INCREASING SHARING TENDENCY WITHOUT REDUCING SATISFACTION: FINDING THE BEST PRIVACY-SETTINGS USER INTERFACE FOR SOCIAL NETWORKS

Completed Research Paper

Bart P. Knijnenburg
University of California, Irvine
Donald Bren School of Information and Computer Sciences
5051 Donald Bren Hall
Irvine, CA 92697-3440
bart.k@uci.edu

Alfred Kobsa
University of California, Irvine
Donald Bren School of Information and Computer Sciences
5092 Donald Bren Hall
Irvine, CA 92697-3440
kobsa@uci.edu

Abstract

Privacy is a major concern of SNS (social networking site) users. Users' profiles contain a large amount of personal information, and most users want to control who has access to this information. However, many SNS users report difficulties in managing their privacy settings. We conducted an online user experiment to systematically evaluate the behavioral and attitudinal effects of several design parameters of a SNS privacy settings interface. We show that the granularity of categories, the possibility to make exceptions, the default setting, and the order in which categories are being presented have strong effects on users' evaluation of the system as well as their sharing behavior. Particularly, an interface that allows users to categorize their contacts into a small number of categories, that is set to shared-by-default, but that allows users to make exceptions for specific contacts results in the highest level of sharing and the highest user satisfaction.

Keywords: privacy, social networking sites, information disclosure, decision making, default effect, user experiment, user interface, human-computer interaction, structural equation modeling.

Introduction

In recent years, social networking sites (SNS) like Facebook and Google+ have come to play an increasingly important role in our social lives. The use of SNS is both pervasive and substantial (73% of online adults are on one or more SNSs, and 63% of Facebook's users log in at least once a day (Duggan and Smith 2014)). Privacy is a major concern of SNS users (Madden 2012). The profiles of most adult SNS users contain a large amount of personal information, such as their photos (66%), email address (46%), home address (30%), or phone number (24%) (Rainie et al. 2013), and the numbers for teens are even higher (Madden et al. 2013). People feel strongly about controlling who has access to this information (Acquisti and Gross 2006; Dey et al. 2012; Rainie et al. 2013; Stutzman et al. 2013), and 80% of SNS users indeed restrict access to their profiles (Madden 2012) by selectively sharing certain information only with certain people (Kairam et al. 2012; Young and Quan-Haase 2009).

At the same time, though, 48% of SNS users report difficulties in managing their SNS privacy settings (Lipford et al. 2008; Madden 2012). With a median of 200 contacts per user (Smith 2014) and users averaging seven new contacts a month (Hampton et al. 2012), the existence of these difficulties is not surprising. To allow users to manage this vast number of connections, SNSs like Facebook have to resort to "labyrinthian" privacy controls (Consumer Reports 2012). As a result, most Facebook users do not seem to know the implications of their own privacy settings (Liu et al. 2011; Strater and Lipford 2008), and share postings in a manner that is often inconsistent with their own disclosure intentions (Madejski et al. 2012).

Several suggestions have been made to simplify SNS privacy-settings user interfaces. For example, most SNS allow users to categorize their contacts as a means to more efficiently determine what to share with whom (Kairam et al. 2012; Watson et al. 2012). But the optimal management of such categories is not straightforward. For example, how granular should the categorization be? And should users be allowed to make exceptions for specific members of a category? Increasing the category granularity and/or allowing users to make exceptions may increase their ability to accurately manage their privacy settings, but this comes at the cost of increasing the complexity of the settings user interface (as well as users' decision process), and may thus decrease its perceived ease of use. Managers of SNS have to carefully navigate this tradeoff between expressiveness and ease of use in developing such interfaces.

Moreover, SNS managers have to decide whether the system should be set to share users' information with their contacts by default or not (Gross and Acquisti 2005). This default setting may not only have consequences for users who refrain from changing their settings: even when users actively engage in setting their sharing preferences, this default setting may unconsciously influence their sharing decisions (John et al. 2011; Johnson et al. 2002; Lai and Hui 2006). Since such default effects are practically unavoidable, Sunstein and Thaler (2003) argue that managers have a moral obligation to choose these defaults wisely.

This work is arguably the first to explore these design parameters of SNS privacy-settings user interfaces in a comprehensive manner. In an online user experiment, we show that interface aspects such as the granularity of categories, the possibility to make exceptions, the default setting, and even the order in which categories are being presented have strong effects on users' evaluation of the system as well as their sharing behavior. Particularly, we find that an interface that allows users to categorize their contacts into a small number of categories (family, friends, classmates, colleagues, and acquaintances), and that is set to shared-by-default but allows users to make exceptions for specific contacts, results in the highest level of sharing as well as the highest user satisfaction.

Related work and hypothesis development

Several researchers in the field of Human-Computer Interaction have developed innovative user interfaces for setting privacy preferences in SNS (Church et al. 2009; Egelman et al. 2011; Lipford et al. 2008, 2010), but this work has mainly been exploratory. Specifically, their evaluations have typically been limited to a comparison of a radically different new interface against a standard interface, with a small number of participants (max. N = 40). In the current work, we systematically explore a number of design parameters instead, and evaluate the effects of these parameters on users' subjective evaluation of the system (namely perceived privacy threat, ease of use, and satisfaction with the system) as well as their objective specification of what to share with whom. With respect to the latter, we define **sharing tendency** as the amount of information users are sharing with their contacts. Sharing tendency is the consequence of users' privacy-setting behavior.

Category granularity and exceptions

Early research on social sharing has shown that people want to share their personal information with their social connections selectively (Deuker 2012; Krasnova, Hildebrand, et al. 2009; Lederer et al. 2003; Olson et al. 2005; Patil and Lai 2005). Indeed, SNS users tend to restrict access to their profiles by sharing certain information with certain people only (Kairam et al. 2012; Madden 2012; Young and Quan-Haase 2009). To simplify this process, most SNSs allow users to categorize their contacts (Kairam et al. 2012; Watson et al. 2012). Facebook has three categories by default (Close Friends, Friends, and Acquaintances), while Google+ has four (Friends, Family, Following and Acquaintances)[1]. Both SNSs allow users to share specific posts and specific profile items with a subset of these categories.

Some researchers suggest that more granular categories are needed to foster information sharing in SNS (Benisch et al. 2011; Brandimarte et al. 2013; Sadeh et al. 2009; Sousa 2009; Tang et al. 2012). These researchers argue that when users are confronted with privacy options that are too coarse, they will "err on the safe side" and restrict their sharing to the level that is appropriate for the worst-case scenario (e.g. the level that is appropriate for the least trusted contact in a given category). Therefore, we hypothesize the following:

H1. Increasing the category granularity increases users' sharing tendency.

Since more granular controls allow users to set their privacy settings to a level that better reflects their sharing preferences, this additional control may decrease the perceived threat of over-sharing (Toch et al. 2010; Tsai et al. 2009; Wang et al. 2011). For example, Brandimarte et al. (2013) demonstrate that users perceive more control when privacy controls are more granular, and Tang et al. (2012) found that users of a finer-grained settings interface were more comfortable with their privacy settings. We therefore hypothesize the following effects of category granularity on users' over-sharing threat:

H2. Increasing the category granularity reduces users' perceived over-sharing threat.

Privacy decisions are among the hardest decisions to make, because they have delayed and uncertain repercussions that are difficult to tradeoff with the possible immediate gratification of disclosure (Acquisti and Grossklags 2008; Acquisti 2004). Due to the complexity of privacy decisions and users' bounded rationality (Acquisti and Grossklags 2005, 2008), an increase in control often just aggravates the problem by introducing choice overload (Iyengar and Lepper 2000; Scheibehenne et al. 2010; Schwartz 2003). Indeed, several researchers have noted that users of fine-grained interfaces find it difficult and time-consuming to accurately set their privacy settings (Korff and Böhme 2014; Madejski et al. 2012; Sadeh et al. 2009; Strater and Richter 2007). Too much control may thus have a detrimental effect on users' perceived ease of use of the privacy settings interface. We therefore hypothesize the following:

H3. Increasing the category granularity reduces the perceived ease of use of the settings interface.

Defaults

Another important parameter of privacy setting interfaces is their default setting. Traditionally, SNS have been set to make users' profile information shared-by-default rather than to keep this information private-by-default. This default setting can have far-reaching consequences for users' sharing behavior. First of all, although most people claim to *want* full control over their personal information, they often avoid the hassle of actually *exploiting* this control (Compañó and Lusoli 2010). This was initially true for Facebook—Gross and Acquisti (2005) indeed found that "only a small number of members change the default privacy preferences, which are set to maximize the visibility of users profiles" (p. 79). Over time,

[1] Facebook also has a "Friends of Friends" and a "Public" category, and Google+ also has a "Followers" category. These categories describe people who are *not* the user's contacts, though. The scope of our study is limited to recipients who are the user's contacts.

though, Facebook users have become increasingly more likely to make their settings more private (Dey et al. 2012; Stutzman et al. 2013).

Research in human decision-making has however demonstrated that default settings influence even those users who actively review their settings (John et al. 2011; Johnson et al. 2002; Knijnenburg et al. 2013a; Lai and Hui 2006). This "default effect" has been explained behaviorally, cognitively, and socially. Behaviorally, the default effect relates to the *status quo bias*: people tend to maintain the status quo because it avoids the effort and stress of making an active decision (Baron and Ritov 1994; Kahneman et al. 1991; Samuelson and Zeckhauser 1988). Since sharing is the status quo in a shared-by-default interface, consumers tend to maintain this setting and consequently end up sharing more information (Johnson and Goldstein 2003).

