

It's Creepy, But It Doesn't Bother Me

Chanda Phelan

Cliff Lampe

Paul Resnick

University of Michigan School of Information
Ann Arbor, Michigan

ABSTRACT

Undergraduates interviewed about privacy concerns related to online data collection made apparently contradictory statements. The same issue could evoke concern or not in the span of an interview, sometimes even a single sentence. Drawing on dual-process theories from psychology, we argue that some of the apparent contradictions can be resolved if privacy concern is divided into two components we call *intuitive concern*, a “gut feeling,” and *considered concern*, produced by a weighing of risks and benefits. Consistent with previous explanations of the so-called privacy paradox, we argue that people may express high considered concern when prompted, but in practice act on low intuitive concern without a considered assessment. We also suggest a new explanation: a considered assessment can override an intuitive assessment of high concern without eliminating it. Here, people may choose rationally to accept a privacy risk but still express intuitive concern when prompted.

Author Keywords

Privacy; social media; privacy paradox; trust; social awareness

ACM Classification Keywords

H.5.3 Group and Organization Interfaces – Web-based Interaction

INTRODUCTION

Sharing personal information is an everyday occurrence online, and weighing the tradeoffs between convenience and privacy is an inescapable reality of using the internet. For designers of online systems, much of the difficulty in establishing the acceptability of data collection practices lies in the uncertainty surrounding consumers’ preferences: people’s views of online privacy are full of contradictions. One such contradiction is the commonly observed mismatch between people’s stated privacy concern and behavior, a phenomenon often referred to as

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI’16, May 07-12, 2016, San Jose, CA, USA
© 2016 ACM. ISBN 978-1-4503-3362-7/16/05...\$15.00
DOI: <http://dx.doi.org/10.1145/2858036.2858381>

the *privacy paradox* [29], but contradictions exist even in the way people speak about their concern.

To explore the source of contradictions in privacy concern, we conducted semi-structured interviews with undergraduates at a large university to study how individuals make sense of privacy concern in specific disclosure decisions. We find that people often seem to contradict themselves when talking about their concern, e.g. describing an intrusion as creepy but denying that it bothers them. The apparent incoherence vanishes when we decompose privacy concern into two parts: an initial “gut feeling,” which we call *intuitive concern*, and a subsequent deliberate weighing of costs and benefits, resulting in what we call *considered concern*.

We argue that the two ways that people talk about concern correspond well to the two systems proposed by the dual process model of cognition popularized by Daniel Kahneman [17], making the two-system view a useful analysis framework for privacy concern. System 1 is the fast, automatic process that generates impressions and feelings; System 2 is the deliberate, reasoning process that one identifies as the conscious self. System 1 is efficient but vulnerable to influences that should be irrelevant to a rational decision, and strongly influences System 2 despite System 2’s outward appearance of rationality. We find that the qualities of intuitive concern correspond to a System 1 process, and those of considered concern to a System 2 process.

Additionally, we identify three factors that commonly affected participants’ privacy concern: trust, social presence of agents collecting data, and using the belief that nothing is private online as a reference frame to assess marginal risk. The three factors have different impacts on the two types of concern. While trust in the organization collecting data affects both, the collector’s social presence affects only intuitive concern and marginal risk affects only considered concern. We additionally conclude that a dual process model of privacy concern sheds light on the privacy paradox. There are two ways in which a privacy paradox can occur in this model. In one, individuals act on their intuitive concern without assessing considered concern; they are usually unaware of the paradox because it arises from an incomplete understanding of the relevant risks. In the other, individuals ignore their high intuitive concern because they are able to establish low considered concern; in this case, they usually recognize the paradox but find it

unproblematic because they view it as irrelevant to their decision, even though it can lead to lingering discomfort.

BACKGROUND

Privacy is a broad subject studied in a number of fields, many of which understand privacy differently. There is no universally agreed-upon direct measure of privacy. A common proxy measure in empirical research is *privacy concern* [37], specifically as related to the collection, unauthorized secondary use, errors in, and improper access of personal data [38]. Using privacy concern as a proxy measure for disclosure intentions or behavior has proven problematic. First, privacy concern has been often measured as an individual's tendency to worry about privacy *in general*, which ignores the effect of context cues like a website's trustworthiness. Second, privacy concern is not predictive of disclosure behavior or even disclosure intention [29]; though often conflated in the past, they are three separate constructs.

Treating privacy concern, disclosure intention, and disclosure behavior as interchangeable assumes that privacy decision-making is an assessment of risks and benefits connected to disclosure, called the *privacy calculus* [12, 13]. This model assumes that people are rational, meaning they "have stable individual preferences for privacy and disclosure, and...will consistently react to changes in the objective risks and benefits" [6]. Under this model, the difference between intentions and actual behavior is negligible, making privacy concern a perfectly acceptable substitute for actual behavior.

More recent research has found that people's behavior is only partly a result of a privacy calculus, introducing a central puzzle of privacy research: people tend to say they have higher levels of privacy concern than is indicated by their behavior. This contradiction is referred to as the *privacy paradox* [29]. Ackerman et al. [1] found that though 39% of survey respondents claimed to be "very concerned" about online privacy, only half of those respondents reported behavior that classified them as such. Other studies have found that people who claim high privacy concern are willing to sell their personal information for small incentives, such as coupons [5].

Several possible sources of the privacy paradox have been proposed. Some have argued that the paradox is simply a methodological quirk, a result of measuring general privacy concern and ignoring how costs and benefits can change in specific disclosure decisions [3]. This explanation is unsatisfactory, as individuals' behavior can diverge from their stated concern even when asked about their concern in specific scenarios [4]. Other explanations of the privacy paradox argue that privacy decisions are affected by bounded rationality, meaning that a person's ability to perform the privacy calculus is bounded by incomplete information and cognitive limitations [3, 21]. This makes it unrealistic to expect that people will behave rationally when making privacy-related decisions [2].

Bounded rationality presents two other possible sources of the privacy paradox. First, privacy decisions are nearly always clouded by uncertainty about risk, making it all but impossible for individuals to understand their own privacy preferences, let alone report them with any accuracy [3]. Second is the use of heuristics, or mental shortcuts, as a means of avoiding effortful cognitive processing [3]. For example, consumers use the presence of a privacy seal on a website as a safety heuristic and fail to assess the objective risks of disclosure [33]. Heuristics are efficient and often lead to satisfactory outcomes, but they are also susceptible to factors that would be discarded as irrelevant in the risk/benefit calculation of the privacy calculus model, such as website design [16], affect [21], or others' observed behavior [27]. For example, people tend to disclose less when they are assured of confidentiality because it draws their attention to and thus heightens privacy concern [16].

