribution of altruistic choices does not seem to be different across conditions on an aggregate level. Rather, selfish behavior seems to be different at the lower bound, i.e. it mainly affects the behavior of extremely selfish dictators. In sum, we find some support for H1, as the proportion of individuals sharing part of their endowment is higher when they are previously given a consent option.

5.2. Privacy valuations

Our measure of privacy valuations was the *WTA* that participants stated in the BDM mechanism to give their consent. Participants did not differ statistically in their mean privacy valuations across treatment conditions ($M = 4.90$, $SD = 3.27$ in the chilling effects treatment; and $M = 4.28$, $SD = 3.02$ in the baseline treatment, $z = 1.00$, $p = .32$). Furthermore, we did not find any evidence for a bimodal distribution of privacy valuations. An Epps-Singleton test showed that distributions were not significantly different ($p = .19$). Therefore, our results do not support H2. We also examined whether privacy valuations could be predicted by the amount shared in the dictator game. Results of a linear regression analysis are presented in Table 3 (columns 5 and 6). They show that lower amounts shared predicted higher privacy valuations, which offers some support for the notion that privacy valuations are related to participants' behavior in the dictator game (see also H3). A non-

significant interaction shows that this result did not depend on whether participants were in the baseline or chilling effects treatment.

5.3. Consenting and non-consenting participants

Out of 122 participants, 19 (15.6%) had a valuation of privacy that exceeded 9.00 € and, thus, did not give their consent.[12] In line with H3, consenting participants ($M = 24.44$, $SD = 20.17$) shared higher amounts in the dictator game than non-consenting participants ($M = 12$, $SD = 17.82$, $z = 2.48$, $p = .01$). A closer look at differences between treatments, however, reveals an interesting result.

In fact, the relationship between the amount given in the dictator game and the propensity to give consent was significantly different for the chilling effects treatment and the baseline treatment. This result is supported by a significant interaction between the amount shared in the dictator game and the treatment condition in a logistic regression on the propensity to give consent ($b = 0.065$, $z = 2.01$, $p = .045$). In the chilling effects treatment, the amount shared in the dictator game predicted the propensity to give consent ($b = 0.073$, $z = 2.79$, $p = .005$), whereas in the baseline treatment, this was not the case ($b = 0.008$, $z = 0.43$, $p = .67$).

Similarly, looking at the means of amounts given illustrates that consenting participants shared significantly higher amounts ($M = 29.27$, $SD = 17.93$) than non-consenting participants ($M = 8.64$, $SD = 15.82$, $z = 3.25$, $p = .001$) in the chilling effects treatment. Conversely, in the baseline treatment the means were not significantly different for consenting ($M = 20.06$, $SD = 21.22$) and non-consenting participants ($M = 16.63$, $SD = 20.40$, $z = 0.26$, $p = .79$). Moreover, consenting participants shared lower amounts in the baseline treatment than in the chilling effects treatment ($z = 2.46$, $p = .01$), whereas non-consenting participants did not share significantly different amounts across treatments ($z = 1.16$, $p = .25$).

A two-part regression model that first tests the propensity to give anything in the dictator game and, in a second step, tests participants who gave something corroborated this result. The significant interaction between the chilling effects treatment and consent (yes/no) indicates that the chilling effect (i.e. consent before dictator game) was stronger for participants who gave consent than people who did not give consent (Table 2, column 3). After controlling for participants who kept the entire endowment

---

[12] Out of the 19 participants who did not consent, 8 refused consent in the baseline treatment, while 11 refused consent in the chilling effects treatment.

Table 3
Treatment effects.

	(5)	(6)	(7)	(8)	(9)
	Privacy valuation	Privacy valuation	Give something (yes = 1/no = 0)	Privacy valuation	Used the right to be forgotten
Amount shared	−0.034[*]	−0.031[+]			−0.002
	(−2.43)	(−1.65)			(−0.06)
Chilling effect (1 = consent before DG)		1.121			
		(1.31)			
Amount shared × Chilling effect		−0.012			
		(−0.44)			
Automatic deletion (1 = automatic)			−0.693	−0.220	
			(−0.73)	(−0.38)	
Consent (1 = gave consent)			0.888		
			(1.28)		
Automatic deletion × Chilling effect			0.628		
			(0.60)		
Privacy valuation					−0.067
					(−0.36)
Amount shared × Privacy valuation					0.008
					(0.82)
Constant	5.34[**]	4.89[**]	0.182	4.699[**]	0.236
	(12.76)	(9.09)	(0.30)	(11.32)	(0.19)
N	122	122	122	122	26

Note: *t*- and *z*-values in parentheses.

[+] $p < .10$.
[*] $p < .05$.
[**] $p < .01$.

for themselves, the amounts were not significantly different by conditions (Table 2, column 4).[13] To sum up, we find some evidence supporting H3. However, consenting individuals only share higher amounts than non-consenting individuals when being offered a consent option before the dictator game.

5.4. Impact of the right to be forgotten

We do not find support for H4. Results revealed no differences in the amount shared by the participants in the automatic deletion and the non-automatic deletion condition ($z = 0.53$, $p = .59$). Moreover, out of 26 participants whose information was made public in the non-automatic deletion treatment, only 6 requested the deletion of their information published on the Google website (23.08%). Additional results pertaining to the right to be forgotten are reported in Appendix A.