Cognitively, the default effect relates to the *anchoring bias*: the default option becomes the reference point in people's decisions (Chapman and Johnson 1994; Jacowitz and Kahneman 1995). People regard this reference point as the endowed option, and evaluate the alternative options in terms of losses and gains compared to this endowed option. This gives the endowed option an advantage (the *endowment effect*, cf. (Kahneman et al. 1991; Park et al. 2000)), because losses tend to loom larger than gains (*loss aversion*, cf. (Kahneman and Tversky 2000; Kahneman et al. 1991). Regarding a privacy setting interface, a shared-by-default interface makes sharing the endowed option and puts users in a "reject frame" (i.e. they have to think of reasons for not sharing the information). A private-by-default interface instead makes "not sharing" the endowed option and puts users in an "accept frame" (i.e. they have to think of reasons to share the information). Consequently, the sharing tendency will be lower in the private-by-default case, because loss aversion dictates that decision-makers need to feel more committed to make an "accept" decision than to forego a "reject" decision (Ganzach 1995; Meloy and Russo 2004; Wedell 1997).

Finally, a social explanation of the default effect puts the effect in a normative *information leakage* framework (Sher and McKenzie 2006). In this interpretation, defaults act as an implied endorsement of the default value by the system (McKenzie et al. 2006), and users tend to comply with this endorsement if they have a positive attitude towards the system (Sher and McKenzie 2006).

Default privacy settings may thus "nudge" SNS users in the direction of more sharing or more privacy (Thaler and Sunstein 2008), and a recent development in the privacy literature is to use such nudges as a means to protect SNS users' privacy (Acquisti 2009; Balebako et al. 2011; Hull et al. 2011; Wang et al. 2014). However, other researchers argue that defaults may threaten consumer autonomy, especially when they cause behavioral or cognitive biases (Smith et al. 2013; Solove 2013). These researchers argue for "smart default" settings that match the preferences of most users. To implement smart defaults, SNS managers would have to analyze users' privacy settings and make the most common setting the default. Such smart defaults would arguably most closely match users' true privacy preferences. Users' sharing tendency in a smart default setting should thus fall between the private-by-default setting (which may cause under-sharing) and the shared-by-default setting (which may cause over-sharing). We therefore hypothesize:

H4. A shared-by-default setting increases users' sharing tendency compared to a private-by-default setting. A smart default setting falls between these two settings.

The default effect may not influence all users' behavior equally. In a study on newsletter signup rates, Lai and Hui (2006) showed that consumers with a high level of privacy concerns were less influenced by the default effect, arguably because they were more careful in deciding what to share (Chapman and Johnson 1994; Connolly and Zeelenberg 2002; Wilson et al. 1996). Similarly, Brown and Krishna (2004) show that when a system uses self-serving defaults, the more skeptic consumers are likely to show reactance (behavior that counters the suggested action). In the case of SNS, this suggests that users with high interpersonal privacy concerns may be more careful in deciding what to share with whom, and therefore be less influenced by the default setting. We thus qualify H4 with a moderating effect of interpersonal privacy concerns:

H4a. The effect of the default setting on users' sharing tendency is smaller for users with a high level of interpersonal privacy concerns than for users with a low level of interpersonal privacy concerns.

Interpersonal privacy concerns may also have a main effect on users' sharing tendency. The existence of such an effect is not as straightforward as it may seem: so many privacy studies have found a gap between privacy attitudes and behavior (Acquisti and Grossklags 2005; Acquisti 2004; van de Garde-Perik et al. 2008; Metzger 2006; Norberg et al. 2007; Spiekermann et al. 2001) that it has been labeled the "privacy paradox" (Norberg et al. 2007). Despite this, many studies found that SNS users with high interpersonal privacy concerns share significantly less personal information with their contacts (Krasnova, Kolesnikova, et al. 2009; Lankton and Tripp 2013; Lo 2010; Posey and Ellis 2007). We therefore hypothesize:

H5. Users with a high level of interpersonal privacy concerns have a lower sharing tendency than users with a low level of interpersonal privacy concerns.

Previous work on default effects in privacy has mainly looked at their behavioral consequences (John et al. 2011; Johnson et al. 2002; Lai and Hui 2006), but defaults may also have an effect on users' perception of privacy threat. Particularly, users who over-share due to exploited decision biases may later come to regret their decisions (Patil et al. 2014; Wang et al. 2011), causing fear of over-sharing. But even users who demonstrate reactance towards the default setting (see H4a) may feel threatened by the default setting: Brown and Krishna (2004) argue that an unwanted default may make skeptical users even more skeptical. Hence we argue:

H6. A shared-by-default setting increases users' perceived over-sharing threat, as compared to a private-by-default or smart default setting.

Users' perception of threat from a default setting is in line with the normative *information leakage* interpretation of the default effect (Sher and McKenzie 2006): Users may perceive the shared-by-default condition as an implied endorsement by the system to over-share, and they may feel threatened by this. Is there anything an SNS can do to reduce this perceived threat? In a study of form auto-completion tools, Knijnenburg et al. (2013a) show that a control (i.e. a button) in the opposite direction of the default softens the endorsement of the default. In a social network, this control could present itself in the form of an ability to make exceptions to the chosen settings not just for the entire category, but for specific members of the category[2]. This option to make exceptions softens tells users that it is okay to deviate from the default-implied norm of disclosure. In other words, the effect of the default on perceived over-sharing threat depends on the availability of the control feature to make exceptions for specific contacts:

H6a. A shared-by-default setting increases users' perceived over-sharing threat, but *not* when the interface allows users to make exceptions for specific contacts.

Category order

Another default effect pertains to the order in which the privacy settings interface presents the contact categories to users. Specifically, the system could present the settings for stronger ties first (e.g. family, friends), or weaker ties first (e.g. acquaintances). Acquisti et al. (2012) demonstrated that asking privacy-sensitive questions in a decreasing order of intrusiveness could increase overall levels of disclosure, because subsequent requests compare favorably to the previous more intrusive requests, and users will therefore be more likely to answer them positively (this is called the "door in the face" technique, cf. Cialdini et al. (1975)). Regarding a privacy-settings user interface, we therefore postulate the following hypothesis:

H7. Presenting the settings for weaker ties first increases users' sharing tendency compared to presenting stronger ties first.

Arguably, though, this effect may be stronger in the accept frame (i.e. the private-by-default condition) than in the reject frame (i.e. the disclosed-by-default condition), because the accept frame highlights the perceptual contrast of subsequent requests with earlier more intrusive requests (Cantrill and Seibold 1986; Shanab and O'Neill 1979). In fact, Acquisti et al. (2012) demonstrate that an accept frame is a *precondition* for the "door in the face" technique in privacy research. Hence we qualify H7 with the following moderating effect:

[2] Both Google+ and Facebook have such an "advanced" option that allows users to make exceptions for specific group members.

H7a. Presenting the settings for weaker ties first increases users' sharing tendency only when the interface uses a private-by-default setting.

Over-sharing threat, perceived ease of use, and system satisfaction

As we have mentioned, the inability to set privacy preferences at the desired level may cause users to perceive over-sharing threat, while increasing the complexity of the interface may result in reduced ease of use. We now turn to the interrelation between these constructs, as well as their ultimate consequence of influencing users' satisfaction with the SNS.

Over-sharing threat can be seen as a lack of comfort or confidence regarding the attained level of sharing (Tang et al. 2012), which results in a system-specific concern of unwanted data collection or loss of control (Knijnenburg and Jin 2013). As people may have certain expectations regarding collection and control, they may use these expectations as key evaluation standards in determining their satisfaction with the SNS (Chen et al. 2012; Fox et al. 2000). To date, SNS researchers have not tested this effect of over-sharing threat on users' satisfaction with the SNS. Outside the realm of social networks, though, researchers have shown that people who feel that personal information collection is intrusive or uncomfortable are less likely to be satisfied with a service (Knijnenburg and Kobsa 2013; Lukaszewski et al. 2008). We thus hypothesize a similar effect for SNS:

H8. Over-sharing threat will be negatively associated with users' satisfaction with the system.

The effect of over-sharing threat on system satisfaction may be moderated by users' interpersonal privacy concerns, though. Specifically, for people with low privacy concerns, the increased threat may not result in a reduced level of satisfaction. In fact, research has consistently found that unconcerned people do not regard threats to their privacy as a significant reason to dislike or abandon a system (Hann et al. 2007; Jensen et al. 2005; Krasnova, Hildebrand, et al. 2009). Therefore, we hypothesize the following moderating effect:

H8a. The negative association between over-sharing threat and system satisfaction will be stronger for users with a high level of interpersonal privacy concern.

Over-sharing threat may influence system satisfaction in yet another way, namely mediated by the perceived ease of using the system. Users who experience over-sharing threat resulting from a mismatch between real and desired level of sharing are likely to want to avoid further over-sharing (Hey Tow et al. 2008; Rainie et al. 2013) and may subsequently feel forced to spend more time and effort "fixing" their settings (Krasnova, Günther, et al. 2009). As users typically avoid such effort as much as possible (Deuker 2012; Sleeper et al. 2013; Strater and Lipford 2008), this will make them more frustrated with the settings interface (Collins et al. 2012; Strater and Lipford 2008). Moreover, these users will feel like they are working "against the tide" when they are changing their settings; in its most extreme case, this is what causes reactance behavior (Brown and Krishna 2004). We therefore hypothesize the following:

H9. Over-sharing threat will be negatively associated with the perceived ease of use of the settings interface.

Subsequently, the perceived ease of use of the settings interface may influence users' satisfaction with the system. This causal relationship is codified in both the Technology Acceptance Model (TAM) (Davis 1989) and the Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh et al. 2003), and has been confirmed in numerous SNS privacy studies (Brandyberry et al. 2010; Bulgurcu et al. 2010; Dhillon and Chowdhuri 2013; Phillips and Shipps 2012; Qin et al. 2011). A recent study even demonstrated that SNS non-adopters cite a lack of ease of use as a main reason for avoiding SNS (Hu et al. 2011). We therefore hypothesize:

H10. Perceived ease of use of the settings interface is positively associated with users' satisfaction with the system.

Experimental setup

Figure 1 summarizes the hypothesized effects. The left side of the model shows the effects on users' sharing tendency, while the right side shows the effects on users' subjective evaluations. We tested these hypotheses in an online user experiment. Participants were asked to imagine using a Facebook-like social network, and were told to set their privacy settings to indicate which of eight personal information items they would share with whom. The privacy-settings user interface was manipulated between-subjects (see Manipulations). Finally users' evaluations of the system and their interpersonal privacy concerns were measured using a post-experimental questionnaire.

