SOCIAL INTERVENTIONS TO REDUCE

THE PRIVACY PRARADOX

By

ISHA GHOSH

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Communication, Information, and Media

Written under the direction of

Vivek K. Singh

And approved by

_____

_____

_____

_____

New Brunswick, New Jersey

October, 2020

ABSTRACT OF THE DISSERTATION

Social Interventions to Reduce the Privacy Paradox

by Isha Ghosh

Dissertation Director: Vivek K. Singh

Individuals often demonstrate privacy behaviors that are contrary to their concerns about information sharing and use. Literature has termed this phenomenon the "privacy paradox." In this dissertation we seek to bridge the gap between information privacy concerns and demonstrated disclosure behavior using social interventions. We examined the essential elements of privacy concerns, information disclosure, social influence, privacy interventions, and individual interactions with interventions to study their complex relationships and the daily information disclosure challenges faced by individuals. The major purposes of this dissertation are to revisit the privacy paradox phenomenon, examine the relationships among privacy concerns and information disclosure, use these relationships to design novel interventions, and explore the role of social interventions in bridging the gap between information privacy concerns and behaviors.

This goal was realized by conducting a 20 day between-subject field study comparing the effectiveness of interventions based on social-proof and information inconsistency with a baseline to identify the most efficient way of reducing the concern-behavior gap. Findings show that knowledge about peer behavior caused individuals to rethink their own disclosure behavior. Individuals tended to believe that their privacy concerns and information disclosure behaviors are similar to many of their peers. When they received a reminder that their concern-behavior gap was higher than their peers, they were inclined to re-evaluate their privacy concerns and behavior.

Findings further showed that while any reminder about information privacy resulted in more privacy sensitive behavior, changing overall privacy concerns was a more nuanced and subtle process often influenced by external or contextual factors.

We also examined the effect that different interventions had on the cognitive processes guiding privacy decision making. From a thematic analysis of participant interviews, we were able to delve deeper into understanding how participants interacted with the interventions, the different ways in which each intervention affected privacy concerns and disclosure, and the different design elements that participants reacted to. We discussed shifts in privacy concerns and disclosure in detail breaking apart different elements of the interventions (textual, visual, numeric) and why participants found some of these elements to be more important than other. Finally, we discuss the implications of the concern-behavior gap, and the effectiveness of interventions in achieving this goal.

Acknowledgements

First, I would like to thank my advisor, Dr. Vivek K. Singh for his guidance, advice, encouragement, and support during my dissertation and beyond. I could not have travelled this journey, nor reached this milestone, without his great help. I am very grateful to him for being a great mentor and a role model of how to give a hand to others in need, and how to be an awesome advisor.

People often said I had a strong committee, and I was very proud of it. I would like to thank my committee members, Dr. Nicholas J. Belkin, Dr. Katherine Ongyanova, and Dr. Jessica Vitak for always motivating me and pushing me to think of the "So…what?" question. The time, energy, and efforts of each member in providing me with insightful comments and suggestions have helped strengthen my dissertation and are greatly appreciated. I would also like to thank my Master's Capstone advisor Dr. Jennifer Gibbs for introducing my to academic research and mentoring me during my initial experiments.

A few faculty members also deserve special acknowledgement for their generous support. I would especially like to thank Dr. Marie Radford, Dr. Charles Sentio, and Dr. Jeff Lane for their help and advice with different stages of the dissertation process. Thanks also to all my colleagues and behavioral informatics lab members: Abdul Aziz, Ishaan Singh, Gautam Sikka, Alex Finsch, and many others for their endless patience while brainstorming and designing this research. I would also like to thank Katie Park who helped me a lot with qualitative coding used in this dissertation.

Last but the most, I express my sincere gratitude to my parents. Drs. Pradyut and Smita Ghosh for their unconditional love and support. They have always believed in and supported me, not only for my research work but in every decision I have made. Without their patience, guidance, and wisdom, my doctoral journey would not be completed. My sister Bidisha for her humor and light-

hearted banter that helped me stay sane during graduate school. Finally, I would like to acknowledge my partner Anabil for always being there and believing in me even when I did not believe in myself. I would never have made it this far without the love, sacrifice, and encouragement from all these people.

Table of Contents

"To follow knowledge like a sinking star, beyond the utmost bound of human thought" –
Ulysees, Alfred, Lord Tennyson.

Chapter 1: Introduction

1.1. The Privacy Paradox

Anecdotal and empirical evidence have suggested that individuals often behave in ways that disclose their personal, perhaps sensitive, information online in spite of feeling concerned about the risks of online disclosure (Barnes, 2006; Acquisti & Grossklags, 2003). This dissertation aims to reduce this gap between privacy concerns and privacy behavior. Early research on this phenomenon by Acquisti and Grossklags (2003) found that the online disclosure of personal information is paradoxical. "*Surveys report that most individuals are concerned about the security of their personal information and are willing to act to protect it. Experiments reveal that very few individuals actually take any action to protect their personal information, even when doing so involves limited costs.*" (p.1). Barnes (2006) specifically discussed the behavior of young people using Social Networking Sites (SNS) and coined the term "*Privacy Paradox*" to describe the counter – attitudinal behavior of these users when disclosing information. Subsequent research analyzed the privacy paradox more explicitly, focusing on, the concern-behavior gap. These studies reported that privacy concerns exert only a weak effect on information disclosure or protection behavior (for a review, see Kokolakis, 2017). While data protection and privacy are found abstractly important, the actual consequences of disclosure at an individual or societal level, are poorly understood and often invisible. Therefore, even though individuals had a clear notion of the act of self-disclosure, they tended to come up with cognitive rationalizations or justifications to cope with the challenges of privacy threats.

Carrascal et al. (2013), for instance, found that internet users tend to trade in their browsing information for the relatively low reward of €7. A web survey by Taddicken (2014) also showed that privacy concerns hardly affect self-disclosure. In a more detailed survey, Spiekermann et al.

(2001) conducted a study attempting to understand the relationship between privacy concerns and online shopping behavior. Their work simulated an online shopping exercise during which they compared self-reported privacy preferences with actual disclosure behavior. Participants were first asked to complete a questionnaire on privacy attitudes and preferences and then, perform transactions in an online store. Spiekermann et al., (2001) found that participants revealed a large volume of highly personal information while conducting online transactions. Analysis by Acquisti and Gross (2006) also found evidence of the privacy paradox. In their work studying the influence of privacy concerns on disclosure behavior, they found that while most of the subjects (approx. 89%) reported to be either moderately or extremely concerned about privacy, more than 21% of that number sample admitted to having sharing their sensitive personal or financial information in return for discounts or raffle tickets. Similarly, Beresford et al., (2012) conducted a field experiment, during which participants were asked to buy a DVD from one of two online websites. The two sites were almost identical except that the first asked for income and date of birth to access the site, while the second store asked for favorite color and year of birth to give access. Participants clearly differentiated that the first store asked for more sensitive information than the second. However, when the price of the DVD was lowered on the first site, participants chose the cheaper option, although it asked for more sensitive information. A post-experimental questionnaire tested if subjects were unconcerned about privacy issues. 75% of participants indicated that they had a strong interest in data protection and 95% said that they were interested in the protection of their personal information.

The research work presented above provides evidence supporting the notion of a dichotomy between privacy concerns and disclosure behaviors. However, there is also a school of

thought, whose researchers who have raised doubts about the existence of a privacy paradox. According to research disputing the privacy paradox, while individuals disclose personal information if it will yield significant benefits, they are, at the same time, worried about how this information is handled. These concerns, however, increase or decrease on a case-by-case basis and are different for different individuals based on their personalities, backgrounds, previous experiences, and a number of other factors. This school of thought claims, every information disclosure exchange is influenced by different privacy concerns in different individuals. Therefore, what is perceived as paradoxical disclosure is rather a lack of understanding of the individuals' actual privacy concerns and motivations for disclosure. This concern about the use of their personal information by third parties is influenced leads to a lowering of information disclosure. Several studies investigating information disclosure in online environments have challenged the assumption that individuals do not attempt to protect their private information (boyd & Hargittai, 2010; Young & Quan – Haase, 2013). Research has found that social media users tended to use a variety of strategies to protect their personal information, such as using pseudonyms and falsifying information (Miltgen & Peyrat-Guillard, 2014), restricting access to their profiles and changing privacy settings (boyd & Hargittai, 2010), limiting friendship requests, and deleting tags and photos (Young & Quan-Haase, 2013)

Research on the privacy paradox phenomenon has therefore, produced contradictory results. Several studies show that a dichotomy exists between privacy concerns and disclosure behavior while other studies indicate that privacy concerns do have an effect on privacy protective behavior. However, if a dichotomy exists then it is especially troubling in today's digital world where sophisticated technologies have made the effective collection, storage, and analysis of vast amounts of personal information a common occurrence. Hence, a tendency to disclose information contrary

to their privacy concerns could have a number of ramifications such as being profiled, personal information being sold or shared to third parties, being stalked, or cyber-bullied. We therefore argue that testing if such a discrepancy exists between individual privacy concerns and disclosure behaviors and identifying effective methods to reduce this gap (if found) is of paramount importance.

1.2. Implications of Counter-Attitudinal Disclosure

Despite the debate about its existence, the privacy paradox continues to be of considerable interest to researchers in the information science community. In spite of the volume of research devoted to understanding the various factors influencing counter-attitudinal behavior, it continues to present a challenge to researchers. The existence of a paradox between privacy concerns and behavior is troublesome for obvious reasons. It may lead to troublesome or regrettable experiences, for example, while mobile coupons based on user's location information can provide highly personalized services, they also produced strong feelings of intrusion and (Sutanto, Palme, Tan, & Phang, 2013). Early research investigating the privacy paradox has found that the information disclosure practices of Internet users are problematic: Although many people are concerned about their online privacy, they still tended to share plenty of personal information on the web (Acquisti & Grossklags, 2005; Barnes, 2006). A paradoxical relationship between concerns and behavior suggests that online information disclosure behavior is irrational and that people are revealing an alarming amount of personal information to unknown audiences including third party institutions, resulting in potentially unintended consequences such as regretting shared posts (Wang, Norcie, Komanduri, Acquisti, Leon, and Cranor, 2011), being profiled or surveilled (Tufecki, 2008). This disclosed personal information could take the form of information about their demographic, social, financial, or medical characteristics.

The discrepancy between information privacy concerns and demonstrated behaviors also has significant implications for e-commerce, online social networking, and government privacy regulation. Privacy policy makers often rely on public opinion surveys measuring privacy attitudes or concerns when creating privacy legislations. However, if these concerns do not accurately reflect enacted behavior, these policies are weakened. This disparity between privacy concerns and actual behavior could also be disadvantageous for service providers. Consumers confronted with their paradoxical behavior, i.e. finding out about their personal data being tracked or shared, might react with resentment, leading to a loss of trust and causing damage to customer relationships. It is therefore of crucial importance not just to understand why people engage in counter-attitudinal behavior but also to identify effective strategies to reduce this discrepancy.

1.3. Theoretical Explanation of the Privacy Paradox

Researchers have put forward multiple theories explaining the privacy paradox. A popular theoretical approach for studying the relationship between attitudes and behaviors refer to the "*privacy calculus*" (Dinev & Hart, 2006). The privacy calculus theory states that when making a disclosure decision, individuals attempt a balance between the cost or privacy risk associated with disclosure and the potential rewards or services gained as a result of disclosure. Their final behavior is determined by the outcome of the privacy trade-off (Dinev & Hart, 2006; Tsai, Egelman, Cranor, & Acquisti, 2011; Acquisti et al., 2015). During this negotiation, individuals' often attribute a higher value to the benefits of disclosure and minimize the risks associated with disclosure. This implies that if the rewards gained by disclosure are momentarily perceived as being greater than the cost associated with disclosure, it may override a user's general concerns about privacy and induce disclosure. Multiple studies have used the privacy calculus model in order to better understand the existence of the privacy paradox (Acquisti et al., 2015, Dinev and Hart, 2006). This explanation is heavily dependent upon the users' level of awareness and understanding of the benefits and risks

associated with any online action, or transaction. It does however seem rather unlikely that an individual can accurately calculate the risks associated with disclosing data in online contexts. These risks are multiple, subjective to individual preferences, and vary depending on a number of random factors. For instance, a person might not mind disclosing their locations or preferences to their online friends but experience privacy intrusion when this information makes them the target for an ad campaign. This implies that an explanation based on purely rational cost-benefit calculations over-simplifies the phenomenon of counter-attitudinal behavior and can therefore not lead to a substantially better understanding of the privacy paradox. Furthermore, such a cognitive rational choice approach neglects the emotional and contextual factors influencing disclosure behavior. Many information disclosure activities – in both online and offline spheres – are routinely performed or driven by irrational affective factors (John, Acquisti, & Lowenstein, 2010).

A second theoretical approach to explaining online self-disclosure despite privacy concerns focuses on trust in the entity to whom disclosure is being made and the norms guiding information disclosure in that space (Nissenbaum, 2010). In this view, users do not necessarily consider and evaluate the specific risks and benefits of an online transaction. Rather, they depend on contextual factors such as the type of information being shared; the trust placed in the recipient; and norms of the environment in which disclosure is made (Taddiken, 2014; Nissenbaum, 2004; Nissenbaum, 2010). This model stresses on the importance of context and trust in information privacy models and contextual integrity is viewed as a key determinant of privacy behavior. Nissenbaum (2010) explains that these norms are largely dependent on the type of information being shared; the trust placed in the recipient; and the medium through which information is transmitted. For instance, an individual might be comfortable disclosing highly personal and sensitive details about their addiction to a doctor due to the formal and legal information norms of information disclosure that govern

a doctor's office. However, if this information is shared by the recipient, i.e. the doctor in an informal gathering among friends or colleagues, it changes the context and audience of that information. This sensitive information has been transmitted from a context with one set of information norms (the office) to another (an informal gathering), and the individual perceives a privacy violation.

However, this work also presumes that the individual in the doctor's office is aware and able to understand the range and social context within which information is disclosed. Although the flow of information to another context is where the privacy violation is experienced, individuals' disclosures depend upon their skills to read a social situation and their perception of context. This can be challenging in computer-mediated environments where the norms of the medium are skewed to encourage information disclosure.

In the case of online information disclosure, if an individual trusts the social network or e-commerce institutions when it comes to their data, their concerns over information disclosure would be lowered. Here trust could emerge as a rational calculation (e.g. decision to trust an institution, like Google, because the benefits of trusting this institution outweigh the costs) or as a result of emotional, intuitive, or contextual factors (e.g. decision to trust a service because it is popular or looks professional). In either case, the positive feeling towards information disclosure can be described as a result of specific beliefs in the entity's trustworthiness (Dwyer, Hilz, & Passerini, 2007). However, this work also presumes that while engaging in self-disclosure individuals' are aware and able to understand the range of disclosure and social context of the environment in which disclosure is made. This ability to accurately estimate the range and magnitude of disclosure is severely challenged in online environments where information is tracked, stored, and shared with large, often invisible, audiences.

Finally, incomplete information or the lack of transparency termed "*information asymmetry*" about data sharing and usage practices within website privacy policies has also been studied as a factor influencing disclosure. According to this theory, people often do not have access to all the necessary information about the risks and benefits of information disclosure, and therefore are unable to make informed privacy decisions. Simply put, individuals make privacy decisions in limited time having incomplete information about risks and benefits.

In order to educate users about the collection and use of their personal information, online website or apps often rely on the use of privacy policies. These services require users to signal their acceptance of the privacy policy or terms of service, typically by clicking "*I Agree*", in order to use the service. These sites argue that this ability to consciously accept or reject the terms governing the collection and use of personal data allows individuals to make an informed and explicit choice between disclosing personal information and protecting their privacy. However, in most cases there is no real choice at all. Either an individual agrees to the terms of data use described by the service provider or they choose not to use the service at all. In addition, online institutions do not make it easy for users to make this choice. Privacy policies often span several pages and are written in unwieldy "*legalese*." Several previous researchers have demonstrated that most people do not read these documents before accepting them (Besmer, Watson, & Lipford, 2010; Milne & Culnan, 2002). Because most users do not take the time to read and understand privacy disclosures, their understanding of what happens to their disclosed information is likely to be low.

Even if users were to take the time to read through the privacy policy, would they be able to understand it? Though a great deal of time is saved by users by choosing not to even read the policies, considerable informational asymmetry exists between users and service providers regarding the collection and processing of personal information online. This asymmetry is compounded

as online services evolve, policies are revised, and users continue to avoid these agreements. These factors result in people often making mistaken assumptions about their meaning: one study found that a majority of Americans who read privacy policies still have an unclear understanding of what it means for the collection and use of their personal information (Bashir, Hayes, Lambert, & Kesan, 2015). In short, individuals who know that a company or organization has a privacy policy may still lack enough information to make informed decision. This makes most users unable to cope with the complexity of privacy management in social media, often resulting in disclosure of sensitive information contrary to their privacy concerns or attitudes.

1.4. Purpose of the Study

As seen from the above work, there is a large body of research from diverse fields such as information science, economics, and social psychology, addressing the privacy decision making phenomenon and attempting explanations for its divergence from information privacy concern. However, there is significantly less work devoted to identifying effective methods to reduce this discrepancy between concerns and behavior. We aim to fulfill this gap in the literature by testing the effectiveness of social interventions reminding users of peer behavior when making privacy decisions. In this work, we describe a 20-day between-subject interventional study testing the effectiveness of different interventions in reducing the gap between privacy concerns and actual information disclosure.

A line of previous research suggests that knowledge of peer behavior can influence behaviors, norms, and preferences in social networks (Christakis & Fowler, 2008; Aral & Walker, 2012; Das, Kramer, Dabbish, & Hong, 2014). The intervention designed in this dissertation is an attempt to reduce the privacy paradox or the concern-behavior gap by reminding users of the actions and concerns of their peers. Our intuition is that showing users these social cues might nudge them to make disclosure decisions in line with their privacy concerns. To the best of our knowledge, this

key idea of explicitly showing and leveraging social cues to reduce the concern-behavior discrepancy is novel.

1.5. Overview of the Dissertation

The organization of the rest of this dissertation is as follows. The introduction chapter has delineated the significance of the privacy paradox and key concepts associated with the discrepancy in information disclosure concerns and behavior. In the chapters to follow, we develop the arguments introduced here and investigate the effects of different interventions on an individual's information disclosure concerns and observed disclosure behavior. The goal of this work is to first identify the difference between information privacy concerns and disclosure behavior and then test the efficiency of social-proof based interventions in reducing this gap. In doing so, we endeavor to contribute to a deeper understanding of the privacy paradox, in particular rationalizations or justifications made by individuals when engaging in behavior contrary to their concerns and the use of interventions as a strategy to help individuals make privacy decisions that result in the alignment of concerns and behavior.

Chapter 2 reviews available theoretical thought and empirical evidence relating to the privacy paradox, the impact of social-proof on individual thoughts and actions, and the use of nudges to communicate privacy risks. In Chapter 3, we present a theoretical framework studying attitudes and behaviors through the lens of the cognitive dissonance theory (CDT). We discuss the different paradigms of this theory and draw connections between the different strategies used to manage dissonance and the coping methods used by individuals to manage threat to their information privacy. We use key concepts of the cognitive dissonance theory to explain the use of interventions as a method to reduce discrepancy between concerns and behavior. Chapter 4 then introduces the role of social-proof based interventions in reducing the gap between concerns and behaviors and presents research questions and hypotheses.

Chapter 5 introduces the methods used in the dissertation to test the research questions and hypotheses presented in the previous chapter. We used a between-subject interventional field study design where we collected quantitative data from surveys and lab sessions and qualitative data from interviews. Chapter 5 details the data collection process and provides discussions about the operationalization of key concepts, such as the measurement of privacy concerns and behavior. It also discusses the design of the different interventions used in this work and outlines the exit interview process.

In Chapter 6, we provide a descriptive analysis of the key variables to present a general picture of the privacy concerns and disclosure behaviors of the sample population. We then present the quantitative and qualitative methods used to answer the research questions and hypotheses. In this section, we focus on teasing out the different effects of the interventions with emphasis on the use of social-proof. This section also delves into the cognitive mechanisms justifying counter-attitudinal disclosure and draws connection between the CDT and disclosure behavior in online settings.

This dissertation concludes in Chapter 7 with a summary and discussion, focusing on: (i) the contradictory relationship between privacy concerns and behavior, (ii) the role of interventions in reducing this discrepancy, (iii) the justifications made by individuals engaging in disclosure contrary to their information privacy concerns, and (iv) an examination of the coping strategies used by individuals when facing privacy threats. Further, the implications of the findings are evaluated in the context of concern – behavior relationships and the role played by the alignment of privacy concerns and disclosure behavior in managing information privacy. In addition, limitations of the study designs and analyses and possible directions for future studies are presented.

Chapter 2: Literature Review

The literature presented here focuses on studies that present theoretical and empirical evidence for the privacy paradox, the use of social influence in guiding behavior, and the use of interventions to help people better manage their information privacy.

2.1. The Privacy Paradox

The discrepancy between privacy concerns and disclosure behavior has been studied in the context of both general internet activities as well as information sharing in SNS and e-commerce transactions. An early study in 2001 explored online shopping popularity and user privacy concerns (Spikerman et al., 2001). Through a series of in-depth interviews with online shoppers, researchers found that while individuals expressed concerns about their privacy being infringed, they were still willing to give their personal details to online retailers as long as they had something to gain in return. Interviewees said they were afraid that too much information about them was collected, but this would not stop them from buying online. Spiekermann et al., (2001) also conducted an experiment to compare self-reported privacy preferences with actual disclosing behavior during online shopping. Participants were first asked to complete a questionnaire on privacy attitudes and preferences and then, to visit an online store. During their shopping in the store they were engaged in a sales dialogue with an online 3D shopping bot. Participants answered most of the questions asked even if these were highly personal. This indicates that even though internet users claim that privacy is a high priority, they do not behave accordingly.

Disclosure has also been studied as a spur-of-the-moment decision based on contextual cues (Knijnenburg & Kobsa, 2013; John, et al., 2010). This research proposes that long-term privacy concerns can be over-ridden by contextual cues, i.e. if an individual believes that the space and context within which they are being asked to disclose information is safe, they will reveal

information irrespective of their long-term or overall privacy concerns. This implies that if an individual is uncertain about disclosing information, a SNS like Facebook, can present certain contextual cues to influence them to disclose. In the social media space, contextual cues can take the form of the websites' privacy policy, the perception of ability to control audiences or information, and the overall perceived trustworthiness of the website.

Privacy Paradox has been measured as a trade-off between risks of disclosure and rewards gained from online information sharing. Acquisti et al., (2015) describe privacy in the age of networking as a "*balance between competing interests – the cost and benefits of sharing or hiding personal information*" (p. 193). This theory of a trade-off implies that privacy decisions follow an economic principle of weighing costs and benefits. It also implies that the privacy paradox seen on SNSs can be explained as a compromise between long-term privacy attitudes or the cost and the advantage gained from building an active social network or benefits.

Norberg et al. (2007) used two experimental studies considering the influence of risk perception and trustworthiness to identify the privacy paradox. Their research attempted to identify the degree to which privacy attitudes or intentions might influence actual disclosure behavior. As opposed to risks, they assume that trust directly influences privacy behavior. Their study found that while risk considerations have an influence on stated preferences, the influence is not strong enough to affect actual behavior. As an environmental factor, trust had stronger effects on actual behavior and outweighed privacy concerns. In contrast, when asked about intentions to provide personal information it is the other way round and risk outweighed trust. Privacy intentions or attitudes and actual data disclosure are paradoxical as risk awareness dominates in abstract decision situations and reliance upon trustworthiness dominates in concrete decision-making processes (Norberg et al., 2007). Contrarily in a study investigating disclosure on social network sites, Dwyer,

Hiltz, and Passerini (2007), found that trust in the website and usage goals did affect people's willingness to disclose personal information. They found that Facebook users expressed greater trust in Facebook than in MySpace and therefore, were more willing to share personal data on Facebook. While trust in the institution represents a key prerequisite for disclosure (Norberg et al., 2007; Dwyer et al., 2007), there is little evidence indicating that users actually trust the SNS (e.g. Facebook, Twitter) where they engage in self-disclosure (Young & Quan-Haase, 2013).

Incomplete information or the lack of transparency termed "information asymmetry" about data sharing and usage practices within website privacy policies has also been studied as a factor influencing disclosure. According to the information asymmetry theory, disclosure of sensitive information can be related to a lack of risk-awareness and missing knowledge about to the potential harms associated with online self-disclosure (Hoofnagle, King, Li, & Turow, 2010; Trepte, Teutsch, Masur, Eicher, Fischer, Hennhöfer, & Lind, 2015). Hargittai and Litt (2013) looked at young users' privacy behavior in the job search context and found that demographic factors such as gender and education were likely to have an effect on an individual's information disclosure behavior. Park (2013) came to similar conclusions and found a positive effect of three literacy dimensions (technical familiarity, surveillance awareness and policy awareness) on privacy protection behavior. Thus, general Internet skills and, more specifically, privacy skills or literacy might be a better predictor of privacy behavior than privacy concerns. Yet, a recent study on privacy literacy (Bartsch & Dienlin, 2016) revealed that most users had low privacy literacy and that the effect of privacy literacy on disclosure was weak.