System 1 and System 2: a model of cognition

A two-system view of thought developed by Daniel Kahneman and Amos Tversky distinguishes intuition and reasoning as separate mental processes. It is often applied as an explanation of apparent contradictions in decisions made under conditions of uncertainty [17]. Given the high levels of uncertainty in privacy decisions, applying this model to information privacy yields useful insights.

The systems, named System 1 and System 2 by Stanovich and West [39], refer to two parallel thought processes. System 1—the intuitive process—is automatic, fast, and often emotionally charged [17]. It generates *impressions* based on associations and heuristics with no sense of voluntary control. System 1 activities include recognizing anger in a voice or knowing the answer to $2+2=?$ [18]. System 2—the reasoning process—is slower, effortful, and may be governed by logic. It generates *judgments* based on consciously controlled mental activity [17]. System 2 activities include monitoring one's behavior in a social setting and evaluating a logical argument [18].

As System 2 is conscious, it is more capable of rationality, and so part of its function is to act as quality control for the impressions generated by System 1. System 2 is costly, though; it takes effort to allocate and sustain attention. Therefore, System 2 monitors System 1 only lightly [19]. During normal operation, System 1's impressions are usually adopted without modification by System 2. Kahneman calls these *intuitive judgments* [43]. If System 1 runs into difficulty—for example, when surprised by a sudden sound—it can call on System 2 to take over. Even when System 2 is fully engaged, it is heavily influenced by System 1's impressions. Further, System 1's impressions are inaccessible to introspection, so System 2 is unaware of its influence. Instead, System 2 fills in the gaps itself: people are able to provide rational explanations for their choices even when their behavior indicates they are following a System 1 heuristic [45].

Previous privacy literature has introduced a particular two-system model called the elaboration likelihood model (ELM). It describes how various persuasive cues are more or less effective depending on whether System 1 (the peripheral route in the ELM) or System 2 (the central route) is engaged. ELM has been used to explain the inconsistent effects of factors such as message framing [7, 23, 47, 48], brand reputation, information quality [47,48] and individual trait anxiety [49] on privacy behavior and/or intentions. In contrast, our focus is on how a two-system model can help us understand different types of privacy concerns, and thus resolve apparent incoherence among statements about concerns as well as inconsistency between stated concerns and behaviors.

Factors contributing to privacy concern

Privacy concern can be influenced by many factors. For example, studies have found that giving users more control over their data can reduce privacy concern [46], while another study found that targeted advertising is considered “creepy” and raises privacy concern [44]. The factors’ effects often interact with each other. One study found that people are concerned about sharing location information from their mobile devices, a concern amplified by the perceived lack of control over access to their location [40]. Further complications arise from the fact that the factors are not always rational; the lighting in the room can affect privacy concern, for instance [11]. Factors may cause an emotional reaction, which can also interfere with rationality: a study of user responses to mobile apps’ data collection found that the users said they were “creeped out” by them; the authors argue creepiness is “an emotional response to a sense of wrongness that is difficult to articulate” and is not necessarily related to objective risks [36], which suggests a System 1 response.

Participants in our study were affected by a complex array of factors found in previous literature. Of these, we chose three of the most common factors to analyze in detail. We examine the background for each here.

Trust

Trust’s effect on privacy concern has been widely studied as a crucial factor in e-commerce and other information disclosure behavior [31]. Research on trust in online environments usually focuses on institutional trust, which refers to a consumer’s belief that a data collector will not misuse personal data [15, 21]. This trust tends to be formed based on the reputation of the company as a whole rather than its privacy policies specifically: Earp and Baumer [14] found that consumers are more likely to intend to disclose information if the site is well-known.

Trust has an extremely powerful influence on both privacy concern and behavior. In a survey, the majority of respondents said that whether a site was run by a trusted company was a very important factor in their general privacy concern [1]. Another study found that building trust was more effective in encouraging information

disclosure than attempting to reduce privacy concern directly [26]. In concrete privacy decisions, trust can overrule an individual’s general privacy concern and increase disclosure behavior [29]. While the importance of trust in online privacy is well-established, its precise role remains unclear: trust has been alternately modeled as an antecedent to privacy concern, an outcome of privacy concern, a mediating variable between privacy concern and disclosure behavior, or a moderator of the influence of privacy concern on disclosure [37].

The importance of trust to online consumers is a rational response, to an extent; trust is most important when there is uncertainty [25], a common characteristic of online privacy decisions. Consumer response to trust can extend beyond what is strictly rational, though. Previous research has found that people can use trust as a heuristic to guide their behavior [35]. This strategy saves time and is less cognitively taxing, but can lead to errors. For example, interventions meant to protect privacy can be too effective in establishing trust: a 2009 survey found that 62% of its respondents incorrectly believed that the simple existence of a privacy policy meant that a site could not share personal information with third parties [42].

Social presence

Social presence describes the feeling of being with another in mediated communication [10]. Whether the other is human or artificial intelligence is immaterial [9]. According to the Computers as Social Actors (CASA) model, humans are predisposed to treat most everything like social entities, including virtual computer agents [27, 28]. The effect is particularly strong if the computer agent possesses characteristics of language and interactivity, but they are not a requirement [22]: targeted advertisements can create social presence [28]. Social presence can also be created by surveillance [41], though it fades if the intrusion does not change over time [24, 30].

Treating computers as social is an automatic response, so people are largely unaware of their behavior. They do not believe that computer agents are literally human [22, 27]. There is some variation in the propensity to feel social presence—gender [34] can have an effect, as can mood [22]—but even so, social presence is pervasive [28].

In research on online learning [41] or remote collaboration [8], social presence is often studied as a desirable occurrence, but the expectations of CASA suggest that social presence may increase privacy concern. Treating computers as social entities entails more than simply feeling presence. It also means holding the agent to the expectations of a social being [22]. People are polite to computers, and expect computer agents to be polite to them [32]. A machine that violates social norms is perceived to be socially—not technologically—deficient. Unsolicited social presence online is likely to be met with the same negative response as another person looking over their shoulder as they browse.

Experimental evidence of the effect of social presence on privacy behavior has produced conflicting results. One study, which used interactivity and vivid visual cues in banner advertisements to induce social presence, found human-seeming agents caused over-disclosure of personal information [3]. In contrast, evidence from another study showed that agents meant to create social presence caused decreased disclosure of personal information if they were perceived to be violating social norms [14].