Participants

485 participants were recruited via Amazon Mechanical Turk, a popular recruitment tool for user experiments (Kittur et al. 2008; Mason and Suri 2012; Paolacci et al. 2010). Participation was restricted to US residents with a high "worker reputation". In previous work we found that Mechanical Turk workers generally match the US Internet population demographics. 98 participants did not meet our stringent comprehension and data quality checks, leaving 387 valid participants (162 males), aged between 18 and 68 (median 29).

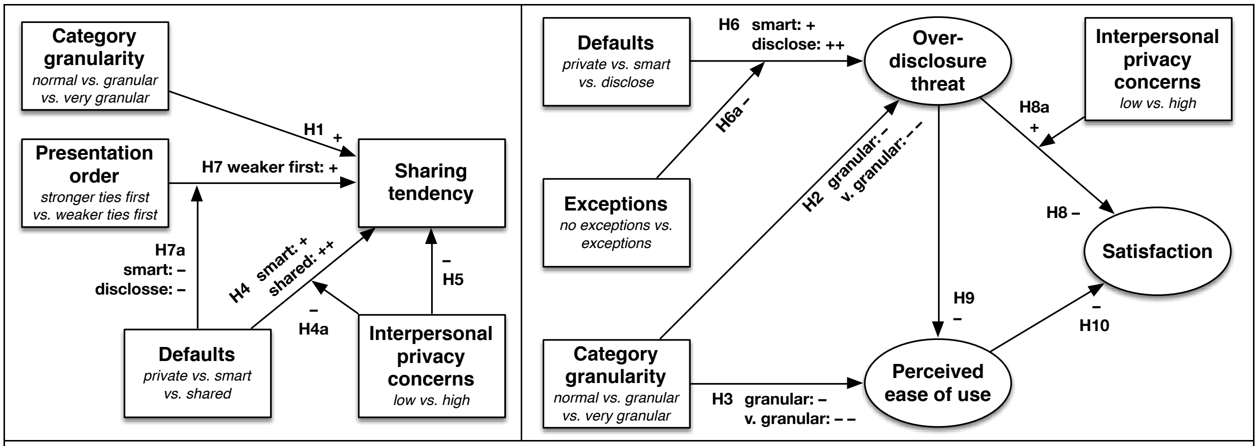

Figure 1. Hypothesized effects on users' sharing tendency and subjective evaluations.

System and procedure

Participants were recruited to test the settings interface of a new SNS. They were promised a US$1.50 payment for their efforts. Upon accepting the task, participants were randomly assigned to one of the experimental conditions (see Manipulations) and given an overview of the experimental steps.

To "seed" the new SNS with a list of real-life contacts, participants were given a list of 50 different "person descriptions" (e.g. "distant relatives", "female friends", "colleagues on your team", "high school classmates", "people you've lost contact with") and were asked to enter the name of one person matching that description, taking care not to provide any duplicate names[3]. Since participants would likely not know a person for every single category, they were allowed to skip up to 10 descriptions[4].

Subsequently, participants were told to "[i]magine that one day in the near future everyone you know has abandoned Facebook and Google+ for a new social network site named Mundo. Mundo is very similar to current social networks in that you have a *profile page* with some basic information about you and a *wall* on which you can share things like status updates and photos." They were then briefed on the two tasks they would be performing on Mundo: categorizing their contacts and setting their sharing preferences.

[3] Participants who provided bogus or duplicate names were removed from the analysis.

[4] 48% of participants did not skip any descriptions, and 96% skipped 5 or fewer descriptions.

To categorize their contacts, participants were given the ~50 names they provided earlier, and asked to categorize them into one of the 5, 10, or 14 categories (see Manipulations) using a dropdown box. They were asked to choose the best-matching category for each contact[5].

Figure 2 shows the settings interface used in the subsequent task of setting the sharing preferences. The default setting, presentation order of categories, and ability to make exceptions for specific contacts varied according to the experimental condition (see Manipulations). Participants were asked to carefully consider what to share with whom, and to set their settings accordingly[6]. Their decisions were saved to our database.

Finally, participants were asked to fill out a questionnaire measuring their perceptions of over-sharing threat, the perceived ease of use of the settings interface, and their anticipated satisfaction with the SNS. The survey concluded by measuring participants' interpersonal privacy concerns.

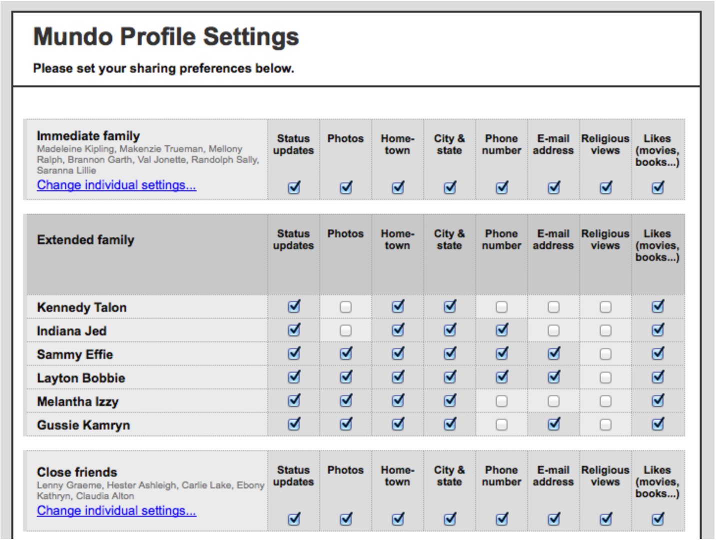

Mundo Profile Settings

Please set your sharing preferences below.

Immediate family Madeleine Kipling, Makenzie Trueman, Mellony Ralph, Brannon Garth, Val Jonette, Randolph Sally, Saranna Lillie Change individual settings...	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
	☑	☑	☑	☑	☑	☑	☑	☑

Extended family	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
Kennedy Talon	☑	☐	☑	☑	☐	☐	☐	☑
Indiana Jed	☑	☐	☑	☑	☑	☐	☐	☑
Sammy Effie	☑	☑	☑	☑	☑	☑	☐	☑
Layton Bobbie	☑	☑	☑	☑	☑	☑	☐	☑
Melantha Izzy	☑	☑	☑	☑	☐	☐	☐	☑
Gussie Kamryn	☑	☑	☑	☑	☐	☑	☐	☑

Close friends Lenny Graeme, Hester Ashleigh, Carlie Lake, Ebony Kathryn, Claudia Alton Change individual settings...	Status updates	Photos	Home-town	City & state	Phone number	E-mail address	Religious views	Likes (movies, books...)
	☑	☑	☑	☑	☑	☑	☑	☑

Figure 2. Mundo profile settings interface.

Manipulations

The study was implemented as a 3x2x3x2 between-subjects experiment. The **category granularity** was manipulated at three levels: users categorized their contacts into 5 (**normal**), 10 (**granular**) or 14 (**very granular**) categories. The categories, presented in Table 1, are based on the results of a previous study (Knijnenburg et al. 2014) that explored the optimal privacy-relevant contact categorization. In that study we determined three candidate categorizations (at different levels of granularity) that optimized their discriminant and convergent validity. In other words, these categories were developed in such a way that users would most likely have similar sharing preferences for all the contacts in a particular category, thereby minimizing the potential for a misspecification of their preferences.

We also manipulated the **exceptions** feature: participants in the **exceptions** condition had the ability to click on a link for each category that would expand it with a row of checkboxes for each individual contact (see Figure 2). In the **no exceptions** condition, this link was not available.

[5] A comprehension task preceded the categorization task to ascertain that participants indeed categorized their contacts carefully. Participants who did not carefully categorize in the comprehension task were removed from the analysis.

[6] Again, participants who spent too little time in this step to have carefully reviewed their settings were removed from the analysis.

The **defaults** of the settings interface were manipulated at three levels: no checkboxes would be checked in the **private-by-default** condition, all checkboxes would be checked in the **shared-by-default** condition, and a subset of the checkboxes would be checked in the **smart default** condition. This subset, presented in Table 1, was determined based on the results of our categorization study (Knijnenburg et al. 2014). Specifically, we checked a box if participants of the previous study shared that item with members of that category at a rate of at least 70% (we used this conservative threshold because potential over-sharing is arguably more problematic than potential under-sharing; cf. Benisch et al. (2011) use this assumption in modeling users' privacy settings for a location-sharing service).

Finally, we manipulated the **presentation order** of the categories. The **stronger ties first** order started with family and friends (typically stronger ties; in our categorization study participants shared on average 88% and 87% of their information to these categories, respectively), then classmates and colleagues (typically weaker ties; 81% and 66% sharing in our categorization study), and then (various types of) acquaintances (typically the weakest ties; 58% sharing in our categorization study). The **weaker ties first** condition used the opposite order. The exact order of the categories in both conditions is listed in Table 1 (SF = stronger ties first, WF = weaker ties first).

Table 1. The categories at different levels of granularity.