This incomplete understanding of the risks of disclosure has been compounded by privacy policies. While it seems intuitive that a company's privacy policy would increase transparency and help users understand the norms governing disclosure, research has found that the opposite is true.

More and more users report difficulties understanding privacy policies and confusion about how their data will be used by the institution (Bashir et al., 2015). Privacy policies are typically presented as lengthy textual documents describing details such as data collection, processing, disclosure and management. Websites collecting personal data commonly use these policies to inform individuals about how their personal data will be used by the site. Users are required to read these policies and decide whether they accept the conditions. In practice, however, privacy policies are often hard to read and of very little help to the user (Jensen, Potts, & Jensen, 2005). Therefore, a privacy policy often compounds the information asymmetry problem giving users the mistaken assumption that they have control over their personal information and thereby encouraging the disclosure of sensitive information contrary to privacy concerns.

As seen from the literature presented here, a large body of work has been devoted to understanding the continued existence of the privacy paradox, however, comparatively fewer works have attempted to reduce or minimize this inconsistency. A recent work by Jackson and Wang (2018) proposed a personalized privacy notification interface that attempted to reduce the discrepancy between attitudes and behaviors in the context of mobile-app downloads. In this study, the researchers present participants with just-in-time notifications that show how changes in granting/rejecting permissions impact privacy risk, and how this privacy risk of currently granted permissions compares to users' general privacy attitudes. We aim to further contribute to the understanding and reduction of discrepancies between reported privacy concerns and disclosure behavior by conducting a field study that (i) operationalized individual privacy concerns and disclosure behavior, (ii) tested the effectiveness of different interventions in aligning concerns and behavior and (iii) examined seemingly erratic disclosure behavior through the lens of the cognitive dissonance theory (CDT). To the best of our knowledge, this is the first work that (i) reduces the privacy

paradox as observed in carrying out daily information sharing behaviors and (ii) implements social-proof based interventions to reduce the privacy paradox.

2.2. Social-proof

Researchers from psychology, sociology, and economics study how peers can influence behaviors, norms, and preferences in social networks. For example, early work by Milgram, Bickman, and Berkowitz (1969), showed that simply getting a crowd of people to look up at the sky on a busy sidewalk caused others to do the same. Studies have shown campaigns based on normative messages with social-proof have incited changes in a wide range of habits, preferences, and behaviors. This includes quitting smoking (Christakis & Fowler, 2008), adoption of online entertainment products (Aral & Walker, 2012), and sensitivity to Facebook's security features (Das, et al., 2014). More recent studies on online platforms such as Facebook have similarly alluded to the power of social-proof. Kramer (2012) found that users were more likely to share emotional content that matched the emotional valence of content shared by friends in the past few days. Research work studying voter mobilization on Facebook found that simply showing people that their Facebook friends voted was sufficient to increase voter turnout in the 2010 U.S. Congressional elections (Bond et al., 2012).

Within the context of understanding the relationship between people's concerns and behavior, social influence is represented by the concept of subjective norm, which describes the amount of pressure that people perceive they are under from a majority peer group to perform or not to perform a particular behavior. In recent years, several researchers have begun to re-examine the role of social influence and normative factors in the concern-behavior relationship. Social injunctive norms reflect perceptions of actions that a majority approves of or thinks is "right". These norms motivate action by highlighting the potential social rewards and punishments for engage-

ment or non-engagement in an approved behavior. In contrast, descriptive norms reflect the perception of whether other people are actually performing the behavior in question. Descriptive norms describe what is typical or normal and motivate action by providing evidence as to what is likely to be effective, adaptive, and appropriate action. Descriptive norms are defined by how likely others are to (dis)approve of a particular behavior (Azjen, 1991). Research has demonstrated that both descriptive and injunctive norms exert an influence over individual behavior and have been used as a prediction model of various behaviors related to trustworthiness and privacy in Internet purchasing behavior (George, 2004) and SNS usage characteristics (Mendel & Toch 2017).

This tendency to look to a peer group in order to ascertain how to act has been termed as social-proof (Cialdini, 2001; Cialdini & Trost, 1998) shown to be effective in driving human behavior. Existing literature shows the presence of a social component that influences peoples' perceptions about the adoption and use of tools to maintain privacy and security. Rader, Wash and Brooks (2012) find that people tend to learn about privacy enhancing practices from informal stories told by others in their network. Research has also found that demonstrating privacy-enhanced behaviors are often driven by social processes. For instance, observing the adoption of a new security feature within an individuals' online social network was a key component in them adopting the same features (Das et al., 2014).

Significant research has also been devoted to understanding how social network users exchange and share information about privacy. Lewis, Kraufman, and Christakis (2008), show that having a roommate with a private profile is a strong predictor for having a private profile. Patil, Page, and Kobsa (2011) studied the effect of reminding an individual of the privacy settings used by their social circle within the context of an instant messaging application. In a similar study, Besmer et al., (2010) tested the effect of informing users about third-party access permissions

granted by a percentage of other users and found that this information impacted user decisions to grant access permissions. Das et al., (2014) studied the impact of social-proof notifications on Facebook security settings. They found that if a majority of friends adopted a particular security feature, users were encouraged to adopt the same feature. Conversely, information about a minority adopting a feature may bias users away from adoption.

Researchers have shown that reliance on social influence increases as the uncertainty in individuals' judgments and decisions rises (Spottswood & Hancock, 2017). As previous research shows individuals often experience confusion over data collection and sharing practices in a phenomenon termed as "information asymmetry" (Acquisti et al., 2015). This implies that people rarely have a clear knowledge of what information other people, third-party sites, corporates, or governments have about them or the policies governing how and with whom that information is used and with what consequences. As this uncertainty persists individuals are likely to look for social cues to guide their information sharing behaviors. We, therefore, used interventions that notified individuals of an "ideal" behavior demonstrated by a majority of the population. We expect receiving social-proof based interventions will significantly impact information sharing behaviors and result in privacy behaviors aligned with concerns. While we recognize the effectiveness of social-proof as a mechanism to guide behavior, it is also important to keep in mind the ethical considerations associated with using social influence. The nudges were hence designed to simply present social cues to the individual but leave the ultimate choice of acting on that information up to the user.

2.3. Privacy Nudges

Communicating privacy risks to individuals via nudges has been an active line of research. Researchers have proposed soft paternalistic interventions that nudge individuals toward certain behaviors (Thaler & Sunstein, 2009). In their work studying user regrets over content shared on

social media, Acquisti (2009) discusses the uses of nudges to influence privacy decision-making and decrease users' regret. Interventions have also been studied in the fields of human – computer interaction (HCI) and persuasive technology as a mechanism to assist users with privacy and security decision making. Forget et al. (2008) designed a system that used persuasive techniques to nudge users to create stronger passwords. In a related study, Ur et al. (2012) used different password meters as an intervention to encourage individuals to create stronger passwords. In the context of SNS, a longitudinal Facebook study highlights how changes in Facebook's interface can impact users' information disclosure behaviors (Stutzman, Gross, & Acquisti, 2012). A similar study investigating the effectiveness of modifications to Facebook's interface tested the efficiency of three different nudging mechanisms (audience, timer, and sentiment) to help users manage privacy when posting on Facebook (Wang et al., 2014).

In the context of increasing privacy awareness in mobile applications, Almuhimedi et al. (2015) designed an intervention that reminds users of the permission settings of various apps installed on their smartphones. In this study, participants were given alerts consisting of messages that described the number of apps accessing location information and frequency of access in a given period. The results suggest that most participants re-evaluated or changed permission settings after receiving nudges showing them how often some of their sensitive data was being accessed by apps.

The use of interventions to help individuals better manage location privacy was also tested in a recent study by the author (Ghosh & Singh, *Under Preparation*). In this exploratory work, the authors investigate the effect of "audience-group" based interventions on Facebook check-in behavior of participants. These "audience-group" based nudges help close the gap between the users' perception of expected audiences and those that actually have access to their data. The nudges

remind participants that their real-time location information may be visible to a larger group of friends than they expect. This work was designed as a 6-week between-subject interventional study where participants were randomly assigned to one of three experimental conditions and location disclosure was measured before and after the interventions. The interventions were designed based on the idea that individuals often have an "expected" audience when they are sharing location information on Facebook. This expected or ideal audience is often much smaller than the actual audience comprising of *all* their Facebook friends. Ghosh and Singh (*Under Preparation*) found that nudges reminding individuals of the actual (and perhaps unexpected) audience to their disclosure was an effective method of helping them manage their location privacy. Based on statistical tests and qualitative analysis, the authors gained several insights into location disclosure on SNS. Further, they identified recommendations for app designers and privacy researchers to better design and evaluate location sharing in online social networks. Based on this prior work, we conducted a field study using interventions to test the efficiency of social – proof nudges in reducing the privacy paradox.

Chapter 3: Theoretical Framework

3.1. The Cognitive Dissonance Theory

Festinger's (1957) theory of cognitive dissonance has been the subject of diverse research studies since its formulation in the mid-1950s. Festinger (1957) theorized that, when an individual holds two or more beliefs that are relevant but inconsistent with one another, a state of discomfort or dissonance is created. In this state, individuals are motivated to re-evaluate either their beliefs or their behavior until a state of mental consonance is reached. According to this theory the relationship between a person's attitudes and behaviors is driven by the need to reduce this state of dissonance (Festinger, 1957). In order to resolve dissonance, individuals could add consonant cognitions (i.e. information or beliefs supporting their behavior), subtract dissonant cognitions, increase the importance of consonant cognitions, or decrease the importance of dissonant cognitions. Scholarly research has frequently studied a change in attitudes as a way of reducing dissonance (Elliot & Devine, 1994; Harmon-Jones, 2000). For instance, a smoker who knew or learned about negative health effects, either stopped smoking, changed his attitudes towards smoking, or added explanations about smoking to reduce dissonance, e.g. car driving is more dangerous than smoking.

Similarly, an individual disclosing information during an online interaction can be supposed to experience two conflicting situations. On the one hand, they appreciate the risks undertaken by disclosing personal information, however, when confronted with the benefits gained by the online interaction (e.g. convenience, network building, or financial discounts) they might be tempted to change the level of privacy concern and generously disclose sensitive information about their preferences, habits, and social networks. Theoretically, the cognitive dissonance theory offers a new explanation of why people disclose data contrary to their information privacy concerns by arguing that stated privacy concerns can be shifted to suit an individuals' disclosure preferences. We now explain the various tenets of the cognitive dissonance theory.

3.1.1. Free Choice Paradigm

The free-choice paradigm (Brehm, 1956; Festinger, 1957), examined dissonance in the post decision making phase. After an individual has chosen between two alternatives (e.g. disclosing or protecting sensitive information), they might encounter dissonance if they have behaved counter-attitudinally. In such a scenario, dissonance can be reduced by viewing the chosen alternative as more attractive and/or viewing the rejected alternative as less attractive. Such dissonance reduction serves to further separate the choice alternatives in terms of their desirability. Brehm (1956) conducted an experiment in which participants were asked to rate eight different small household appliances (Brehm, 1956). The participants were then given either a difficult choice (i.e., between two alternatives that had both been rated high) or an easy choice (i.e., between one alternative that had been rated high and another alternative that had been rated low). Participants were asked to evaluate the decision options before and after the decision. Brehm (1956) found that, after persons made a difficult decision i.e. where both alternatives were valuable, they changed their attitudes to become more negative toward the rejected alternative. After an easy decision i.e. where one alternative was clearly more valued than the other, participants did not change their attitudes.

During our study, participants will be faced with several opportunities to protect or disclose information with costs and benefits associated with each choice. We expect that when individuals choose to perform certain tasks instead of others, there may be post-hoc change in valuation of the chosen alternative and the free-choice paradigm will help us gain a theoretical understanding of how these decisions are made.

3.1.2. Induced Compliance Paradigm

The induced compliance paradigm placed people in a situation where they were persuaded to behave in a manner contrary to their personal beliefs by offering an incentive to do so. The research shows that large incentives provided external justification for the desired action and therefore,

individuals did not experience dissonance in demonstrating a behavior that contradicts their beliefs (Harmon-Jones & Mills, 1999). In order to reduce the dissonance between desired action and internal beliefs, people shifted their beliefs (often based on external justifications) in the direction of the desired action. To test this prediction, Harmon et al. (1999) brought participants into the laboratory and asked them to perform a boring task. Then, participants were paid either $1 or $20 to tell "another participant" that the task was interesting. According to dissonance theory, lying for a payment of $20 should not arouse much dissonance, because $20 provides sufficient justification for the counter-attitudinal behavior (i.e., it adds 20 cognitions consonant with the behavior). However, being paid $1 for performing the same behavior should arouse much dissonance, because $1 was just enough justification for the behavior (i.e., it adds only one consonant cognition). As expected, participants in the $1 (low-justification) condition expressed a belief that the task was actually interesting, whereas participants in the $20 (high-justification) condition did not change their beliefs.

The induced compliance paradigm points to the phenomenon of individuals' using a perception of high benefits as an external justification for exhibiting behavior contrary to their attitudes or concerns. When attempting to understand privacy decision making, the induced compliance paradigm might provide a useful theoretical lens to examine the influence of contextual cues and biases on the disclosure of sensitive information.

3.1.3. Social Conformity

Festinger, Riecken, and Schachter (1956) studied knowledge of social behaviors as a method for reducing dissonance. An initial study by Festinger et al., (1956) addressed the support provided by members of a doomsday group in reducing each other's dissonance when their group's predictions of the apocalypse failed. Since then other studies investigating the role of social information (i.e.

knowledge of peer behavior) on dissonance reduction have found that knowing that others have behaved in the same manner (i.e., counter-attitudinally) acts as a consonant cognition, thereby reducing dissonance (Strobe & Diehl, 1981). According to Stone and Cooper's (2001) model, dissonance occurred when people evaluated their attitudes or actions and found it different from some accepted standard. This standard could be created from personal beliefs of "good" or "acceptable" behavior, or from social factors such as the normative rules and beliefs held by most people in a culture. Research on dissonance identifies group-level activities such as gaining social consensus for a particular belief or attributing dissonance-producing acts/beliefs to members of out-groups to resolve dissonance (Festinger et al., 1956).

Classic perspectives on social influence coincide with cognitive dissonance research where counter-attitudinal behavior generates dissonance for social and informational reasons. Within social influence literature, normative influence is conceived as the pressure that people perceive they are under from a majority peer group to perform or not to perform a particular behavior. Engaging in acceptable behavior helps achieve a favorable conception of self as well as establishing positive relations with others (Cialdini & Trost, 1998). In contrast, informational influence pressures originate in people's desire to have a valid understanding of reality and thereby to effectively negotiate their world. However, individuals often depend on their social peers to meet these informational needs. Being part of a social consensus often provides a sense of reality because similar others have the power to define reality. Research on social influence has demonstrated that disagreement from peers can threaten self-esteem and social acceptance crating feelings of isolation and uncertainty in one's beliefs (Christakis & Fowler, 2008; Cialdini & Trost, 1998).

Therefore, dissonance and social influence theories both suggest that disagreement from others in a group produces negative states of dissonance in an individual. Regardless of whether

dissonance is produced directly from others' disagreement or is an indirect product of the normative and informational challenges posed by disagreement, there is a good theoretical reason to believe that dissonance arises from interpersonal inconsistencies in judgments and can be resolved by engaging in peer-supported activities. In the context of the current study, we design interventions based on the cognitive dissonance and social influence theories to reduce the concern – behavior gap. The interventions are designed to create an awareness of dissonance and provide information about peer-supported behavior. We expect that participants will attempt to reduce dissonance by re-evaluating attitudes and behaviors so they again feel a sense of belonging within the social group.

Chapter 4: Research Questions and Hypothesis

Privacy scholars have shown that many social network users are afraid their privacy might be violated online (Acquisti et al., 2015; Stutzman et al., 2012), although few users implement any steps necessary to safeguard sensitive data (Taddicken, 2014; Norberg et al., 2007). While this gap between privacy concerns and behaviors has been well-studied in research literature (Barnes, 2006; Beresford et al., 2012; Brandimarte et al., 2013), identifying effective ways of reducing this misalignment is still under-studied. We designed an interventional field study using social cues to help users reduce the gap between their privacy concerns and disclosure behaviors. The theories of cognitive dissonance and social-proof were used to test the intuition that when individuals engage in counter-attitudinal behavior, information about "good" or "accepted" peer attitudes and behavior, will result in an increased alignment between concerns and behavior. Previous research in social psychology has also pointed to the use of social cues as an effective way of influencing human behavior. Participants in a "music market" experiment were much more likely to download a song if they believed the song was popular among other participants (Salganik, et al., 2006). Hotel guests were similarly motivated to reduce their use of towels by showing them that previous occupants chose to be less wasteful (Goldstein et al., 2008). Social-proof has even been used to eliminate young children's phobia of dogs by showing them film clips of other children playing with dogs (Bandura et al., 1967). In a recent study on the use of social-proof to increase security awareness among Facebook users, Das et al., (2017) found that receiving a notification of their friend's use of security features was a key enabler for security related behavior change among their participants.

Alerting users to privacy risks via notifications has also been an active line of research. Soft paternalism (or nudging) interventions do not aim to restrict choice but attempt to account for

bounded rationality in decision making. (Thaler & Sunstein, 2009). In the context of mobile app installation, a recent study used interventions to reduce the discrepancy between privacy attitudes and behaviors (Jackson & Wang, 2018). In this study, the researchers presented participants with an interface that simulated the app installation screen in Android phones. They then gave participants just-in-time notifications that showed how changes in granting/rejecting permissions influenced privacy risk, and how the privacy risk undertaken by currently granting permissions compared to participants general privacy concerns. We follow their lead in utilizing nudges for privacy management, but our study is different in some important ways. Jackson & Wang's, (2018) work focused on privacy paradox in the context of permissions granted during mobile app installations. Despite the fact that information disclosure via mobile apps is an important area of investigation, our research focuses on the broader question of information sharing behaviors demonstrated in daily life activities. We do not restrict the medium through which information disclosure occurs rather we concentrate on commonly performed tasks in daily life through which sensitive information can often (inadvertently) be disclosed. Further, our focus is on "*social*" interventions, which have never been tested as a method for reducing the privacy paradox.

Interventions have also been used as a mechanism to increase privacy sensitive behavior in individuals (Acquisti et al., 2017). A recent work by the author (Ghosh & Singh, *In Preparation*) used audience-group based interventions to help people better manage their location privacy on Facebook. The audience-group based nudges were designed to remind participants that their real-time location information may be visible to a larger group of friends than they expect. Ghosh and Singh (*In Preparation*) found that reminding participants of different audiences within their Facebook friends group nudged them towards privacy enhanced location sharing. In a previous study, Almuhimedi, et al., (2015) designed a system to inform users about location data collected by apps

installed on their devices. After downloading an app, users were periodically nudged to review and adjust (i.e., restrict or permit) permission settings. Similarly, another study used Privacy Facts, a "just-in-time" privacy display warning users when sensitive information such as location is being requested (Kelley et al., 2013).

Based on the outlined research, this work used interventions based on social – proof to nudge users to demonstrate behavior consistent with their privacy concerns. We also tested the effectiveness of the social – proof nudge by comparing it to a nudge that reminded people of the inconsistency between their privacy concerns and behaviors and a baseline nudge that reminded people of their information disclosure. We therefore used the following research question and hypotheses in this study:

RQ1: Can social-proof based nudges be an effective way to reduce the privacy paradox?

H1a: The reduction in discrepancy between concerns and behaviors will be higher in the social-proof intervention group when compared to the other groups.

H1b: The reduction in discrepancy between concerns and behaviors will be higher in the information inconsistency intervention group when compared to the baseline group.

H2a: The number of people who had a reduction in discrepancy between concerns and behaviors will be higher in the social condition compared to other groups.

H2b: The number of people who had a reduction in discrepancy between concerns and behaviors will be higher in the inconsistency condition compared to the baseline group.

As individuals navigate online websites and apps in pursuance of their daily activities, they are faced with numerous information disclosure decisions. These decisions range from understanding and configuring privacy settings on SNS, to deciding whether to download an app based on the permission it requires, to deciding whether a website offering discounts can be trusted. All the

activities mentioned require individuals to disclose some personal information, which can often be contrary to their privacy concerns. This decision-making process is further complicated by cost-benefit trade-offs individuals make to justify disclosure decisions. Research shows that individuals often view the effects of disclosure as uncertain and distant while the benefits are viewed as immediate and concrete (Hallam & Zenella, 2017). A previous study investigating the use of location sharing apps on Smartphones found evidence that individuals often make rationalizations (e.g. the ubiquity of information tracking apps) to justify their own location sharing decisions (Ghosh & Singh, 2017).

Different people attribute dissimilar values to different types of information and perceptions of cost and benefits associated with information disclosure vary from person to person. Research has also shown the existence of cognitive biases and heuristics that govern privacy decision making (Acquisti et al., 2015). In the context of using interventions in privacy decision making, scholars have investigated the efficiency of social and non-social interventions on privacy behaviors (Acquisti et al., 2017; Wang et al., 2014). In this dissertation, we focused on the use of different interventions in reducing the gap between information privacy concerns and disclosure behaviors. These interventions were based on social-proof, information inconsistency and a baseline. RQ1 tests the effectiveness of the interventions and whether the social-proof intervention is the most effective way of reducing the privacy paradox. However, it is also important to understand *underlying mechanisms* as to why these interventions work (or do not). We wanted to further investigate the differences in how participants reacted to each intervention and the different cognitive factors that led to changes (if any) in privacy concerns and behaviors. We therefore used a second RQ designed to understand the mechanisms of each intervention as follows:

RQ2: How do the interventions change counter-attitudinal disclosure across the three experiment groups?

We used interventions designed to guide users towards choosing behaviors aligned with their privacy attitude. Specifically, we aimed to test the effect of social-proof based interventions in reducing the gap between privacy attitudes and behaviors. In order to accomplish this goal, we conducted a 20-day between – subject study with three experiment conditions: (i) social-proof based interventions: based on the idea that individuals will mimic actions performed by a majority of the population, (ii) privacy inconsistency interventions: based on highlighting the inconsistency between the user's privacy attitude and disclosure behaviors, and (iii) baseline interventions: based on reminding users of their information sharing behaviors. This work used a mixed-method study design using quantitative data gathered from surveys and tasks and qualitative data gathered from follow-up interviews.

Chapter 5: Study Methodology

The process of conducting a research study requires a careful understanding of what we are attempting to measure and how. Designing a research study requires making decisions about the type of data to be collected, identifying the right instruments and measures, and what research techniques need to be employed, and how this data will address the research questions (Neuman, 2006). In this section, we focus on the various methodological approaches that could be useful and insightful for this study. We also compared these methods and in doing so, identified the relevant approach for data collection and analysis. We then present a detailed description of the methods used as well as the instruments used in the measurement of different types of data.

5.1. Methods in Previous Work

The work proposed here attempts to measure gap between privacy attitudes and behavior and investigate the efficacy of social-proof based nudges in reducing this gap. Several researchers have attempted to provide explanations for this misalignment; however, identifying strategies that can reduce this gap is an under-studied area. To identify proper methods for collecting and analyzing data, we review research methods used in some key research works investigating attitudes and behaviors as well as important research investigating behavioral change.

5.1.1. Survey Data

A survey is probably the most common data collection method used in social science research. It is a flexible approach which can be used to investigate a wide range of topics. Surveys often employ set questionnaires as a tool for data collection. This method has the advantage of being relatively inexpensive and easy to deploy while being able to gather data from larger populations. It is therefore useful in collecting data that can be generalized. It can also be useful when collecting demographic information and gathering participant opinions on multiple topics. It is also easier to replicate and compare and can be used to identify some initial patterns that could benefit from

further investigation. However, a major drawback of this method is that it depends wholly on what people report they do or think which might be different from what they actually do or think. If a particular word or phrase is misinterpreted by participants, it does not allow for further explanations or probes. Long and complicated surveys can often become tiring to complete and lead to errors (McNeely, 2012). It is therefore important to compare the advantages and drawbacks of surveys as pertaining to the research study and perhaps consider using it as a complement to other research methods.

When measuring information privacy concerns previous studies have used large-scale surveys to measure concerns over information exchange in online and offline settings. Early work measuring information privacy concerns involved surveying a large number of participants using questionnaires. Seminal work by Acquisti and Gross (2006) surveyed 506 students, faculty, and staff affiliated with a particular institution and asked to answer questions about their usage, knowledge, and concerns over information posted on social networking sites. Other work attempting to understand motivations for information disclosure in social media sites surveyed 704 undergraduate students to understand their disclosure behavior on two social network sites (Tufecki, 2008). This survey combined questions on concerns over information privacy as well as questions on practices enacted by students to protect their information (e.g. using their real name, changing privacy settings, etc.).