5.5. Additional results: privacy claims correlate with selfish behavior, not with privacy valuations

Finally, we examined whether individuals consider privacy as an *instrument* to avoid damage to their social image but not as a *value* per se. Our results suggest that different privacy types categorized on the basis of a variant of the Westin Privacy Index do not differ on the dimension of privacy valuations but on the dimension of decisions in the dictator game (Fig. 3).[14] A regression analysis

shows that *privacy fundamentalists* shared lower amounts in the dictator game ($M = 17.1$, $SD = 19.62$) than *privacy realists* ($M = 23.24$, $SD = 20.45$) and *privacy ignorants* ($M = 32.4$, $SD = 18.44$, $b = 15.32$, $t(121) = 2.51$, $p = .014$). Conversely, no differences emerged for privacy valuations ($ps > .358$). These results indicate that expressed privacy preferences are correlated with deviations from social norms (i.e. in this case: antisocial behavior). Additional results are reported in Appendix B.

6. Discussion

Combining elements of a lab and a field experiment, we set out to test whether a salient option to consent to the publication of personal information may trigger chilling effects (i.e. adherence to the 50–50 norm in a dictator game). We also wanted to investigate whether automatic deletion of personal information about the choice made in the dictator game facilitates selfish choices (deviations from the 50–50 norm) and how it affects privacy valuations. In addition, we study whether privacy valuations are negatively correlated with prosocial behavior and how revealing a refusal of consent actually is. Finally, we explore whether behavior in the dictator game can be matched to specific privacy types.

---

[13] In line with these results, a two-sample Kolmogorov-Smirnov test rejects the equality of distributions for participants who consented (p = 0.043, exact = 0.033).
[14] Our categorization is based on the *Privacy Segmentation Index* (Kumaraguru & Cranor 2005). The *privacy fundamentalists* and the *privacy unconcerned* were coded in

accordance with the original Privacy Segmentation Index. We gave participants the opportunity to not know the response to the questions of the Privacy Segmentation Index. Hence, the *privacy pragmatists* were subdivided in *privacy realists* (who differ from the first two categories but claim that they do not know the response only once) and *privacy ignorants* (who differ from the first two categories but claim that they do not know the response more than once). In our experiment, 30.33% are *privacy fundamentalists*, 54.92% are *privacy realists*, 12.30% are *privacy ignorants*, and 2.46% are *privacy unconcerned*.
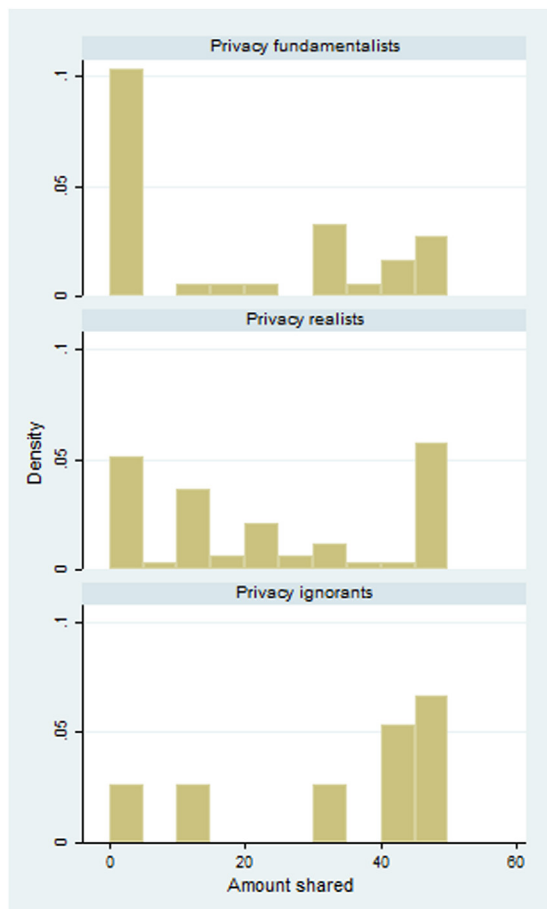
Fig. 3. Amounts shared in the dictator game by privacy type.

Our results show that salient consent options significantly increase both the average of amounts shared and the propensity to give something in the dictator game. However, as our econometric analysis also shows, the treatment effect on the average amount shared is mainly driven by the different proportions of dictators sharing nothing in the dictator game. Furthermore, while the mode shifts from selfish behavior in the baseline treatment to the 50-50 norm in the chilling effects treatment, we do not observe a significant difference in the distribution of altruism. We can therefore conclude that our chilling effects treatment is sufficiently strong to deter extremely selfish behavior. This result is in line with our (conservative) hypotheses, which only bore on the propensity of giving something. Nevertheless, it is worth noting that this effect is also strong enough to entail a significant upward shift of average social behavior. As discussed in the hypotheses section, the most plausible explanation for this effect may be social image concerns. Our chilling effects treatment can be seen as a nudge to the risk of publicity. Observing this cue, participants can establish a cognitive connection between giving and consent. This eventually triggers social image concerns. Hence, participants become aware that their behavior in the dictator game may be a public signal when deciding whether to give consent. Thus, our manipulation may have activated self-awareness and prompted participants to make a moral commitment through consent without incurring high monetary losses. This explanation is in line with the fact that the main effect of our chilling effects treatment is driven by consenting participants. When facing salient consent options, participants seem to be tempted by the prospect of having their (slightly) prosocial choices published. In this situation, giving in to the temptation of boasting with not entirely selfish choices is not really costly.