Granularity: normal			**Smart default settings**							
SF	*WF*	*Label*	*Status updates*	*Photos*	*Home-town*	*City & state*	*Phone number*	*E-mail address*	*Religious views*	*Likes*
1	5	Family members	✓	✓	✓	✓	✓	✓	✓	✓
2	4	Friends	✓	✓	✓	✓	✓	✓	✓	✓
3	3	Classmates	✓	✓	✓	✓	•	•	•	✓
4	2	Colleagues	•	•	✓	✓	•	✓	•	•
5	1	Acquaintances	•	•	✓	•	•	•	•	•
Granularity: granular			**Smart default settings**							
SF	*WF*	*Label*	*Status updates*	*Photos*	*Home-town*	*City & state*	*Phone number*	*E-mail address*	*Religious views*	*Likes*
1	9	Immediate family	✓	✓	✓	✓	✓	✓	✓	✓
2	10	Extended family	✓	✓	✓	✓	✓	✓	•	✓
3	6	Close friends	✓	✓	✓	✓	✓	✓	✓	✓
4	7	Regular friends	✓	✓	✓	✓	✓	✓	✓	✓
5	8	Best behavior friends	✓	✓	✓	✓	✓	✓	•	✓
6	4	Peer colleagues on my team	•	•	✓	✓	✓	✓	•	✓
7	5	Other colleagues	•	•	✓	✓	•	✓	•	•
8	3	Classmates	✓	✓	✓	✓	•	•	•	✓
9	1	Infrequent contacts	✓	✓	✓	✓	•	•	•	✓
10	2	People I hardly know / don't trust	•	•	•	•	•	•	•	•
Granularity: very granular			**Smart default settings**							
SF	*WF*	*Label*	*Status updates*	*Photos*	*Home-town*	*City & state*	*Phone number*	*E-mail address*	*Religious views*	*Likes*
1	12	Immediate family	✓	✓	✓	✓	✓	✓	✓	✓
2	13	Relatives younger than me	✓	✓	✓	✓	✓	✓	✓	✓
3	14	Extended family	✓	✓	✓	✓	✓	✓	•	✓
4	8	Close friends	✓	✓	✓	✓	✓	✓	✓	✓
5	9	Regular friends	✓	✓	✓	✓	✓	✓	✓	✓
6	10	Best-behavior friends	✓	✓	✓	✓	✓	✓	•	✓
7	11	Friends from my past	✓	✓	✓	✓	•	✓	•	✓
8	5	Peer colleagues on my team	•	•	✓	✓	✓	✓	•	✓
9	6	Colleagues on other teams	•	•	✓	✓	•	•	•	•
10	7	Superiors/Subordinates	•	•	✓	✓	•	✓	•	•
11	4	Classmates	✓	✓	✓	✓	•	•	•	✓
12	1	Infrequent contacts	✓	✓	✓	✓	•	•	•	✓
13	2	People I hardly know	✓	•	✓	•	•	•	•	✓
14	3	People I don't trust	•	•	•	•	•	•	•	•

Table 2. The CFA outcome of the post-experimental questionnaire.

Factor	Item	Loading
System satisfaction	I would use Mundo if it was available	0.868
Alpha: 0.93, AVE: 0.747	Using Mundo is annoying	
	Based on what I have seen, Mundo is useful	0.872
Correlations with:	Using Mundo makes me happy	0.854
PEOU: 0.541	So far, I am satisfied with Mundo	0.902
PODT: −0.336	I would recommend Mundo to others	0.913
SPC: −0.157	I would quickly abandon using Mundo	−0.768
Perceived ease of use	Setting my preferences in Mundo is convenient	
Alpha: 0.79, AVE: 0.689	It takes many mouse-clicks to set my preferences with Mundo	
	I was able to quickly set my sharing preferences	0.781
Correlations with:	I find the Mundo application easy to use	
SAT: 0.541	Setting my preferences in Mundo was unnecessarily complex	−0.818
PODT: −0.660	I felt lost using the Mundo profile settings	−0.888
SPC: −0.223		
Perceived over-sharing threat	I am afraid that due to my Mundo settings, I am sharing my profile information too freely	0.882
Alpha: 0.88, AVE: 0.676	I am comfortable with the amount of profile sharing I chose	−0.859
	Due to Mundo, people will know too much about me	0.867
Correlations with: SAT: −0.336	I made sure that nobody gets to see more information about me than I am comfortable with	−0.732
PEOU: −0.660	I fear that I have been too liberal in selecting my profile settings	0.798
SPC: 0.313	My profile settings are spot on; I am not disclosing too much to anyone	−0.783
Interpersonal	It usually bothers me when people ask me something personal	0.817
privacy concerns	I will tell people anything they want to know about me	−0.706
Alpha: 0.90, AVE: 0.622	I have nothing to hide from other people	
	I am concerned that people know too many personal things about me	
Correlations with: SAT: −0.157	Compared to others, I am more sensitive about sharing personal information with other people	0.799
PEOU: −0.223	To me, it is the most important thing to keep things private from others	0.706
PODT: 0.313	I am not bothered that other people know personal things about me	
	Most of the personal things I share on are publicly available anyway	
	I worry about personal information being viewed by other people than those I wanted to see it	
	I worry that others could take things I reveal about myself out of context	
	When people ask me something personal, I sometimes think twice before telling them	0.815
	I think it is risky to tell people personal things about myself	0.872
	Sharing personal information with others may involve many unexpected problems	
	I feel safe telling people personal things about me	−0.892
	I feel comfortable sharing my private thoughts and feelings with others	−0.774
	I usually discuss my problems and concerns with others	

Dependent variables

This work considers both the subjective and behavioral consequences of the manipulated variables. The main behavioral dependent variable was **sharing tendency**, a dichotomous (yes/no) variable that was measured for each item (i.e. the 8 items listed in Table 1) and each recipient. Additionally, at the end of the study, we measured the following subjective constructs (listed in order of presentation):

- **System satisfaction**: This construct is based on benefit-related constructs that have been used in privacy research, such as "preference for benefits" (Hui et al. 2006), "disclosure-privacy benefits" (Xu et al. 2009) and "perceived benefits of info disclosure" (Xu et al. 2011). The items are taken from Knijnenburg and Kobsa (2013).

- **Perceived ease of use of the settings interface**: This construct is related to the "perceived ease of use" construct in TAM (Davis 1989; Davis et al. 1989) and the "effort expectancy" construct in UTAUT (Venkatesh et al. 2003). Krasnova et al. (2010) adapt the ease of use construct to SNS privacy settings with their "convenience in relationship maintenance" construct.
- **Perceived over-sharing threat**: This construct represents users' fear of unwanted collection and loss of control resulting from a lack of comfort or confidence regarding the system's privacy settings. In the human-computer interaction field it has generally been measured by a single item (e.g. Church et al. 2009; Lipford et al. 2010; Tang et al. 2012). Our multi-item construct is a social adaptation from the "perceived privacy threats" factor of Knijnenburg and Kobsa (2013), which is in itself a system-specific analogy to the "collection" and "control" factors of the Internet Users Information Privacy Concerns scale (Malhotra et al. 2004).
- **Interpersonal privacy concerns**: Several scales have been developed to measure people's privacy concerns (Dinev and Hart 2004; Malhotra et al. 2004; Smith et al. 1996; Stewart and Segars 2002), but these scales primarily focus on information privacy rather than interpersonal privacy (Page et al. 2013). Our interpersonal privacy concerns scale is based on the "risk" construct of Malhotra et al. (2004) and the items of the General Information Privacy Concerns scale (Smith et al. 1996). Several new items were inspired by scales developed at the CSCW2013 workshop on Measuring Networked Social Privacy (Page et al. 2013).

Measurement validity of these constructs was tested with a confirmatory factor analysis (CFA, Table 2). Since the items were measured on a 5-point scale, we used a WLS estimator that treats them as ordered-categorical. We iteratively removed items with a communality < 0.450 or high cross-loadings; these items have no loading in Table 2. The results show adequate convergent and discriminant validity. Since the interpersonal privacy concerns construct serves as a moderator in several hypotheses, its items were summed as an equal-weight index and dichotomized. This approach allows a "simple slope analysis" of the moderation results (Aiken and West 1991), which is easier to interpret and more robust to outliers.

Results

Behavioral Hypotheses (H1, H4-5, and H7)

The hypothesized effects of the experimental conditions and interpersonal privacy concerns on users' sharing tendency were tested as a repeated measures generalized linear model (glmer) with a logit link function and random intercepts for participant, recipient and item[7]. Category granularity (H1) did not have significant effects on users' sharing tendency, so this effect was trimmed from the model. The regression coefficients of the final model are presented in Table 3. The condition with the lowest hypothesized level of sharing (weaker ties first, private-by-default, high privacy concerns) was chosen as the baseline condition. For reference, Figure 3 presents the average observed sharing levels (across all items and recipients) in the relevant experimental conditions. These observed levels match the estimated levels rather well. Below we discuss the individual hypothesized effects.

As said, there is no difference in sharing tendency between the normal, granular (p = .316) and very granular (p = .134) conditions. H1 is thus not supported.

H4 argues that users would have a higher sharing tendency in the shared-by-default condition than in the private-by-default condition, and that the smart default condition would be somewhere in between these two. Indeed, compared to the private-by-default condition, the odds of sharing are estimated to be 2.1 times as high in the smart default condition, and 3.9 times as high in the shared-by-default-condition. H4 is thus supported.

[7] The analysis treats the sharing tendency to each individual contact as a separate decision. This ignores the fact that participants were mainly (or in some conditions: exclusively) making decision on a *per category* basis rather than a *per contact* basis. We also performed the presented analysis at the category level (necessarily ignoring the exceptions some participants made for specific contacts). This alternative analysis produced essentially the same results as those presented here.

Interestingly, though, the effect is *weaker* (rather than stronger) for participants with low interpersonal privacy concerns: the interaction effect more than halves the odds ratio for the smart default and shared-by-default condition. This is opposite to what we hypothesized in H4a. H5 is supported, though: the odds of sharing for participants with low interpersonal privacy concerns are estimated to be 2.7 times as high as for participant with high interpersonal privacy concerns.

Ordering the categories weaker ties first results in 1.8 times higher odds of sharing than ordering them stronger ties first (H7 supported). However, this effect only occurs in the private-by-default condition: the interaction effect again just about halves the odds ratio for the smart default and shared-by-default condition (H7a supported).

Consequently, Figure 3 shows that there is indeed a strong default effect, but mainly for participants with high privacy concerns in the stronger ties first condition. The default effect is less apparent in the weaker ties first condition and for participants with low privacy concerns. Nevertheless, the shared-by-default condition results in significantly higher sharing tendencies than the other two conditions under all circumstances, except for participants with low privacy concerns in the weaker ties first condition. Under these specific circumstances, the private-by-default condition shows a sharing tendency that is significantly *higher* than the smart default condition, and about equal to the shared-by-default condition.

Table 3. Regression of sharing tendency.