In attempting to understand Facebook users' awareness of privacy issues and perceived benefits and risks of disclosure, Debatin et al., (2009) surveyed 119 undergraduate students from a mid-western university. The online questionnaire consisted of 36 multiple-choice questions. Participants were asked for basic information regarding Facebook habits (e.g. duration of account), including the amount, personal information in their profile (e.g. descriptors, contact information),

as well as yes or no questions asking about any unpleasant experiences on Facebook (e.g. being trolled or harassed). In order to gain contextual information about privacy invasion experiences, a subset of eight participants from the online pool were selected for in-person interviews. As seen from this review, surveys have been used to gather general information about attitudes or behaviors from large samples. However, in cases where more precise information was required the survey was used in conjunction with other research methods (e.g. interviews).

5.1.2. Semi-structured Interviews

There are three fundamental types of research interviews: structured, semi-structured and unstructured. Structured interviews are, essentially, verbally administered questionnaires, in which a list of predetermined questions are asked, with little or no variation (similar to a questionnaire). Unstructured interviews on the other hand do not reflect any preconceived theories or ideas and are performed with little or no organization. In this format, participants are asked open-ended questions about their experiences and are encouraged to simply "talk" to the interviewer. Semi-structured interviews stand in between these two formats and is the most commonly used interview technique. This method consists of several key questions that help to define the areas to be explored, but also allows the interviewer or interviewee to diverge in order to pursue an idea or response in more detail. The semi-structured interview allows researchers to explore the views of homogenous as well as diverse groups of people in detail and help unpack these differing perspectives within a community. However, such a research design is often time and labor intensive and often subject to research biases. This technique is ideally suited for exploratory research within a small group where the main goal of the research is to gain a subjective understanding of social interactions.

Research investigating privacy concerns and behaviors, especially in cross-cultural contexts, have used semi-structured interviews to gain a rich, layered, and in-depth understanding of the contextual nature of privacy and security. For instance, a study by Ghosh and Singh (2017)

investigating the cognitive dissonance theory as an explanation for location-sharing information via phone-based apps interviewed fourteen participants who had already been part of a previous study gathering logged data. The added interviews allowed researchers to ask in-depth questions about motivations for location sharing and the seemingly paradoxical behavior observed via phone logs. In this work, researchers asked participants questions about the number of location-sharing apps installed, knowledge of tracking, as well as concerns over location privacy. The questions were framed as open-ended and allowed participants to talk in-depth about their own or friends' experiences and the meanings and significance associated with location data.

Other work investigating regrets experienced due to social media posts also used semi-structured interviews in connection with surveys to understand why such regrets are experienced as well as ways of helping social media users avoid these regrets. Interviews were conducted with nineteen college students where researchers spent 1 to 1.5 hours understating their experiences with Facebook. Questions ranged from frequency of Facebook activity to negative incidents on Facebook. The interviewers used a template of three basic questions that allowed them to ask more probing or detailed questions as required and encourage participants to share details about their feelings, perceptions, and thoughts along with actions.

Recent research has discussed the importance of gaining an understanding of privacy in non-western contexts. In an effort to incorporate other contexts, Nissenbaum (2004) argues that notions of privacy change with place, people, culture, and context. Semi-structured interviews have been frequently used to gather data for this exploratory research area. Notably, researchers investigating privacy and surveillance in Bangladesh conducted semi-structured interviews with thirty families from different socioeconomic classes (Ahmed Haque, Guha, Rifat, & Dell, 2017). Each of the home visits lasted approximately an hour and discussed the participants' backgrounds,

mobile phone use, and their experiences (if any) with government surveillance. These in-depth interviews allowed researchers to understand the meanings and significance attributed to information, ownership, and identity. As seen in these research works, semi-structured interviews are often used as a follow up to survey or logged data or during exit interviews in order to gain a deeper understanding of the data.

5.1.3. Logged Data

Logged data refers to data gathered from one or more devices or sensors that can help researchers observe human behavior. These sensors are the everyday smart devices (e.g. phones) that capture and record a range of activities and have recently been used study and predict human behavior (de Montjoye, Quoidbach, Robic, & Pentland, 2013). As more and more of our daily activities are mediated by smart devices, the metadata gathered by these devices present huge opportunities for researchers. Sensors allow for large-scale, longitudinal, and unobtrusive data collection eliminating many of the biases seen in survey based studies. Logs of data gathered by devices collect natural observations of people as they perform daily activities, uninfluenced by experimenters or observers. However, log studies are best suited when the research aims to gain an abstract, high-level picture of user behavior rather an understanding of individual intentions or goals, or the contexts in which these behaviors occur.

Research studying changes in disclosure behavior over time have used longitudinal social network logs where OSN profile creations and updates were studied for a large sample of participants (>5000) from 2005-2011. This logged data allowed researchers to gain a high-level understanding of understand information disclosure and Facebook usage over time (Stutzman et al., 2012). Another study aimed at inferring individual privacy attitudes used phone meta-data gathered over a period of 10-weeks (Ghosh & Singh, 2018). This work allowed researchers to gain important insights into individual privacy attitudes without having to rely on self-reported data.

These works point to the advantages of using logged data in order to gain a high-level understanding of a particular behavior. However, relying solely on logs might not be the most effective way of understanding human behavior. Other data collection techniques (surveys or interviews) described in this section can be used in addition to logged data to help confirm and provide insights into what is learned from log data.

5.1.4. Field Studies

When interested in understanding how the manipulation of a variable can explain specific outcomes on another variable, some researchers find it useful to conduct experiments. A field study typically comprises of 1) taking an action and 2) observing the consequences of that action. As with all research methods, field studies are more appropriate for some topics and research purposes than others. A classic field study design typically involves comparing two groups, one called the experimental group, the other the control group, to both of which respondents have been randomly assigned (Neuman, 2006). In the experimental group, the researcher conducts some treatment (or interventions) on the subjects and measures its effects either to a group that does not receive the treatment or to a group receiving a different kind of treatment. Field studies are particularly effective when attempting to capture the relationship between two variables (Neuman, 2006).

The use of field studies using interventions has been used in multiple works to understand privacy behavior. Early work examining privacy behaviors and disclosure decisions have argued for the need to understand the effects and influencers behind disclosure decisions (Acquisti, 2009). As such, researchers from diverse disciplines like behavioral economics, social psychology, etc. have used field studies to identify complementary and rich tools to understanding and guide privacy decision-making. In a field-study testing the efficiency of different interventions, Wang et al., proposed different interventional mechanisms (audience, timer, and sentiment) that could be integrated into Facebook. In this 6-week within-subject study, researchers observed information

disclosure behavior during a 3-week control period and compared it to a 3-week treatment period during which participants were exposed to different types of interventions. Another work by Ghosh and Singh (*In Preparation*) used a between-subject field study design where participants were randomly divided into three groups and disclosure behavior across these groups was compared to understand the efficacy of different interventions. This work used a before-after design so that disclosure behavior in the period before interventions can be compared to disclosure behavior after interventions. Another work focusing on location disclosure, also opted for a within-subject study design where the location disclosure behavior of 23 participants was compared with and without interventions. As seen from this brief outline, field studies are often effective when trying to understand subjective and context-dependent behaviors like information disclosure. We, therefore, chose a between – subject interventional field study design for this dissertation.

The objective of this dissertation is to reduce the gap between privacy concerns and behaviors through the use of social interventions. We do not intend to make privacy decisions for users, but rather to develop a complementary approach to support users so that they can make disclosure decisions that better align with their information privacy concerns. To achieve this goal we conducted a 20 day between – subject interventional study consisting of an online entry session, in person lab session, three consecutive interventions delivered over a 3-day period, a second online session, a second lab session, and an exit interview. A detailed timeline of the study is presented in the Figure 5.1.

We used a mixed method approach combining quantitative data gathered before and after the interventions with semi – structured interview data. Mixed method is defined as *"a method [which] focuses on collecting, analyzing and mixing both quantitative and qualitative data in a single study. Its central premise is that the use of [both] approaches in combination provides a*

better understanding of research problems than either approach alone" (Creswell & Clark, 2011, p. 5). In this dissertation, we aim to test the effect of interventions on reducing the gap between privacy concerns and disclosure. In order to better understand the results obtained from the quantitative analysis we included an interview stage that allowed us to gather rich, contextual data on the diverse factors affecting individual privacy concerns and behavior. A mixed method analysis allowed us to identify a set of themes from the quantitative data analysis that were used as an analytical framework to analyze the qualitative data.

Figure 5.1. Timeline of study.

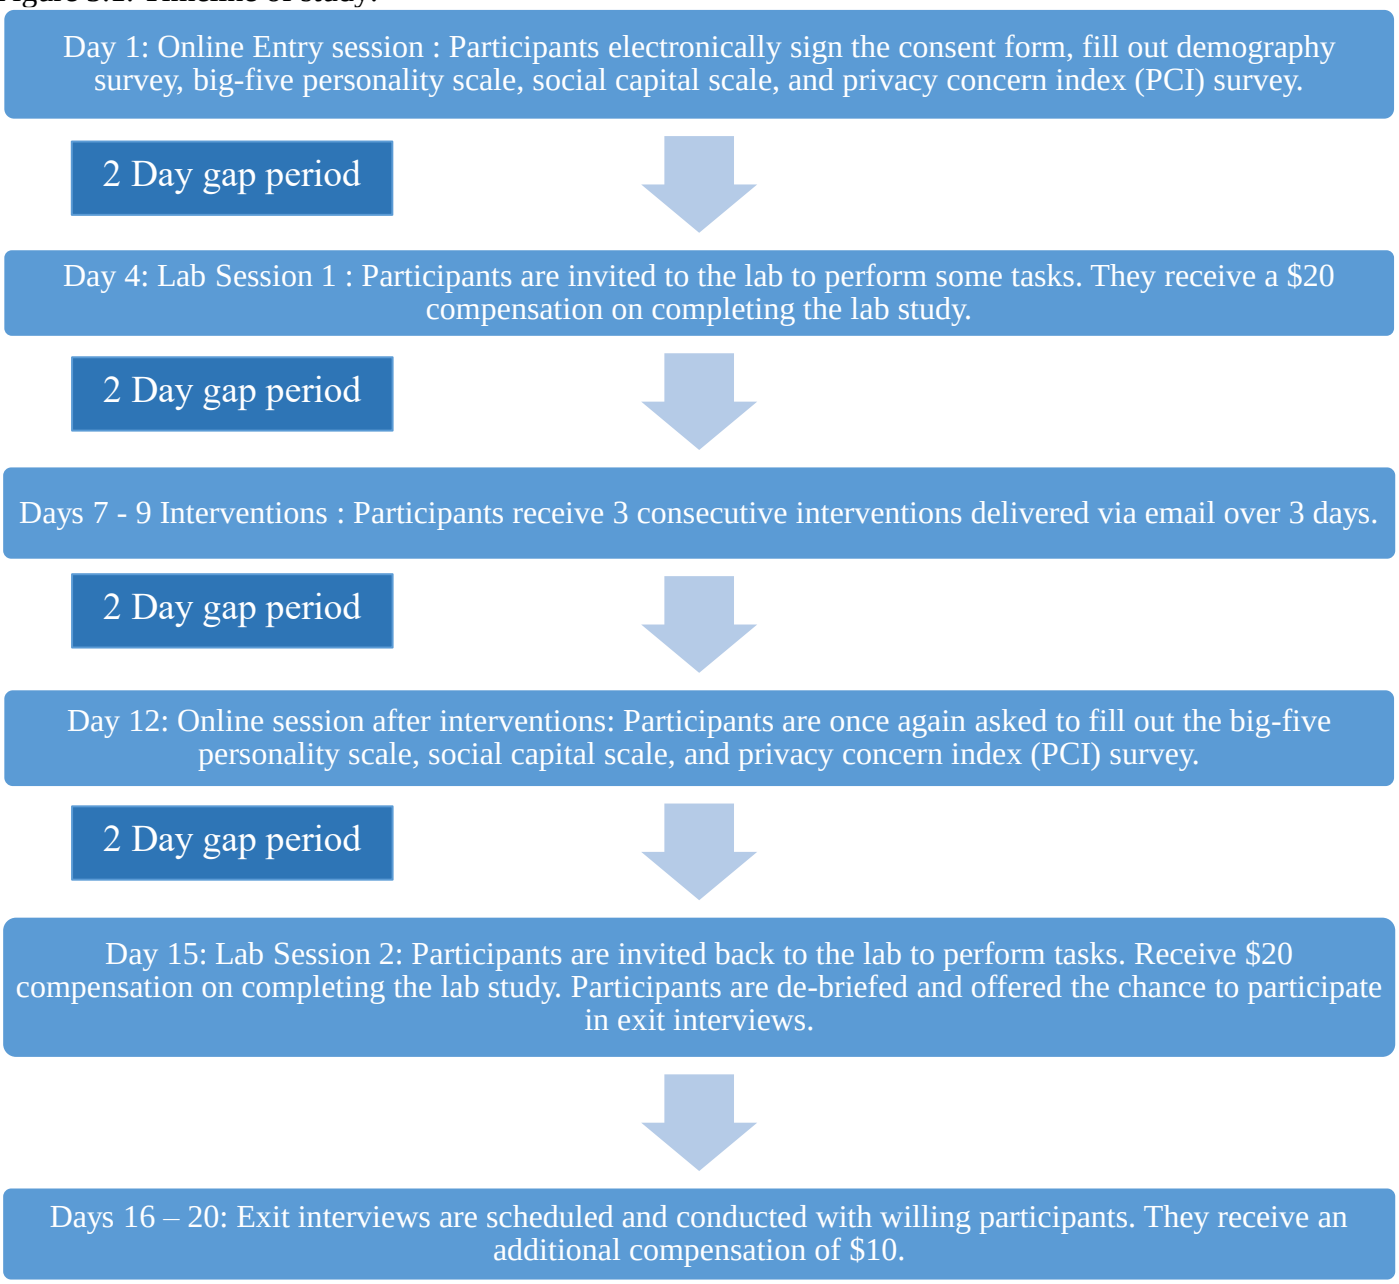

We now provide a detailed description of the research design including recruitment process, concern and behavior measures, interventions, and exit interviews.

5.2. Recruitment

Given the goal of our work to test the effectiveness of interventions in reducing the gap between privacy concerns and behaviors, therefore, we tried to use various recruitment methods (Section

5.2.2) to reach participants who meet the research criteria. In this study, we focused on different forms of information disclosure in online settings and the motivations guiding disclosure decisions. We therefore collected self-reported data on privacy concerns as well as captured actual disclosure behaviors via activities performed in the lab session. The information disclosure activities were designed as an online web form where participants could upload plain text, images, or videos. The privacy concern scale and disclosure activities were both well-suited to measuring disclosure in online environments. In order to participate in the study, individuals needed to: (i) have a smartphone with file sharing and Internet capabilities (ii) be active social media users (iii) have a bank account and credit card (iv) be conversant in written and spoken English and (v) be able to travel to the test site at New Brunswick at least two times.

We measured privacy concerns and information disclosure as they occur in everyday transactions. Hence, participants did not require additional expertise in a specific area to be a part of the study. We focused on an adult population (over 18 years) that can travel to the lab for at least two lab sessions. Some of the activities listed in Appendices E and F require the disclosure of financial or medical information (e.g. credit card details, medical insurance, etc.). Due to age restrictions in creating bank accounts, an adult population was considered most appropriate for this study. We attempted to simulate information disclosure that occurs while performing various online activities, for example, submitting health information via an online web form. As these activities are usually performed by adults, we thought it valuable to sample from an over 18 population. We wanted to recruit a diverse demographic group in order to gain a rich and in-depth understanding of privacy concerns and behaviors and perhaps yield additional insights about the effect of social and non-social interventions on different population groups.

5.2.1. Recruitment Strategy

Among nonprobability sampling techniques, we used convenience sampling. Convenience sampling is where the sample is selected from members of the target population who meet certain practical criteria, such as easy accessibility, geographical proximity, availability at a given time, or the willingness to participate are included for the purpose of the study (Creswell & Clark, 2011). The main characteristics of the population for this study is knowledge and frequent use of online services and the ability to commute to the study site. We used convenience sampling methods to recruit individuals who meet the criteria described in the previous section.

We referred to similar interventional field studies to identify the proper number of participants for this study. A brief review is presented below:

- Almuhimedi et al., (2015) recruited 23 participants via Craigslist and a city-wide participant pool for a within-subject Interventional Study measuring the effect of notifications on location disclosure.

- Wang et al., (2014) recruited 21 participants from Pittsburgh and Syracuse using Craigslist, flyers, email distribution lists, and social media posts. Participants came from a variety of occupations and age groups (19 to 51) including medical staff, students, managers, retired, and unemployed. They used a within subject research design where the same participants were observed during a control period and a treatment period.

- Norberg et al., (2007) recruited 23 part-time, evening program graduate students at a university in the Northeast United States to examine longitudinal disclosure behavior. Participants were mostly college students recruited via the University. The age range of the sample was 22–40, and the gender split was even.

The studies described here use a within subject study design where the participants are not divided into separate groups. The entire sample population (~20) is observed during a control period and subsequent treatments period. We also conducted an *apriori* power analysis to calculate the required sample size for observing a large effect size between groups with $\alpha = .05$, power standard to .80 (O'Keefe, 2007), the minimum required sample size is 36 participants. However, prior research has also shown that there is often a drop-off rate in participants during field studies, therefore we purposely over sampled to recruit 54 participants in the study. Due to some participants quitting the study halfway due to not completing online components or lab sessions we obtained results from 42 participants out of whom 20 agreed to a follow-up interview.

5.2.2. Recruitment Procedure

Participants were recruited by spreading information about the study through word-of-mouth communication, as well as from online ads and from personal and social networks. Recruitment methods included posting on social media sites, recruitment ads on Craigslist, Facebook, and Google, emails, and flyers distributed in the New Brunswick area. We also contacted directors of undergraduate and graduate programs as well as instructors of several classes asking them to inform students about the research study, students thus recruited were offered a small amount of extra credit along with monetary compensation for research participation. Participants who successfully completed the study were compensated with a sum of $40 ($20 at the end of lab session 1 and $20 at the end of lab session 2) with the option to earn an additional $10 by participating in an exit interview.

The key idea of this study was to first identify the gap between privacy concerns and behaviors and then use social-proof based interventions to reduce this inconsistency. This meant that it was necessary to first measure privacy concerns over the disclosure of personal information and compare it to actual information disclosure. We then needed to measure the gap (if any) between

concerns and behaviors. The next step was to repeat this process after the intervention to test if there was any change in the concern-behavior gap before and after the intervention. However, participating in a privacy-centric study may cause individuals to think harder about disclosure decisions than they would in their daily lives.

Previous research has also highlighted the effect of contextual factors on privacy decision making. Trust in the entity that information is being disclosed to, for instance, has been shown to lead to higher disclosure behavior (Malheiros, Preibusch, & Sasse, 2013). It is therefore possible that participants might feel a sense of security from knowing that the data is being gathered for academic research purposes and is protected from abuse. This optimism may lead to participants sharing more information than they normally would. On the other hand, research has established the presence of social desirability bias among individuals (Mendel & Toch, 2017). This implies that when participating in a research study, participants try to sub-consciously give the "right" answers. Therefore, knowledge that this is a privacy related study may result in participants enacting unnaturally privacy sensitive behavior.

Previous research investigating the privacy paradox has used deception and undertaken the guise of marketing agencies (Norberg et al., 2007), an online store (Spikermann et al., 2001) or a marketing study for a credit card company (Malheiros et al., 2013). These studies while highlighting the importance of studying privacy concerns and behaviors also caution against giving participants the opportunity to thinking too deeply about either. We used similar deceptive techniques when communicating the purpose and goal of this study in order to gather data that is reflective of everyday attitudes and behaviors.

We informed participants that this study is being conducted as a preliminary step towards launching a new credit scoring method in the United States and explained the importance of credit

scores as well as current methods for building these scores (materials included in Appendix B). Some participants were unsure about the credit score process, we provided these participants with a handout explaining what a credit score is, how it is calculated, and its importance in a person's financial life. We expect that obfuscating the true purpose of the study will allow us to observe actual disclosure behavior in a real-world scenario where participants believe that their information is being disclosed to a third-party entity. We now describe each of the sessions in detail.

5.3. Stage 1: Before Intervention

In the first online session (day 1), participants recruited for the study were sent an email welcoming them to the study and asking them to electronically sign the consent form (Appendix A). Once they signed the form, participants were asked to fill in three surveys, which asked for information about their demography, personality, privacy concerns, and social capital (Appendices C, and D). Williams (2006) Internet social capital scales was used as a measure of social capital while personality was measured using the five-factor model of personality (often termed the "Big Five") (John, Naumann, & Soto, 2008). The social capital and personality scales were administered to shift the focus away from privacy concerns. We wanted participants to believe that we were collecting information about different aspects of their personality so that the true purpose of the study remained hidden. Participants were informed that these surveys were part of a "screening process"; however, all participants who completed the surveys were invited to the lab session.

5.3.1. Measuring Privacy Concerns

To operationalize and measure people's general privacy concerns about information disclosed over the Internet, we adopted the Internet Users' Information Privacy Concerns (IUIPC), a validated scale designed for online contexts (Malhotra, Kim, and Agarwal, 2004). (Appendix C). This ten-

question scale measures privacy concerns along the dimensions of control, collection, and awareness of privacy practices. The questions are based on a 5-point Likert Scale, ranging from 1 (strongly disagree) to 5 (strongly agree). The three dimensions are explained as follows:

- *Collection* refers to the concern an individual feels about the amount of data collected (knowingly or unknowingly) about them in online contexts. This refers to data gathered by institutions as well as individuals.
- *Control* refers to the specific permissions that individuals give to institutions or other individuals allowing them to collect information about themselves.
- *Awareness of privacy practices* which is the third dimension of the IUIPC scale, refers to the concern felt by individuals about their awareness (or lack thereof) of institutional data collection practices.

Previous research investigating privacy concerns has discussed the influence of feeling vulnerable about shared information and the ability to control access to information in shaping individual privacy concerns (Brandimarte et al., 2013). We therefore used the IUIPC scale as a measure of privacy concerns about information disclosed over the Internet. As there were no reverse loading questions and each question was measured on a scale of 1 to 5, the mean score from this survey quantified a person's privacy concern. Finally, we converted the raw score into standardized z-scores, i.e. a measure of how many standard deviations below or above the population mean a raw score is and referred to it as the Privacy Concern Index (PCI).

Participants were asked to complete this survey both before and after the interventions (described in Section 5.4). The interventions were designed to remind participants of their privacy

concerns as well as behaviors, this implies that in some cases the intervention may cause participants to re-think their overall privacy concerns. We therefore asked participants to complete the IUIPC scale in the after-intervention stage as well.

5.3.2. Measuring Privacy Behavior

In order to observe actual information disclosure behavior in online contexts, we invited participants to attend an in-person lab session (day 4). During the lab session, we informed participants that the surveys were part of a screening process and they were now "eligible" to participate in the project. While we did not actually screen participants (all participants who signed the consent form were invited to the lab session), we used this term to discourage participants from focusing on their survey responses. Lab sessions lasted approximately 60 minutes and there were no more than 8 participants in each session. The low number of participants allowed us to verify that all participants were correctly completing the activities. Participants were compensated $20 after completing the first lab session.

During the lab session, participants were asked to choose from a set of activities that required them to disclose personal information over an online web form. The information disclosure activities undertaken during this session were treated as a baseline measure of their behavior. The activities ranged from information disclosure about their preferences and hobbies to describing their opinions on potentially sensitive or controversial topics to disclosing sensitive social, medical, or financial information. Research investigating disclosure has found that an impulsive or emotional response to social media posts often leads to regret (Wang et al., 2011). We therefore consciously used "trigger-words" (e.g. pro-life or pro-choice) designed to encourage participants to make impetuous disclosure decisions.

During the lab session, participants were informed that any information they provided would be analyzed and used to create a credit-scoring algorithm. We explained the significance of

credit scores (Appendix J) and informed participants that while information was being gathered for research purposes, there are no special protections or anonymity granted to their information. The consent form used in this study (Appendix A) was modified to exclude confidentiality and anonymity provisions.

We told participants that any information disclosed during this study would be shared and used by unknown third parties; they should therefore use their own judgement when performing tasks and not assume that their data would be protected. Individuals' sharing information online are often subject to the optimism bias, i.e. they believe that even if there is a security breach, their information will not be stolen or disclosed resulting in the disclosure of (sensitive) personal information that is contrary to their privacy attitudes. We included this step to simulate information disclosure over online web forms.

Research investigating disclosure behaviors has found that different forms of personal information are valued differently by individuals (Staiano et al., 2014). Some empirical studies have attempted to quantify subjective privacy valuations of personal information in different contexts, such as personal information revealed online (Beresford et al., 2012), access to location data (Almuhimedi et al., 2015), or data about an individual's friends or family (Statiano et al., 2014). In a study investigating monetary values that people assign to different kinds of personal information (contact information, social information, photos and videos, and location information), Staiano et al., (2014) find that people assign higher value to information bundles than individual bits of information. Similarly, Horne and Horne (1998) found that consumers were much more concerned about the use of medical, financial, and family information by third parties than they were about their product preferences or daily habits. A similar experiment by Carrascal et al., (2013) investi-

gating the monetizing of personal information finds that individuals placed a high value on financial and medical information when compared to social network interactions. Participants were most comfortable sharing information about their hobbies, activities, and preferences (favorite sport, music, etc.). A related study by Horne, Norberg, and Cemal Ekin (2007) explores the extent to which people lie when disclosing personal information and finds that participants were likely to falsify information that could be perceived as sensitive (e.g. alcohol consumption). Researchers have also studied concerns over socially risky information which when disclosed causes embarrassment. In their study, White (2004) showed a decrease in willingness to discuss attitudes or beliefs (e.g. religious or political) that could lead to stigma.