It merely requires some hedging: what consenting dictators lose by making a small sacrifice in the dictator game may, to a certain extent, be recovered by the monetary reward associated with consent. An explanation for the mere shift in selfish behaviors may be – in behavioral economics parlance – 'over-justification' (Bénabou and Tirole, 2006, 2010). Accordingly, participants in the chilling effects treatment may have discounted the meaning of prosocial behavior by ascribing their motivation to a favorable social image rather than to altruism. This may have partially crowded out sharing incentives and could explain why the chilling effects treatment only pulled participants away from extremely selfish choices. In sum, the trade-off is relatively easy to operate when the consent option is cognitively retrievable before making a choice in the dictator game.

By contrast, the cognitive connection between giving and consent is weaker in our baseline treatment where the consent option is given after the dictator game. Participants are not aware of the consent option when playing the dictator game. Hence, the only possibility to adjust to social image concerns is to increase privacy valuations. However, this is not what we observe. While privacy valuations are indeed somewhat higher in the baseline treatment, this effect is not statistically significant. We only find a statistically significant negative relationship between privacy valuations and amounts shared in the dictator game: the less participants share in the dictator game, the higher their privacy valuations. This result holds across treatments and is in line with our predictions. A potential explanation for observing constant privacy valuations across treatments is that the social signal provided through our website may not have been 'public' enough. Previous findings indicate that giving is distorted to highly visible, glamorous or fashionable recipients (Bénabou and Tirole, 2006, 2010). The Google website where participants' information was published may not have been sufficiently visible or memorable to trigger a quest for social prestige and decrease privacy valuations. An alternative explanation is 'mental accounting' (see Thaler, 1999). Participants carry their earnings from the dictator game in a mental account when making their consent choice in the baseline treatment. By contrast, participants in the chilling effects treatment have to anticipate earnings from the dictator game when making their choice in the consent stage. Hence, the marginal benefit from earnings through consent is lower for participants in the baseline treatment. This might have induced a decrease of privacy valuations, a reaction caused by something close to an income effect. This explanation also ties in with 'immediate gratification' (see Acquisti, 2004; Acquisti and Grossklags, 2004). Participants in the chilling effects treatment may have opted for the immediate monetary reward (or expected value) of consent. As a result, they may have been more inclined to forego additional earnings in the subsequent dictator game. This is a typical feature of privacy choices, which usually require thinking about future choices conditional on disclosing personal information and thus involve mental accounting (Brandimarte et al., 2013).

Our analysis of consenting and non-consenting participants yields a surprising result that only partly squares with our predictions. While a refusal of consent does not signal that individuals have behaved selfishly in the past, it seems to be a good predictor of selfish behavior in the future. On the one hand, this finding provides further evidence that salient ex ante consent options may serve as a commitment device and foster prosocial behavior (Shu et al., 2012). On the other hand, it suggests that a refusal to give consent (and disclose information about one's type) may be less revealing than assumed in the literature on unraveling models (Baird et al., 1994; Peppet, 2011; Benndorf et al., 2015). A refusal of consent may reveal more about intentions and future behavior than about past behavior. This asymmetry could be the starting point for future research on unraveling. A closer investigation of the sequence of choices, the endogeneity of types and social or norma-

tive commitment will hopefully provide insights about the forces underlying this asymmetry. In sum, the results on consenting and non-consenting participants cast doubt on the claim that concealing personal information is a guarantee of privacy (Bull, 2011). It might sometimes be the opposite.

Finally, our results indicate that a more cautious interpretation of the so-called privacy paradox is needed. Many scholars have conceived the discrepancy between high privacy concerns and low privacy valuations as a paradox (Acquisti and Grossklags, 2005a; Berendt et al., 2005; Strandburg, 2006). The problem of this paradox is that it fails to explain why stated privacy preferences or concerns should influence privacy valuations and information revelation rather than any other type of behavior. While privacy concerns do not necessarily concur with a high willingness to pay, they may be expressed through anti-social or socially non-compliant behavior. Our chilling effects treatment corroborates this conjecture: privacy choice architectures, such as salient consent options, may have an impact on behaviors not directly connected to privacy choices. Cuing a risk of publicity is likely to enhance conformity instead of privacy valuations. This dovetails with earlier findings suggesting that revealed preferences might thus not always reflect expressed preferences (Sunstein and Thaler, 2009; Acquisti et al., 2013). Hence, our findings suggest a different interpretation of the privacy paradox. When people say that they value their privacy, this may well mean that they value their freedom to make self-centered or non-conforming choices. In sum, this hints at an important deontological and behavioral distinction: privacy seems to be an instrumental value to evade social or normative pressure rather than a value per se.

Finally, we find no effect of the default designs of the right to be forgotten on privacy valuations. Moreover, our data do not support the claim that automatic expiry dates for personal information reduce chilling effects (see Mayer-Schönberger, 2009). Taking due caution in interpreting this null result, it may well be that people do not anticipate how personal information will be deleted when exposed to the temptation to make a selfish choice. Furthermore, despite the low opportunity costs of requesting deletion via Email only few participants actually use their right to be forgotten. While we cannot make any statistically relevant claim due to the small number of deletion requests, publicity on the Google website might not have been strong enough to pose an immediate threat to the social image of participants; the disadvantages of "losing" personal information from the website might have loomed larger than the advantages of inertia (see Kahneman et al., 1991); or participants might simply have forgotten their right to be forgotten. In future work, these factors potentially driving deletion behavior will have to be explored and distinguished more clearly.