Marginal risk

Individuals tend to assess outcomes as losses or gains relative to a reference frame, rather than the absolute level of the outcome [20]. In the case of online privacy, this would mean that individuals judge the risk of a new privacy intrusion by comparing it to their understanding of how exposed they already are. Studies examining the effect of marginal risk on privacy concern and behavior are recent and still relatively sparse, but there is evidence of marginal risk assessments affecting online privacy behavior. For example, Acquisti et al. [5] found that people will pay much less to protect information that would otherwise be public than the amount they will accept to disclose information that would otherwise be private.

A closely related but distinct concept is relative risk, judgment of the risk of a privacy intrusion relative to a specific similar intrusion that is close to mind. For example, a recent study found that subjects assessing the intrusiveness of two surveys had lower privacy concern when the level of protection increased from the first to the second survey, rather than when it decreased [6]. The same study also found that relative risk had an effect on both disclosure intention and actual privacy behavior, but was more pronounced in actual behavior.

METHOD

Recruitment

Participants were recruited from an introductory undergraduate course on information studies at Large University. Anyone enrolled in the class was eligible to participate, and received credit in the course for doing so. Students were invited by email to sign up for an interview. Students had access to alternative exercises for points. The study was approved by the university institutional review board.

Procedure

Three weeks before the invitations to participate in the current study were sent out, students were invited to sign up for a separate research study about social media use, called the Media Monitor (MM) for the purposes of this narrative, in exchange for class credit. To participate, students visited MM's recruiting website and downloaded a Chrome browser extension that could be perceived as privacy invasive: the extension logged visits to social media and other popular websites, collecting the URL, time, and duration of their visit. The data was collected

pseudonymously: participants created accounts with their email addresses, which were stored separately. All gave their informed consent to install MM. Of the 230 students in the class, 181 (79%) signed up.

Later, students received another email inviting them to participate in this privacy study, again in exchange for class credit. Students who had not signed up for MM were particularly encouraged to participate.

Semi-structured interviews with 37 undergraduates were conducted in person over a period of three weeks, each lasting 23-56 minutes (median 35 minutes). During the first stage of the interview, the study was disguised as a user experience interview for MM; the true focus of the study, to elicit views on privacy, was concealed. In this first stage, participants were not asked about privacy directly, and were instead asked to talk about their experience signing up for and using MM, as well as social media habits more generally. When participants brought up privacy in one of their responses, the interviewer asked exploratory questions but allowed the participant to direct the conversation. The concealment accomplished two things: it discouraged priming participants to over-report privacy concern [29], and it anchored the discussion to a real, specific privacy choice, which is important in capturing privacy concern accurately [37].

In the second stage of the interview, participants were informed of the concealment and given the opportunity to retract their consent (none did). With privacy revealed as a topic, participants were asked to talk about their general privacy concern, a common measure of privacy that has been demonstrated to be inaccurate [3]. Its inaccuracy was what made general privacy concern important to the study design: when contrasted with the more reliable method of questioning in the first stage, it encouraged contradictory responses and created an opportunity to observe if and how the participants resolved those contradictions for themselves.

The interviewer never defined online privacy in either stage. This was essential for us, because we wanted to see what factors would emerge as important to them, rather than forcing a focus on factors that we had pre-selected. A disadvantage of this strategy was that it was difficult to gauge how privacy concern may have been affected by an individual's understanding of privacy issues.

Audio of the interview was recorded and later transcribed. Interview transcripts were analyzed using inductive coding to iteratively identify themes. We first coded for anything that affected a participant's perception of a privacy intrusion, subsequently identifying a number of factors that commonly affected participants. These factors included the type of data being collected, trust in the entity collecting data, and their perceived level of control over their data. During this stage we noticed that participants differentiated between factors that affected

how *comfortable* they were, and how *acceptable* they considered the intrusion to be. We coded for these two themes. We then refined the codes for each factor contributing to privacy concern over several iterations.

Once we had refined the contributing factors, we delved more deeply into the characteristics of the themes *acceptable* and *comfortable* and found they corresponded to the characteristics of System 1 and System 2. Three of the most common factors were chosen to be analyzed in detail and were catalogued according to the type of privacy concern they affected. In a final round of analysis, transcripts were examined for counter evidence for our interpretation of each of these factors.

Participants

Interviews were conducted with 37 undergraduates (18 female). The majority, 62%, were studying a technology-related field, though all were enrolled in an introduction to information studies course. Nearly all reported having an account on multiple social media platforms, though most said they were primarily active on Facebook. One participant had recently deactivated his social media accounts to reduce distractions. All of the participants had signed up for MM. One reported uninstalling it before the interview, and three more said they planned to uninstall it.

RESULTS

Most participants were fairly content with their level of privacy on Facebook and other social media. They were largely aware that nothing on social media was truly private but found this unproblematic because *“I keep everything to myself that I don't want people [on Facebook] to know in the first place”* (S28). Most were similarly unworried by government surveillance, accepting it as an inevitability they believed was unlikely to ever affect them personally. In contrast, there was a lot of variation in participants' level of concern regarding the collection of browsing data by entities like MM or private companies such as Google. There was no common standard of what constituted a comfortable level of data collection or which risks were most salient.

Intuitive and considered privacy concern

The way participants spoke about their privacy concern sometimes appeared contradictory, simultaneously acknowledging and distancing themselves from their concern when reflecting on it. For example, S35's statement, *“The more I think about [data collection], the more disturbing it is, I guess. Personally, I'm not that worried right now,”* or S08's reaction to Facebook's data collection practices, *“I don't think it is that big of a deal. [...] I just think it's the fact that they're able to do that [...] is what scares me a little bit.”*

These statements were not simply evidence of confusion. Rather, they were examples of participants' tendency to divide their overall privacy concern into two separate components: whether a given data collection practice was

“creepy,” and whether they were *“bothered”* by it. This is evident in S17 speaking of targeted advertisements: *“It's kind of strange for me, [but] it doesn't necessarily bother me. [...] it's just kind of like, ‘This is weird.’”* A number of other participants phrased their concern in a similar way, such as S34: *“It doesn't really bother me. [...] It gets you thinking, ‘Wow, that's annoying that's weird, but whatever, I guess. That's fine.’”*

Though participants used particular words differently, when examined in context it appeared that participants were referring to two distinct components to their concern. The first component was the result of an emotional response to an intrusion—most frequently described using such words as *“annoying,” “weird,” “creepy,”* or *“disturbing”*—while the second was the result of a more deliberate assessment of how problematic the intrusion was—most frequently described as the participant being *“bothered”* or *“worried.”* The distinction can be observed in S23's response to targeted ads:

“I don't like the outcome ‘cause that's annoying. But once you see the outcome, you realize what they're doing, and then that starts to bother you. If you [...] don't know that they're manipulating you, you're just annoyed with them [...] But once I started learning about the manipulation, that started to bother me.”