Finally, researchers have found an inverse relationship between effort and disclosure. Specifically, if a website requires information that is difficult to remember or requires higher cognitive effort, or in cases where a large number of data items are requested, individuals will perceive this interaction as requiring higher effort. *The higher the perceived effort the more likely an individual is to withhold data (Malheiros et al., 2013). Contrarily, in cases where the effort of disclosure was perceived as low (e.g. location disclosure via Facebook check-in) individuals will often share highly sensitive information, often resulting in regrettable disclosure* (Noulas et al., 2011). This research shows the effort required to disclose information is directly connected to the individuals' propensity to share or protect the information. Therefore, when measuring information disclosure behavior we take into account *both the effort required to disclose information as well as the sensitivity of information disclosed*.

We created a list of information disclosure tasks that vary in the sensitivity of information disclosed as well as the effort (cognitive or physical) required to complete the task, and valued at a given number of points (Appendix E and F). The list was distributed to participants as a handout

during the lab session. Participants were asked to read the entire list and select (check off) the tasks they would like to complete to reach a total of 20 points. We then verified that a sufficient number of tasks (i.e. totaling 20 points) were selected and gave participants a link a web form where they could complete the tasks. Participants could access the web form on their smartphones or computers and fill in the required information. The number of tasks an individual chose was used to create a behavioral index for that person.

In order to assign points to tasks, we evaluated tasks on *sensitivity (low, medium, or high)* of information disclosed and *effort (low and high)* required to complete the task. For instance, revealing age, gender, favorite music etc. were considered as low on effort and sensitivity of information and therefore carried 1 point each. Tasks that asked participant for descriptive responses like writing 3-5 sentences about their favorite college campus, or to call and speak with a friend for 3 minutes were considered as low in terms of sensitive information disclosure but required a higher effort and were therefore valued at 3 points. On the other hand, tasks that asked for information in visual formats (e.g. video recording self or others) or medical information (screen shot of insurance page) were considered moderately sensitive and worth 5 points. Highly sensitive information like social security number or email and social media passwords were valued at 10 points. We included multiple tasks of the same value, so that participants had the option of choosing a combination of high and low value tasks, only high value tasks, or only low value tasks.

In order to confirm that the points assigned to the tasks make sense, we ran a pilot study on Amazon Mechanical Turk and asked 35 Turkers to rank each task on effort and sensitivity. The Turkers were shown a list of all tasks in Appendices E and F and asked to assign a low, medium, or high score for sensitivity of information and a low or high score for effort on each task. We clarified in the survey introduction that Turkers should assign ranking based on how they would

feel if a third-party organization asked them to provide the information in an online setting. For instance, if an organization asked them about their favorite musician or preferred hobbies, they should assign a low or high ranking based on the *effort (cognitive or physical)* required to answer and a low, medium, or high rank based on the *sensitivity* they assign to that information. We then tallied the results and divided the tasks into different sections shown in Appendices E and F. A task that more than 20 participants (~60%) ranked high was considered worth more points than tasks ranked low or medium. As a post-hoc measure, we compared the average effort and sensitivity scores assigned to 1-point, 3-point, 5-point, and 10-point tasks (Table 5.1). As the table shows, there is an increasing trend in terms of sensitivity for the points assigned. Also, note that the 3-point and 5-point tasks have a much higher effort than the 10-point tasks which is lower in terms of effort but has a significantly high sensitivity score. The 10-point tasks are designed to cost the least in terms of effort to answer in the study even though they ask for highly sensitive information (e.g. social security number, credit card pin, etc.). Participants can potentially finish the study by just answering two questions instead of going through the time-consuming and effortful process of choosing a selection of different tasks, thinking back to past event or memories, composing and writing lengthy answers, or recording and uploading videos. The turkers score reflected this reduction of effort for the 10-point tasks as they were ranked easier (purely in terms of effort) to complete that the 3-point or 5-point tasks.

Table 5.1

Average effort and sensitivity values as assigned by Turkers.

Task Value	Effort	Sensitivity
1 Point	5.58	5.25
3 Point	23.22	16.61
5 Point	24.90	25.70
10 Point	9.50	34.00

Once participants finished tasks, we looked over the completed online form to verify that questions worth 20 points had indeed been answered, and the lab session was considered complete. Participants were compensated $20 and thanked for their time. Recording the number of tasks participants choose to perform allowed us to compute a behavioral index (BI) for each participant, which was their z-score based on the number of tasks undertaken. Therefore, if a participant chose to complete high value tasks to get to the desired 20 point total, their behavior showed a lack of concern for privacy. Conversely, if they performed more low value tasks to get to 20 points, they demonstrated a greater privacy concern. The BI was then compared to PCI to check for a gap between concerns and behavior.

In this work, we compare overall privacy concerns with general information disclosure behavior. While previous studies examining the privacy paradox have compared (un)willingness to disclose specific information (e.g. medical history) with actual disclosure of that information (Norberg et al., 2007) we take a more holistic approach to comparing concerns and beliefs. Given the wide-ranging nature of information sharing and disclosure (SNS, Smart devices, third-party agencies, governments, etc.) in today's digitized world, a 1:1 comparison of concerns and behaviors would be limiting in its scope as well as fail to account for contextual factors or trade-off negotiations that influence information disclosure decisions. A more generalized measurement of privacy concerns and behaviors on the other hand, allowed us to gauge general privacy concerns over information sharing and compare this to information disclosure via a broad range of daily life activities. We perform a relative comparison of concerns and behaviors across the sample population. For instance, if an individual had higher privacy concerns than 90% of the population but their actions demonstrated privacy sensitivity that is only higher than 20% of the population there is a clear gap and one that can perhaps be addressed by timely reminders or interventions.

5.4. Stage 2: Interventions

Participants who completed stage 1 of the study were divided into three groups. We aimed for a similar demographic representation across the groups. In the current study, there were no participants whose concern (PCI) and behavior (BI) indices were aligned. Therefore, all participants were considered eligible to receive the interventions. The interventions were designed to nudge participants to re-think their privacy concerns and behaviors in order to reduce the concern-behavior gap. It was possible for this gap to occur in two directions. Participants could either have behaviors that reflect a lower concern for privacy than their beliefs suggest *or* demonstrated behavior that was more privacy sensitive than their concerns suggested. We argue that a misalignment in either direction would benefit from an intervention that reduces the gap between concerns and behaviors.

While it seems intuitive that an increased privacy sensitive behavior (even with lower privacy concerns) is a good thing, such counter-attitudinal behavior demonstrates a departure from privacy concerns. We argue that concerns over information privacy are a personal matter and there is no prejudice attached to having a lower concern for privacy. However, this concern should be reflected in their daily behaviors. By demonstrating behavior that is more privacy sensitive than their beliefs, individuals are again falling prey to biases (heuristic, spur-of-the-moment decision making) and perhaps losing out on opportunities to leverage their information for different kinds of financial or social opportunities. On a societal level, a misalignment between attitudes and behaviors in either direction can have similar consequences and result in similar challenges when framing privacy policy. We therefore designed this study to help individuals synchronize their concerns and behaviors without passing judgement on "ideal" privacy concerns or behaviors.

Participants were emailed interventions according to their experimental condition (social, inconsistency, or control) for a period of three days (days 8-10). The first intervention (day 8) told

participants about their privacy concerns and disclosure behaviors. This intervention differed across the three groups and was explained in detail in Sections 5.4.1-5.4.3. The second and third intervention contained information about the participant's personality and social capital based on the information collected from online surveys. These interventions (days 9 and 10) were the same for all three groups.

In order to make sure that participants saw the intervention, we included a unique code in each email. Participants needed to use one of the three codes to access the online sessions in stage 3 and were asked for all three codes when they came in for the lab session.

5.4.1. Social-proof Intervention

The first intervention was designed based on the concept of *social-proof*, i.e. the idea that knowledge of actions performed by a large group of people will influence individuals to act accordingly (Cialdini, 2001). Previous research has highlighted the value of using social-proof as a method for increasing security awareness in individuals (Das et al., 2014). In their work, Norberg et al., (2004) compared people's willingness to disclose information with their perception of *information disclosure by others* and found no significant difference. This implies that people believe that their own disclosure behavior is no different from others. We, therefore, hypothesize that reminding individuals of different behaviors enacted by a majority would cause them to re-think their own information disclosure behavior. The social conformity paradigm of the cognitive dissonance theory also points to the importance of gaining social consensus and looking to group norms when engaging in risky behavior (Zanna & Sande, 1987). Stone and Cooper's (2001) model on the influence of group norms on dissonance note the importance of maintaining membership within a group as a necessary condition for conformity. This implies that if membership within the group is not valued by the individual, the need to maintain norms is also weakened. Although

longstanding relationships with family, friends, and colleagues can affect motivation, there is also research showing that minimal social cues can create a sense of social connection with even unfamiliar others and result in people adopting the interests and goals of these others as their own (Walton, Cohen, Cwir, & Spencer, 2012). This phenomenon has been called "mere belonging" and defined as "an entryway to a social relationship-a small cue of social connection to another person or group in a performance domain" (Walton et al., 2012, pp. 514). Through a series of experiments investigating this theory, Walton et al., (2012) found that people internalize goals and motivation even from unfamiliar others automatically, as a result of small or minimal cues of social connectedness. When designing the social-proof intervention, we wanted participants to experience a sense of belonging within a group. We therefore use the term "peers" to create an impression of a group that shares certain characteristics with the participant.

Based on the theoretical insights gained from social-proof and cognitive dissonance theory, the first set of interventions were designed to show participants that for the majority, privacy behaviors are aligned with concerns. We included visual and textual components highlighting this departure of participant behavior from the majority group. We expect that seeing this notification will prompt individuals to mimic their peer group and reduce the discrepancy between concerns and behaviors. The social-proof based interventions used are shown below:

Figure 5.2. Social-proof interventions over three days for Group 1 participants

Dear xyz,

Thank you for participating in the "*Power of Information*" study. Here is some feedback based on an analysis of your activities:

You have a higher privacy concern than 80% of your peers.

You shared more personal information than 40% of your peers.

It might be interesting for you to know that most of the participants **did not have this gap** between their privacy concerns and disclosure behavior. In fact, almost all of your peers had a smaller gap between their privacy concerns and disclosure behavior.

Your code value: RE5H4

Dear xyz,

We are glad you chose to be a part of the "*Power of Information*" study. We analyzed your data and found the following results that you may be interested in:

Your openness score is **High**.
Your conscientiousness score is **Low**.
Your extraversion score is **High**.
Your agreeableness score is **High**.
Your neuroticism score is **Low**.

This means that you are spontaneous, open to new experiences, love meeting new people and making connections, and have a wide variety of interests. You are friendly, cooperative, and positive by nature. These traits make you likeable and prone to success in your chosen field.

Your code value: TF9G2

Dear xyz,

As a valued member of the Power of Information study, we gathered some additional insights from your data. Here are the results:

You have a **Low** bridging score.
Your bonding score is **High.**
You have **High** social capital.

Social capital is the value gained from an informal network of relationships. Your scores indicate that you are someone who prefers to form close, long-lasting meaningful bonds within a relatively small network. The time and energy you have invested in building and maintaining these close ties make you a great friend and an extremely valuable member of your network.

Your code value: UH3E6

5.4.2. Information Inconsistency Intervention

The second intervention is designed to remind individuals of the gap between their concerns and behavior. This intervention contrasted the PCI and BI for each individual. Here, the intuition is that showing participants a message that explicitly highlights the inconsistencies in their concerns and actions would nudge them to make disclosure decisions more aligned with their privacy concerns. A recent study by Jackson and Wang (2018) used a similar intervention in the context of mobile app installation. We designed an intervention along these lines and tested if highlighting the inconsistency between concerns and behavior could result in a reduced discrepancy in real life information sharing behaviors. We expect this intervention to perform better than the baseline in reducing the privacy paradox. The interventions for Group 2 are shown below:

Figure 5.3. Privacy Inconsistency interventions over three days for Group 2 participants.

Dear xyz,

Thank you for participating in the "*Power of Information*" study. Here is some feedback based on an analysis of your activities:

We found that your information sharing practices **do not match** your privacy concerns.

Your level of information sharing was **higher** than your stated comfort level. This means that you **shared more information** than you would want to.

Your code value: RE5H4

Dear xyz,

We are glad you chose to be a part of the "*Power of Information*" study. We analyzed your data and found the following results that you may be interested in:

Your openness score is **High**.
Your conscientiousness score is **Low**.
Your extraversion score is **High**.
Your agreeableness score is **High**.
Your neuroticism score is **Low**.

This means that you are spontaneous, open to new experiences, love meeting new people and making connections, and have a wide variety of interests. You are friendly, cooperative, and positive by nature. These traits make you likeable and prone to success in your chosen field.

Your code value: TF9G2

Dear xyz,

As a valued member of the Power of Information study, we gathered some additional insights from your data. Here are the results:

You have a **Low** bridging score.
Your bonding score is **High.**
You have **High** social capital.

Social capital is the value gained from an informal network of relationships. Your scores indicate that you are someone who prefers to form close, long-lasting meaningful bonds within a relatively small network. The time and energy you have invested in building and maintaining these close ties make you a great friend and an extremely valuable member of your network.

Your code value: UH3E6

5.4.3. Baseline Intervention

We also included a baseline group who received daily emails similar to participants in Groups 1 and 2. However, notification sent to this group showed participants their information disclosure behavior without any additional information on their privacy concerns. We included Group 3 to verify that it is indeed the social cues presented or inconsistency highlighted that results in reducing the gap between concerns and behaviors rather than simply receiving daily notifications. The baseline interventions are modelled to mimic the design of interventions for Groups 1 and 2 as closely as possible. A sample is presented below:

Figure 5.4. Baseline interventions over 3 days for Group 3 participants.

Dear xyz,

Thank you for participating in the "*Power of Information*" study. Here is some feedback based on an analysis of your activities:

During the previous in-lab session you performed 10 tasks and which resulted in some **personal information disclosure**.

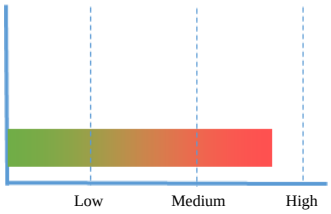

Overall, you have a **High** information sharing level.

Your code value: DG78H

Dear xyz,

As a valued member of the Power of Information study, we gathered some additional insights from your data. Here are the results:

You have a **Low** bridging score.
Your bonding score is **High.**
You have **High** social capital.

Social capital is the value gained from an informal network of relationships. Your scores indicate that you are someone who prefers to form close, long-lasting meaningful bonds within a relatively small network. The time and energy you have invested in building and maintaining these close ties make you a great friend and an extremely valuable member of your network.

Your code value: QT2P1

Dear xyz,

We are glad you chose to be a part of the "*Power of Information*" study. We analyzed your data and found the following results that you may be interested in:

Your openness score is **High**.
Your conscientiousness score is **Low**.
Your extraversion score is **High**.
Your agreeableness score is **High**.
Your neuroticism score is **Low**.

This means that you are spontaneous, open to new experiences, love meeting new people and making connections, and have a wide variety of interests. You are friendly, cooperative, and positive by nature. These traits make you likeable and prone to success in your chosen field.

Your code value: VX391

A two day "gap period" after the interventions was built into the study design to reduce chances of capturing artificially inflated concerns or behavior in the post-intervention stage. This work tested the effectiveness of interventions to help users reduce any gap between their privacy concerns and disclosure behaviors. However, measuring concerns and behaviors immediately after the intervention may show an artificially inflated concern for privacy that is not demonstrative of actual everyday behavior. Therefore, we gathered data on privacy concerns and behavior data two days after the interventions.

5.5. Stage 3: After Intervention

This work uses a before/after study design, it is therefore important for Stage 3 to replicate actions performed in Stage 1 as closely as possible. Hence, participants were again asked to fill in the personality, IUIPC, and social capital scales online before the second lab session (day 13). The

interventions described in section 5.3 were designed to remind participants of the gap between their privacy concern and disclosure behavior. It was however possible for some participants to rethink their concerns over information privacy while participating in the study. We, therefore, measured privacy concerns, using the same scale as in Stage 1, after the interventions and computed a new PCI for each participant.

5.5.1. Lab Session 2

Participants were then invited to the second lab session (day 13) two days after completing the online surveys similar to the Stage 1 process. In the second lab session, we thanked participants for their continued participation and asked them for the three codes they received over email to confirm eligibility. We continued with the deception telling participants that we require some further information in order to fine-tune our "algorithm."

Participants were again given a handout with a list of tasks comparable to the tasks performed in Stage 1 (full list in Appendix F), asked to select the tasks they would like to perform in this stage, and get it signed off by staff before they start completing the web form. For instance, in Stage 1 we asked participants to share a video recording of themselves logging in and checking their email for 5 points. The rationale here, was that the participant would disclose their email preferences, login screen names, and perhaps some social information. In Stage 3, we asked participants to share a similar video recording of themselves logging into and browsing their social media sites resulting in similar disclosure. We gave participants the same instructions as we did in the before intervention step in order to replicate Stage 1 conditions as closely as possible. Once the tasks were completed and we verified that sufficient questions (worth 20 points) had been correctly answered, participants were thanked and paid the remaining $20.

We included a de-briefing session at the end of lab session 2, during which we revealed the true purpose of the study (Appendix H). Most participants expressed some surprise at the deception but we did not observe any significant expression of shock and anger. The second lab session was planned to last 90 instead of 60 minutes so that we could discuss the need for deception in detail with each participant. We explained that when measuring privacy concerns and behaviors it was important for them to perform activities while believing that no special protections would be given to their data. We also informed participants about the effects of "priming" and our need to use deception in order to gather data that accurately reflects privacy concerns. We assured participants that the information disclosed by them in the lab sessions had been destroyed and only the task logs (i.e. number of tasks completed) would be used in analysis and reporting. We then offered participants the chance to earn an additional $10 by engaging in an exit interview session.

5.6. Exit Interview

The relationship between privacy attitudes and behaviors is a complex and subjective one often dependent on contextual and heuristic cues (Acquisti et al., 2017). It was therefore important to understand the underlying motivations governing changes in privacy concerns and disclosure rather than simply measuring the difference. The second research question in this work attempted to delve deeper into the cognitive processes guiding privacy concerns and disclosure behaviors and tried to understand the effect of different interventions on counter-attitudinal disclosure.

We therefore conducted exit interviews during which we discussed actions undertaken by participants in the before and after intervention stages and attempted to understand what prompted them to make these decisions. At the end of the interview, participants were compensated an additional $10 and thanked for their participation.

5.6.1. Interview Analysis

Participants who successfully completed before and after intervention sessions were invited back to do a standardized open-ended interview to better understand the effect of interventions on counter-attitudinal disclosure. All participants answered identical standardized questions; however, follow-up questions were asked wherever a need for greater clarification arose (Appendix G). Asking open-ended questions allowed participants to fully express their viewpoints and experiences with the interventions and study designs.

We wanted to gain a deeper understanding of the effect of social-proof based interventions on counter-attitudinal disclosure. We therefore asked participants across the three conditions to describe their thought processes during the experiment. All interviews lasted approximately 30-45 minutes and focused on the effect of interventions on privacy concerns and disclosure behaviors. Thematic analyses (Braun & Clarke, 2006) were then conducted on the transcriptions to reveal major themes and issues. Thematic analysis is a method for identifying, analyzing, and reporting themes within data and is frequently used in mixed method studies.

We used an analytical approach where a thematic framework was developed based on the results derived from the initial quantitative analysis. Once the thematic factors were identified, all interviews were thoroughly read through and all parts relevant to specific themes were collected, and finally a coherent narrative of the different ways in which participants engaged with the interventions was produced. A list of themes and quotes illustrating each theme can be found in Appendix K.

We summarize the approach taken to answer the identified research questions in Table 5.2. RQ1 is studied using quantitative and qualitative analysis before and after interventions. A number of related hypotheses are tested using quantitative analysis. The interpretation of the results is

facilitated by standardized open-ended exit interviews. RQ2 is studied based on the qualitative analysis of the data obtained via standardized open-ended interviews (Appendix G).

Table 5.2

List of research questions, methods and instruments used, and analysis styles.

Research Questions	**Method**	**Instruments**	**Analyses**
RQ1. Can social-proof based nudges be an effective way to reduce the privacy paradox?	1. Testing hypotheses based on the effectiveness of different interventions. 2. Interpreting associations found based on qualitative interviews.	1. Gap between PCI and BI before and after interventions. 2. Standardized open-ended interviews	Quantitative and qualitative analysis.
RQ2. What effect did the interventions have on the gap between individual privacy concerns and disclosure behaviors?	1. Qualitative analysis based on interviews.	1. Standardized open-ended interviews	Qualitative analysis.

Chapter 6: Results

6.1. Demographic Description

A total of 42 participants completed the study out of whom 20 participants agreed to the exit interviews. In this study, we used a between-subject interventional design where participants were randomly assigned to one of three experimental groups. There were 15 participants in the social intervention group, 13 in the information inconsistency group, and 14 in the control group. 21 participants self-reported their gender as males and 21 as females. Most participants (73%) were in the 18-24 year range. Table 6.1 gives a more detailed description of the demographic characteristics of the participants.

Table 6.1

Demographic profile of participants across all three experiment conditions.

	Social (n = 15)	**Inconsistency (n = 13)**	**Control (n = 14)**	**Total (n = 42)**
Gender	8 Male 7 Female	6 Male 7 Female	7 Male 7 Female	21 Male 21 Female
Age	11 between 18-24 yrs 4 between 24-34 yrs	11 between 18-24 yrs 1 between 24-34 yrs 1 between 35-44 yrs	9 between 18-24 yrs 3 between 24-34 yrs 2 between 45-54 yrs	31 between 18-24 yrs 8 between 24-34 yrs 1 between 35-44 yrs 2 between 45-54 yrs

6.2. Aligning Concerns and Behavior

The main purpose of this work was to test the effectiveness of different interventions in reducing the discrepancy between privacy concerns and behaviors. In this chapter we describe findings from the statistical analyses for hypotheses (H) and research questions (RQ). Findings will be described in the order of hypotheses and research questions as presented in Chapter 4. For H1, H2 and RQ1, which examine the relationship between interventions and the concern-behavior gap, we used a mixed method analysis gathering data from statistical difference of means tests (while controlling for demographics gender, age, and education) as well as qualitative interviews. RQ2 investigated

how interventions changed counter-attitudinal disclosure across the experiment groups using a qualitative approach. This RQ2 used thematic analyses on transcriptions from exit interviews. The next section presents a detailed analysis of RQs and hypotheses.

6.2.1. Social Interventions and the Concern-Behavior Discrepancy

The first RQ in this work asked about the effectiveness of social-proof based interventions in reducing the discrepancy between information privacy concerns and disclosure behaviors. This meant that we needed to first identify the difference (Δ) between concerns and behaviors before the intervention (Gap_1) and the Δ after interventions (Gap_2). We then need to check if the average (ΔGap_2, Gap_1) was higher in the social-proof intervention group when compared to other groups. In order to effectively compare privacy concerns and behavior, we first needed to identify a standard measurement scale. As described in the previous sections, privacy concerns were measured using the IUIPC scale while privacy behavior was measured based on the number of disclosure activities undertaken by an individual. Since concerns and behaviors were measured using different instruments (i.e. a survey and task list), we created standardized z-scores for concerns (PCI) and behavior (BI) before and after the intervention. Standardized z-scores are calculated as a function of the mean and standard deviation of the population. In this case, the raw scores from IUIPC surveys and tasks completed were both different for participants before and after interventions. Hence, we calculated PCI and BI before interventions using the mean and standard deviation before intervention and a PCI and BI after intervention based on the mean and standard deviation after intervention. This brought both PCI and BI into a comparable range (as per the central limit theorem (Abdi, 2007), 99% of all z-scores will lie between +3 and -3) so that we could perform further analysis. Table 6.2 shows the preliminary analysis of concern-behavior discrepancy across the three experiment groups before and after interventions.

Table 6.2

Average scores for privacy concerns, disclosure behaviors, and the concern-behavior discrepancy (based on standardized z - scores) for individuals across the three experiment conditions.

	Social (n = 15)	**Inconsistency (n = 13)**	**Control (n = 14)**	**Total (n = 42)**
Average Concerns_Before	-0.21	0.05	0.18	0
Average Concerns_After	-0.14	-0.19	0.32	0
Average Behavior_Before	-0.20	0.31	-0.07	0
Average Behavior_After	-0.35	0.41	0.00	0
Average Gap_1	1.44	1.01	0.70	1.06
Average Gap_2	0.83	1.21	0.99	1.01
Percentage reduction in discrepancy	61%	(-)20%	(-)29%	5%
Percentage of people for whom $Gap_2 < Gap_1$	73.33% (11 out of 15)	30.77% (4 out of 13)	28.57% (4 out of 14)	45.23% (19 out of 42)

* Gap_1 = concern behavior discrepancy before intervention. Gap_2 = concern behavior discrepancy after intervention.