As any experiment, our experiment relies on stylized facts and context. Therefore external validity has to be scrutinized. First, in complex digital environments, the kind of information people disclose depends on various contextual factors (John et al., 2011). Internet users are subject to a much richer decision-making environment than participants in our setting. We acknowledge this but point to the problem that a richer decision context may have added noise to our results. Hence, we keep the type of information to be disclosed and the experimental frame narrow in our study. This allows us to make interpersonal comparisons and identify the impact of salient consent options. On a structural view, however, our experimental setup comes quite close to social reality in the world of Google and Facebook: people can publicly share personal information, but the purposes of information processing remain rather vague; and people can decide how social they would like to be and appear. Second, our results require a careful interpretation due to the relatively small sample size. This is a limitation of our study that may be addressed in future studies with bigger samples. Third, in our setting, identification of the specific causes driving

social norm compliance is a stiff task, since publicity is endogenous to consent. Participants might have self-selected into publicity _because_ of their inclination to behave fairly and comply with social norms. However, this is precisely part of our research question. In our experiment, just like on digital platforms, salient consent options are designed to push people into consent and behavioral commitment. Self-selection into publicity is thus intended by our design. The latter calls for a precise distinction between two different factors: consent and consent options. We do not and cannot identify the impact of consent. Rather, our aim is to identify the impact of salient consent options as one element of choice architectures. To accomplish this, we vary the timing of the consent option to trigger self-awareness in the chilling effects treatment. Against this backdrop, a choice architecture that allows for self-selection into publicity before making a socially relevant choice entails a lower proportion of extremely selfish choices (compared to a situation where the same choice architecture is presented after the socially relevant choice). This is what our experiment was supposed to examine. It is up to future empirical studies to explore the reasons for giving consent in depth. On this basis, it might be possible to further disentangle and identify the forces that characterize our experiment: material incentives for consent, social image-seeking and the salience of consent options.

7. Conclusion

The European Charter of Fundamental Rights, the EU-GDPR and other privacy laws aim at empowering people to make free and informed privacy choices through consent. Conversely, one of the prime purposes of these rules is to mitigate chilling effects even when this implies an increase of selfish choices. Our study hints at a regulatory dilemma, which arises from the fact that current privacy laws are designed to steer consent choices through salient information and notice: instead of empowering people to make a free and informed choice over consent, salient information and consent options may push people into conformity. Lawmakers and lawyers might want to consider this risk of backfire effects in the implementation of information and notice policies. To be sure, the problem should not be overstated through dystopian accounts of a uniform society. On the one hand, publicity in small web communities may not trigger chilling effects because normative expectations are commonly shared by community members due to selection effects. On the other hand, according to our results chilling effects seem to mainly affect extreme deviations.

Nevertheless, legal scholars should be aware of the paradox associated with the autonomous decision to share personal information publicly: as people freely decide to give their consent, they may be subject to a chilling effect. This effect is akin to a self-imposed cost or internality (Loewenstein and Haisley, 2008). Should privacy law counter these internalities by decreasing the salience of notice? This will ultimately depend on the strength of chilling effects and on whether the suppressed behaviors are considered as socially beneficial. Our results should not prematurely justify restrictions of "the autonomy of one person for the sake of the greater autonomy [. . .] of that person himself in the future" (see Raz, 1986). However, there is a risk that salient and incentivized consent architectures will systematically push people towards consent with short-term monetary benefits and long-term costs to liberty. Lawmakers and lawyers should consider that societies need an adequate portion of dissent and non-conformity to flourish (Sunstein, 2003). In doing so, they should also take into account how revealing a refusal of consent might be. In the end, not sharing may be just as daring as sharing.

Appendix A. : Additional results on the right to be forgotten

In line with our results summarized above, a logistic regression did not reveal any significant effects on participants' propensity to share anything in the dictator game based on automatic vs. non-automatic deletion (Table 3, column 7). Moreover, participants did not differ in their privacy valuations ($M = 4.70$, $SD = 3.27$ in the non-automatic deletion treatment; and $M = 4.48$, $SD = 3.06$ in the automatic deletion treatment, $z = 0.31$, $p = .76$) (see also Table 3, column 8). It seems that the manipulation of the deletion modality influenced neither participants' privacy valuations nor their behavior in the dictator game. A logistic regression analysis of participants who requested deletion did not reveal any significant effect (Table 3, column 9), which may be due to the small number of independent observations. While we cannot make any inferential claim with respect to the reasons for this "stickiness" of the retention default, our results descriptively suggest that users are likely to incur higher risks of information abuse or damages to their social image when they have to claim deletion actively.

Appendix B. : Additional results on the impact of experience and privacy attitudes

The experience of regret does not seem to have a significant impact on privacy valuations in our experiment. Participants who stated to have experienced regret after posting personal information in networked publics did not have higher privacy valuations ($M = 4.59$) than participants who had never experienced regret ($M = 4.58$, $z = 0.008$, $p = .99$). In a similar vein, bad experiences resulting from the publication of personal information did not entail higher privacy valuations. Participants who had a bad experience with publicity did not have significantly higher privacy valuations ($M = 4.87$) than participants who never had a bad experience ($M = 4.52$, $z = 0.44$, $p = .66$). However, the 51 participants who reported having previously requested deletion of their information on public websites had significantly higher privacy valuations ($M = 5.46$) than the 71 participants who had never requested deletion ($M = 3.96$, $z = 2.45$, $p = .01$). Finally, while participants who had read the privacy policy of their preferred social network expressed significantly higher privacy valuations ($M = 5.35$) than those who had not done so ($M = 4.02$, $z = 2.48$, $p = .01$), they were not less inclined to give their consent ($z = 1.46$, $p = .14$). One possible explanation for these results is that people do not learn from bad experiences. An alternative explanation might be that the privacy implications of our experiment were not perceived as strong enough by participants who had a bad experience with sharing personal information.