As tended to be the case with participants generally, whether the intrusion was bothersome was the most important in explaining behavior when reflecting on past choices; the intrusion's creepiness was largely dismissed.

Another difference between the two components was time. When presented with new, specific privacy intrusions during the interviews, participants tended to respond immediately with an emotional reaction, with more deliberate consideration following after. This exchange between the interviewer and S05 illustrates:

Interviewer: [...] Would [reading your Facebook messages] change how you felt about [MM]?

S05: Oh, definitely. That's pretty invasive and, or at the very least someone who agrees that would probably be not, would probably censor themselves a bit and act kind of differently.

Interviewer: [...] What do you think is different?

S05: [pause] Good question. I don't... [know] how to explain it. It's just... I guess it's a matter of knowing who is going to see it. [...] I just feel like I would need to censor myself. It would be kind of, just like... I don't know, it just kinda makes me less comfortable.

S05 knew his feelings immediately but had trouble articulating the reasons for his concern, an indication that he moved to the second stage of consideration only after being prompted. The sequence was generally the same for all participants when establishing overall privacy concern.

Upon encountering a new privacy intrusion, a “gut feeling” formed immediately, based on a first impression often established before learning many details about objective privacy risks. As a result, it could be affected by factors that should have been irrelevant, and the participant did not always appear to understand which factors affected her gut feeling. When discussing past choices as well, participants named factors that should have been irrelevant, even though they were sometimes unsure as to why those factors mattered; this suggests the gut feeling was also formed in real privacy choices. We identified this gut feeling using the participants’ emotional responses to an intrusion, i.e. how the intrusion made them feel. We call this *intuitive concern*.

After establishing intuitive concern, a participant sometimes engaged in a second stage of deliberate assessment, resulting in what we call *considered concern*. In this stage, participants identified possible privacy risks and considered whether they were worth taking, given the intrusion’s benefits (e.g. class credit for installing MM).

Participants only assessed considered concern under certain conditions, namely when intuitive concern was high but there were benefits to accepting the intrusion anyway. In other cases, participants skipped assessing considered concern entirely. For privacy choices presented during the interview, participants usually had to be prompted to engage in a considered evaluation of concern. When intuitive concern was high and no benefits were apparent (e.g. a version of MM run by Facebook with no class credit) participants followed their intuition and rejected the intrusion; if intuitive concern was sufficiently low, participants accepted, again without further consideration. Participants who made a choice based only on intuitive concern usually had clear feelings but were unable to articulate the reasons with confidence.

Many factors contributed to participants’ concern about their privacy; some affected assessments of both intuitive and considered concern, while others primarily affected just one. Therefore, intuitive concern and considered concern were related but ultimately independent components of overall concern, which explained the apparent contradictions in some responses. Participants could achieve low considered concern even if intuitive concern was high, but if they were unable to address the factors contributing to their intuitive concern, their intuitive concern remained high. This resulted in responses such as S08 earlier referring to data collection on Facebook as both “scary” and “[not] that big of a deal” in the same thought.

Factors contributing to privacy concern

The way a factor tended to affect participants provided clues as to whether it had an impact on intuitive concern, considered concern, or both. In the following sections, we identify three of the most common factors and how they contribute to each concern component.

Trust

Trust appeared to affect both intuitive and considered concern; it also appeared to have the strongest effect on participants of all the identified factors. Trust’s proposed effect on concern is modeled in Figure 1; see Figure 4 for the full model with all factors included. Both institutional trust and interpersonal trust were discussed in the interviews, primarily in the context of the MM sign-up decision. Institutional trust in Large University was mentioned more frequently than interpersonal trust in their professor, but it is not possible to untangle the two entirely. Consequently, we refer to trust generally here.

Effect of trust on intuitive concern: Trust had a powerful effect on keeping participants’ intuitive concern about MM low. Nearly all participants felt comfortable sharing their data with MM even if they disliked the idea in other contexts, like S36:

“I already think [academic department] is really cool. So anything that they’re putting out [...] it’s really easy to just be like, ‘Oh, well, you’re cool. [...] Whereas I feel like I’d maybe just put more thought into it if it was coming from someone I don’t know.”

A number of participants failed to recognize MM as a potential privacy intrusion, an indication that trust had kept their intuitive assessment of concern so low that they had bypassed assessing considered concern before signing up. About half of the participants said they had no hesitation about signing up for MM (even when they were otherwise concerned about their privacy) and they usually did not have a complete understanding of MM’s data collection practices. S08 said he did not realize that MM collected data at all, initially:

“I was just flipping through, yay, whatever, install, and then when I went and looked back [...] I was like, ‘Wow. They must be collecting something in my computer.’ [...] I feel like that’s not their motive, to collect personal information from me. [...] Especially when it’s coming from professors from the university, they’re trustworthy people” (S08).

In S08’s case, he eventually investigated what data was being collected. Many others did not, and even though they had gone through MM’s informed consent process,

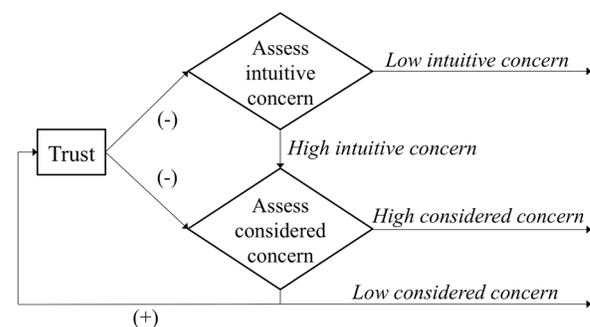


Figure 1: Effect of trust on privacy concern.

they were surprised to learn about MM’s data collection. These participants never evaluated the objective privacy risks of MM, suggesting they did not engage in the second stage of assessing considered concern.

Effect of trust on considered concern: Trust also appeared to have a major impact on participants’ considered assessment of concern. It was the most commonly-cited reason for why participants were comfortable sharing their data with MM. The ease with which most participants were able to explain their lack of intuitive concern was more evidence trust affected both components of concern.