We then performed a difference of means test to check if the sample populations were similar in terms of differences in concerns and behaviors across the three intervention groups. We do not find significant differences in concerns (section 6.2.2) and behaviors (section 6.2.3) before and after the interventions. We do however find that participants in the social-proof intervention condition had a significant difference in Gap_1 (*F-value = 2.71; p-value < 0.05*) when compared to participants in the control condition. There was no significant difference in Gap_2 across the three experiment conditions. This implies that participants in the social-proof condition had a higher gap between their concerns and behaviors, before interventions, than participants in the other two conditions. When creating intervention groups, we attempted to maintain a similar distribution of age and gender across the three groups. While participants in all groups had similar levels of privacy concerns and exhibited similar disclosure behaviors, the *gap* between concerns and disclosure was significantly higher for some participants.

Table 6.2 shows that participants in the social-proof intervention experienced the most reduction in discrepancy as well as in the number of people who had a reduction in the concern-behavior gap. We also saw an increase in discrepancy among the inconsistency and control experimental groups. While this result is contrary to our expectations, an in-depth analysis of the variation in privacy concerns and behaviors discussed in Sections 6.2.2 and 6.2.3 can provide an explanation for this increase in discrepancy. Some participants in the inconsistency and control conditions experienced a drop in their privacy concern levels (Table 6.4) after the interventions while demonstrating higher privacy sensitive behavior (Table 6.5). This implies that any interventions focused on information sharing will result in an increase in privacy-sensitive behavior however, a more nuanced intervention, (e.g. showing comparative percentages of privacy concern and behavior), is required to achieve the desired goal of aligning concerns with behavior.

Further, our results showed that simply pointing out a gap in concerns and behavior, or reminding individuals' of their disclosure behavior is not as effective as showing this discrepancy along with information on peer behavior. This was borne out during the exit interviews when participants in the social condition used terms like "*felt isolated*", "*singled out*", and "*felt weird about acting differently*" when describing their reactions on seeing the interventions. The group norm paradigm of the Cognitive Dissonance Theory (CDT) discusses the tendency of individuals to look to peer behavior as a way of reducing dissonance. In the context of the current study, the social-proof intervention provided two critical pieces of information to participants: (i) that there was a gap between their concerns and behavior, and (ii) this gap did not exist for most of their peers. From the analysis shown in Table 6.2, it is clear that participants used this knowledge of peer actions to model their own concerns and behaviors and therefore, improve the alignment between their privacy concerns and disclosure behaviors.

Participants in the information inconsistency condition received an intervention that said their privacy concerns and disclosure behavior did not match (Figure 5.3). As seen in Table 6.2 participants in this condition actually had an increased discrepancy between their privacy concerns and disclosure behaviors. The qualitative interviews pointed out that for most participants in this condition being made aware of the gap between their privacy concerns and disclosure behaviors did cause them to feel dissonance but that was not sufficient to make them consciously try to reduce this gap ("*It [the email intervention] said my information sharing didn't match my concerns…I felt like so what? I mean was it supposed to be the same?*" – P32). The CDT looks at changes in attitude (or concerns) *after* an individual has performed an action. According to the CDT, when an individual faces dissonance due to engaging in counter-attitudinal behavior, *they tend to shift their attitudes to reflect their behavior rather than the other way around*. It is usually not possible for a person to change an action they have already performed; therefore, dissonance is reduced by shifting the attitude to align with the exhibited behavior (Festinger, 1957). In the current study, people in the information inconsistency condition, did not have the advantage of being told how their peers had performed. Therefore, when individuals received notification that their disclosure was higher than their comfort level, rather than try to align their concerns and disclosure, participants may have reported lowered concerns to avoid dissonance. From this analysis, it is clear that when investigating a gap between privacy concerns and disclosure behaviors, it is not enough to simply do measure the discrepancies. One must also investigate individual changes in privacy concerns as well as disclosure behaviors to identify where and how these discrepancies occur.

In the next two sub-sections (6.2.2 and 6.2.3) we tease apart participant privacy concerns and observed disclosure during the study and examine how these two factors changed across the different experiment groups.

6.2.2. Measuring Concern for Information Privacy

Individual privacy concerns for participants were measured using the IUIPC scale (Malhotra et al., 2004) both before and after the interventions. Mean scores from this survey were converted into standardized z – scores to create a privacy concerns index (PCI) for each participant before and after the intervention. The IUIPC scale measures privacy concerns along the dimensions of control, collection, and awareness of privacy practices. In order to check if any of these dimensions could be an accurate predictor of disclosure behavior, we calculated z-scores separately for each dimension (zControl, zCollection, and zAwareness). We then performed a regression analysis with z-score of the behavioral index as the dependent variable and zControl, zCollection, and zAwareness as the independent variables. The results $F\ (3, 38) = 1.017$, $p\text{-value} > 0.05$ shows a non-significant relationship between the dimensions of IUIPC survey and disclosure behavior. We also performed multiple correlation analyses between the behavioral index and each dimension i.e. zBehavior vs. z Control, zBehavior vs. z Collection, and zBehavior vs. z Awareness. These comparisons yielded non-significant ($p > 0.05$) results as well. We now proceed to further testing with the overall privacy concern score.

We use Cronbach's alpha to assess the internal consistency of the scale both before and after the interventions. Cronbach's alpha is a commonly used method to assess the reliability of a multi-item scale. Simply put, it describes the extent to which all the items in a test measure the same concept or construct and hence it is connected to the inter-relatedness of the items within the test (Cronbach, 1951). Reliability for the scale was good in both the before ($\alpha = 0.86$, $mean = 4.10$, $SD = 0.64$) and after ($\alpha = 0.87$, $mean = 3.99$, $SD = 0.70$) intervention stages.

When comparing the concern-behavior gap across different groups, it is important to verify that participants in all three conditions have similar privacy concerns. We therefore perform a difference of means test and find differences to be non-significant in the before (*F (2, 40) = 0.56, p-value = 0.58*) and after (*F (2, 40) = 0.40, p-value = 0.68*) intervention stages. We now perform a detailed analysis of the changes in privacy concerns before and after interventions across the three experiments conditions (Table 6.3).

Table 6.3

Average privacy concern scores (based on raw survey scores) for individuals across the three experiment conditions.

	Social (n = 15)	**Inconsistency (n = 13)**	**Control (n = 14)**	**Total (n = 42)**
Average IUIPC Score_Before	3.93	4.13	4.18	4.06
Standard Deviation_Before	0.78	0.59	0.53	0.64
Average IUIPC Score After	4.08	3.90	4.24	4.07
Standard Deviation_After	0.64	0.69	0.51	0.61
Average Increase in Concerns	0.15	(-) 0.23	0.06	0.01
Percentage of people who had higher concerns after intervention	40% (6 out of 15)	30% (4 out of 13)	42% (6 out of 14)	38% (16 out of 42)

As seen in Table 6.3, participants in the social condition had the highest increase in privacy concern followed by participants in the baseline condition while participants in the information inconsistency condition had reduced privacy concerns after the intervention. The introductory analysis in Section 6.2.1 suggested that while participants in the information inconsistency condition received the notification about a gap between their concerns and behaviors, they attempted to resolve this gap by lowering concerns for information privacy to match their disclosure levels. Participants in the baseline condition show a slight increase in concerns for privacy, but this increase was not enough to reduce the gap between concerns and behaviors. In fact, Table 6.2 shows

that the concern-behavior gap increased for participants in the information inconsistency and baseline conditions.

When we asked participants about this change in concerns during the exit interviews, we found that individuals in the social condition found the intervention design clear ("*The percentages were pretty easy to understand*" – P1) and were motivated to model their behavior on the "ideal" majority behavior of having aligned privacy concerns and disclosure behaviors ("*I didn't like being the only one with a 'gap*" – P34). This implies that participants in the social-proof condition thought about their privacy concerns and disclosure behaviors as well as the concerns and behavior of their peers and demonstrated a proportional increase in privacy concerns and disclosure behavior resulting in a reduced discrepancy.

Table 6.3 also shows that participants in the information inconsistency condition had lowered privacy concerns in the after intervention stage. During the qualitative interviews, we asked participants to describe their information privacy concerns during the study. Most participants in the information inconsistency condition reported feeling more or similar levels of concern in the after intervention stage. Participant P26 mentioned, "*I don't think my concerns changed during the study. I still think as much about privacy as I always did*." However, from an examination of survey responses, we found P26's privacy concern survey score dropped from 5.00 in the before intervention stage to 3.33 in the after intervention stage. This means that there was a drop of 1.67 points on the IUIPC scale in the after intervention stage. While P26 did not have an explanation for why this drop could have occurred, examining this change in privacy concerns through the lens of the CDT provides an explanation for decreased privacy concerns. When participants in the information inconsistency stage encountered dissonance, they might *subconsciously* have lowered their privacy concerns to reflect their disclosure behavior.

For participants, in the baseline condition there was a small increase in information privacy concerns (Table 6.3). While the baseline intervention did not contain any information on privacy concerns, during qualitative interviews a few participants mentioned that simply being a part of the experiment caused them to think more about information privacy, which was reflected in the survey scores. Participant P27 who had a slight (0.22) increase in privacy concerns in the after intervention stage noted "*I was definitely thinking a lot more about privacy just from all the activities I did in the lab and the emails and everything*." Table 6.3 showed that a similar number of participants in the social – proof and baseline conditions showed increased privacy concerns. However, the overall increase in privacy concerns is vastly different for these two groups. Even though 6 out of 14 participants in the baseline condition had increased privacy concerns the average increase was only 0.06. The main purpose of this work, however, is to find an effective way to reduce the discrepancy between privacy concerns and behavior. Therefore, we now need to examine disclosure behaviors across the three intervention groups in order to make sense of the results shown in Table 6.2.

6.2.3. Measuring Disclosure Behavior

In order to measure information disclosure in online contexts we use an online web form where participants could choose to share or protect information by completing some activities. Institutions commonly use web forms to ask for personal information for marketing purposes – for instance, when registering to access a website or joining an online community or to provide more personalized services to consumers. Researchers have also found that entering information in online web forms often encourages increased disclosure and candid responses (Preibusch, Krol, & Beresford, 2013) resulting in a heightened sharing of personal information often contrary to information privacy concerns. In order to measure if there was a discrepancy between information pri-

vacy concerns and disclosure behavior, we asked participants to choose which information disclosure activities they would like to perform and enter the information in an online web form (full list in Appendix E and F). The number of tasks participants chose to perform was treated as a measure of their disclosure behavior, i.e. an individual who completed 20 tasks was considered to have *higher* privacy sensitivity compared to an individual who completed 10 tasks. A more detailed distribution of the average number of tasks in each category that participants performed before and after the interventions is available in Appendix L. Similar to the measurement of privacy concerns, information disclosure was measured both before and after the interventions. These raw scores were then converted to standardized z-scores in order to create the behavioral index (BI).

The first step is to verify that there are no significant differences in disclosure behavior across the experiment groups i.e. no single group had a cluster of highly concerned (or not concerned) participants. We again tested for the difference of means and found insignificant differences in the before ($F\ (2, 40) = 1.29$, $p\text{-value} = 0.29$) and after ($F\ (2, 40) = 2.78$, $p\text{-value} = 0.08$) intervention stages. This meant that disclosure behavior was more or less equally distributed across the three experimental conditions. A detailed analysis of the changes in disclosure behavior before and after interventions across the three experiments conditions is shown in Table 6.4.

Table 6.4

Average disclosure behavior for individuals (based on number of tasks) across the three experiment conditions.

	Social (n = 15)	**Inconsistency (n = 13)**	**Control (n = 14)**	**Total (n = 42)**
Average number of activities_Before	13.07	15.62	13.71	14.07
Standard Deviation_Before	5.26	3.32	5.40	4.97
Average number of activities_After	14.60	18.31	16.29	16.31
Standard Deviation_After	5.24	2.81	5.12	4.88
Average increase in activities	1.53	2.69	2.57	2.23
Percentage of people who did more activities after intervention	33% (5 out of 15)	69% (9 out of 13)	57% (8 out of 14)	52% (22 out of 42)

As seen from Table 6.4 all three interventions were effective in increasing privacy sensitive behavior. This table also shows that participants in the information inconsistency had the highest increase in privacy protective behaviors followed by participants in the baseline condition and then participants in the social-proof condition. Comparing these with the changes in overall privacy concerns (Table 6.3) through the study, we see that people in the social-proof and baseline experiment groups both had an increase in both privacy concerns and disclosure behavior. People in the information inconsistency group however, had lowered privacy concerns but displayed more privacy protective behaviors. A further examination between the social-proof and baseline conditions shows that only participants in the social-proof conditions had a similar increase in both privacy concerns and disclosure behaviors. While participants in the baseline condition had a minor increase in privacy concerns (0.06), they had a much higher increase in privacy protective behaviors (2.57) resulting in an increased gap between privacy concerns and disclosure behavior.

This further examination of variations in concerns and behaviors across the three experiment conditions, helps us make sense of the results shown in Table 6.2 as well as answer RQ1 i.e.

social interventions were more effective than the other two interventions in reducing the gap between privacy concerns and disclosure behaviors. We now move to testing the hypotheses using a quantitative analyses of changes in Gap_1 and Gap_2 across the three experiment groups.

6.3. The effect of interventions in reducing concern-behavior gap

Hypotheses H1 and H2 were used to compare the effectiveness of different interventions in reducing the gap between privacy concerns and disclosure behaviors. We performed a difference of means test to check if the type of intervention had a significant effect on reducing the discrepancy between privacy concerns and behaviors. The results of each hypothesis test are reported below.

6.3.1. Hypothesis H1a

Hypothesis H1a states that people receiving the social-proof intervention will have a greater reduction in discrepancy between their concerns and behaviors when compared to the other two groups. A difference of means test (ANCOVA) was used to test this hypothesis. Findings showed that the type of intervention had a significant effect on reducing the discrepancy between attitudes and behaviors after controlling for age, gender, and education $F\ (5, 36) = 5.852$, $p\text{-value} < 0.01$, $\beta = 0.41$ (table 6.5). Pairwise tests (using Bonferroni's correction for multiple comparisons) revealed that the social intervention significantly decreased the discrepancy between concerns and behaviors when compared to the inconsistency group ($p = 0.044$) and the baseline group ($p = 0.018$). Therefore, the hypothesis H1a was supported (Table 6.6). Among control variables, gender and education had a statistically significant effect on reduction in discrepancy. Women were more likely than men to align concerns and behaviors, $\beta = .20$, $p < .05$. The greater the level of education, the more likely people were to have a lower concern-behavior discrepancy, $\beta = .18$, $p < .05$.

6.3.2. Hypothesis H1b

Hypothesis 1b stated that people receiving the information inconsistency nudge will have a greater reduction in discrepancy between their concerns and behavior than people in the baseline condition.

While the difference of means test (ANCOVA) found that the type of intervention had a significant effect on reducing the discrepancy between concerns and behaviors (Table 6.5), identifying where these differences lay required further analysis. We performed pairwise tests (using Bonferroni's correction for multiple comparisons) to check if the people in the information inconsistency condition experienced a significantly higher reduction in concern-behavior discrepancy when compared to people in the baseline condition (Table 6.6). Our analysis showed that there was no significant difference ($p > 0.05$) in the reduction in discrepancy between people in the information inconsistency and baseline condition (Table 6.6). Therefore, the hypothesis was not supported.

Table 6.5

Test of between-subject effects controlling for age, gender, and education.

	Sum Squares	F value	Significance
Intervention	7.51	6.16	0.005**
Gender	2.82	4.62	0.04*
Age	0.28	0.44	0.51
Education	2.67	4.39	0.04*

***. The mean difference is significant at the 0.01 level.*
*$R^2 = 0.36$ Adjusted $R^2 = 0.27$, $F(5, 36) = 4.089$, p-value = 0.005***

Table 6.6

Post-hoc test (using Bonferroni correction) comparing intervention groups.

	Estimate	95% Confidence Interval for difference		Significance
		Lower Bound	Upper Bound	
Social - Inconsistency	0.80	0.15	1.46	0.02*
Social – Control	0.90	0.25	1.54	0.04*
Inconsistency – Control	-0.09	-0.73	0.55	1.00

**. The mean difference is significant at the 0.05 level.*

6.3.3. Hypothesis H2a

Hypothesis H2a states that the number of people who had a reduction in discrepancy between concerns and behaviors will be more in the social condition compared to other groups. In order to

compare the number of people who experienced a reduction in the discrepancy between their concerns and behaviors, we first coded participants who for whom $Gap_2 < Gap_1$[1] as *I (Increase)* and participants for whom $Gap_2 >= Gap_1$ as *D (Decrease)*. (There were no participants for whom the gap remained the same.) We now had two categorical variables (i.e. intervention and reduction in Gap), and needed to test if these two variables were related. We performed a Chi-square test of independence which showed that there was a significant relationship between the number of people who experienced reduction in discrepancy and the type of intervention ($\chi^2 (2) = 9.82, p < 0.01$). In order to test if the social intervention had a stronger relationship with reduction in discrepancy when compared to the information inconsistency and baseline groups, we performed a post – hoc test comparing the adjusted residuals and the odds ratios.

Standardized residuals were then used to assess significance of the effect, any residual that lies outside ± 1.96, is considered significant at $p < 0.05$ (Haberman, 1973). The odds ratio on the other hand is used to quantify the strength of the association. Findings showed that the social intervention was significantly related with the number of people who had a reduction in discrepancy (Table 6.7). We also found that the odds of an individual having a lower discrepancy between concerns and behavior after receiving the social intervention is 6.25 times higher than after receiving the information inconsistency intervention and 6.87 times higher than after receiving the baseline intervention. Hypothesis 2a is hence supported.

[1] Gap_1 = concern behavior discrepancy before intervention. Gap_2 = concern behavior discrepancy after intervention.

Table 6.7

Relationship between number of people with a reduced discrepancy and type of intervention.

			Type of Intervention		
			Social – proof (N = 15)	Information In-consistency (N = 13)	Baseline (N = 14)
$Gap_2 < Gap_1$	I	Count	12	4	4
		Adjusted Residual	3.13	-1.47	-1.75
	D	Count	3	9	10
		Adjusted Residual	-3.13	1.47	1.75
Total		Count	15	13	14

I = Increase in concern-behavior discrepancy; D = Decrease in concern-behavior discrepancy.

6.3.4. Hypothesis H2b

Hypothesis 2b states that the number of people who had a reduction in discrepancy between concerns and behaviors will be more in the information inconsistency condition than in the baseline condition. While the chi-square test found a significant relationship between the number of people who had a reduction in discrepancy and the type of intervention, we needed further analysis to identify the difference between each group. We therefore performed post-hoc testing looking at adjusted residuals and the odds ratio to check if the number of people who experienced reduction in concern – behavior discrepancy were significantly higher in the information inconsistency condition when compared to the baseline condition. Table 6.7 shows that the adjusted residuals in the information inconsistency and baseline condition lie within ± *1.96*, therefore we concluded that the information inconsistency and baseline interventions were not significantly related ($p > 0.05$) to the number of people who had a reduction in discrepancy between their concerns and behavior. Further, the odds of an individual having a reduction in concern-behavior discrepancy was only 1.11 times higher after receiving the inconsistency intervention rather than the baseline intervention. Therefore, hypothesis H2b was not supported.

Based on the hypothesis testing we found that the social-proof intervention outperformed both the information inconsistency and baseline interventions in reducing the gap between privacy concerns and disclosure behaviors. Further, the information inconsistency intervention and baseline intervention produced similar effects (i.e. an increased display of privacy sensitive behaviors without a similar change in privacy concerns) in the after intervention stage. Taken together these results make a clear case for the effectiveness of social-proof as a strategy to help individuals better align their information privacy concerns and disclosure behaviors.

We now delve deeper into the cognitive processes guiding the changes described above through a qualitative analysis of the exit interview data in order to answer RQ2.

6.4. Variations in counter-attitudinal disclosure by experiment condition

In recent years, several scholars from different streams of research have attempted to understand and help individuals make online privacy and security decisions (Acquisti et al., 2017). Previous studies examining the discrepancy between information privacy concerns and behavior have suggested several hypotheses to explain disclosure on social network or e-commerce sites. These include contextual disclosure i.e. the influence of situational norms or contextual factors on disclosure decisions (Nissenbaum, 2010; John et. al, 2011) and information asymmetry or the gap in understanding institutional privacy practices (Milne & Culnan, 2002; Hoofnagle et al., 2010). Another theory described in the literature is the privacy calculus theory that views disclosure decision making as a trade-off between the costs and benefits associated with disclosure. The construal level theory (CLT) (Hallam & Zenella, 2017), provided an explanation of online self-disclosure where users minimize the risks of disclosure as abstract or less applicable to them than their peers while viewing the benefits of disclosure as immediate and concrete.

A growing line of research has also been devoted to the use of interventions or nudges to help users make "good" or at least less regretful decisions online (Almuhimedi et al., 2015; Wang et al., 2013; Das et al., 2014). While there is a large body of research on the use of nudges to influence privacy and security behavior (for a review see Acquisti et al., 2017), so far little attention has focused understanding the effect of nudges on general information privacy *concerns*. Our investigation of RQ1 clearly shows that both information privacy concerns as well as disclosure practices are influences by interventions. We therefore use RQ2 to gain a deeper understanding of the effect of different interventions on the cognitive processes guiding the privacy decision making process.

All participants who completed the study were given the option of participating in an exit interview. There were a total of 20 participants with 8 participants from the social-proof condition, 8 from the information inconsistency condition, and 4 participants from the baseline condition who completed the exit interview process. Interviews were conducted over the phone and in-person and each interview lasted approximately 45 minutes. We used the questions described in Appendix G as a guide to interviews with additional questions asked where a clarification or more details were required. In this study, we use the insights gained from an analysis of RQ1 to identify the major themes that the qualitative data could be used to illuminate. Once these themes were identified, all transcriptions were thoroughly read, and participant quotes relevant to specific themes were collected (Appendix K). During the exit interviews, we only had access to a sub-sample of the population, we therefore present the results obtained from this analysis with the caution that it may differ from the overall population results. We now present a detailed analysis of each theme in order to gain a richer and deeper understanding of the different ways in which interventions affected information privacy concerns and disclosure behaviors.

6.4.1. Shift in concerns and behavior

An examination of the quantitative data shows the clear effectiveness of social-proof based interventions in reducing the discrepancy between information privacy concerns and disclosure behaviors. In order to gain a deeper understanding of this phenomenon we asked participants across different conditions about their reactions to the email interventions and whether these emails had an influence on their responses to the IUIPC survey or information disclosure activities.

We found that participants across the three experiment conditions were easily able to remember the information contained in the interventions, however, almost all interviewees (7 out of 8) from the social-proof condition brought up feelings of alarm and worry about the *gap* between their concerns and behaviors while only 2 out of the 8 interviewees in the information inconsistency condition mentioned the gap between privacy concerns and disclosure behaviors. Interviewees in the baseline condition, on the other hand, only discussed their information disclosure in the lab sessions and the influence of interventions on choosing tasks in the second lab session.

During the interviews, we asked participants about their reactions on receiving the email interventions and found that interviewee reactions were very different across the different intervention groups. For instance, P22 who was part of the social-proof experiment group mentioned "*I thought I was much more careful than people in my generation, so being told that I was acting differently really made me think if I was putting everyone at risk because of what information I had shared*". Similarly, participant P16 stated "*It was very scary to be told that the gap between my privacy concerns and behavior is higher than normal. I think privacy is a really important and sensitive topic and I don't want people to think that I don't care about it*." Examining PCI and BI before and after interventions, we found that both these participants showed an increased alignment between privacy concerns and behaviors in the after intervention stage. This implies that

communicating information about majority behaviors created a sense of being left out resulting in an increased motivation to demonstrate aligned privacy concerns and behaviors. According to the social-proof theory (Cialdini & Trost, 1998), individual's attitudes and behavior are often influenced by the attitudes and behavior of others. Research has also found that feelings of membership and belonging can be strong even within a group of unfamiliar others (Walter et al., 2012). Any communication that an individual does not belong within their group can undermine motivation and reduce self-worth (Walton et al., 2012). From interviews with participants from the social-proof condition, we found that the knowledge of acting differently from peers exerted a strong influence on the individuals' own privacy concerns and information disclosure behaviors.

Another participant P1, who had a large discrepancy between their concerns and behaviors before the intervention reported feelings of confusion on seeing the intervention. P1 mentioned "*I always felt I was similar to everyone else, but then I thought about it and I realized I'm probably much more open than other people. I hear everyone being really scared about what will happen to their information and I just don't feel that way*." This social comparison of information privacy concerns was reflected in the after intervention stage where P1 tried to align their privacy behavior and concerns to get closer to their peers.