Independent variable	*odds ratio*	*95% CI*	*p*
Intercept	1.079	(0.766, 1.521)	
Weaker ties first	1.827	(1.480, 2.255)	.004
Low concerns	2.699	(2.186, 3.332)	< .001
Smart default	2.138	(1.663, 2.749)	.002
Shared-by-default	3.940	(3.050, 5.090)	< .001
Smart, weaker first	0.545	(0.407, 0.730)	.038
Share, weaker first	0.554	(0.411, 0.746)	.047
Smart, low concerns	0.483	(0.360, 0.649)	.014
Share, low concerns	0.462	(0.343, 0.623)	.010

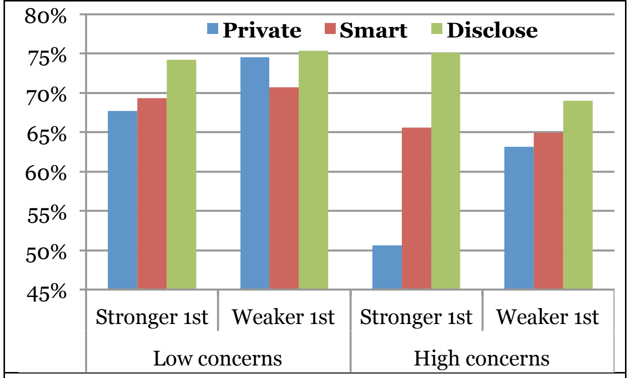

Figure 3. Observed sharing levels.

Subjective Hypotheses (H2-3, H6, and H8-10)

We tested the hypothesized effects of the experimental conditions on the measured subjective constructs using structural equation modeling (SEM) with a weighted least squares estimator. To allow a moderating relationship, interpersonal privacy concerns was not modeled as a latent construct but calculated as a sum of relevant items based on the CFA results and then dichotomized. Paths that were non-significant ($p > .05$) were removed from the model, and a post-hoc interaction effect (over-sharing threat on defaults × exceptions) was added based on high modification indices. In general, models can be trimmed or built based on theoretical and/or empirical standards (Kline 2004). Figure 3 shows the resulting model, which has a good[8] fit ($\chi^2(433) = 552.883$, p = .0001; CFI = 0.988, TLI = 0.989; RMSEA = .038, 90% CI: [.028, .047]). Below we discuss the individual hypothesized effects.

[8] A good model has a χ^2 that is not statistically different from a saturated model ($p > .05$). However, this statistic is regarded as too sensitive, and researchers have proposed other fit indices (Bentler and Bonett 1980). Hu and Bentler (1999) propose cut-off values for these indices to be: *CFI* > .96, *TLI* > .95, and *RMSEA* < .05, with the upper bound of its 90% CI falling below 0.10.

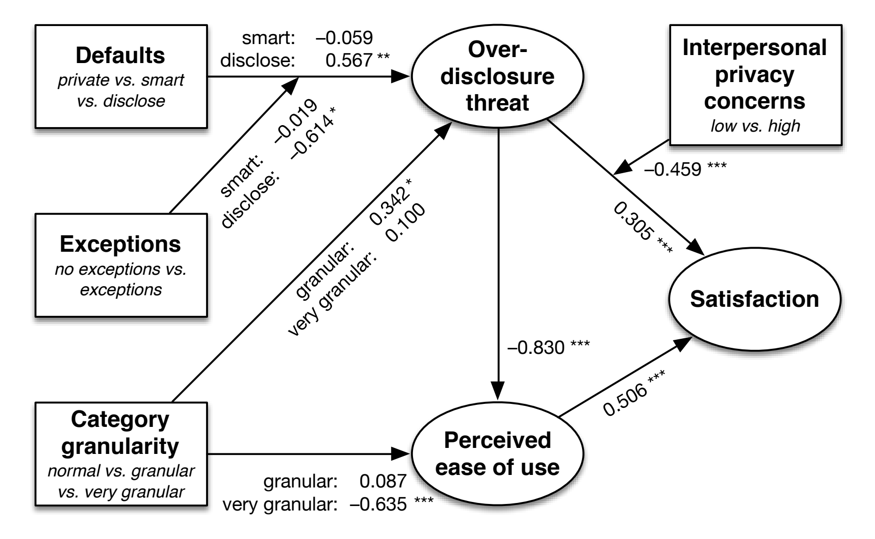

Figure 2. Trimmed structural equation model.
*** p < .05, ** p < .01, *** p < .001.**

H2 argues that the more granular categorizations would reduce users' perceived over-sharing threat. However, compared to the normal granularity condition, *participants perceive higher (rather than lower) over-sharing threat in the more granular categorization conditions* (yet interestingly not significantly in the very granular categorization condition). This is opposite to what we hypothesized in H2. H3 argues that the more granular categorizations reduce the perceived ease of use of the settings interface. The perceived ease of use is indeed significantly lower in the very granular categorization condition.

In line with H6 we find that users of the shared-by-default condition perceive a higher level of over-sharing threat than users of the private-by-default and smart default conditions. As hypothesized, this effect disappears when participants can make exceptions; in that case shared-by-default results in no higher over-sharing threat than the private-by-default condition (p = .852). H6a is thus also supported.

H8 and H8a argue that over-sharing threat is negatively related to satisfaction, and that this effect is stronger for users with a high level of interpersonal privacy concerns. Interestingly, the effect of over-sharing threat on satisfaction is *positive* for participants with low interpersonal privacy concerns (H8 not supported), and only slightly negative for participants with high interpersonal privacy concerns ($\beta_{\text{contrast}} = -0.154$, p = .052; the interaction effect is strongly significant, so H8a is supported). This unexpected effect is likely due to the strong mediated effect of over-sharing threat on satisfaction via perceived ease of use (hence, H9 and H10 are supported). The total effect of over-sharing on satisfaction is slightly negative for participants with low interpersonal privacy concerns ($\beta_{\text{total}} = -0.115$, p = .102) and strongly negative for participants with high interpersonal privacy concerns ($\beta_{\text{total}} = -0.574$, p < .001).

Discussion

Table 4 shows an overview of the support for our hypothesized effects. We discuss some of our surprising findings below. As hypothesized, defaults do have an effect on users' sharing tendency, but this effect is strongest when the stronger ties are presented first. Conversely, presenting weaker ties first increases sharing, but only in the private-by-default condition. This is in line with Acquisti et al. (2012), who demonstrate that an "accept frame" is a precondition for the "door-in-the-face" effect to occur. Alternatively, it may be that the lack of an effect of category order on sharing tendency in the shared-by-default condition is simply due to a ceiling effect.

Table 4. Overview of supported and rejected hypotheses.

Hypothesis		Support	Reason
H1.	Granularity → sharing tendency	No	No effect found
H2.	Granularity → threat	No	Opposite effect; granular *increases* threat
H3.	Granularity → ease of use	Partial	Only for the very granular categorization
H4.	Defaults → sharing tendency	Yes	
H4a.	Defaults × concerns → sharing tendency	No	Opposite effect: default effect *increases* with higher concerns
H5.	Concerns → sharing tendency	Yes	
H6.	Defaults → threat	Yes	
H6a.	Defaults × exceptions → threat	Yes	
H7.	Order → sharing tendency	Yes	
H7a.	Order × defaults → sharing tendency	Yes	
H8.	Threat → satisfaction	No/Yes[9]	Opposite effect with low concerns, small negative effect with high concerns
H8a.	Threat × concerns → satisfaction	Yes	
H9.	Threat → ease of use	Yes	
H10.	Ease of use → satisfaction	Yes	

Unexpectedly, the default effect is *stronger* for people with *high* interpersonal privacy concerns than for people with low concerns (contrary to Lai and Hui's (2006) findings and our own hypothesis). This finding is surprising: typically people with low levels of concern are less motivated to make accurate privacy decisions, and are therefore more (rather than less) amenable to the default effect. Again, this finding may be due to a ceiling effect: people with low privacy concerns share at very high rates (around 70%) in any condition. A more provocative explanation would be that this is a case of reactance (McKenzie et al. 2006) in the *opposite* direction: unconcerned individuals disagree with the more conservative default policy of the smart defaults and private-by-default condition, and end up overcompensating by specifying their settings to share even more.

We find that granularity has no effect on users' sharing tendency, so evidently participants do not "err-on-the-safe-side" as was claimed by Benisch et al. (2011), Tang et al. (2012) and Sadeh et al. (2009). The "err-on-the-safe-side" suggestion seems to be based on the implicit assumption that users find over-sharing more problematic than under-sharing[10]. If that is true, users will tend to under-share whenever they encounter misspecifications. However, given that our study found no effect of granularity, we may argue that our participants found under-sharing and over-sharing roughly equally problematic. In fact, some recent studies have found similar effects (Knijnenburg et al. 2013b; Lin et al. 2012). Another explanation of the lack of an effect of granularity on sharing would be that even our coarsest category granularity is rather accurate, and leads to very few misspecifications. Indeed, very few participants in our study actually made any exceptions (see Table 5), and these participants were evenly distributed over the categorization granularity conditions, which suggests that even our coarsest categorization was very accurate.

Note that contrary to our expectations, users perceive more (rather than less) threat in the granular condition. This suggests that users may not evaluate threats objectively, but rather superficially: the larger number of categories in the more granular conditions may make users *feel* like they are sharing more, thereby increasing their perceived over-sharing threat. Moreover, in line with our expectations, users find the very granular condition harder to use. Consequently, the "normal" granularity 5-category condition (family, friends, classmates, colleagues and acquaintances) eventually results in the highest user

[9] The unexpected positive effect is due to a strong mediation effect of threat on satisfaction via ease of use (see H9 and H10). The total effect of threat on satisfaction is indeed negative as expected.

[10] Indeed, in predicting users' privacy preferences, Benisch et al. explicitly define a "cost ratio" that varies from c = 1 (over- and under-sharing are equally penalized) to c = 100 (over-sharing is considered one-hundred times as bad as under-sharing).

satisfaction. This reflects an increasing trend in privacy research that shows that more control is not always better (Brandimarte et al. 2013; Knijnenburg et al. 2013b; Korff and Böhme 2014).