Trust was important in dampening considered concern primarily because participants felt they could trust that the creators’ intentions were good and the application truthfully described exactly what it did. Together, these significantly reduced the perceived risks of signing up. Participants felt they could reasonably infer what they were agreeing to from the university context without investing the time to investigate MM itself in much detail. Consequently, most participants who reported initial reservations about the app did not express any regret about their choice to install MM.

Social presence of data collectors

Social presence was the most difficult factor to identify; it affected primarily intuitive concern, so participants were unlikely to cite it directly as a factor. Instead, its existence was inferred through participants’ descriptions of their emotional response to data collectors. Though not all participants felt social presence, it seemed to have a powerful effect on those who did, and led to lingering privacy concern even if their considered concern overruled their intuitive concern and they chose to accept an intrusion. The proposed effect of social presence is modeled in Figure 2.

Identifying social presence: Data collectors were generally unpopular among participants, but some participants had an appreciably higher level of concern than others about data collectors and tended to have a strong negative emotional reaction to them. Even these participants commonly had trouble identifying their issue with data collectors. For example, S27 had trouble thinking of what data she wanted to be kept private:

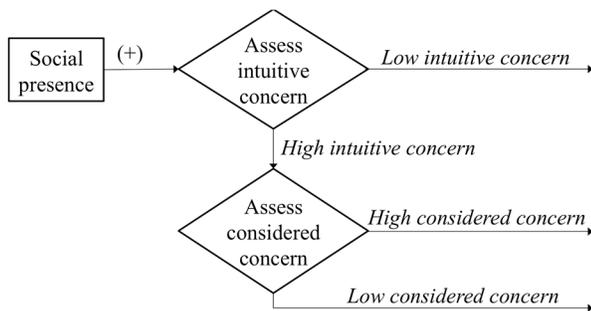


Figure 2: Effect of social presence on privacy concern.

“That’s a good question. I want to say location but... Like I’m always on campus, so where would I go? I don’t know. That’s a good question. It’s hard to pinpoint.”

As a result, in most cases it was not possible to directly identify what made these participants more concerned about data collectors than others. Instead, social presence was inferred by participants describing data collectors as making them *feel watched*. Participants who felt watched brought up a specific browsing experience in which they became aware of a data collector’s presence in real time (usually through targeted advertising) and used language of surveillance to describe the experience:

“I do still use Google and stuff, but it’s definitely opened my eyes to see, okay, they are pretty much watching, they’re watching.” (S35)

“[I don’t like] the fact that people know where I’ve been to [...] the fact that there’s somebody behind me, trailing me, it’s just a little scary.” (S27)

No one described feeling watched constantly. Most reported feeling negatively about the data collectors only intermittently while browsing; for example, after being reminded by targeted advertisements. Many participants, including those who showed no signs of feeling social presence generally, reported an initial feeling of being under surveillance after installing MM that faded quickly. For those who tended to feel social presence, the feeling returned more frequently and intensely. In addition, when otherwise comfortable participants were asked to describe their “nightmare” privacy scenario, it usually involved not feeling alone while browsing. For similar reasons, many participants said they would be more concerned about more invasive versions of MM, like one that tracked all browsing: *“It just creates a level of transparency where we feel like we’re being watched... where people are so scared of the level of surveillance they’re under that they are afraid to do certain things”* (S30).

Effect of social presence on intuitive concern: There were a number of indications that social presence primarily influenced the intuitive assessment of concern. Take S04 describing his response to social presence:

“[The collection of browsing data is] a weird subject for me. I don’t know. Sometimes, when I’m researching things that [...] would seem sketchy then I have concerns about privacy, but I would say that that’s a very small way I use the internet. [...] But it’s just like a weird thing to think about that someone’s sort of watching you, whatever you’re doing.”

This implies intuitive concern because his concern was rooted in an emotional reaction (the discomfort he felt about a data collector watching him). This would explain why he did not state social presence as a factor outright: as the concern stemmed from the sensation itself, rather than any concrete risks, he was unable or unwilling to use

social presence as a factor during his assessment of considered concern and struggled to find an alternative explanation (“[It’s] a weird subject for me. I don’t know”). He attempted to minimize his concern (“that’s a very small way I use the internet”) but remained unsettled, unable to resolve his intuitive concern.

Social presence had a pronounced effect on overall concern and was one of the few factors that convinced participants to take protective action. S35 reported switching to the no-tracking search engine DuckDuckGo when she felt uncomfortable, one of the few instances where participants described accepting inconvenience for more privacy. Similarly, the one participant who reported uninstalling MM did so because “I didn’t want to have something checking up on me” (S24) and said he “definitely” felt less watched after uninstalling it. Several other participants said they planned to uninstall MM, all because the app made them feel watched.

Effect of social presence on considered concern: Social presence was not frequently discussed as a factor contributing to considered concern. Instead, some participants discussed their strategies for ignoring or avoiding intuitive concern caused by social presence. A number of participants said that their lack of concern about data collectors was contingent on not being reminded of them. S15 was one: “I knew that they could [target ads], but the fact that I keep seeing it keeps reminding me of it. As long as it’s out of sight, out of mind.” S06 was another: “The Facebook advertising thing creeps me out [...] I don’t like to think about it. That’s just having it right there and [I think to myself], ‘Oh, don’t look at it!’” These participants tended to have lower overall privacy concern than participants who were affected by social presence; it may be they simply had better strategies for ignoring feelings of social presence.

Marginal risk

Some participants gauged their concern by evaluating the additional risk of a new intrusion compared to their current perceived level of exposure, namely the belief that nothing on the Internet is private. S30 used this to justify his lack of concern about MM: “All you guys were asking for was monitoring my sites and my hits, and basically a lot of other sites already do that without my permission.” Figure 3 models the proposed effect of marginal risk.

Effect of marginal risk on intuitive concern: There was little evidence that marginal risk reduced intuitive concern. Indeed, some subjects explicitly described feeling intuitive concern despite believing that marginal risk was minimal. For example, S06 said:

“I just hate the idea of...I don’t know, just having [browsing activity] recorded, although, I’m sure it’s recording somewhere else right now, but just I guess visualizing it and giving it to someone else is just weird to me. It just freaks me out.”

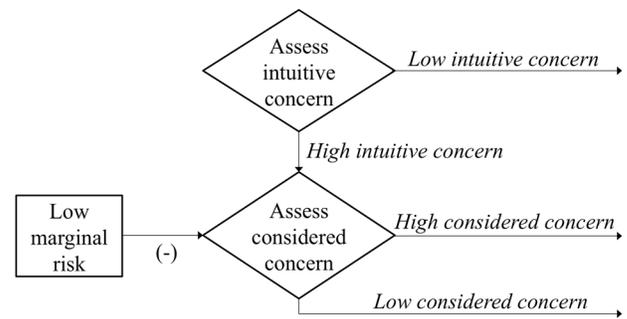


Figure 3: Effect of marginal risk on privacy concern.