While participants in the information inconsistency condition did note feeling worried that their privacy concerns and behavior were not aligned, we did not find a strong motivation within participants to reduce this gap. For instance, during the exit interview P8 stated "*It said my information sharing didn't match my comfort level, so I started thinking about what tasks I did and I thought sharing the call log was risky*" Another participant P11 mentioned, "*The email said that my information sharing was higher than my comfort level, I agreed with that result. I think I do tend to give out more information and regret it later, so I think it was accurate*." The information

inconsistency intervention was designed to simply create an awareness of the gap between concerns and behaviors. Both these participants understood that there was a gap between their privacy concerns and disclosure behaviors and they focused on demonstrating privacy sensitive behavior to reduce this gap. Similarly, participants P26 and P41 discussed their worry about information disclosed during lab activities they performed when we asked them to describe their feelings after receiving the intervention. Most participants in the information inconsistency condition did not discuss their general information privacy concerns. In fact, as the analysis in Section 6.2.2 shows, a number of participants had reduced privacy concerns. When we specifically asked interviewees from the information inconsistency group if their concerns over information privacy may have reduced, most interviewees thought it stayed the same. However, this was not reflected in the IUIPC survey scores as there was a noticeable decrease in the concern levels. We interpret these results to mean that while the intervention resulted in participants consciously displaying privacy sensitive behavior, information about the concern-behavior gap resulted in participants subconsciously lowering their privacy concerns.

Participants in the baseline condition received an intervention that simply reminded them of their information disclosure without additional information about their privacy concerns or the gap between their concerns and behaviors (Figure 5.4). Interviewees from the baseline group, unsurprisingly, focused mainly on the tasks they had performed in the lab sessions. None of the four interviewees mentioned their privacy concerns or connected concern and disclosure in any way. For instance, participant P42 mentioned "*I didn't think I had shared too much information but then I remembered sharing a screenshot of the call log so maybe that's what it [email intervention] was referring too.*" Another participant P9 stated "*I remembered uploading a lot of images in the first*

lab session, there was one that I shared my Facebook friends list, another my insurance information and some others. I didn't think about it at the time, but when I got the email I was like yeah I probably did share a lot of information so I was more careful the second time." Both these participants showed an increased privacy sensitivity in terms of behavior in the after intervention phase. However, privacy concerns shifted in different directions for these two participants. Participant P42 showed an increased privacy concern in the after intervention stage and therefore lowered the discrepancy between concerns and behaviors, while P9 showed no difference in privacy concern after intervention, leading to an *increase* in the gap between their privacy concerns and behaviors. While for some participants like P42 simply participating in the study may have caused an increase in privacy concerns, for most participants in the control group there was little or no increase in privacy concerns. Combined with an increase in privacy sensitive behavior, this resulted in increasing rather than reducing the gap between privacy concerns and disclosure behaviors.

Interviewees from the information inconsistency and baseline groups also mentioned feeling "worried" about information privacy but were also unsure about how this concern translated into day to day behavior. For instance, when we asked participant P38 (who scored high on the IUIPC survey) to describe their general information privacy concerns, they reported, "*I know it's something I should worry about, and I do, most of the time…but there's usually so many other things going on that I don't really think about it unless something bad happens*." Another participant P45 mentions "*Of course, I'm concerned about privacy but I'm also concerned about so many other things. Just being worried about privacy should not stop me from going about my daily life right?*" Similarly, other interviewees who mentioned "*thinking a lot*" about their information privacy also seemed to disassociate it from actual behavior.

From this analysis, it is clear when thinking about general concerns over information privacy, most individuals tend to view it as an abstract concept and connecting general privacy concerns with actual disclosure is not an automatic process. Therefore, when attempting to reduce the gap between privacy concerns and disclosure behaviors it is useful to provide individuals with an example on which to model their own behavior.

6.4.2. Interactions with Interventions

A second major theme from the interviews was the usefulness of different aspects of the interventions. In this section, we delve deeper into investigating how participants interacted with different aspects of the interventions and what they liked or disliked about each intervention. We asked participants questions about the overall design of interventions, textual and visual information contained in the intervention, and participants understanding of the interventions. The following sub-sections address each topic that emerged from the thematic analysis of exit interview questions.

Simple Design: Most participants across the three intervention groups had positive feelings about the interventions. Many participants said they liked the integration of images along with the textual information. Participant P16 in social-proof condition states "*The first thing I saw when I opened the email was the image. It was quite a large graphic and seeing a red figure standing apart made me curious about the email. I read the email and the image and text together was really useful to tell me that I am out of the group*" Similarly, P40 (social-proof condition) focused on colors displayed in the visual part of the intervention saying "*The green and red was really effective. Like everyone is green and I am the only red that actually conveyed to me that I was singled out.*" Another participant in the information inconsistency condition mentioned "*I liked the way the im-*

age was in the center of the mail. The arrow in the red part made it clear that I need to be concerned" – P38. Participants in the baseline condition also noted the usefulness of the graph which allowed them to quickly understand that they had disclosed sensitive information.

A number of participants across all conditions also mentioned the simplicity in design and the use of red bolded text to draw attention towards the concern-behavior discrepancy. Participant P45 stated "*The email was simple, it wasn't too cluttered with paras and paras of information. It just said what was going on in 2 – 3 lines that was easy to understand*." Interviewees in the social – proof condition also mentioned the effectiveness of numeric information P22 stated "*I liked seeing where I stood and where everyone else was…it was like a S.A.T score. Seeing the comparison really opened my eyes to how different I was from everyone else*." This implies that P22 was able to leverage both the numeric information and as well as the underlying information highlighting peer concerns and behavior. From the interviews, we found that the most favored elements in the interventions, were the comparative percentages and the dramatic visuals included in the email.

Information Interpretation: We wanted to verify that participants were not just reacting to the images contained in the emails but were actually paying attention to the privacy information contained in the emails. We therefore asked participants what the interventions meant according to them and whether they agreed with the information contained in the interventions. Most interviewees (7 out of 8) in the social-proof condition interpreted the email as saying that they were doing something different from a majority group. We did find a difference in the perception of who this majority group comprised of. P22 interpreted peers to mean people of the same age group while P1 had a more ambiguous definition characterizing peers as "*people around me or people that know me.*" P16 on the other hand interpreted peers as other participants of the study. While the

interpretation of peers varied the feeling of discomfort at being separate from the group was common for these participants. We also asked these interviewees if they agreed with the information contained in the intervention. Most interviewees (7 out of 8) in the social-proof experiment group, mentioned having shifting thoughts about the accuracy of the intervention. For instance, P40 mentions, "*I always thought most people were like me, so I initially didn't agree with the results, but you were looking at the data; so you probably knew better than me. Then, I thought maybe I was unconsciously doing something different.*" Another participant P22 mentioned "*I guess the email was accurate because it had all the numbers, but I didn't think I was behaving differently from everyone else, but maybe I was.*" For both these interviewees, it was hard to believe that they were acting differently from their peers, however, due to contextual factors (participating in research, presence of numeric information) they perceived the information to be accurate and were therefore motivated to model themselves based on peer thoughts and action.

On the other hand, we found that participants in the information inconsistency group focused on the information sharing aspect of the intervention. While interviewees grasped the idea that there was a gap between their concerns and behaviors, for most interviewees (6 out of 8) this meant that they had disclosed sensitive information. While these interviewees did realize that they had disclosed more information than their comfort level, they did not connect this comfort level to the IUIPC survey. We also found that not all interviewees in the information inconsistency condition agreed with the information contained in the intervention. P38 mentions "*It said that I share more information than I want to, but I'm pretty careful about my data sharing practices so I don't think that result was right.*" Other interviewees however, did find the information accurate and reported reflecting on the activities performed in the lab session on seeing the email "*It [The intervention] said my information sharing didn't match my comfort level, so I started thinking about*

what tasks I did and I thought sharing the call log was risky" – P8. Some participants in the information inconsistency condition (3 out of 8) also reported wanting additional information about how the discrepancy between concerns and behaviors was calculated.

Interviewees in the baseline condition found the email intervention clear, succinct and easy to understand. The baseline condition simply alerts participants about the number of tasks they have performed and the information they disclosed (Figure 5.4). All interviewees (4 out of 4) were easily able to understand this information and also agreed with it.

Chapter 7: Discussion

This chapter of the dissertation will summarize key findings, provide a discussion of contributions and implications of this work, identify the limitations and suggests directions for future studies.

7.1. Summary of Findings

Privacy literature has established strong support for the contradictory relationship between privacy concerns and disclosure behavior (Barnes, 2006; Norberg et al., 2007). Scholars investigating the privacy paradox have put forward a number of theoretical and empirical explanations to explain the privacy paradox (Acquisti and Grossklags 2005; Brandimarte et al., 2013; Hallam & Zenella, 2017). Researchers have also reviewed the literature on the privacy paradox in an attempt to provide a systematic understanding of the various factors affecting the discrepancy between information privacy concerns and behaviors (Barth & de Jong 2017; Kokolakis 2017). Prior literature has linked the disclosure of sensitive information contrary to information privacy concerns to a lack of risk awareness, presence of contextual cues, and a skewed comparison of the risks and benefits of disclosure (Acquisti & Grossklags, 2005; Bashir et al., 2015; Hallam & Zenella, 2017). However, there is a lack of research identifying strategies to correct this gap between privacy concerns and behaviors. In this work, we test the effectiveness of social-proof based interventions on reducing the discrepancy between information privacy concerns and disclosure behaviors and explore the different ways in which participants interacted with the interventions.

When measuring the gap between privacy concerns and disclosure behavior, we found that while most participants had a higher disclosure level than their concerns, for some participants this gap was *reversed*. That is, some participants were more conservative in their information disclosure while reporting lower concern levels. While it could be argued that an increased privacy sensitive behavior (even with lower privacy concerns) is a good thing, we would posit that even pri-

vacy protective counter-attitudinal behavior could have negative consequences. For instance, lowered privacy scores could be perceived as evidence that people in general no longer care about protecting their personal information. While this notion would be false, it could still be used by corporations to block privacy protective legislation or push for legislation allowing them further access to user data. Lowered levels of privacy concerns could convey an impression that an individual is likely to share sensitive personal, medical, or financial information. Even though, this impression would not reflect their actual behavior, it could result in the individual being increasingly targeted by phishing email, spam calls, and even identity theft calls. We would therefore argue that an increased alignment of privacy concerns and behaviors would allow individuals, corporations, and governments to more accurately understand overall information privacy needs and design mechanisms that can match these needs.

The goal of this work was to compare social-proof based interventions with an information inconsistency and baseline intervention to check which condition resulted in privacy concerns and disclosure behaviors being more aligned. In this work, we first operationalized and measured privacy concerns and disclosure behavior and then compared the gap between the two before and after interventions. The changes in concern-behavior gap across different experimental conditions were then examined through hypotheses H1, H2 and RQ1. The results of our study suggested that the social-proof intervention was the most effective in aligning individual privacy concerns and disclosure behaviors. By bringing awareness of peer concern-behavior alignment to the forefront, participants were motivated to pay attention to and reduce the discrepancies between their own information privacy concerns and disclosure behaviors. Table 7.1 shows results of the hypotheses and research questions.

The second RQ was used to gain a better understanding of the role played by the interventions in aligning privacy concerns and disclosure behavior. We conducted semi-structured interviews with 20 participants and performed a thematic analysis to better understand why the social-proof intervention was the most effective. We found that while both the social-proof and information inconsistency conditions created an awareness of concern-behavior gap, the added knowledge that this gap was *not commonly* found among peers motivated social-proof intervention participants to re-evaluate their concerns and behaviors and therefore reduce the discrepancy between their own privacy concerns and disclosure behaviors.

Table 7.1

Results of Hypotheses/Research Questions

Hypothesis and Research Questions	**Testing**	**Significance**
H1a: People receiving the social nudge will demonstrate privacy behaviors more aligned with their attitudes than other groups.	Significant	$p < 0.05$
H1b: People receiving the inconsistency nudge will demonstrate privacy behaviors more aligned with their attitudes than people in the baseline group.	**Not** Significant	$p > 0.05$
H2a: The number of people who had a reduction in discrepancy between attitudes and behaviors will be more in the social condition compared to other groups.	Significant	$p < 0.05$
H2b: The number of people who had a reduction in discrepancy between attitudes and behaviors will be more in the inconsistency condition compared to the baseline group.	**Not** Significant	$p > 0.05$
RQ1: Can social-proof based nudges be an effective way to reduce the privacy paradox?		
RQ2: How do the interventions change counter-attitudinal disclosure across the three experiment groups?		

In this work we use the concepts of nudges, which refers to the idea of designs that guide rather than force users to make privacy decisions in certain directions (Thaler & Sunstein, 2009). Previous research using nudges to improve information privacy management has focused on the use of nudges to guide disclosure behavior, for example, interventions to increase interactions with pri-

vacy settings (Das et al., 2014) or interventions to improve location privacy management on Facebook (Ghosh & Singh, *in preparation*). While this approach works in improving privacy behaviors, it does not take into account the privacy concerns of individuals. Nudging a user with a low concern for privacy towards privacy sensitive behaviors may result in the individual losing out on the benefits gained by information disclosure (e.g. discounts or building a large network). In this work, we therefore, first measure information privacy concerns and disclosure behaviors and then design nudges to guide users towards their own preferences.

A second question is however, whether this is enough? That is, is simply making an individual aware of the discrepancy between their privacy concerns and behavior enough to create an alignment between the two. Prior research investigating concern-behavior discrepancy in the case of mobile app installation has found that users were more inclined to refuse to install mobile apps when they were shown an interface comparing their privacy concerns with the privacy risk they would be exposed to by installing the app (Jackson & Wang, 2018). Jackson and Wang's (2018) work however, does not investigate whether this discrepancy existed before the intervention was given to users. It presumes the presence of a paradoxical relationship and implements an intervention to reduce the misalignment. The measurement of privacy concerns and behaviors before and after the intervention in our study, allowed us to see the effect of interventions on changes in both concerns and behavior. In fact, we found adjustment of privacy concern to be an important strategy used by participants to reduce the perceived gap between their concerns and behavior. We found that the information inconsistency intervention, i.e. where participants were informed about the gap between their concerns and behavior, was not as effective as the social-proof intervention, where participants were told that this gap separated them from a majority, when reducing the concern – behavior gap.

Additionally, the qualitative analyses allowed us to understand how these interventions affected the participants' knowledge on their concern-behavior gap and why the social-proof intervention was effective in reducing the gap. We found that for most participants the knowledge of having a gap between their concerns and behaviors was not as alarming as realizing that this gap separated them from their peers. The visuals and numeric information contained in the social – proof intervention conveyed a strong impression of isolation from the group that in turn motivated participants to adopt peer behavior and reduce the gap between their own concerns and behaviors.

7.2. Discussion for Hypotheses and Research Question

This section of the chapter will focus on interpreting findings of the current project regarding the association between information privacy concerns and behavior, the effectiveness of social-proof intervention, and information disclosure in online SNS drawing on the cognitive dissonance theory (CDT) as a theoretical framework.

Social-proof interventions (RQ1, H1&2): Findings from this work show the effectiveness of the social-proof based interventions in reducing the privacy paradox. Comparing the design of the three interventions (Figures 5.2 – 5.4), we see that the social-proof and information inconsistency interventions both informed participants about the discrepancy between their information privacy concerns and disclosure behavior. However, the social-proof intervention also contained additional information about peer behavior, which led participants to believe that it was "normal" to have an alignment between concerns and behaviors. This perception was especially brought out during the qualitative interviews (Section 6.4) when multiple interviewees from the social-proof condition mentioned experiencing shock or surprise upon realizing that they were separate from a majority group. The interviews also brought out nuances about how participants shifted their concerns and behaviors to be more aligned. Some participants, for instance, engaged in cognitive work to think

more about their general privacy concerns resulting in them lowering their concerns to match their disclosure behavior. For other participants, it was important to increase their privacy protective behaviors so that it matched their privacy concerns.

According to theories of social influence, the perception of actions that a majority group approves of become the norm within that group (Azjen 1991). Actions that follow or break these norms are associated with potential social rewards or risks of membership or non-membership within the group. Cialdini and Trost (1998) state that norms guiding daily activities have evolved from behaviors that are performed and reinforced through repeated interactions with others. These norms then become preferred responses to certain situations, in the context of privacy behaviors knowledge of close friends changing privacy settings acts as a trigger for an individual to perform similar actions. Once these norms are established members of the social group discourage any deviation by stating what others "should do." Research examining inter-group dynamics has found that people tend to use minimal cues of social connections (e.g. being arbitrarily placed in the same group) as a way to create group membership and collectively work toward the goals and interests of the whole group (Walton et al., 2012). This implies that while long-standing trusted ties can influence an individual's choices, the influence exerted by a minimal group cannot be discounted. In our study, participants did not know each other beforehand, however, the knowledge of being part of an experiment group could be enough to create a feeling of group membership. The behavior of other group members therefore, became an important standard against which participants compared their own thoughts and actions. Dissonance was caused when participants realized that they were different from the majority group. In order to reduce dissonance, individuals attempted to reconcile their thoughts and actions with the majority.

The interviews also helped us understand the need for individuals to be perceived as part of a larger group. For instance, participants from the social-proof condition were worried about other people getting the wrong impression about their information privacy concerns. This implies that individuals often try to present themselves in a way that matches how they want to be perceived by other people. Goffman's (1959) seminal work on impression management explains that individuals express themselves differently based on the impression they want to create within an audience. For example, we may look or behave quite differently in a business meeting than at dinner with close friends based on the impression we want to create within each group (Goffman, 1959). This impression management theory can be used to understand the problem of managing disclosure in order to maintain membership within a group. An inaccurate self-presentation results in group members gaining a different impression of the individual rather than one they wanted to produce. In the context of our study, when participants received the notification that the mismatch between their privacy concerns and behaviors was not commonly observed among their peers, they felt worried about the impression that was being conveyed to a larger group. As one participant mentioned, "*privacy is an important and sensitive topic and I don't want people to think that I don't care about it*" – P16. Another participant P34 mentioned, "*I have a Master's in Computer Sciences, I know all about information tracking and big data…in fact I mostly advise my friends on what to do and what not to do so of course I don't want people to think that I don't know what I'm doing*." Further, P34 mentions advising "friends" about information privacy. Here, P34 was worried about their competency as an information technology expert being questioned because of their concern-behavior gap. Hence mimicking peer behavior allowed P34 to maintain the impression of being someone who not only cares about privacy but is also knowledgeable about the subject.

The issue of impression management is further compounded by the notion of "networked privacy" (boyd, 2012) i.e. the idea that privacy concerns and behaviors are not only influenced by the individual's perceptions, actions, and beliefs, but also by interpersonal and group – level actions and perceptions. The notion of contextual integrity (Nissenbaum, 2004) also addresses privacy management at a network level. A central tenet of this framework is that maintaining privacy occurs at a collective rather than individual level and the norms of information protection or information disclosure are collectively shared, understood, and practiced by different people in the network. This implies that engaging in "risky" disclosure impacts not just an individuals privacy but also affects the collective privacy of the network. For participants in the social-proof condition, this notion of risking collective privacy was brought to the forefront by the intervention. For instance, P22 mentions "*putting everyone at risk*" by their own information disclosure while P34 talks about "*advising friends on what to do or what not to do*". This implies that along with the fear of being excluded from the group, the knowledge of acting and thinking differently from the majority also caused participants to think about the effect they had on the network as a whole. While they wanted to belong to the network, they also accepted the responsibility of maintaining the information privacy norms of the network. Therefore, the knowledge that they were putting collective privacy at risk provided an increased stimulus to fix their own concern – behavior gap.

Information Inconsistency Intervention: In this dissertation, we tested the use of an information inconsistency intervention to reduce the concern-behavior gap. The idea behind this intervention was to create dissonance by providing information about the gap between information privacy concerns and behaviors. According to the CDT when an individual holds two conflicting elements

of knowledge (e.g. knowledge of their information privacy concerns and knowledge of their disclosure behavior) a state of dissonance is created. Since a state of dissonance causes discomfort and individuals are motivated to engage in "psychological work" so as to reduce this dissonance (Festinger 1957). We expect creating an awareness of the discrepancy between information privacy concerns and disclosure behavior will result in dissonance and cause individuals to resolve this dissonance by demonstrating behaviors more aligned with their information privacy concerns. However, Table 6.6 shows that the information inconsistency intervention did not have a significant effect of the reduction of discrepancy. A further investigation into the changes in privacy concerns and behaviors finds that while a majority of participants in the information inconsistency condition did demonstrate more privacy sensitive behavior after the intervention, they also shifted their privacy concerns to be *lower concerns* over information privacy. Participants in the baseline condition demonstrated similarly heightened privacy behavior while having almost no change in their privacy concerns (Tables 6.2, 6.3) resulting in an increase in the discrepancy between concerns and behaviors.

According to the CDT, when an individual experiences dissonance by gaining two or more conflicting elements of information, dissonance is resolved by supporting the information most resistant to change (Festinger, 1957). In the context of our study, when an individual received information that their disclosure behavior did not match their information privacy concerns, they might have experienced discomfort caused by these two dissonant knowledge elements. However, changing knowledge of their disclosure behavior (e.g. denying that they had disclosed sensitive personal information) was not possible as the activity had been completed. In order to resolve dissonance participants, therefore, shifted their privacy concerns to align with this heightened disclosure resulting in a lower concern for privacy.

Participants in both the social-proof and information inconsistency condition experienced dissonance brought by the knowledge of a gap between their concerns and behaviors. However, participants in the social-proof condition had the additional knowledge of peer behavior to help resolve this dissonance. These participants therefore shifted concerns as well as behavior closer to each other. A quantitative analysis of participants in the information inconsistency condition shows a lower concerns for privacy with an increased privacy sensitive behavior. From interviews with participants in the information inconsistency condition, we found that the intervention did create an awareness of the concern-behavior gap, but did not give an additional norm-inducing nudge to reduce this gap. Some participants viewed the information inconsistency intervention as simply increasing their knowledge about information privacy without finding the necessity to take any action. For instance, P11 mentioned that they understood and agreed that they tended to share more information than they were comfortable with. An analysis of the after intervention session for this participant shows an increase in privacy sensitive behavior but a lower concern for privacy. According to the CDT, these lowered concerns could be due to a post-hoc rationalization where once participants were told that their concerns and behavior did not match each other, they lowered their concerns to reflect information disclosure. On the other hand, the information about disclosure habits contained in the intervention, caused participants to become sensitized towards information disclosure. This resulted in participants consciously choosing tasks that protected their information privacy and thus displayed more privacy sensitive behavior.

Baseline Intervention: Participants in the baseline intervention also showed similar results in terms of privacy behaviors. As seen in Figure 5.4, participants in the baseline condition were only provided information about their disclosure behaviors without giving them any information about privacy concerns or a misalignment between privacy concerns and behaviors. These participants

therefore may not have experienced any dissonance and simply focused on information disclosure behaviors. A quantitative analysis of the after intervention session showed that participants in the baseline condition had an increase in privacy sensitive behaviors (Table 6.4). Interestingly, there was a slight increase (0.06) in average privacy concerns among participants in the baseline condition in the after intervention session (Table 6.3). From qualitative interviews we gathered that for some participants just the experience of participating in the study caused them to think more about information privacy resulting increased privacy sensitivity. Participant P9 mentioned, "*I was thinking more about privacy after the email, initially I was thinking more about the financial aspect but after getting emails and the lab session I started thinking more about what information I'm sharing and who is looking at it*." For participant P9 this heightened focus on privacy resulted in a small increase in their privacy concerns from 4.11 to 4.22 out of 5. However, their privacy sensitive behavior (i.e. number of tasks) had a much greater increase from 11 to 20 resulting in an increase in the discrepancy between their information privacy concerns and disclosure behavior.

From these results, it is clear that in order to align privacy concerns and disclosure behaviors, it is not enough to create an awareness of the discrepancy between concerns and behaviors. Individuals require more information about normative behaviors (in this case aligning concerns and attitudes) in order to make the effort to reduce this discrepancy. An argument could be made that increase in privacy sensitive behaviors is a worthy and desirable outcome that is more valuable than reducing the misalignment between information privacy concerns and behaviors. However, this argument implies that disclosure behaviors can be examined as separate from concerns over information privacy. Understanding the different factors that influence individual privacy concerns is of critical importance in digital contexts where data collection policies often rely upon user privacy concerns. In these contexts, simply guiding individuals towards privacy sensitive behavior

rather than behavior that reflects information privacy concerns would result in data policies that reflect a lowered concern for privacy contrary to what individuals actually want. Encouraging individuals to align their information privacy concerns and behaviors might also help direct individuals away from shortcut heuristic decision making by guiding disclosure towards their own preferences. At the same time, increased awareness of the relationship between privacy concerns and disclosure behaviors could also lead to a higher (more realistic) valuation of personal information forcing institutions to improve their offers involving the exchange of personal information.

7.3. Ethical Considerations of Nudging

In this section, we focus on some ethical considerations that may arise from the use of interventions or nudges. Arguments have been made that the use of interventions, reduces an individual's freedom to choose. However, we would argue that any system or design influences the individual in some way. As Thaler and Sunstein's (2009) work on interventions highlights, at some point the user has to make a choice between the available options. They use the example of a cafeteria, where the order in which different food items are organized influences a person's choice. The food however, must be set out. This implies that some sort of a nudge of is inevitable. Any system irrespective of whether or not it was designed to influence an individual's behavior, will impact how the person interacts with a system. Thaler and Sunstein (2009) define nudging as "*any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives*." Therefore, a nudging approach does not force people to do things, but rather provides or clarifies information about the different options available to them. We use a nudging approach because we recognize the difficulty people have in effectively comparing information privacy concerns and disclosure behaviors, and we seek to help users with reducing the gap between the two.