We expected that exceptions would be harder to use, but this hypothesis was not supported, arguably because very few participants made any exceptions. Exceptions do erase the over-sharing threat caused by the shared-by-default condition, though. This is a very important finding, as it allows a system to have a default that increases sharing without increasing users' perceived over-sharing threat. This finding is all the more surprising given that only very few participants made exceptions: it turns out that *merely the presence of the exception feature is enough to reduce users' perception of over-sharing threat!*

Table 5. Exceptions made by participants.

# of exceptions *(max: 400)*	**# of participants** *(total: 200)*
0	173
1-5	3
6-10	3
11-20	4
21-30	2
>30	2

Conclusion and Future work

In this paper we conducted a systematic evaluation of the effect of several design parameters of an SNS privacy settings interface on users' evaluation of the system and their sharing behavior. In terms of managerial implications, we find that it is important to beware of perceived over-sharing: over-sharing causes a threat that makes using the system seem more of a hassle (reduced ease of use) and that will—for people with high interpersonal privacy concerns—reduce system satisfaction.

Luckily, we find that it is possible to increase users' sharing tendency without inducing such negative subjective consequences: the shared-by-default setting increases sharing, and giving users the option to make exceptions for specific contacts reduces the over-sharing threat that may ensue from this default setting. Interestingly, this strategy works despite the fact that most users will not even use this exception feature at all. Alternatively, managers could employ the private-by-default setting and present the contact categories of weaker ties (i.e. acquaintances) first, but this only works for users with low interpersonal privacy concerns.

In terms of category granularity, we find no effect on sharing tendency, but the granular (10-category) categorization increases over-sharing threat, and the very granular (14-category) categorization is harder to use. Managers should thus rather employ the 5-category categorization.

There are several limitations to our work. First of all, our study design is rather complex, and while our sample has sufficient power to test the hypothesized main and 2-way interaction effects, it is not large enough to carefully examine 3- or 4-way interaction effects. A larger sample is needed to test these effects and assure the robustness of our results.

Moreover, in our study the categories were static and "system-defined", whereas users of existing systems such as Google+ and Facebook can create their own categories. Note, though, that Google+ users only have a median of five categories (just one more than the default four; Watson et al. (2012)), and most Facebook users only use the Friends category (Carr 2010; Deuker 2012; Strater and Lipford 2008). Kelley et al. (2011) provide a reason for this limited success of user-defined categories: they found that when people are prompted to categorize their friends into semantically meaningful categories, they often create categories that make sense from a social perspective, but that are inadequate for making privacy decisions. Consequently, system-defined categories based on an analysis of users' actual tendency to share information with the cateogry (such as those used in the current paper) may be better suited for the purpose of selectively sharing personal information.

Future work should further investigate some of our surprising results, such as why default effects are more prominent for people with high (rather than low) interpersonal privacy concerns, and why exceptions counteract the over-sharing threat of shared-by-default systems. One could argue that the availability of exceptions demonstrates the system's intentions to allow users to set their privacy settings in a flexible manner, making the system seem benevolent, thereby reducing the over-sharing threat resulting from the shared-by-default setting.

The smart default setting, as proposed by Smith et al. (2013) did not stand out as the superior solution: it did not result in higher sharing tendencies than the private-by-default setting (except for people with high privacy concerns in the stronger ties first condition), and it did not reduce over-sharing threat or increase ease of use. Although we based the smart default on previous data using a somewhat arbitrary threshold, this threshold seemed to be quite accurate: analyzing users' deviations from the default, we find that the number of "shares" turned into "not shares" (median: 11, mean: 21.9) is about equal to the number of "not shares" turned into "shares" (median: 10, mean: 15.9). Also, the total number of changes in the smart default condition (median: 25, mean: 37.7) is much lower than the number of changes in both the private-by-default condition (median: 257, mean: 251.9) and the shared-by-default condition (median: 91, mean: 103.3). This means that users in the smart default condition indeed used far fewer clicks to optimally set their privacy settings, but that this reduction in physical effort was not accompanied by a reduction in cognitive effort (i.e. ease of use). Smith et al. (2013) argue that "Smart defaults can become even smarter by adapting to information provided by the consumer as part of the decision-making process" (p. 167). Such adaptive defaults are an interesting venue to explore in future work (cf. Knijnenburg 2014; Wisniewski et al. 2014).

More generally, we encourage privacy researchers, policy-makers, and industry executives to consider the effects of privacy settings interfaces on privacy outcomes. This paper shows that subtle changes in the design of such interfaces can have important subjective and behavioral consequences. Careful design of these systems is thus very important to make users share optimally without feeling like they over-share.

References

Acquisti, A. 2004. "Privacy in electronic commerce and the economics of immediate gratification," in *Proceedings of the 5th ACM conference on Electronic commerce*, New York, NY: ACM, pp. 21–29.

Acquisti, A. 2009. "Nudging Privacy: The Behavioral Economics of Personal Information," *IEEE Security and Privacy* (7), pp. 82–85.

Acquisti, A., and Gross, R. 2006. "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," in *Privacy Enhancing Technologies*, Lecture Notes in Computer Science, G. Danezis and P. Golle (eds.), (Vol. 4258) Springer Berlin / Heidelberg, pp. 36–58.

Acquisti, A., and Grossklags, J. 2005. "Privacy and Rationality in Individual Decision Making," *IEEE Security & Privacy* (3:1), pp. 26–33.

Acquisti, A., and Grossklags, J. 2008. "What Can Behavioral Economics Teach Us About Privacy?," in *Digital Privacy: Theory, Technologies, and Practices*, A. Acquisti, S. De Capitani di Vimercati, S. Gritzalis, and C. Lambrinoudakis (eds.), New York/London: Auerbach Publications, pp. 363–377.

Acquisti, A., John, L. K., and Loewenstein, G. 2012. "The Impact of Relative Standards on the Propensity to Disclose," *Journal of Marketing Research* (49:2), pp. 160–174.

Aiken, L. S., and West, S. G. 1991. *Multiple Regression: Testing and Interpreting Interactions*, Thousand Oaks, CA: SAGE Publications, Inc.

Balebako, R., Leon, P. G., Mugan, J., Acquisti, A., Cranor, L. F., and Sadeh, N. 2011. "Nudging users towards privacy on mobile devices," in *CHI 2011 workshop on Persuasion, Influence, Nudge and Coercion Through Mobile Devices*, Vancouver, Canada, pp. 23–26.

Baron, J., and Ritov, I. 1994. "Reference Points and Omission Bias," *Organizational Behavior and Human Decision Processes* (59:3), pp. 475–498.

Benisch, M., Kelley, P. G., Sadeh, N., and Cranor, L. F. 2011. "Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs," *Personal Ubiquitous Computing* (15:7), pp. 679–694.

Bentler, P. M., and Bonett, D. G. 1980. "Significance Tests and Goodness of Fit in the Analysis of Covariance Structures," *Psychological Bulletin* (88:3), pp. 588–606.

Brandimarte, L., Acquisti, A., and Loewenstein, G. 2013. "Misplaced Confidences: Privacy and the Control Paradox," *Social Psychological and Personality Science* (4:3), pp. 340–347.

Brandyberry, A., Li, X., and Lin, L. 2010. "Determinants of Perceived Usefulness and Perceived Ease of Use in Individual Adoption of Social Network Sites," in *AMCIS 2010 Proceedings*, Lima, Peru.

Brown, C. L., and Krishna, A. 2004. "The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice," *Journal of Consumer Research* (31:3), pp. 529–539.

Bulgurcu, B., Cavusoglu, H., and Benbasat, I. 2010. "Understanding Emergence and Outcomes of Information Privacy Concerns: A Case of Facebook," in *ICIS 2010 Proceedings*, St. Louis, MO.

Cantrill, J. G., and Seibold, D. R. 1986. "The Perceptual Contrast Explanation of Sequential Request Strategy Effectiveness," *Human Communication Research* (13:2), pp. 253–267.

Carr, A. 2010. "Facebook's New Groups, Dashboards, and Downloads Explained," *Fast Company*.

Chapman, G. B., and Johnson, E. J. 1994. "The limits of anchoring," *Journal of Behavioral Decision Making* (7:4), pp. 223–242.

Chen, J.-J. V., Huang, A. H., and Muzzerall, A. 2012. "Privacy concerns and expectation of control," *Human Systems Management* (31:2), pp. 123–131.

Church, L., Anderson, J., Bonneau, J., and Stajano, F. 2009. "Privacy Stories: Confidence in Privacy Behaviors Through End User Programming," in *Proceedings of the 5th Symposium on Usable Privacy and Security*, Mountain View, CA: ACM, pp. 20:1–20:1.

Cialdini, R. B., Vincent, J. E., Lewis, S. K., Catalan, J., Wheeler, D., and Darby, B. L. 1975. "Reciprocal concessions procedure for inducing compliance: The door-in-the-face technique," *Journal of Personality and Social Psychology* (31:2), pp. 206–215.

Collins, R., Dwyer, C., Hiltz, S., and Shrivastav, H. 2012. "Do I Know What You Can See? Social Networking Sites and Privacy Management," in *AMCIS 2012 Proceedings*, Seattle, WA.

Compañó, R., and Lusoli, W. 2010. "The Policy Maker's Anguish: Regulating Personal Data Behavior Between Paradoxes and Dilemmas," in *Economics of Information Security and Privacy*, T. Moore, D. Pym, and C. Ioannidis (eds.), New York, NY: Springer US, pp. 169–185.

Connolly, T., and Zeelenberg, M. 2002. "Regret in decision making," *Current directions in psychological science* (11:6), pp. 212–216.

Consumer Reports. 2012. "Facebook & your privacy: Who sees the data you share on the biggest social network?," *Consumer Reports*.

Davis, F. D. 1989. "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly* (13:3), p. 319.

Davis, F. D., Bagozzi, R. P., and Warshaw, P. R. 1989. "User Acceptance of Computer Technology: A Comparison of Two Theoretical Models," *Management Science* (35:8), pp. 982–1003.

Deuker, A. 2012. "Friend-to-Friend Privacy Protection on Social Networking Sites: A Grounded Theory Study," in *AMCIS 2012 Proceedings*, Seattle, WA.