S33 was similarly unsatisfied with low marginal risk as a reason to be less concerned when appraising a hypothetical version of MM run by Google: “I might have been a little bit more skeptical [...] It’s like, ‘I already give Google so much of my information, so why am I really concerned?’ I don’t know. It’s a good question.” Not being able to explain the residual concern suggests that the residual concern was intuitive, not considered.

Effect of marginal risk on considered concern: Some participants used marginal risk arguments to justify discarding intuitive concern upon further deliberation, signaling considered concern. For example, S11:

“I’m just numb to the fact that people can get information about me. I guess, it did occur to me like, ‘Oh, what if they can see my Facebook and see my interactions with other people?’ [...] [but in the end] I just signed up for it.”

S23 expressed a similar sentiment:

“I probably wouldn’t prefer to have a big corporation have all that information stored about me, [...] but it’s not like I would be in an uproar about it. I’m pretty content that they already have a lot of information, so it wouldn’t be that much more.”

Further evidence that marginal risk affected considered assessments of concern was that it was often employed in an explicit cost-benefit analysis. S25 was one participant who did this: “I think just with the way that times have changed, it’s impossible to really [keep personal information offline] and stay in touch with everything. So I think you just have to accept that reality.”

Interestingly, S33, who denied minimal marginal risk was a satisfactory reason to discard intuitive concern in a specific choice, *did* use low marginal risk to justify lower considered concern: “People get so angry [about privacy][...] it’s like, whatever, I’m already giving the Internet all this information about me. I can’t even think of what they can do to that really, really make me angry.”

Unlike trust, low considered concern based on assessment of minimal marginal risk did not have a mechanism for reducing intuitive concern. Intuitive concern remained even after a considered assessment of low concern.

DISCUSSION

Based on the results of the interviews, we developed a model of privacy concern with two semi-independent pieces, which we call intuitive concern and considered concern. The model explains apparent contradictions in our student respondents’ expressed privacy concerns.

As we analyzed the characteristics of intuitive and considered concern, we found that they correspond to the results of assessments by the two systems of a dual process model of cognition. Intuitive concern results from a process of the intuitive System 1: emotional, formed immediately, and affected by factors that are irrelevant to objective privacy risk. Considered concern results from a process of the reasoning System 2: conscious, slower, and more likely to engage in an explicit cost-benefit analysis using factors that are relevant to objective risk.

A two-system view can explain why individuals only assess considered concern under certain conditions. System 2 processes take up more cognitive resources, so they do not engage unless System 1 raises a warning; otherwise, System 2 adopts the intuitive judgment without modification. As this would predict, in situations where there is conflict in an intuitive assessment of concern—namely, if intuitive concern is high but there is some apparent benefit to be gained by accepting the intrusion anyway—individuals assess considered concern. Otherwise, when intuitive concern is very low or no obvious benefit exists that is worth the extra cognitive effort, an evaluation of considered concern is bypassed for the sake of efficiency. The impression generated by the intuitive assessment of concern is then adopted without modification. Switching between the two types of assessment occurs invisibly, usually without the person being aware they are moving between the two systems.

The two-system view also predicts that intuitive concern, as a product of a System 1 process, would be more influential in determining overall concern than considered concern, the product of a System 2 process. As Kahneman put it, though System 2 “believes itself to be where the action is,” it is System 1 that is really in charge [18]. This is consistent with the findings of the interviews: it would explain why factors affecting considered concern are most effective when they also address intuitive concern, as was the case with the trustworthiness of MM. Participants who were intuitively concerned about MM’s data collection were able to alleviate their intuitive concern through considered concern because once they decided that MM’s creators were trustworthy, subsequent intuitive assessments produced lower intuitive concern. The two-system view also explains why factors *unable* to address intuitive concern were likely to lead to residual intuitive concern even if considered concern was low, as was the case with marginal risk.

Further, because the impressions generated by System 1 processes usually cannot be articulated, individuals evaluating considered concern may not understand the influence of intuitive concern. Therefore, if the benefits of disclosure are sufficiently large or they can achieve low considered concern, they may disregard their lingering intuitive concern. However, as a System 2 process, an assessment of low considered concern can override System 1’s conclusion of high intuitive concern, but cannot disregard its influence; they therefore remain concerned. This can have long-term consequences on an individual’s privacy attitudes and behavior: every participant who discussed uninstalling MM did so because trust in Large U was insufficient to allay the intuitive concern caused by social presence. The complete model is summarized in Figure 4.

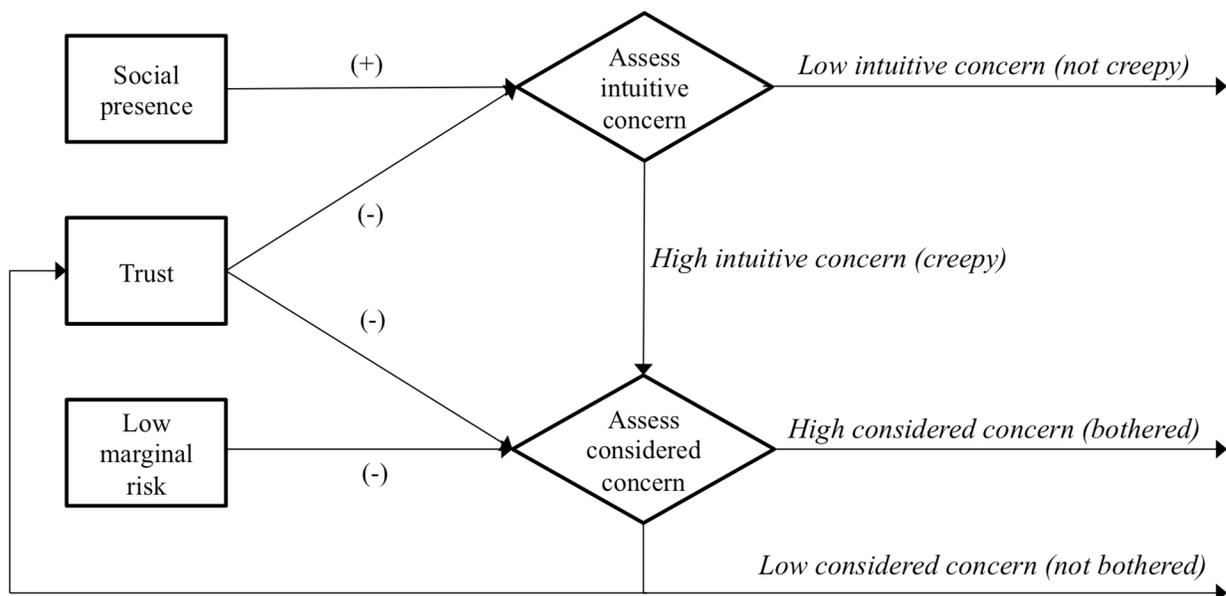


Figure 4: Model of privacy concern with components intuitive concern and considered concern.