Unintended Consequences: The possibility of unintended consequences has already been studied in HCI literature (Brown, Weilenmann, McMillan, & Lampinen, 2016), however interventions are especially fraught with ethical repercussions (Acquisti et al., 2017). It is therefore necessary to be especially careful about possible unintended consequences when designing an intervention. For instance, a nudge designed to increase a person's awareness of the different audiences that had access to their location disclosure resulted in some participants culling their Facebook friend networks (Ghosh & Singh, in preparation). In the current study, we found that participants in the social – proof condition felt scared or isolated when they were informed that the gap between their privacy concerns and behavior separated them from a majority. While we intended to create an awareness of divergent behavior, we did not intend for participants to experience negative emotions due to the intervention. While this may not seem significant at an individual level, viewed over time and across a large sample population, an intervention using social-proof could also result in an increased perception of isolation. While none of the participants mentioned lasting feelings of loneliness, it reminds us to be mindful of any unintended social effects of our interventions in future. It has also been argued that an intervention maybe seen as stigmatizing certain choices. For instance, an intervention designed to improve an individual's privacy sensitivity may unintentionally pass judgement about people who are unconcerned with privacy. While these effects are, by definition, difficult to foresee or predict, it is important for system designers and researchers to be aware of the possibilities of unintended side-effects and always consider the effects of the intervention at a. individual as well as societal level.

Individual Choices: When designing an intervention, it is important to respect the individual's right to choose. Even if there are great benefits to be gained from the intervention (or terrible consequences from ignoring it), it must remain secondary to the individuals' decision. If a person

decides to disregard a nudge, we must assume that they have a reason for doing so. For example, it may be appropriate to warn social media users that their information disclosure habits may be exposing them to theft or harassment, but an intervention that limits their ability to interact or build connections with others would be unethical. Intervention designers and researcher should also consider making the intervention more transparent and customizable. For instance, letting the individual decide under what circumstances or particular time frames the intervention will be triggered or allowing the individual to learn more about why they received the intervention.

Use of Deception in Nudging: Nudges should also be designed to respect the user's expectation of truthful information. Glaeser (2005) recommends thinking carefully about mis-representing information and only using deception to nudge the user when there is a strong rationale for doing so. For example, Kumaraguru et al., (2007) sent fake phishing emails to users in order to nudge them to adopt stronger security settings, which was deemed permissible given increasing concerns over phishing. This approach has since been adopted by many organizations to teach users to better protect themselves from phishing attacks (Gartner Group, 2014).

The social-proof intervention told individuals that their concern-behavior gap was more than their peers *without* having a measurement of peer concerns and behaviors. The widening gap between privacy concerns and behavior has been identified as a worrisome phenomenon by many research scholars (Jackson & Wang, 2018; Acquisti et al., 2017; Brandimarte et al., 2013). Hence, similar to the abovementioned works we considered the use of deception permissible in order to encourage users to reduce this worrisome privacy gap.

7.4. Limitations

As with any research study, this work exploring the use of social-proof interventions in reducing the discrepancy between information privacy concerns and disclosure behavior is also not free

from limitations. First, we cannot be sure that the disclosure behavior demonstrated in the study can be generalized to other settings. While we use deception to inform participants that information collected will be used by non-academic third-parties and no confidentiality protections are granted for any information provided, it is possible that performing activities within a University may give subjects confidence that their information would not be misused and caused them to disclose more information than they would in real-world settings.

Secondly, even though we cautioned subjects against falsifying data and provide them the option of not performing a disclosure task when they were not comfortable supplying personal answers, there exists the possibility that falsification may occur. We did include certain checks like inspecting data for blatant misrepresentations (e.g. special characters or extra spaces) or missing descriptive length, and asked participants if they falsified information during the exit interviews. However, it is possible that falsification occurred resulting in a mis-calculation of disclosure behaviors.

Finally, our interventions were designed to notify participants after the activities were completed. While some literature has made the case for just-in-time notifications (Acquisti et al., 2017; Wang et al., 2014), the results from this work show a clear effect of the nudges on individual privacy concerns and behavior. The design of the interventions is flexible enough to be adapted to runtime settings. We believe that the concept of social-proof interventions can be adapted to apriori, just-in-time, and post activity notifications and complement each other.

7.5. Implications of the study

In this study, we investigated the privacy paradox phenomenon and examined ways to reduce it. Previous research has documented the persistence of this privacy paradox as well as proposed multiple explanations for its existence (Norberg et al., 2007; Acquisti et al., 2015; Debatin et al., 2009). However, significantly lesser research has been devoted to reducing this paradox. In this

work, we tested the effectiveness of social-proof based interventions in reducing the privacy paradox. The privacy paradox occurs in situations where individual's information disclosure does not match their privacy attitudes. These inconsistencies can sometime lead to the disclosure of sensitive information resulting in troublesome or regrettable incidents. For instance, while people can reap the benefits of using discounts received based on disclosing personal information, they often do not know or underestimate the full implications of disclosing their personal information to companies (Ghosh & Singh, 2017; Tufecki, 2008). Additionally, while people are aware that they are being tracked, they rarely have a clear and accurate idea of what information other people (within or outside their network), SNS like Facebook, or government agencies have about them or how that information will be used. As long as individuals remain unaware of these factors, they are likely to experience uncertainty about their disclosure behavior and therefore demonstrate behavior unaligned with their concerns.

The various factors influencing concerns over information privacy, variety in motivations of information disclosure, effects of interventions on re-evaluating concerns and behavior, and the different ways in which participants interacted with the interventions suggest that there exists a complex and non-linear relationship between information privacy concerns and disclosure behaviors. We find that measuring changes in information privacy concerns, occurring over time and influenced by multiple contextual factors, is vitally important for an understanding of the privacy paradox. Previous studies examining the privacy paradox have either used a post-experimental questionnaire about privacy concerns (Beresford et al., 2012) or used participant valuations of personally identifiable information as a measure of privacy concerns (Carrascal et al., 2013; Spiekermann et al., 2001). Another work using just-in-time interventions to reduce the privacy paradox measures general privacy concerns about mobile app installation and compares it to installation

behavior after different interventions (Jackson & Wang, 2018). In their work, Jackson and Wang (2018) also measure the privacy concerns only in the before intervention stage. While these works did find evidence for the privacy paradox, privacy concerns or preferences were only measured at a *single point* in the study. Multiple instances of information disclosure under different contexts (e.g. entering information into a form versus disclosing information to a chat bot with an avatar) were compared to these privacy concerns. However, as the results from Section 6.3 point out information privacy concerns are just as likely to be influenced by external and contextual factors as disclosure behavior. It is therefore important for future research in privacy paradox to take into account this possibility of shifting privacy concerns. An increased focus on the change in privacy concerns at different points in the research can significantly help researchers and scholars pin down the factors causing gap between privacy concerns and disclosure behaviors and identify better and more efficient ways of reducing this gap.

A better understanding of the privacy paradox and ways to reduce it could also help create a new perspective on the legal and ethical framework of information privacy. For instance, information privacy policy is often dependent on large-scale privacy concern surveys collected from a general population. If these concerns do not reflect actual disclosure behavior, then any policies that are based on these surveys would not account for actual disclosure behavior. Fixing this gap between concerns and behaviors can have significant benefits for organizations. Information disclosure has an "opportunity cost" associated with it. In economic terms an "opportunity cost" is the worth of the alternative option that could replace the activity under consideration. In case of the privacy paradox, an individual might incur a certain opportunity cost (cognitive effort + time spent) by never disclosing any personal information on the internet. They may not receive many discount coupons, may not be able to use many apps, and not be introduced to many potential

social connections. However, if this person did not feel strongly about information privacy and knew the level of their privacy concerns, they could avoid this cost incurred by displaying privacy protective behavior that better aligns with their information privacy concerns. Opportunity costs can impact businesses as well as individuals. An inaccurate estimation of the willingness to disclose personal information based on artificially inflated privacy concerns can result in businesses losing out on the opportunity to provide personalized services to their customers.

A misalignment resulting in more information sharing than one is comfortable can also result in negative outcomes for organizations. Building and maintaining customer trust can be of significant value for an organization. If an organization bases its understanding of individual privacy concerns simply based on information users are willing to share on their websites and not try to understand their concerns and attitudes, this might lead to negative outcomes in the future. For instance, this may lead to an erosion of trust once the individual realizes that any personal information they have shared will be used for unwanted marketing phone calls or past purchase orders will serve as input for price discrimination. Ultimately, a lack of trust might result in the user refusing to use the service completely. An increased alignment of privacy attitudes and behaviors would allow individuals' to more accurately value their personal information and use SNS settings and tools in ways that help them protect their privacy while taking advantage of the opportunities provided by these networks.

7.6. Conclusion

As more and more of our daily activities and interactions are converted into online contexts, a considerable proportion of individuals tend to disclose sensitive personal information contrary to individual privacy concerns. Our study of interventions to reduce the privacy paradox showed that individual privacy concerns and disclosure habits are influenced by a number of subjective and

contextual factors that impede privacy decision making. This counter-attitudinal disclosure sometimes carries significant consequences, such as identity theft, loss of trust, and loss of friendships. Drawing on behavioral and social influence research, we designed three interventions that attempted to guide individuals to better align their privacy concerns and disclosures. In this study, we used a relatively simple set of information disclosure activities as a measure of disclosure behavior, however, in real world settings smart sensors (e.g. mobile phones, wearable devices, etc.) have the ability to measure information disclosure in much more nuanced settings and in a much larger scale. These devices can also be used to collect information on individual privacy concerns. In our study, the social-group interventions used deceptive information telling participants that their concern-behavior gap separated them from a majority of their peers. When implementing the social-group intervention as a mobile or smart device app notifying individuals' about paradoxical behavior, the app designers would need to use similar messaging i.e. use a deceptive message about peer concerns and behaviors to help participants align their own concerns and behaviors. While we believe that interventions must provide honest and truthful feedback to users, in this case, we would argue that the risk of misaligned concerns and behaviors (regretful posts, erosion of trust, etc.) balances the need for deceptive messaging.

The results from our study suggested that social-proof based interventions can be a powerful mechanism to help some people avoid counter-attitudinal disclosure. Although we use an online web form to simulate disclosure, this idea of social-proof based interventions can be extended to social network sites services such as Facebook or Twitter, or to other types of e-commerce, location sharing, and smart phone applications. Finally, we advocate the social nudging approach to researchers, designers, and policy-makers to help people's privacy decision-making.

References

Abdi, H. (2007). Z-scores. *Encyclopedia of measurement and statistics*, 3, 1055-1058.

Acquisti, A. (2009). Nudging privacy: The behavioral economics of personal information. *IEEE security & privacy*, 7(6), 82-85.

Acquisti, A., Adjerid, I., Balebako, R., Brandimarte, L., Cranor, L. F., Komanduri, S., ... & Wang, Y. (2017). Nudges for privacy and security: understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3), 44.

Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509-514.

Acquisti, A., & Gross, R. (2006). Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies* (pp. 36-58). Springer, Berlin, Heidelberg.

Acquisti, A., & Grossklags, J. (2005). Privacy and rationality in individual decision making. *IEEE security & privacy*, 3(1), 26-33.

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., ... & Agarwal, Y. (2015). Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (pp. 787-796). ACM.

Aral, S., & Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, 1215842.

Aronson, E., & Mills, J. (1959). The effect of severity of initiation on liking for a group. The *Journal of Abnormal and Social Psychology*, 59(2), 177.

Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2):179–211.

Baek, Y. M., Kim, E. M., & Bae, Y. (2014). My privacy is okay, but theirs is endangered: Why comparative optimism matters in online privacy concerns. *Computers in Human Behavior*, 31, 48-56.

Barnes, S. B. (2006). A privacy paradox: Social networking in the United States. *First Monday*, 11(9).

Bartsch, M., & Dienlin, T. (2016). Control your Facebook: An analysis of online privacy literacy. *Computers in Human Behavior*, 56, 147-154.

Bashir, M., Hayes, C., Lambert, A. D., & Kesan, J. P. (2015). Online privacy and informed consent: The dilemma of information asymmetry. *Proceedings of the Association for Information Science and Technology*, 52(1), 1-10.

Beresford, A. R., Kübler, D., & Preibusch, S. (2012). Unwillingness to pay for privacy: A field experiment. *Economics letters*, 117(1), 25-27.

Besmer, A., Watson, J., & Lipford, H. R. (2010, July). The impact of social navigation on privacy policy configuration. In *Proceedings of the Sixth Symposium on Usable Privacy and Security* (p. 7). ACM.

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415), 295.

Bougie, R., Pieters, R., & Zeelenberg, M. (2003). Angry customers don't come back, they get back: The experience and behavioral implications of anger and dissatisfaction in services. *Journal of the Academy of Marketing Science*, 31 (4), 377-693

boyd, D. (2012). Networked privacy. *Surveillance & society*, 10(3/4), 348.

Brandimarte, L., Acquisti, A., & Loewenstein, G. (2013). Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3), 340-347.

Braun, V., & Clarke, V. (2006). "Using Thematic Analysis in Psychology," *Qualitative Research in Psychology* 3(2), pp. 77-101.

Brehm, J. W. (1956). Post – decision changes in the desirability of alternatives. *Journal of Abnormal Social Psychology*, 52, 384–389. doi: 10.1037/h0041006

Brown, B., Weilenmann, A., McMillan, D., & Lampinen, A. (2016). Five provocations for ethical HCI research. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 852-863).

Burke, M., Marlow, C., & Lento, T. (2009, April). Feed me: motivating newcomer contribution in social network sites. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 945-954). ACM.

Carrascal, J. P., Riederer, C., Erramilli, V., Cherubini, M., & de Oliveira, R. (2013). Your browsing behavior for a big mac: Economics of personal information online. In *Proceedings of the 22nd international conference on World Wide Web* (pp. 189-200). ACM.

Christakis, N. A., & Fowler, J. H. (2007). The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4), 370-379.

Christakis, N. A., & Fowler, J. H. (2008). The collective dynamics of smoking in a large social network. *New England journal of medicine*, 358(21), 2249-2258.

Cialdini, R. B. (2001). *Influence: Science and practice* (Vol. 4). Boston, MA: Pearson education.

Cialdini, R. B., & Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. In *The Handbook of Social Psychology*, ed. DT Gilbert, ST Fiske, G Lindzey, 2:151–92. Boston: McGraw-Hill. 4th ed.

Cresswell, J. W., & Plano Clark, V. L. (2011). *Designing and Conducting mixed method research* (2nd ed.). Thousand Oaks, CA: Sage.

Cronbach, L. (1951). Coefficient alpha and the internal structure of tests. *Psychomerika*.16:297-334. 10.1007/BF02310555

Curtis, S., Gesler, W., Smith, G., & Washburn, S. (2000). Approaches to sampling and case selection in qualitative research: Examples in the geography of health. *Social Science and Medicine*, 50(7–8), 1001–1014.

Das, S., Kramer, A. D., Dabbish, L. A., & Hong, J. I. (2014). Increasing security sensitivity with social-proof: A large-scale experimental confirmation. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security* (pp. 739-749). ACM.

Debatin, B., Lovejoy, J. P., Horn, A. K., & Hughes, B. N. (2009). Facebook and online privacy: Attitudes, behaviors, and unintended consequences. *Journal of computer-mediated communication*, 15(1), 83-108.

Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information systems research*, 17(1), 61-80.

Dwyer, C., Hiltz, S., & Passerini, K. (2007). Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace. *AMCIS 2007 proceedings*, 339.

Elliot, A.J., & Devine, P.G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology*, 67, 382–394.

Ellison, N. B., Steinfield, C., & Lampe, C. (2007). The benefits of Facebook "friends:" Social capital and college students' use of online social network sites. *Journal of computer-mediated communication*, 12(4), 1143-1168.

Festinger, L. (1957). A theory of cognitive dissonance. Stanford, CA: Stanford University Press.

Festinger, L., Riecken, H., & Schachter, S. (1956). *When prophecy fails*. Minneapolis: University of Minnesota Press.

Forget, A., Chiasson, S., Van Oorschot, P. C., & Biddle, R. (2008). Improving text passwords through persuasion. In *Proceedings of the 4th symposium on Usable privacy and security* (pp. 1-12). ACM.

Gambino, A., Kim, J., Sundar, S. S., Ge, J., & Rosson, M. B. (2016). User disbelief in privacy paradox: Heuristics that determine disclosure. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (pp. 2837-2843). ACM.

Gartner Group. (2014). Magic Quadrant for Security Awareness Computer-based Training Vendors. *Technical Report. Gartner Group*.

George, J. F., (2004). The theory of planned behavior and internet purchasing. *Internet research*, 14(3):198–212.

Ghosh, I., & Singh, V. (2017). Using cognitive dissonance theory to understand privacy behavior. *Proceedings of the Association for Information Science and Technology*, 54(1), 679-681.

Ghosh, I., & Singh, V. (*Under Review*). *"Not all my friends are friends": Audience-group based nudges for managing location privacy*.

Ghosh, I., & Singh, V. (2018). Phones, privacy, and predictions. *Online Information Review*. https://doi.org/10.1108/OIR-03-2018-0112

Glaeser, E. L. (2005). *Paternalism and psychology* (No. w11789). National Bureau of Economic Research..

Goffman, E. (1967). *On face-work*. Interaction ritual, (pp. 5–45). New York: Doubleday Anchor.

Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of consumer Research*, 35(3), 472-482.

Haberman, S. J. (1973). The analysis of residuals in cross-classified tables. *Biometrics*, 205-220.

Hallam, C., & Zanella, G. (2017). Online self-disclosure: The privacy paradox explained as a temporally discounted balance between concerns and rewards. *Computers in Human Behavior*, 68, 217-227.

Hargittai, E., & Litt, E. (2013). New strategies for employment? Internet skills and online privacy practices during people's job search. *IEEE security & privacy*, 11(3), 38-45.

Harmon-Jones, E. (2000). Cognitive dissonance and experienced negative affect: Evidence that dissonance increases experienced negative affect even in the absence of aversive consequences. *Personality and Social Psychology Bulletin*, 26, 1490–1501.

Harmon-Jones, E. E., & Mills, J. E. (1999). Cognitive dissonance: Progress on a pivotal theory in social psychology. In *Scientific Conferences Program, 1997, U Texas, Arlington, TX, US; This volume is based on papers presented at a 2-day conference at the University of Texas at Arlington, winter 1997.*. American Psychological Association.

Hoofnagle, C. J., King, J., Li, S., & Turow, J. (2010). How different are young adults from older adults when it comes to information privacy attitudes and policies? Berkeley, CA: University of California, Berkeley. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1589864

Horne, D. R., and Horne, D.A. (1998). Domains of Privacy: Toward an Understanding of Underlying Factors. *Presented at the Direct Marketing Educators' Conference*, October 11. San Francisco, CA.

Horne, D. R., Norberg, P. A., & Cemal Ekin, A. (2007). Exploring consumer lying in information-based exchanges. *Journal of Consumer Marketing*, 24(2), 90-99.

Jackson, C. B., & Wang, Y. (2018). Addressing the Privacy Paradox through personalized privacy notifications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(2), 68.

John, L. K., Acquisti, A., & Loewenstein, G. (2010). Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of consumer research*, 37(5), 858-873.

John, O. P., Naumann, L. P., & Soto, C. J. (2008). Paradigm shift to the integrative big five trait taxonomy. *Handbook of personality: Theory and research*, 3, 114-158

Kelley, P. G., Cranor, L. F., & Sadeh, N. (2013). Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 3393-3402). ACM.

Knijnenburg, B. P., & Kobsa, A. (2013). Making decisions about privacy: information disclosure in context-aware recommender systems. *ACM Transactions on Interactive Intelligent Systems* (TiiS), 3(3), 20.

Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & security*, 64, 122-134.

Kramer, A. D. (2012). The spread of emotion via Facebook. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 767-770). ACM.

Kumaraguru, P., Rhee, Y., Acquisti, A., Cranor, L. F., Hong, J., & Nunge, E. (2007). Protecting people from phishing: the design and evaluation of an embedded training email system. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 905-914).

Lewis, K., Kaufman, J., & Christakis, N. (2008). The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication*, 14(1), 79-100.

Malheiros, M., Preibusch, S., & Sasse, M. A. (2013, June). "Fairly truthful": The impact of perceived effort, fairness, relevance, and sensitivity on personal data disclosure. In *International Conference on Trust and Trustworthy Computing* (pp. 250-266). Springer, Berlin, Heidelberg.

Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information systems research*, 15(4), 336-355.

McNeeley, S. (2012). Sensitive issues in surveys: Reducing refusals while increasing reliability and quality of responses to sensitive survey items. In *Handbook of survey methodology for the social sciences* (pp. 377-396). Springer, New York, NY.

Mendel, T., & Toch, E. (2017, February). Susceptibility to social influence of privacy behaviors: Peer versus authoritative sources. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing* (pp. 581-593).

Metzger, M. J. (2007). Communication privacy management in electronic commerce. *Journal of Computer-Mediated Communication*, 12(2), 335-361.

Milgram, S., Bickman, L., & Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of personality and social psychology*, 13(2), 79.

Milne, G. R., & Culnan, M. J. (2002). Using the content of online privacy notices to inform public policy: A longitudinal analysis of the 1998-2001 US Web surveys. *The Information Society*, 18(5), 345-359.

Neuman, W. L. (2006) *Social Research Methods: Qualitative and Quantitative Approaches,* 6th Edition, Pearson International Edition, USA.

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

Nissenbaum, H.F. (2004). Privacy as contextual integrity. *Washington Law Review*, 79, 119.

Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of consumer affairs*, 41(1), 100-126.

Noulas, A., Scellato, S., Mascolo, C., and Pontil, M. (2011). "An empirical study of geographic user activity patterns in Foursquare". *ICwSM*, 11(70-573), 2.

O'Keefe, D. J. (2007). Brief report: Post hoc power, observed power, a priori power, retrospective power, prospective power, achieved power: Sorting out appropriate uses of statistical power analysis. *Communication Methods and Measures*, 1, 291-299. doi:10.1080/19312450701641375

Park, Y. J. (2013). Digital literacy and privacy behavior online. *Communication Research*, 40(2), 215-236.

Patil, S., Page, X., & Kobsa, A. (2011, March). With a little help from my friends: can social navigation inform interpersonal privacy preferences? In *Proceedings of the ACM 2011 conference on Computer supported cooperative work* (pp. 391-394). ACM.

Preibusch, S., Krol, K., & Beresford, A. R. (2013). The privacy economics of voluntary over-disclosure in Web forms. In *The Economics of Information Security and Privacy* (pp. 183-209). Springer, Berlin, Heidelberg.

Robert D. Putnam. 2000. *Bowling Alone*. New York: Simon & Schuster

Rader, E., Wash, R., & Brooks, B. (2012, July). Stories as informal lessons about security. In *Proceedings of the Eighth Symposium on Usable Privacy and Security* (p. 6). ACM.

Rainie, L. (2018). Americans' complicated feelings about social media in an era of privacy concerns. *The Pew Research Center*.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854-856.

Salmon, S. J., Fennis, B. M., de Ridder, D. T., Adriaanse, M. A., & De Vet, E. (2014). Health on impulse: When low self-control promotes healthy food choices. *Health Psychology*, 33(2), 103.

Spiekermann, S., Grossklags, J., & 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* (pp. 38-47). ACM.

Spottswood, E. L., & Hancock, J. T. (2017). Should I share that? Prompting social norms that influence privacy behaviors on a social networking site. *Journal of Computer-Mediated Communication*, 22(2), 55-70.

Staiano, J., Oliver, N., Lepri, B., de Oliveira, R., Caraviello, M., & Sebe, N. (2014). Money walks: a human-centric study on the economics of personal mobile data. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 583-594). ACM.

Stone, J., & Cooper, J. (2001). A self-standards model of cognitive dissonance. *Journal of experimental social psychology*, 37(3), 228-243.

Stroebe, W., & Diehl, M. (1981). Conformity and counter – attitudinal behavior: The effect of social support on attitude change. *Journal of Personality and Social Psychology*, 41, 876–889.

Stutzman, F., Gross, R., & Acquisti, A. (2012). Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of privacy and confidentiality*, 4(2).

Sutanto, J., Palme, E., Tan, C. H., & Phang, C. W. (2013). Addressing the personalization-privacy paradox: an empirical assessment from a field experiment on smartphone users. *MIS quarterly*, 1141-1164.

Taddicken, M. (2014). The 'privacy paradox' in the social web: The impact of privacy concerns, individual characteristics, and the perceived social relevance on different forms of self-disclosure. *Journal of Computer-Mediated Communication*, 19(2), 248-273.

Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York: New York.

Trepte, S., Teutsch, D., Masur, P. K., Eicher, C., Fischer, M., Hennhöfer, A., & Lind, F. (2015). Do people know about privacy and data protection strategies? Towards the "Online Privacy Literacy Scale"(OPLIS). In *Reforming European data protection law* (pp. 333-365). Springer, Dordrecht.