Our results suggest that participants' privacy attitudes are not directly reflected in their privacy choices, but in their degree of sharing in the dictator game. A linear regression shows that specific concerns about how much information is publicly available in the Internet do not predict participants' privacy valuations ($b = -0.11$, $t(121) = -0.59$, $p = .55$).[15] Similarly, specific concerns about the way in which personal information is used by social networks generally can predict neither consent ($b = 0.36$, $z = 1.40$, $p = .16$)[16] nor privacy valuations ($b = -0.38$, $t(121) = -1.40$, $p = .16$). Moreover, worries about who can access personal information in social networks and how personal information may be used are not reflected in privacy valuations ($b = 0.03$, $t(121) = 0.16$, $p = .87$). However, our results revealed that privacy valuations were predicted by how important control of information was for participants ($b = -1.92$, $t(121) = -3.40$, $p = .001$). More specifically, privacy valuations were significantly higher for those who expressed greater need to control information. Surprisingly, neither risk aversion ($b = -0.07$, $t(121) = -0.53$, $p = .60$) nor income ($b = 0.001$, $t(121) = 1.51$, $p = .13$) predicted participants' privacy valuations. Further regression analyses revealed that women valued private information slightly lower than men ($b = -0.96$, $t(121) = -1.69$, $p = .09$) and that women shared slightly higher amounts in the dictator game ($b = 6.11$, $t(121) = 1.67$, $p = .09$).

Appendix C. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.irle.2017.06.001.

References

Acquisti, Alessandro, Grossklags, Jens, 2004. Privacy attitudes and privacy behavior. losses, gains, and hyperbolic discounting. In: Camp, L. Jean, Lewis, Stephen (Eds.), The Economics of Information Security. Springer, New York, NY, pp. 179–186.

Acquisti, Alessandro, Grossklags, Jens, 2005a. Privacy and rationality in individual decision making. IEEE Secur. Priv. 3, 26–33.

Acquisti, Alessandro, Grossklags, Jens, 2005b. Uncertainty, Ambiguity and Privacy Working Paper. Submitted to the 4th Annual Workshop on Economics and Information Security (WEIS 2005). Carnegie Mellon University, Pittsburgh, PA.

Acquisti, Alessandro, John, Leslie, Loewenstein, George, 2013. What is privacy worth? J. Legal Stud. 42, 249–274.

Acquisti, Alessandro, 2004. Privacy in Electronic Commerce and the Economics of Immediate Gratification Working Paper. Carnegie Mellon University, Pittsburgh, PA.

Akerlof, George A., Kranton, Rachel E., 2000. Economics and identity. Q. J. Econ. 115, 715–753.

Akerlof, George A., Kranton, Rachel E., 2005. Identity and the economics of organizations. J. Econ. Perspect. 19, 9–32.

Andreoni, James, Bernheim, B. Douglas, 2009. Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. Econometrica 77, 1607–1636.

Bénabou, Roland, Tirole, Jean, 2006. Incentives and prosocial behavior. Am. Econ. Rev. 96, 1652–1678.

Bénabou, Roland, Tirole, Jean, 2010. Individual and corporate social responsibility. Economica 77, 1–19.

Backes, Julian, Backes, Michael, Dürmuth, Markus, Gerling, Sebastian, Lorenz, Stefan, 2011. X-Pire! – A Digital Expiration Date for Images in Social Networks Working Paper. Universität des Saarlandes, Saarbrücken, Germany.

Baird, Douglas G., Gertner, Robert H., Picker, Randall C., 1994. Game Theory and the Law. Harvard University Press, Cambridge, MA.

Bannon, Liam J., 2006. Forgetting as a feature:not a bug: the duality of memory and implications for ubiquitous computing. CoDesign 2, 3–15.

Bateson, Melissa, Callow, Luke, Holmes, Jessica R., Redmond Roche, Maximilian L., Nettle, Daniel, 2013. Do images of 'Watching eyes' induce behaviour that is more pro-social or more normative? A field experiment on littering. PLoS One 8, e82055.

Becker, Gordon M., DeGroot, Morris H., Marschak, Jacob, 1964. Measuring utility by a single-response sequential method. Behav. Sci. 9, 226–232.

Ben-Shahar, Omri, Chilton, Adam, 2016. Simplification of privacy disclosures: an experimental test. J. Legal Stud. 45, S41–S67.

---

[15] All regressions used in this section are OLS regressions if not reported otherwise.
[16] These are the results of a logistic regression.

Benndorf, Volker, Normann, Hans-Theo, 2014. The Willingness to Sell Personal Data. Düsseldorf Institute for Competition Economics, Discussion Paper No. 143.

Benndorf, Volker, Kübler, Dorothea, Normann, Hans-Theo, 2015. Privacy concerns, voluntary disclosure of information, and unraveling: an experiment. European Economic Review 75, 43–59.

Berendt, Bettina, Günther, Oliver, Spiekermann, Sarah, 2005. Privacy in E-Commerce: stated preferences vs. actual behavior. Commun. ACM 48, 101–106.

Beresford, Alastair R., Kübler, Dorothea, Preibusch, Sören, 2012. Unwillingness to pay for privacy: a field experiment. Econ. Lett. 117, 25–27.