The privacy paradox in the dual process model

This dual process model of privacy concern additionally contributes to our understanding of the privacy paradox. The model shows there are two ways in which a mismatch in preferences and behavior can be created:

- 1) An individual bypasses considered assessment of concern entirely and acts only on System 1's intuitive assessment. The intuitive concern assessment may be at odds with a considered assessment of concern when that is undertaken.
- 2) An individual accepts an intrusion because of low considered concern, overriding (but not reducing) System 1's assessment of high intuitive concern.

The first type provides insight into privacy paradox findings of existing research. For example, John et al. [16] found that people were more likely to disclose sensitive information when privacy concern was suppressed, even when the website was judged as high in disclosure danger by others. That is, people with low intuitive concern skipped assessing considered concern and were thus susceptible to errors caused by heuristic thinking. The same study also found that people are less vulnerable to the effect of interface design on disclosure if their privacy concern is activated, an indication that considered assessment of concern will override the mistaken impressions of intuitive concern more often under those conditions.

The second source of the paradox described above appears to be a new finding. It is of special interest because the intuitive concern, though overruled, is not erased. It may have lingering effects on behavior and should therefore be a special concern for designers.

Design implications

System designers naturally rely on people to *tell* them what privacy practices users want; as a result, they design for considered concern. The dual process model demonstrates that this is problematic: people may be rationally okay with an intrusion, according to their considered concern, but still be intuitively uncomfortable. Ensuring low marginal risk is one example of a factor that looks like an attractive target to lower privacy concern, since people say it is important. However, it is unable to make people *feel* good about a privacy intrusion, weakening its effect on overall concern.

People are less likely to think of or state their intuitive concerns when asked to think about their concerns. This makes it harder for designers to identify factors that affect assessments of intuitive concern. Nonetheless, it is this type of concern that designers should target to help people make privacy choices they can be happy with when they are no longer rationally considering them. This poses a difficulty, as strategies that target considered concern directly—such as emphasizing confidentiality before a disclosure decision—are often ineffective or even

counterproductive in addressing intuitive concern [16]. An alternate approach would be to offset factors that raise intuitive concern with factors that lower it. MM achieved this by offsetting the social presence it caused with high levels of institutional and interpersonal trust, for example; the same effect might be achieved by improving users' mood whenever they are making disclosure decisions [21] or using website design that is perceived as trustworthy [16].

Limitations and future work

This study has several limitations. The participants were drawn from a population that was non-representative even of undergraduate students; it is likely that they were more knowledgeable than average about privacy issues. Further, the only specific, real privacy choice discussed was the MM sign-up decision. It may therefore be the case that participants' reflections on other privacy choices were inaccurate, which would limit the strength and generalizability of our findings. In addition, all participants had installed MM, even though we encouraged participation by students who had not consented to MM. Those students may have had higher privacy concern, been less trusting, or otherwise differed significantly from the students who participated in the interviews. However, qualitative work like this does not claim generalizability as a contribution, and we see the themes uncovered as important.

As our evidence is preliminary, further work is necessary to test the validity of our proposed model. Future research can investigate how well a dual process model of privacy concern predicts the impact of other factors. It would be particularly useful to test if one type of concern is more closely connected to behavior than the other, and if one type of concern is particularly effective in producing preference-consistent behavior.

CONCLUSION

We propose a dual process model of privacy concern, where privacy concern can be decomposed into intuitive concern, one's "gut feeling," and considered concern, a deliberate calculation of risks and benefits. In this model, high intuitive concern can be overruled by low considered concern, but residual concern can remain. Consequently, there are two ways that a privacy paradox can occur: if an individual fails to engage in considered concern, or if an individual engages in considered concern and is unable to adequately address the factors contributing to intuitive concern. The challenge for future designers is to differentiate between the two causes of a privacy paradox in particular situations and respond with appropriate strategies that help people make better privacy choices.

ACKNOWLEDGMENT

We gratefully acknowledge financial support from Google's Social Interactions Focused Program.