Treiblmaier, H. (2005). Antecedents of the quality of online customer information. In the proceedings of: *the 2005 International Conference on Information Quality (MIT IQ Conference)*, Cambridge, MA.

TRUSTe, T. R. U. S. T. (2014). US consumer confidence privacy report: consumer opinion and business impact. *Research Report*, TRUSTe Inc.

Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2011). The effect of online privacy information on purchasing behavior: An experimental study. *Information systems research*, 22(2), 254-268.

Tufekci, Z. (2008). Can you see me now? Audience and disclosure regulation in online social network sites. *Bulletin of Science, Technology & Society*, 28(1), 20-36.

Tufekci, Z. (2012). Facebook, youth and privacy in networked publics. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*. 2012.

Ur, B., Kelley, P., Komanduri, S., Lee, J., Maass, M., Mazurek, M., Passaro, T., Shay, R., Vidas, T., Bauer, L., Christin, N., and Cranor, L. (2012). How does your password measure up? the effect of strength meters on password creation. In *Presented as part of the 21st {USENIX} Security Symposium ({USENIX} Security 12)* (pp. 65-80).

Walton, G. M., Cohen, G. L., Cwir, D., & Spencer, S. J. (2012). Mere belonging: The power of social connections. *Journal of personality and social psychology*, 102(3), 513.

Wang, Y., Leon, P. G., Acquisti, A., Cranor, L. F., Forget, A., & Sadeh, N. (2014, April). A field trial of privacy nudges for facebook. *In Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (pp. 2367-2376). ACM.

Wang, Y., Norcie, G., Komanduri, S., Acquisti, A., Leon, P. G., & 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* (p. 10). ACM.

White, T. B. (2004). Consumer disclosure and disclosure avoidance: A motivational framework. *Journal of Consumer Psychology*, 14(1-2), 41-51.

Williams, D. (2006). On and off the 'net: scales for social capital in an online era. *Journal of Computer Mediated Communication*, 11(2), 593–628

Yang, H. (2012). Young American consumers' prior negative experience of online disclosure, online privacy concerns, and privacy protection behavioral intent. *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior*, 25, 179-202.

Yao, M. Z., & Linz, D. G. (2008). Predicting self-protections of online privacy. *CyberPsychology & Behavior*, 11(5), 615-617.

Young, A. L., & Quan-Haase, A. (2013). Privacy protection strategies on Facebook: The Internet privacy paradox revisited. *Information, Communication & Society*, 16(4), 479-500.

Zafeiropoulou, A. M., Millard, D. E., Webber, C., & O'Hara, K. (2013). Unpicking the privacy paradox: can structuration theory help to explain location-based privacy decisions? In *Proceedings of the 5th Annual ACM Web Science Conference* (pp. 463-472). ACM.

Appendix A: Informed Consent Form

You are invited to participate in a research study that is being conducted by students at the School of Communication and Information at Rutgers University in conjunction with a large credit reporting agency. The purpose of this research is to identify newer ways of building a credit score using social, behavioral, and financial data. This research will help us understand if it is possible to use different modes of data to generate more accurate credit scores for individuals. The participation will last a total of 15 days during which participants have to attend 2 in-person lab sessions.

Participation in this study will involve the following:

- Subjects will be asked to complete a survey comprising of questions on demography and attitudes online. We expect each survey to take anywhere between 10-15 minutes to complete.
- Subjects will also need to attend 2 in-person lab sessions during which they will be asked to complete a survey (personality traits) and perform some activities.
- Subjects will also receive email notifications during the study. These notifications will be for information only and should not take much time.
- Subjects can also choose to participate in an exit interview in order to discuss their experiences with the overall study. We expect the interview to take approximately 30 minutes.

Participants will be compensated up to $50 in total for completing the study. Participants will receive $20 on completing lab session 1, $20 on completing lab session 2, and an additional $10 for participating in a follow-up interview. If subjects decide to stop partway through the study, they will be compensated on a pro-rated basis.

In order to be eligible for this study, subjects must own a smartphone with camera and file-sharing capabilities which they need to bring to the lab, regularly use an email and social media account,

own at least one credit card, be between 18-60 years, be comfortable with spoken and written English, and be able to travel to the test site for lab sessions.

Participation in this study is voluntary. You may choose not to participate, and you may withdraw at any time during the interview without any penalty to you. In addition, you may choose not to answer any questions with which you are not comfortable.

If you have any questions about the study or study procedures, you may contact:

Isha Ghosh, Principal Investigator, 848-932-7588, isha,ghosh@rutgers.edu,

4 Huntington St., New Brunswick, NJ 08901.

If you have any questions about your rights as a research subject, please contact an IRB Administrator at the Rutgers University, Arts and Sciences IRB:

Institutional Review Board
Rutgers University, the State University of New Jersey
Liberty Plaza / Suite 3200
335 George Street, 3rd Floor
New Brunswick, NJ 08901
Phone: 732-235-9806
Email: humansubjects@orsp.rutgers.edu

You will be given a copy of this consent form for your records.

Sign below if you agree to participate in this research study:

Subject (Print) _____

Subject Signature _____ Date _____

Principal Investigator Signature _____ Date _____

Appendix B: Explanation of Research

**Communicated to participants in Lab session 1 after online entry session is completed*

We are researchers working with a credit reporting agency, whose name we cannot reveal at this point, to refine a revolutionary idea that will change the way credit scoring is performed. We have done some previous studies and focus groups, and most people are dissatisfied with the current way credit scores are calculated. The majority view was that credit scoring is too limited and generalized and does not address individual personalities. This agency is investigating a new way of building a credit score that takes into account all aspects of a person's activities rather than just financial data. The agency wants to build an algorithmic model that uses social and behavioral data about a person to predict the best credit rate and offers for them. This is a testing project where our research team, working with the agency, is trying to identify the quickest and most efficient way of building this score. In order to do this, we have divided this project into two phases during which we will ask you for different types of information. The first phase will be today and the second after a week. You will receive compensation of $20 after completing today's activities, another $20 after the second round of data collection and have the option to receive an additional $10 by participating in an interview. Data collection will take place in 2 phases, today, we collect first round of data and do some preliminary processing. To start with please use the following link to complete a survey (Big 5 personality quiz – Appendix D)"

After participants complete the Big5 personality quiz they will be given a handout with a list of tasks (Appendix E) and asked to choose a selection of tasks that they want to complete.

Communication to participants after completing personality survey:

"Thank you for completing the first round, to start with the second round please select the activities you would like to perform from the activity sheet provided. The idea is for you to gain 20

points by completing a selection of these tasks. As you see each task has a number of points associated with it, you need to perform enough tasks to gather 20 points. Once you check off the tasks you will be doing, you will receive a link and can start answering the questions.

A few things to keep in mind:

- You must be completely honest when providing this information. All your info will be cross-referenced with credit reference agency (similar to TransUnion) and any deception will disqualify you from the study.
- We cannot disclose the name of the corporation we are hired by or the fact-checking agency we will use as we are still in product development phase and we want to keep our algorithm as confidential as possible.
- If responses are incomplete or inaccurate, you may be disqualified from the study or will have to restart the survey from the beginning.

Note: As this is in the design phase, we cannot guarantee security or confidentiality of any data you provide.

Appendix C: Demography and Privacy Concern Surveys

**Administered online at the start of in Phase 1 (Day 1) and 2 (Day 15)*

What is your applicant id: _____

State your first and Last name: _____

What is your age?

√ Under 18 √ 18-21 √ 22-24 √ 25-34 √45-54 √55-64 √Age 65 or older

What is your gender?

√ Male √ Female √ Other _____

What is your marital status?

√Single (Never Married) √ Married √ Separated √ Divorced

What is your highest grade or year of school you completed?

√ Grade 9-11 (Some High School) √ Grade 12 or GED (High school equivalent)

√ College 1-3 years (Some College) √ College 4 years (College Grad)

√ Grad School (Advance Degrees)

How would you describe yourself? Select more than one option if needed.

√American Indian or Alaskan Native √ Hawaiian or Other Pacific Islander

√ Asian or Asian American √ Black or African American

√ Hispanic or Latino √ White

What is your annual household income from all sources?

√ Less than $25,000 √ $25,000 to $34,999 √ 35,000 to $49,000

√ $50,000 to $74,999 √ $75,000 to $99,000$ √ $100,000 to $149,000

√ $150,000 or more

For the following questions, choose an option from 1 to 5 where 1 is Strongly Disagree and 5 is Strongly Agree

1. Online privacy is really a matter of user's right to exercise control and autonomy over decisions about how their information is collected, used, and shared.
2. Control of personal information lies at the heart of consumer privacy.
3. I believe that privacy is invaded when control is lost or unwillingly reduced as a result of an online transaction.
4. Websites seeking information online should disclose the way the data are collected, processed, and used.
5. A good online privacy policy should have a clear and conspicuous disclosure.
6. It is very important to me that I am aware and knowledgeable about how my personal information will be used.
7. It usually bothers me when websites or apps ask me for personal information.
8. When websites ask me for personal information, I sometimes think twice before providing it.
9. It bothers me to give personal information to so many websites.
10. I'm concerned that websites are collecting too much personal information about me.

Appendix D: Big-Five Personality Index

**Administered online at the start of in Phase 1 (Day 1) and 2 (Day 15)*

What is your applicant id: _____

State your first and Last name: _____

Here are a number of characteristics that may or may not apply to you. Please respond to each statement to indicate the extent to which you agree or disagree with that statement.

For the following questions, choose an option from 1 to 5 where 1 is Strongly Disagree and 5 is Strongly Agree

I am someone who...

Is talkative	Is inventive
Tends to find fault with others	Has an assertive personality
Does a thorough job	Can be cold and aloof
Is depressed, blue	Preserves until the task is finished
Is original, comes up with new ideas	Can be moody
Is reserved	Values artistic, aesthetic experiences
Is helpful and unselfish with others	Is sometimes shy, inhibited
Can be somewhat careless	Is considerate and kind to almost everyone
Is relaxed and handles stress well	Does things efficiently
Is curious about many things	Remains calm in tense situations
Is full of energy	Prefers work that is routine
Starts quarrels with others	Is outgoing, sociable
Is a reliable worker	Is sometimes rude to others
Can be tense	Makes plans and follows through with them
Is ingenious, a deep thinker	Gets nervous easily
Generates a lot of enthusiasm	Likes to reflect, play with ideas
has a forgiving nature	Has few artistic interests
Tends to be disorganized	Likes to cooperate with others
Worries a lot	Is easily distracted
Has an active imagination	Is sophisticated in art, music, or literature
Tends to be quiet	Is talkative
Is generally trusting	Tends to find fault with others
Tends to be lazy	Does a thorough job
Is emotionally stable, not easily upset	Is depressed, blue

Here are a number of characteristics that may or may not apply to you. Please respond to each statement to indicate the extent to which you agree or disagree with that statement. Choose an option from 1 to 5 where 1 is Strongly Disagree and 5 is Strongly Agree

1. There are several people online/offline I trust to help solve my problems.
2. There is someone online/offline I can turn to for advice about making very important decisions.
3. There is no one online/offline that I feel comfortable talking to about intimate personal problems. (reversed)
4. When I feel lonely, there are several people online/offline I can talk to.
5. If I needed an emergency loan of $500, I know someone online/offline I can turn to.
6. The people I interact with online/offline would put their reputation on the line for me.
7. The people I interact with online/offline would be good job references for me.
8. The people I interact with online/offline would share their last dollar with me.
9. I do not know people online/offline well enough to get them to do anything important. (reversed)
10. The people I interact with online/offline would help me fight an injustice.

11. Interacting with people online/offline makes me interested in things that happen outside of my town.
12. Interacting with people online/offline makes me want to try new things.
13. Interacting with people online/offline makes me interested in what people unlike me are thinking.
14. Talking with people online/offline makes me curious about other places in the world.
15. Interacting with people online/offline makes me feel like part of a larger community.
16. Interacting with people online/offline makes me feel connected to the bigger picture.
17. Interacting with people online/offline reminds me that everyone in the world is connected.
18. I am willing to spend time to support general online/offline community activities.
19. Interacting with people online/offline gives me new people to talk to.
20. Online/Offline, I come in contact with new people all the time.

Appendix E: Behavioral Index 1

**Hand out provided to participants in lab session 1 (before interventions)*

	1 Point Tasks: Low Effort Low Sensitivity
1	Do you have one or more favorite books? If yes, which one(s)
2	What would you most likely sing at Karaoke night?
3	Do you have any siblings? How many?
4	Which is your dream car?
5	What is your favorite music style?
6	What is your favorite sport?
7	How often have you moved in the last 5 years?
8	Approximately how long is your daily commute?
9	If you could live anywhere where would it be?
10	How many pillows do you sleep with?
11	Do you love or hate rollercoasters?
12	What's your favorite fast food chain?
13	Is your glass half full or half empty?
14	Where is the next place on your travel bucket list?
15	What would you say is your favorite season (Spring, Summer, Fall, or Winter)
16	Do you have a favorite (non-alcoholic) beverage? What is it?
17	What is your favorite genre of book or movie?
18	Which topic could you give a 5-minute presentation on with absolutely no preparation?
19	What is your favorite sitcom?
20	What is your favorite band or music artist?
	3Point Tasks: High Effort Low Sensitivity
1	Describe in 5-6 sentences (at least 500 characters) a place you visited, worked at, or lived in that has remained memorable to you?
2	Describe in 5-6 sentences (at least 500 characters) a recurring dream you have?
3	Would you rather be alone for the rest of your life or always surrounded by annoying people? Use 5-6 sentences (at least 500 characters) to explain your reasons.
4	Which candidate would you support in the upcoming elections? Use 5-6 sentences (at least 500 characters) to explain your reasons.
5	Describe in 5-6 sentences (at least 500 characters) 3 positive and 3 negative qualities your closest friend/significant other would say you have.
6	Describe in 5-6 sentences (at least 500 characters) what you would like the inscription on your gravestone to be and why?
7	Do you believe organized religion is relevant in today's world? Use 5-6 sentences (at least 500 characters) to explain your reasons.
8	Should animals be used to make skin and hair products safer for humans? Use 5-6 sentences (at least 500 characters) to explain your reasons.
9	Should vaccinations be made mandatory for children (irrespective of religious/cultural beliefs)? Use 5-6 sentences (at least 500 characters) to explain your reasons.
	5 Point Tasks: High Effort Medium Sensitivity
1	Open the Facebook app on your phone, navigate to your profile page and scroll down to the snapshot of your friends list. Take a screenshot and upload the file.
2	Call a friend and speak with them for 3 mins, upload a screenshot of the call log
3	Ask someone to video record you using your email account (logging in/ reading/answering email) for 45 seconds.
4	Video record a 45 second conversation with a friend (Video conference or face to face) where you exchange daily routines and activities
5	Navigate to your medical insurance provider page and upload a screenshot showing your 5 most recent claims
	10 Point Tasks: Low Effort High Sensitivity
1	Enter your social security number
2	Enter your Facebook (or favorite social media site) id and password

Appendix F: Behavioral Index 2

**Hand out provided to participants in lab session 2 (After Interventions)*

	1 Point Tasks: Low Effort Low Sensitivity
1	Do you have one or more favorite movies? If yes, which one(s)
2	What radio station are you most likely to tune into? (a) Classic Rock (b) Top Hits (c) Regional Channel (d) News/NPR (e) Other _____
3	How many people are there in your immediate family
4	What form of transportation do you prefer to use during trips? (a) Air (b) Train (c) Bus (d) Ferries (e) Other_____
5	What is your favorite cuisine?
6	What is your favorite hobby?
7	How many jobs have you had in the past 5 years?
8	How often do you exercise?
9	Which is your favorite holiday destination?
10	What's the longest you've gone without sleep?
11	Do you enjoy adventure sports? Which is your favorite?
12	What is your go-to comfort food?
13	Are you a clean or messy person?
14	Do you have a favorite holiday? Which is it? (E.g.: Christmas/Hanukkah, Thanksgiving, Fourth of July etc.)
15	Which part of the continental U.S. would you like to live in (a) North-East (b) Mid-West (c) South (d) Pacific Coast (e) North-West
16	What part of the day is your favorite? (a) Morning (b) Afternoon (c) Evening (d) Night
17	What quote or one-liner do you most frequently use (E.g.: Life's not fair; Tomorrow is another day, etc.)?
18	What is the one thing you take for granted?
19	Which show would you be most likely to binge-watch?
20	What is your favorite dance style?
	3 Point Tasks: High Effort Low Sensitivity
1	Describe in 5-6 sentences (at least 500 characters) an experience or incident that has remained memorable to you?
2	Describe in 5-6 sentences (at least 500 characters) what you think is the ugliest vegetable and your reasons?
3	Describe in 5-6 sentences (at least 500 characters) if you would rather be unable to use search engines or social media?
4	Are you pro-choice or pro-life? Give at least 3 reasons (500 characters) justifying your answer.
5	Describe in 5-6 sentences (at least 500 characters), what quality do you think you have that is most valued by your family?
6	If you could put your brain in a computer and live indefinitely would you? Describe in 5-6 sentences (at least 500 characters).
7	Do you believe capital punishment, or the death penalty is unconstitutional? Use 5-6 sentences (at least 500 characters) to explain your reasons.
8	Genetic human cloning has become a very real prospect in today's world. Do you believe this practice should be encouraged?
9	The availability of guns directly influences the crime rate in the US. Do you agree with this statement? Use 5-6 sentences (at least 500 characters) to explain your reasons.
	5 Point Tasks: High Effort Medium Sensitivity
1	From the list of contacts on your phone, pick 3 that can act as your references. Open the contact page and upload a screenshot of their contact information.
2	Open the Facebook messenger app and upload a screenshot of the page.
3	Ask someone to video record you logging into your social media accounts and visiting your friends' profile pages for 45 seconds and upload the file.
4	Video record a 45 second conversation with someone on the street in which you exchange names, phone numbers, house address, and favorite movie. Upload the video file
5	Enter your medical insurance number and all medical ailments in the last year.
	10 Point Tasks: Low Effort High Sensitivity
1	Upload a picture of your credit card front and back
2	Enter your Gmail id and Password

Appendix G: Exit Interview

Questions for exit interview conducted after lab sessions:

1. Thank you for agreeing to the interview. To start with can you describe your overall experiences of the study from the start of the study to today?
2. In the first step, we asked you to fill out a survey about your information privacy concerns. So is privacy something you think about often?
3. Have you ever had a negative experience related to information privacy? How did that make you feel?
4. Did you take any steps to protect or hide your information?
5. Okay, now coming to the first lab session. Can you describe how you felt in the first lab session? Any tasks that stood out to you? Any tasks that you felt more concerned about?
6. How did you choose which tasks to do?
7. Did a concern for information privacy play a role in the tasks you selected?
8. After the lab sessions, we sent you a few emails, what did you feel when you saw the email? Did you agree with the information given?
9. What did you think about the design of the email? Was there anything that stood out to you or anything that you found really helpful?
10. There was a second round of surveys, where we again asked questions about your privacy concerns, did you feel your answers changed in the second round?
11. Did you choose any different tasks in the second phase? Why?
12. Were you thinking about the emails you received while answering the survey or choosing tasks?
13. Did you think of giving any false information to finish the tasks quicker?
14. Did the study influence any other aspects of your daily life (i.e. those occurring outside of the lab sessions)?

Appendix H: Debriefing Statement

Debriefing Statement

Thank you for participating in our study. In privacy research, it is sometimes necessary to conceal our hypotheses because when people know what is being studied, they often alter their information disclosure behavior. However, we do not want you to leave misinformed, so we will now tell you what we were actually studying.

During this study, you were asked to fill in two surveys and complete selected activities. You were told that the purpose of the study was to change the way credit scoring is performed. The purpose of this study is to actually study whether a person's disclosure behavior matches their privacy concerns and whether an intervention can help reduce the inconsistency in cases where there is a mis-match. In order to test whether social interventions are more effective than non-social interventions in reducing inconsistencies between privacy concerns and disclosure behaviors.

We apologize that we could not reveal our true hypotheses to you at the beginning of the study, but we hope you can see why it was necessary to keep this information from you. When people know exactly what the researcher is studying, they often change their behavior, thus making their responses unusable for drawing conclusions about human nature and experiences. For this reason, we ask that you please not discuss this study with others who might participate any time after you.

If you have any questions about this study, feel free to ask the researcher, Isha Ghosh, isha.ghosh@rutgers.edu, 848-932-7588.

If you have any questions about the study or study procedures, you may contact myself at Isha Ghosh, isha.ghosh@rutgers.edu, 848-932-7588. You may also contact my faculty advisor Vivek K. Singh, vivek.k.singh@rutgers.edu. If you have any questions about your rights as a research

subject, please contact an IRB Administrator at the Arts and Sciences Institutional Review Board, Rutgers University by phone: 732-235-9806 or by email: humansubjects@orsp.rutgers.edu.

Now that you understand the true nature of our study, we would like to give you the chance to refuse the use of your data for our research purposes. You are free to ask us not to use your data in our study analysis. If you have any concerns about your participation or the data you provided in light of this disclosure, please discuss this with us. We will be happy to provide any information we can to help answer questions you have about this study. Please again accept our appreciation for your participation in this study.

You will be given a copy of this form for your records. Please choose one (1) statement below and sign/date:

You have read this debriefing form and you AGREE to allow the use of your data for research purposes:

_____

Agree---Subject's Signature Date

You have read this debriefing form and you DO NOT AGREE to allow the use of your data for research purposes and would like your data to be immediately withdrawn and destroyed (where possible).

_____

Disagree---Subject's Signature Date

Subject Name (Print) _____ Subject ID/# _____ (if applicable)

Principal Investigator Signature _____ Date _____

Appendix J: Credit Score Information

In recent decades, consumers have become increasingly dependent on credit. When you use credit, you are borrowing money that you promise to pay back within a specified period of time. Your credit report and rating compose a financial snapshot that presents you to the business world. Your financial history can affect how easily you can get a mortgage, rent an apartment; make big-ticket purchases; take out loans, and in some industries even get hired. When you apply for a credit card or even a cable hookup, lenders check your credit rating. Your credit rating helps to determine the probability that you could and would pay back the money that you have borrowed; it also indicates the degree of risk that you pose to a lender. Lenders combine your credit score with the information in your credit report to assess your risk as a borrower. If your score is high, you look like less of a risk; if your score is low, lenders may question your ability to pay what you owe.

A good credit score can save you thousands of dollars over the life of a loan. For example, you may get a better mortgage interest rate with a high credit score than you would with a lower score. On a 30-year mortgage for $200,000, the savings can be significant.

The same principle applies whether you are borrowing for a car, an education, or a personal loan: The better your credit score, the more you can save when you decide to borrow.

Many Internet, TV, and cell phone service providers now check your credit before they set you up with service. In some cases, if your credit is poor enough, you might be denied an account.

Even if you aren’t denied service, you might have to pay a security deposit or pay some part of your service up front. This can be frustrating and costly as it can change your monthly cash flow and strain your budget.

Appendix K: Qualitative Coding Themes

Interview Quotes Sorted by Themes

Shift in concerns and behavior

"*I always felt I was similar to everyone else, but then I thought about it and I realized I'm probably much more open than other people. I hear everyone being really scared about what will happen to their information and I just don't feel that way"* – P1 (social-proof condition)

"*It was very scary to be told that the gap between my privacy concerns and behavior is higher than normal. I think privacy is a really important and sensitive topic and I don't want people to think that I don't care about it*." – P16 (social-proof condition)

"It said my information sharing didn't match my comfort level, so I started thinking about what tasks I did and I thought sharing the call log was risky" – P8 (information inconsistency condition)

"*It said my information sharing didn't match my concerns…I felt like so what? Were they supposed to match?*" – P32 (information inconsistency condition)

"I didn't think I had shared too much information but then I remembered sharing a screenshot of the call log so maybe that's what it [email intervention] was referring too" – P42 (baseline condition)

"I remembered uploading a lot of images in the first lab session, there was one that I shared my Facebook friends list, another my insurance information and some others. I didn't think about it at the time, but when I got the email I was like yeah I probably did share a lot of information so I was more careful the second time" – P9 (baseline condition)

Interpretation of the Interventions

"The green and red was really effective. Like everyone is green and I am the only red that actually conveyed to me that I was singled out" – P40 (social-proof condition)

"*I liked seeing where I stood and where everyone else was…it was like a S.A.T score. Seeing the comparison really opened my eyes to how different I was from everyone else*." – P22 (social-proof condition)

"*The email was simple, it wasn't too cluttered with paras and paras of information. It just said what was going on in 2 – 3 lines that was easy to understand*." – P45 (baseline condition)

Appendix L: Tasks Performed by Participants

Average number of tasks in each category performed by participants in each experiment group *before and after interventions*

	Before Intervention					After Intervention				
	10 Point	5 Point	3 Point	1 Point	Total	10 Point	5 Point	3 Point	1 Point	Total
Social-proof Condition	0.27	0.73	0.80	11.27	**13.07**	0.20	0.73	0.33	13.33	**14.60**
Information Incon-sistency Condition	0.08	0.69	0.46	14.38	**15.62**	0.00	0.31	0.23	17.77	**18.31**
Baseline Condition	0.14	0.86	0.79	12.29	**13.71**	0.14	0.64	0.25	15.43	**16.29**