Bernheim, B. Douglas, 1994. A theory of conformity. J. Polit. Econ. 102, 841–877.

Bicchieri, Cristina, 2006. The Grammar of Society: The Nature and Dynamics of Social Norms. Cambridge University Press, Cambridge.

Bolton, Gary E., Ockenfels, Axel, 2000. ERC: a theory of equity, reciprocity, and competition. Am. Econ. Rev. 90, 166–193.

Brandimarte, Laura, Acquisti, Alessandro, Loewenstein, George, 2013. Misplaced confidences: privacy and the control paradox. Soc. Psychol. Pers. Sci. 4, 340–347.

Brandts, Jordi, Charness, Gary, 2011. The strategy versus the direct-response method: a first survey of experimental comparisons. Exp. Econ. 14, 375–398.

Britz, Gabriele, 2010. Informationelle Selbstbestimmung zwischen rechtswissenschaftlicher Grundsatzkritik und Beharren des Bundesverfassungsgerichts. In: Hoffmann-Riem, Wolfgang (Ed.). Offene Rechtswissenschaft, Tübingen: Mohr Siebeck, pp. 561–596.

Bull, Hans Peter, 2011. Informationelle Selbstbestimmung – Vision oder Illusion? Datenschutz im Spannungsverhältnis von Freiheit und Sicherheit, 2nd ed. Mohr Siebeck, Tübingen.

Bursztyn, Leonard, Jensen, Robert, 2015. How does peer pressure affect educational investments? Q. J. Econ. 130, 1329–1367.

Camerer, Colin F., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, NJ.

Charness, Gary, Gneezy, Uri, 2008. What's in a name? Anonymity and social distance in dictator and ultimatum games. J. Econ. Behav. Organ. 68, 29–35.

Chellapa, Ramnath K., Sin, Raymond G., 2005. Personalization versus privacy: an empirical examination of the online consumers' dilemma. Inf. Technol. Manage. 6, 181–202.

Cherry, Todd L., Frykblom, Peter, Shogren, Jason F., 2002. Hardnose the dictator. Am. Econ. Rev. 92, 1218–1221.

Cohen, Julie E., 2012. Configuring the Networked Self. Yale University Press, New Haven, CT.

Cohen, Julie E., 2013. What privacy is for. Harv. Law Rev. 126, 1904–1933.

Dufwenberg, Martin, Gächter, Simon, Hennig-Schmidt, Heike, 2011. The framing of games and the psychology of play. Games Econ. Behav. 73, 459–478.

Duval, T. Shelley, Wicklund, Robert A., 1972. A Theory of Objective Self-awareness. Academic Press Inc., New York, NY.

Elster, Jon, 1989. The Cement of Society. Cambridge University Press, Cambridge.

Engel, Christoph, 2011. Dictator games: a meta study. Exp. Econ. 14, 583–610.

European Data Protection Supervisor, 2011. Opinion on the communication from the Commission to the European Parliament, the Council, the Economic and Social Committee and the Committee of the Regions. A comprehensive approach on personal data protection in the European Union, § 85.

Fathi, Moe, Bateson, Nettle, Daniel, 2014. Effects of watching eyes and norm cues on charitable giving in a surreptitious behavioral experiment. Evol. Psychol. 12, 878–887.

Fehr, Ernst, Schmidt, Klaus M., 1999. A theory of fairness, competition, and cooperation. Q. J Econ. 114, 817–868.

Fischbacher, Urs, 2007. z-Tree. zurich toolbox for ready-made economic experiments. Exp. Econ. 10, 171–178.

Forsythe, Robert, Horowitz, Joel L., Eugene Savin, N., Sefton, Martin, 1994. Fairness in simple bargaining experiments. Games Econ. Behav. 6, 347–369.

Foucault, Michel, 1977. Surveiller Et Punir. Gallimard, Paris.

Frey, Bruno, Bohnet, Iris, 1999. Social distance and other-regarding behavior in dictator games: comment. Am. Econ. Rev. 89, 335–339.

Geambasu, R., Tadayoshi, K., Levy, Admit A., Levy, Henry M., 2009. Vanish: increasing data privacy with self-destructing data. Proceedings of the 18th Usenix Security Symposium.

Goodin, Robert E., Jackson, Frank, 2007. Freedom from fear. Philos. Public Aff. 35, 249–265.

Greiner, Ben, 2004. An online recruitment system for economic experiments. In: Kurt Kremer, Volker Macho (Eds.), Forschung und wissenschaftliches Rechnen. GWDG Bericht 63. Ges. für Wiss. Datenverarbeitung, Göttingen, pp. 79–93.

Grimmelmann, James, 2009. Saving facebook. Iowa Law Rev. 94, 1137–1206.

Haley, Kevin J., Fessler, Daniel M.T., 2005. Nobody's watching? Subtle cues affect generosity in an anonymous economic game. Evol. Hum.Behav. 26, 245–256.

Hermstrüwer, Yoan, 2016. Informationelle Selbstgefährdung. Zur rechstfunktionalen, spieltheoretischen und empirischen Rationalität der datenschutzrechtlichen Einwilligung und des Rechts auf informationelle Selbstbestimmung. Mohr Siebeck, Tübingen.

Hermstrüwer, Yoan, 2017. Contracting around privacy: The (behavioral) law and economics of consent and big data. J. Intellect. Prop. Inf. Technol. Electron. Commer. Law 8, 9–26.