Dey, R., Jelveh, Z., and Ross, K. 2012. "Facebook users have become much more private: A large-scale study," in *2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pp. 346–352.

Dhillon, G., and Chowdhuri, R. 2013. "Individual values for protecting identity in social networks," in *ICIS 2013 Proceedings*, Milan, Italy.

Dinev, T., and Hart, P. 2004. "Internet Privacy Concerns and Their Antecedents: Measurement Validity and a Regression Model," *Behaviour & Information Technology* (23:6), pp. 413–422.

Duggan, M., and Smith, A. 2014. "Social media update 2013," Pew Research Center.

Egelman, S., Oates, A., and Krishnamurthi, S. 2011. "Oops, I Did It Again: Mitigating Repeated Access Control Errors on Facebook," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Vancouver, Canada: ACM, pp. 2295–2304.

Fox, S., Rainie, L., Horrigan, J., Lenhart, A., Spooner, T., and Carter, C. 2000. "Trust and Privacy Online: Why Americans Want to Rewrite the Rules," The Pew Internet & American Life Project.

Ganzach, Y. 1995. "Attribute Scatter and Decision Outcome: Judgment versus Choice," *Organizational Behavior and Human Decision Processes* (62:1), pp. 113–122.

Van de Garde-Perik, E., Markopoulos, P., de Ruyter, B., Eggen, B., and Ijsselsteijn, W. 2008. "Investigating Privacy Attitudes and Behavior in Relation to Personalization," *Social Science Computer Review* (26:1), pp. 20–43.

Gross, R., and Acquisti, A. 2005. "Information revelation and privacy in online social networks," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, Alexandria, VA, USA: ACM, pp. 71–80.

Hampton, K., Goulet, L. S., Marlow, C., and Rainie, L. 2012. "Why most Facebook users get more than they give," Pew Internet & American Life Project.

Hann, I.-H., Hui, K.-L., Lee, S.-Y., and Png, I. 2007. "Overcoming Online Information Privacy Concerns: An Information-Processing Theory Approach," *Journal of Management Information Systems* (24:2), pp. 13–42.

Hey Tow, W., Dell, P., and Venable, J. 2008. "Understanding Information Disclosure Behaviour in Australian Facebook Users," in *ACIS 2008 Proceedings*, Toronto, Canada.

Hu, L., and Bentler, P. M. 1999. "Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives," *Structural Equation Modeling: A Multidisciplinary Journal* (6:1), pp. 1–55.

Hu, T., Poston, R., and Kettinger, W. 2011. "Nonadopters of Online Social Network Services: Is It Easy to Have Fun Yet?," *Communications of the Association for Information Systems* (29:1), pp. 441–458.

Hui, K.-L., Tan, B. C. Y., and Goh, C.-Y. 2006. "Online information disclosure: Motivators and measurements," *ACM Transactions on Internet Technology* (6:4), pp. 415–441.

Hull, G., Lipford, H. R., and Latulipe, C. 2011. "Contextual gaps: privacy issues on Facebook," *Ethics and Information Technology* (13:4), pp. 289–302.

Iyengar, S. S., and Lepper, M. R. 2000. "When choice is demotivating: Can one desire too much of a good thing?," *Journal of Personality and Social Psychology* (79:6), pp. 995–1006.

Jacowitz, K. E., and Kahneman, D. 1995. "Measures of Anchoring in Estimation Tasks," *Personality and Social Psychology Bulletin* (21:11), pp. 1161–1166.

Jensen, C., Potts, C., and Jensen, C. 2005. "Privacy Practices of Internet Users: Self-Reports versus Observed Behavior," *International Journal of Human-Computer Studies* (63:1-2), pp. 203–227.

John, L. K., Acquisti, A., and Loewenstein, G. 2011. "Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information," *Journal of consumer research* (37:5), pp. 858–873.

Johnson, E. J., Bellman, S., and Lohse, G. L. 2002. "Defaults, Framing and Privacy: Why Opting In ≠ Opting Out," *Marketing Letters* (13:1), pp. 5–15.

Johnson, E. J., and Goldstein, D. 2003. "Do Defaults Save Lives?," *Science* (302:5649), pp. 1338–1339.

Kahneman, D., Knetsch, J. L., and Thaler, R. H. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias," *Journal of Economic Perspectives* (5:1), pp. 193–206.

Kahneman, D., and Tversky, A. 2000. *Choices, Values, and Frames*, Cambridge: Cambridge Univ. Press.

Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. 2012. "Talking in circles: Selective Sharing in Google+," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Austin, TX: ACM Press, pp. 1065–1074.

Kelley, P. G., Brewer, R., Mayer, Y., Cranor, L. F., and Sadeh, N. 2011. "An Investigation into Facebook Friend Grouping," in *INTERACT*, Lecture Notes in Computer Science, P. Campos, N. Graham, J. Jorge, N. Nunes, P. Palanque, and M. Winckler (eds.), (Vol. 6948) Lisbon, Portugal: Springer Heidelberg, pp. 216–233.

Kittur, A., Chi, E. H., and Suh, B. 2008. "Crowdsourcing user studies with Mechanical Turk," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Florence, Italy: ACM Press, pp. 453–456.

Kline, R. B. 2004. *Beyond significance testing: reforming data analysis methods in behavioral research*, Washington, DC: American Psychological Association.

Knijnenburg, B. P. 2014. "Information Disclosure Profiles for Segmentation and Recommendation," in *SOUPS2014 Workshop on Privacy Personas and Segmentation*, Menlo Park, CA.

Knijnenburg, B. P., and Jin, H. 2013. "The Persuasive Effect of Privacy Recommendations," in *Twelfth Annual Workshop on HCI Research in MIS*, Milan, Italy.

Knijnenburg, B. P., and Kobsa, A. 2013. "Making Decisions about Privacy: Information Disclosure in Context-Aware Recommender Systems," *ACM Transactions on Interactive Intelligent Systems* (3:3), pp. 20:1–20:23.

Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013a. "Counteracting the Negative Effect of Form Auto-completion on the Privacy Calculus," in *ICIS 2013 Proceedings*, Milan, Italy.

Knijnenburg, B. P., Kobsa, A., and Jin, H. 2013b. "Preference-based location sharing: are more privacy options really better?," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Paris, France: ACM, pp. 2667–2676.

Knijnenburg, B. P., Kobsa, A., and Jin, H. 2014. *Segmenting the Recipients of Personal Information*, unpublished manuscript, UC Irvine.

Korff, S., and Böhme, R. 2014. "Too Much Choice: End-User Privacy Decisions in the Context of Choice Proliferation," in *Symposium on Usable Privacy and Security (SOUPS)*, Menlo Park, CA.

Krasnova, H., Günther, O., Spiekermann, S., and Koroleva, K. 2009. "Privacy concerns and identity in online social networks," *Identity in the Information Society* (2:1), pp. 39–63.

Krasnova, H., Hildebrand, T., and Guenther, O. 2009. "Investigating the Value of Privacy in Online Social Networks: Conjoint Analysis," in *ICIS 2009 Proceedings*, Phoenix, AZ.

Krasnova, H., Kolesnikova, E., and Guenther, O. 2009. "'It Won't Happen To Me!': Self-Disclosure in Online Social Networks," in *AMCIS 2009 Proceedings*, San Francisco, CA.

Krasnova, H., Spiekermann, S., Koroleva, K., and Hildebrand, T. 2010. "Online social networks: why we disclose," *Journal of Information Technology* (25:2), pp. 109–125.

Lai, Y.-L., and Hui, K.-L. 2006. "Internet Opt-In and Opt-Out: Investigating the Roles of Frames, Defaults and Privacy Concerns," in *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research*, Claremont, CA, pp. 253–263.

Lankton, N., and Tripp, J. 2013. "A Quantitative and Qualitative Study of Facebook Privacy using the Antecedent-Privacy Concern-Outcome Macro Model," in *AMCIS 2013 Proceedings*, Chicago, IL.

Lederer, S., Mankoff, J., and Dey, A. K. 2003. "Who wants to know what when? privacy preference determinants in ubiquitous computing," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Ft. Lauderdale, FL: ACM Press, pp. 724–725.

Lin, J., Benisch, M., Sadeh, N., Niu, J., Hong, J., Lu, B., and Guo, S. 2012. "A comparative study of location-sharing privacy preferences in the United States and China," *Personal and Ubiquitous Computing* (17:4), pp. 697–711.

Lipford, H. R., Besmer, A., and Watson, J. 2008. "Understanding Privacy Settings in Facebook with an Audience View," in *Proc. of the 1st Conference on Usability, Psychology, and Security*, Berkeley, CA, USA: USENIX Association.

Lipford, H. R., Watson, J., Whitney, M., Froiland, K., and Reeder, R. W. 2010. "Visual vs. compact: a comparison of privacy policy interfaces," in *Proceedings of the 28th international conference on Human factors in computing systems*, Atlanta, GA: ACM, pp. 1111–1114.

Liu, Y., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. 2011. "Analyzing Facebook privacy settings: user expectations vs. reality," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, Berlin, Germany: ACM, pp. 61–70.

Lo, J. 2010. "Privacy Concern, Locus of Control, and Salience in a Trust-Risk Model of Information Disclosure on Social Networking Sites," in *AMCIS 2010 Proceedings*, Lima, Peru.

Lukaszewski, K. M., Stone, D. L., and Stone-Romero, E. F. 2008. "The Effects of the Ability to Choose the Type of Human Resources System on Perceptions of Invasion of Privacy and System Satisfaction," *Journal of Business and Psychology* (23:3-4), pp. 73–86.

Madden, M. 2012. "Privacy management on social media sites," Washington, DC: Pew Internet & American Life Project, Pew Research Center.

Madden, M., Lenhart, A., Cortesi, S., Gasser, U., Duggan, M., and Smith, A. 2013. "Teens, Social Media, and Privacy," Pew Research Center.

Madejski, M., Johnson, M., and Bellovin, S. M. 2012. "A study of privacy settings errors in an online social network," in *2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, Lugano, Switzerland, pp. 340–345.