REFERENCES

1. Mark S Ackerman, Lorrie Faith Cranor and Joseph Reagle. (1999) Privacy in e-commerce: examining user scenarios and privacy preferences. in *Proceedings of the 1st ACM conference on Electronic commerce*, ACM, 1-8.
2. Alessandro Acquisti. (2004) Privacy in electronic commerce and the economics of immediate gratification. in *Proceedings of the 5th ACM Conference on Electronic Commerce*, ACM, 21-29.
3. Alessandro Acquisti, Laura Brandimarte and George Loewenstein. (2015) Privacy and human behavior in the age of information. *Science*, 347 (6221). 509-514.
4. Alessandro Acquisti and Ralph Gross. (2006) Imagined communities: Awareness, information sharing, and privacy on the Facebook. in *Privacy Enhancing Technologies*, Springer, 36-58.
5. Alessandro Acquisti, Leslie K John and George Loewenstein. (2013) What is privacy worth? *The Journal of Legal Studies*, 42 (2). 249-274.
6. Idris Adjerid, Eyal Peer and Alessandro Acquisti. Beyond the privacy paradox: Objective versus relative risk in privacy decision making. Available at http://www.krannert.purdue.edu/academics/MIS/workshop/Idris%20Adjerid_RDJournalPaper-Purdue_Final.pdf.
7. Corey M Angst and Ritu Agarwal. (2009) Adoption of electronic health records in the presence of privacy concerns: the elaboration likelihood model and individual persuasion. *MIS Quarterly*, 33 (2). 339-370.
8. Mark Billinghurst and Hirokazu Kato. (2002) Collaborative augmented reality. *Communications of the ACM*, 45 (7). 64-70.
9. Frank Biocca. (1997) The cyborg's dilemma: progressive embodiment in virtual environments. *Journal of Computer-Mediated Communication*, 3 (2). 0-0.
10. Frank Biocca, Chad Harms and Judee K Burgoon. (2003) Toward a more robust theory and measure of social presence: Review and suggested criteria. *Presence*, 12 (5). 456-480.
11. Alan L Chaikin, Valerian J Derlega and Sarah J Miller. (1976) Effects of room environment on self-disclosure in a counseling analogue. *Journal of Counseling Psychology*, 23 (5). 479-481.
12. Mary J Culnan and Pamela K Armstrong. (1999) Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization science*, 10 (1). 104-115.
13. Tamara Dinev and Paul Hart. (2006) An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17 (1). 61-80.
14. Julia B Earp and David Baumer. (2003) Innovative web use to learn about consumer behavior and online privacy. *Communications of the ACM*, 46 (4). 81-83.
15. Ellen Garbarino and Olivia F Lee. (2003) Dynamic pricing in internet retail: effects on consumer trust. *Psychology & Marketing*, 20 (6). 495-513.
16. Leslie K John, Alessandro Acquisti and George Loewenstein. (2011) Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of Consumer Research*, 37 (5). 858-873.
17. Daniel Kahneman. (2003) A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, 58 (9). 697-720.
18. Daniel Kahneman. (2011) *Thinking, fast and slow*. Macmillan.
19. Daniel Kahneman and Shane Frederick. (2002) Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, 49-81.
20. Daniel Kahneman and Amos Tversky. (1979) Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*. 263-291.
21. Flavius Kehr, Tobias Kowatsch, Daniel Wentzel and Elgar Fleisch. (2015) Blissfully ignorant: the effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. *Information Systems Journal*. 25 (6). 607-635.
22. Kwan Min Lee and Clifford Nass. (2003) Designing social presence of social actors in human computer interaction. in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 289-296.
23. Paul Benjamin Lowry, Greg Moody, Anthony Vance, Matthew Jensen, Jeff Jenkins and Taylor Wells. (2012) Using an elaboration likelihood approach to better understand the persuasiveness of website privacy assurance cues for online consumers. *Journal of the American Society for Information Science and Technology*, 63 (4). 755-776.
24. David Lyon. (2002) Everyday surveillance: Personal data and social classifications. *Information, Communication & Society*, 5 (2). 242-257.
25. Roger C Mayer, James H Davis and F David Schoorman. (1995) An integrative model of organizational trust. *Academy of Management Review*, 20 (3). 709-734.
26. George R Milne and Maria-Eugenia Boza. (1999) Trust and concern in consumers' perceptions of marketing information management practices. *Journal of Interactive Marketing*, 13 (1). 5-24.
27. Youngme Moon. (2000) Intimate exchanges: Using computers to elicit self-disclosure from consumers. *Journal of Consumer Research*, 26 (4). 323-339.
28. Clifford Nass, Jonathan Steuer and Ellen R Tauber. (1994) Computers are social actors. in *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, ACM, 72-78.

29. Patricia A Norberg, Daniel R Horne and David A Horne. (2007) The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41 (1). 100-126.
30. Antti Oulasvirta, Aurora Pihlajamaa, Jukka Perkiö, Debarshi Ray, Taneli Vähäkangas, Tero Hasu, Niklas Vainio and Petri Myllymäki. (2012) Long-term effects of ubiquitous surveillance in the home. in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, ACM, 41-50.
31. Paul A Pavlou. (2003) Consumer acceptance of electronic commerce: Integrating trust and risk with the technology acceptance model. *International Journal of Electronic Commerce*, 7 (3). 101-134.
32. Byron Reeves and Clifford Nass. (1996) *The media equation*. Cambridge University Press.
33. Nora J Rifon, Robert LaRose and Sejung Choi. (2005) Your privacy is sealed: Effects of web privacy seals on trust and personal disclosures. *Journal of Consumer Affairs*, 39 (2). 339-362.
34. Paul Schermerhorn, Matthias Scheutz and Charles R Crowell. (2008) Robot social presence and gender: Do females view robots differently than males? in *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction*, ACM, 263-270.
35. John T Scholz and Mark Lubell. (1998) Trust and taxpaying: Testing the heuristic approach to collective action. *American Journal of Political Science*. 398-417.
36. Irina Shklovski, Scott D Mainwaring, Halla Hrund Skúladóttir and Höskuldur Borgthorsson. (2014) Leakiness and creepiness in app space: Perceptions of privacy and mobile app use. in *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, ACM, 2347-2356.
37. H Jeff Smith, Tamara Dinev and Heng Xu. (2011) Information privacy research: an interdisciplinary review. *MIS Quarterly*, 35 (4). 989-1016.
38. H Jeff Smith, Sandra J Milberg and Sandra J Burke. (1996) Information privacy: measuring individuals' concerns about organizational practices. *MIS Quarterly*. 20 (2). 167-196.
39. Keith E Stanovich and Richard F West. (2000) Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23 (5). 645-665.
40. Janice Y Tsai, Patrick Gage Kelley, Lorrie Faith Cranor and Norman Sadeh. (2010) Location-sharing technologies: Privacy risks and controls. *ISJLP*, 6. 119.
41. Chih-Hsiung Tu. (2002) The measurement of social presence in an online learning environment. *International Journal on E-learning*, 1 (2). 34-45.
42. Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley and Michael Hennessy. (2009) Americans reject tailored advertising and three activities that enable it. *Available at SSRN 1478214*.
43. Amos Tversky and Daniel Kahneman. (1974) Judgment under uncertainty: Heuristics and biases. *Science*, 185 (4157). 1124-1131.
44. Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay and Yang Wang. (2012) Smart, useful, scary, creepy: perceptions of online behavioral advertising. in *Proceedings of the Eighth Symposium on Usable Privacy and Security*, ACM, 4.
45. PC Wason and J St BT Evans. (1975) Dual processes in reasoning? *Cognition*, 3 (2). 141-154.
46. Heng Xu, Hock-Hai Teo, Bernard CY Tan and Ritu Agarwal. (2009) The role of push-pull technology in privacy calculus: the case of location-based services. *Journal of Management Information Systems*, 26 (3). 135-174.
47. Shu-Chen Yang, Wan-Chiao Hung, Kai Sung and Cheng-Kiang Farn. (2006) Investigating initial trust toward e-tailers from the elaboration likelihood model perspective. *Psychology and Marketing*, 23 (5). 429-445.
48. Tao Zhou. (2012) Understanding users' initial trust in mobile banking: An elaboration likelihood perspective. *Computers in Human Behavior*, 28 (4). 1518-1525.