Hoffman, Elizabeth, McCabe, Kevin, Shachat, Keith, Smith, Vernon L., 1994. Preferences, property rights, and anonymity in bargaining games. Games Econ. Behav. 7, 346–380.

Hoffman, Elizabeth, McCabe, Kevin, Smith, Vernon L., 1996. Social distance and other-regarding behavior in dictator games. Am. Econ. Rev. 86, 653–660.

Holm, Hakan J., Samahita, Margaret, 2016. Curating Social Image: Experimental Evidence on the Value of Actions and Selfies. Lund University, Department of Economics, Working Paper 2016: 8.

Holmes, Oliver W., 1872. The Poet at the Breakfast-Table. University Press, Welch, Bigelow, & Co, Cambridge, MA.

Huberman, Bernardo A., Adar, Eytan, Fine, Leslie R., 2005. Valuating privacy. IEEE Secur. Priv. 3, 22–25.

Hui, Kai-Lung., Png, Ivan P.L., 2006. The economics of privacy. In: Hendershott, Terrence (Ed.). Handbooks in Information Systems: Economics and Information Systems, vol. 1. Emerald, Bingley, pp. 471–498.

Hui, Kai-Lung, Teo, Hock H., Lee, Sang-Yong T., 2007. The value of privacy assurance: an exploratory field experiment. MIS Q. 31, 19–33.

Jentzsch, Nicola, Preibusch, Sören, Harasser, Andreas, 2012. Study on Monetising Privacy: An Economic Model for Pricing Personal Information. European Network and Information Security Agency.

John, Leslie, Acquisti, Alessandro, Loewenstein, George, 2011. Strangers on a plane: context-dependent willingness to divulge personal information. J. Consum. Res. 37, 858–873.

Johnson, Eric, Bellman, Steven, Lohse, Gerald, 2002. Defaults, framing and privacy: why opting in-opting out. Market. Lett. 13, 5–15.

Jolls, Christine, 2010. Rationality and Consent in Privacy Law Working Paper. Yale Law School, New Haven, CT.

Kahneman, Daniel, Ritov, Ilana, 1994. Determinants of stated willingness to pay for public goods: a study in the headline method. J. Risk Uncertain. 9, 5–38.

Kahneman, Daniel, Knetsch, Jack L., Thaler, Richard H., 1991. Anomalies: the endowment effect, loss aversion, and status quo bias. J. Econ. Perspect. 5, 193–206.

Krupka, Erin L., Weber, Roberto A., 2013. Identifying social norms using coordination games: why does dictator game sharing vary? J.Eur. Econ. Assoc. 11, 495–524.

Kumaraguru, Ponnurangam, Cranor, Lorrie F., 2005. Privacy Indexes: A Survey of Westin's Studies. Working Paper. Institute for Software Research International, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Laibson, David, 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112, 443–477.

Lessig, Lawrence, 2007. The code of privacy. Proc. Am. Philos. Soc. 151, 283–290.

Levitt, Steven D., List, John A., 2007. What do laboratory experiments measuring social preferences reveal about the real world? J. Econ. Perspect. 21, 153–174.

Levmore, Saul, Nussbaum, Martha C. (Eds.), 2010. Speech, Privacy, and Reputation. Harvard University Press, Cambridge MA/London.

List, John A., 2007. On the interpretation of giving in dictator games. J. Polit. Econ. 115, 482–493.

Loewenstein, Georg, Haisley, Emily, 2008. The economist as therapist: methodological ramifications of light paternalism. In: Caplin, Andrew, Schotter, Andrew (Eds.), Foundations of Positive and Normative Economics. Oxford University Press, Oxford/New York, NY, pp. 210–245.

Loewenstein, George, 1992. The fall and rise of psychological explanations in the economics of intertemporal choice. In: Loewenstein, George, Elster, Jon (Eds.), Choice Over Time. Russell Sage Foundation, New York, NY, pp. 3–34.

Marthews, Alex, Tucker, Catherine, 2015. Government Surveillance and Internet Search Behavior, Working Paper (April 29, 2015).

Mayer-Schönberger, Viktor, 2009. Delete: The Virtue of Forgetting in the Digital Age. Princeton University Press, Princeton, NJ.

McAdams, Richard H., Rasmusen, Eric B., 2007. Norms and the law. In: Mitchell Polinsky, A., Shavell, Steven (Eds.), Handbook of Law and Economics, vol. 1. Elsevier, Amsterdam, pp. 573–618.

McAdams, Richard H., 1997. The Origin: development and regulation of norms. Mich. Law Rev. 96, 338–433.

McKenzie, Craig R.M., Liersch, Michael J., Finkelstein, Stacey R., 2006. Recommendations implicit in policy defaults. Psychol. Sci. 17, 414–420.

Mitrou, Lilian, 2010. The impact of communications data retention on fundamental rights and democracy. In: Haggerty, κevin D., Samatas, Minas (Eds.), Surveillance and Democracy. Routledge, London, pp. 127–147.

Nettle, Daniel, Harper, Zoe, Kidson, Adam, Stone, Rosie, Penton-Voak, Ian S., Bateson, Melissa, 2013. The watching-eyes effect in the Dictator Game: it's not how much you give, it's being seen to give something. Evol. Hum. Behav. 34, 35–40.

Nissenbaum, Helen, 2010. Privacy in Context. Technology, Policy, and The Integrity of Social Life. Stanford University Press, Stanford, CA.