Malhotra, N. K., Kim, S. S., and Agarwal, J. 2004. "Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Nomological Framework," *Information Systems Research* (15:4), pp. 336–355.

Mason, W., and Suri, S. 2012. "Conducting behavioral research on Amazon's Mechanical Turk," *Behavior Research Methods* (44:1), pp. 1–23.

McKenzie, C. R. M., Liersch, M. J., and Finkelstein, S. R. 2006. "Recommendations Implicit in Policy Defaults," *Psychological Science* (17:5), pp. 414–420.

Meloy, M. G., and Russo, J. E. 2004. "Binary choice under instructions to select versus reject," *Organizational Behavior and Human Decision Processes* (93:2), pp. 114–128.

Metzger, M. J. 2006. "Effects of Site, Vendor, and Consumer Characteristics on Web Site Trust and Disclosure," *Communication Research* (33:3), pp. 155–179.

Norberg, P. A., Horne, D. R., and Horne, D. A. 2007. "The Privacy Paradox: Personal Information Disclosure Intentions versus Behaviors," *Journal of Consumer Affairs* (41:1), pp. 100–126.

Olson, J. S., Grudin, J., and Horvitz, E. 2005. "A study of preferences for sharing and privacy," in *CHI '05 Extended Abstracts*, Portland, OR: ACM, pp. 1985–1988.

Page, X., Tang, K., Stutzman, F., and Lampinen, A. 2013. "Measuring networked social privacy," in *Proceedings of the 2013 conference on Computer supported cooperative work companion*, San Antonio, TX: ACM, pp. 315–320.

Paolacci, G., Chandler, J., and Ipeirotis, P. 2010. "Running experiments on amazon mechanical turk," *Judgment and Decision Making* (5:5), pp. 411–419.

Park, C. W., Jun, S. Y., and MacInnis, D. J. 2000. "Choosing What I Want versus Rejecting What I Do Not Want: An Application of Decision Framing to Product Option Choice Decisions," *Journal of Marketing Research* (37:2), pp. 187–202.

Patil, S., and Lai, J. 2005. "Who Gets to Know What when: Configuring Privacy Permissions in an Awareness Application," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Portland, OR: ACM, pp. 101–110.

Patil, S., Schlegel, R., Kapadia, A., and Lee, A. J. 2014. "Reflection or action?: how feedback and control affect location sharing decisions," in *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, Toronto, Canada: ACM, pp. 101–110.

Phillips, B., and Shipps, B. 2012. "Frequency of Usage: The Impact of Technology Acceptance Factors Versus Social Factors," *International Journal of Virtual Communities and Social Networking* (4:2), pp. 30–45.

Posey, C., and Ellis, S. 2007. "Understanding Self-Disclosure in Electronic Communities: An Exploratory Model of Privacy Risk Beliefs, Reciprocity, and Trust," in *AMCIS 2007 Proceedings*, Keystone, CO.

Qin, L., Kim, Y., Hsu, J., and Tan, X. 2011. "The Effects of Social Influence on User Acceptance of Online Social Networks," *International Journal of Human-Computer Interaction* (27:9), pp. 885–899.

Rainie, L., Kiesler, S., Kang, R., and Madden, M. 2013. "Anonymity, Privacy, and Security Online," Washington, DC: Pew Research Center's Internet & American Life Project.

Sadeh, N., Hong, J., Cranor, L., Fette, I., Kelley, P., Prabaker, M., and Rao, J. 2009. "Understanding and capturing people's privacy policies in a mobile social networking application," *Personal and Ubiquitous Computing* (13:6), pp. 401–412.

Samuelson, W., and Zeckhauser, R. 1988. "Status quo bias in decision making," *Journal of Risk and Uncertainty* (1:1), pp. 7–59.

Scheibehenne, B., Greifeneder, R., and Todd, P. M. 2010. "Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload," *Journal of Consumer Research* (37:3), pp. 409–425.

Schwartz, B. 2003. *The paradox of choice: why more is less*, New York: Harper Perennial.

Shanab, M. E., and O'Neill, P. 1979. "The effects of contrast upon compliance with socially undesirable requests in the door-in-the-face paradigm," *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement* (11:3), pp. 236–244.

Sher, S., and McKenzie, C. R. M. 2006. "Information leakage from logically equivalent frames," *Cognition* (101:3), pp. 467–494.

Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., and Cranor, L. F. 2013. "The post that wasn't: exploring self-censorship on facebook," in *Proceedings of the 2013 conference on Computer supported cooperative work*, San Antonio, TX: ACM, pp. 793–802.

Smith, A. 2014. "6 new facts about Facebook," *Pew Research Center*.

Smith, H. J., Milberg, S. J., and Burke, S. J. 1996. "Information Privacy: Measuring Individuals' Concerns about Organizational Practices," *MIS Quarterly* (20:2), pp. 167–196.

Smith, N. C., Goldstein, D. G., and Johnson, E. J. 2013. "Choice Without Awareness: Ethical and Policy Implications of Defaults," *Journal of Public Policy & Marketing* (32:2), pp. 159–172.

Solove, D. J. 2013. "Privacy Self-Management and the Consent Dilemma," *Harvard Law Review* (126), pp. 1880–1903.

Sousa, L. F. 2009. "Privacy Policy Dynamics in Location Sharing Applications," Master Thesis, Lisbon, Portugal: University of Lisbon, Department of Computer Science.

Spiekermann, S., Grossklags, J., and Berendt, B. 2001. "E-privacy in 2nd Generation E-Commerce: Privacy Preferences versus Actual Behavior," in *Proceedings of the 3rd ACM conference on Electronic Commerce*, Tampa, FL, pp. 38–47.

Stewart, K. A., and Segars, A. H. 2002. "An Empirical Examination of the Concern for Information Privacy Instrument," *Information Systems Research* (13:1), pp. 36–49.

Strater, K., and Lipford, H. R. 2008. "Strategies and struggles with privacy in an online social networking community," in *Proc. of the 22nd British HCI Group Annual Conference on People and Computers*, Swinton, UK: British Computer Society, pp. 111–119.

Strater, K., and Richter, H. 2007. "Examining privacy and disclosure in a social networking community," in *Proceedings of the 3rd symposium on Usable privacy and security*, Pittsburgh, Pennsylvania: ACM, pp. 157–158.

Stutzman, F., Gross, R., and Acquisti, A. 2013. "Silent Listeners: The Evolution of Privacy and Disclosure on Facebook," *Journal of Privacy and Confidentiality* (4:2).

Sunstein, C. R., and Thaler, R. H. 2003. "Libertarian Paternalism Is Not an Oxymoron," *The University of Chicago Law Review* (70:4), pp. 1159–1202.

Tang, K., Hong, J., and Siewiorek, D. 2012. "The implications of offering more disclosure choices for social location sharing," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, Austin, TX: ACM Press, pp. 391–394.

Thaler, R. H., and Sunstein, C. 2008. *Nudge : improving decisions about health, wealth, and happiness*, New Haven, NJ & London, U.K.: Yale University Press.

Toch, E., Cranshaw, J., Drielsma, P. H., Tsai, J. Y., Kelley, P. G., Springfield, J., Cranor, L., Hong, J., and Sadeh, N. 2010. "Empirical models of privacy in location sharing," in *Proc. of the 12th ACM intl. conference on Ubiquitous computing*, Copenhagen, Denmark: ACM Press, pp. 129–138.

Tsai, J. Y., Kelley, P., Drielsma, P., Cranor, L. F., Hong, J., and Sadeh, N. 2009. "Who's viewed you?: the impact of feedback in a mobile location-sharing application," in *Proceedings of the 27th international conference on Human factors in computing systems*, Boston, MA, USA: ACM, pp. 2003–2012.

Venkatesh, V., Morris, M. G., Davis, G. B., and Davis, F. D. 2003. "User Acceptance of Information Technology: Toward a Unified View," *MIS Quarterly* (27:3), pp. 425–478.

Wang, Y., Leon, P. G., Acquisti, A., Cranor, L. F., Forget, A., and Sadeh, N. 2014. "A Field Trial of Privacy Nudges for Facebook," in *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*, Toronto, Canada: ACM, pp. 2367–2376.

Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., and Cranor, L. F. 2011. "'I regretted the minute I pressed share': a qualitative study of regrets on Facebook," in *Proceedings of the Seventh Symposium on Usable Privacy and Security*, Pittsburgh, PA: ACM, pp. 10:1–10:16.

Watson, J., Besmer, A., and Lipford, H. R. 2012. "+Your circles: sharing behavior on Google+," in *Proc. of the 8th Symposium on Usable Privacy and Security*, Pittsburgh, PA: ACM.

Wedell, D. H. 1997. "Another look at reasons for choosing and rejecting," *Memory & Cognition* (25:6), pp. 873–887.

Wilson, T. D., Houston, C. E., Etling, K. M., and Brekke, N. 1996. "A new look at anchoring effects: Basic anchoring and its antecedents," *Journal of Experimental Psychology: General* (125:4), pp. 387–402.

Wisniewski, P., Knijnenburg, B. P., and Richter Lipford, H. 2014. "Profiling Facebook Users' Privacy Behaviors," in *SOUPS2014 Workshop on Privacy Personas and Segmentation*, Presented at the SOUPS2014 Workshop on Privacy Personas and Segmentation, Menlo Park, CA.

Xu, H., Luo, X. (Robert), Carroll, J. M., and Rosson, M. B. 2011. "The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing," *Decision Support Systems* (51:1), pp. 42–52.

Xu, H., Teo, H.-H., Tan, B. C. Y., and Agarwal, R. 2009. "The Role of Push-Pull Technology in Privacy Calculus: The Case of Location-Based Services," *Journal of Management Information Systems* (26:3), pp. 135–174.

Young, A. L., and Quan-Haase, A. 2009. "Information Revelation and Internet Privacy Concerns on Social Network Sites: A Case Study of Facebook," in *Proceedings of the Fourth International Conference on Communities and Technologies*, University Park, PA: ACM, pp. 265–274.