O'Donoghue, Ted, Rabin, Matthew, 2001. Choice and procrastination. Q. J. Econ. 116, 121–160.

Oda, Ryo, Kato, Yuta, Hiraishi, Kai, 2015. The watching-eye effect on prosocial lying. Evol. Psychol., 1–5.

Peppet, Scott R., 2011. Unraveling privacy: the personal prospectus and the threat of a full-disclosure future. Northwestern Univ. Law Rev. 105, 1153–1204.

Posner, Richard A., 1978. The right of privacy. Georgia Law Rev. 12, 393–422.

Posner, Richard A., 1981. The Economics of privacy. Am. Econ. Rev. 71, 405–409.

Posner, Eric A., 2000. Law and Social Norms. Harvard University Press, Cambridge, MA.

Posner, Richard A., 2008. Privacy, surveillance, and law. Univ. Chicago Law Rev. 71, 405–409.

Preibusch, Sören, Kübler, Dorothea, Beresford, Alastair R., 2013. Price versus privacy: an experiment into the competitive advantage of collecting less personal information. Electron. Commer. Res. 13, 423–455.

Raz, Joseph, 1986. The Morality of Freedom. Clarendon Press, Oxford.

Richards, Neil M., 2013. The dangers of surveillance. Harv. Law Rev. 126, 1934–1965.

Rigdon, Mary, Ishii, Keiko, Watabe, Motoki, Kitayama, Shinobu, 2009. Minimal social cues in the dictator game. J. Econ. Psychol. 30, 358–367.

Rivenbark, David R., 2012. Valuing the Risk from Privacy Loss: Experimentally Elicited Beliefs Explain Privacy Behavior Working Paper. University of Central Florida, Orlando, FL.

Rosen, Jeffrey, 2001. The Unwanted Gaze. The Destruction of Privacy in America. Random House, New York, NY.

Rosen, Jeffrey, 2011. Free speech, privacy, and the web that never forgets. J. Telecommun. High Technol. Law 9, 345–356.

Rosen, Jeffrey, 2012. The right to be forgotten. Stanford Law Rev. Online 9, 345–356.

Samuelson, William, Zeckhauser, Richard, 1988. Status Quo Bias in Decision. Mak. J. Risk Uncertain. 1, 7–59.

Schauer, Frederick, 1978. Fear, risk and the first amendment: unraveling the chilling effect. Boston Univ. Law Rev. 58, 685–732.

Schudy, Simeon, Utikal, Verena, 2017. You must not know about me – On the willingness to share personal data. J. Econ. Behav. Org. (in press).

Selten, Reinhard, 1967. Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In: Sauermann, Heinz (Ed.), Beiträge zur experimentellen Wirtschaftsforschung. Mohr Siebeck, Tübingen, pp. 136–168.

Shu, Lisa L., Mazar, Nina, Gino, Francesca, Ariely, Dan, Bazerman, Max H., 2012. Signing at the beginning makes ethics salient and decreases dishonest self-reports in comparison to signing at the end. PNAS 109 (38), 15197–15200.

Slovic, Paul, 2000. The Perception of Risk. Routledge, London.

Stigler, George J., 1980. An introduction to privacy in economics and politics. J. Legal Stud. 9, 623–644.

Strandburg, Katherine J., 2006. Social norms, self-control, and privacy in the online world. In: Strandburg, Katherine J., Raicu, Daniela Stan (Eds.), Privacy and Technologies of Identity: A Cross-disciplinary Conversation. Springer, New York, NY, pp. 31–53.

Sugden, Robert, 1986. The Economics of Rights, Cooperation and Welfare. Palgrave Macmillan, London.

Sunstein, Cass R., Thaler, Richard H., 2009. Nudge: Improving Decisions About Health, Wealth and Happiness. Penguin, London.

Sunstein, Cass R., 2003. Why Societies Need Dissent. Harvard University Press, Cambridge, MA.

Sunstein, Cass R., 2009. On Rumors. How Falsehoods Spread, Why We Believe Them, What Can Be Done. Farrar, Straus and Giroux, New York, NY.

de Terwangne, Cécile, 2012. Internet privacy and the right to be forgotten/right to oblivion, revista de internet. Derecho y Política 13, 109–121.

Thaler, Richard H., 1999. Mental accounting matters. J. Behav. Decis. Mak. 12, 183–206.

Tsai, Janice Y., Egelman, Serge, Cranor, Lorrie, Acquisti, Alessandro, 2011. The effect of online privacy information on purchasing behavior: an experimental study. Inf. Syst. Res. 22, 254–268.

Catherine, Tucker, 2011. Social Networks, Personalized Advertising, and Privacy Controls. MIT Sloan School Working Paper 4851-10. Massachusetts Institute of Technology, Cambridge, MA.

Wathieu, Luc, Friedman, Allan, 2007. An Empirical Approach to Understanding Privacy Valuation Working Paper. Harvard University, Cambridge, MA.

White, Gregory L., Zimbardo, Philip G., 1975. The chilling effects of surveillance: deindividuation and reactance, office of naval research. Tech. Rep., Z-15.

Willis, Lauren E., 2014. Why not privacy by default? Berkeley Technol. Law J. 29, 61–134.

Young, H. Peyton, 2008. Social norms. In: Durlauf, Steven, Blume, Lawrence E. (Eds.), The New Palgrave Dictionary of Economics Online. , 2nd ed. Palgrave Macmillan, London.

Zittrain, Jonathan, 2008. The Future of the Internet – And How to Stop It. Yale University Press, New Haven, CT.