Nudging Users Towards Privacy on Mobile Devices

Rebecca Balebako, Pedro G. Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Faith Cranor and Norman Sadeh
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213 USA

ABSTRACT

By allowing individuals to be permanently connected to the Internet, mobile devices ease the way information can be accessed and shared online, but also raise novel privacy challenges for end users. Recent behavioral research on "soft" or "asymmetric" paternalism has begun exploring ways of helping people make better decisions in different aspects of their lives. We apply that research to privacy decision making, investigating how soft paternalistic solutions (also known as *nudges*) may be used to counter cognitive biases and ameliorate privacy-sensitive behavior. We present the theoretical background of our research, and highlight current industry solutions and research endeavors that could be classified as nudging interventions. We then describe our ongoing work on embedding soft paternalistic mechanisms in location sharing technologies and Twitter privacy agents.

Author Keywords

Nudge, Privacy, Security, Location Sharing, Mobile Devices, Soft Paternalism

ACM Classification Keywords

H.1.2 User/machine systems: Human information processing; J.4 Social and behavioral sciences: psychology

INTRODUCTION

As mobile devices and applications become pervasive, privacy risks to their users also grow. The accessibility and ease of use of these devices make it easy to casually broadcast personal information at any time, from anywhere, to friends and strangers. Without a doubt, users benefit from and enjoy such streams of information sharing. However, they also expose themselves to tangible and intangible risks: from tracking by commercial entities interested in exploiting personal information for profit, to surveillance or even stalking by malicious parties. However, it is difficult for individuals to determine the optimal balance between revealing and hiding personal data. Sometimes we are not even aware that information about us is being broadcast, shared, or monitored; other times, while aware of ongoing information flows, we do not understand their consequences, or properly assess their risks. Such challenges are magnified in mobile scenarios. Therefore, a mobile device user may end up sharing information in a manner that goes against her own long-term self interests.

In recent years, there has been growing interest in using lessons from behavioral economics to influence and ameliorate decision making in situations where cognitive and behavioral biases may adversely affect the individual [11, 16]. This approach is often referred to as soft or asymmetric paternalism, or with the more popular term "nudges." Soft paternalism aims at countering and overcoming those biases, so as to assist individual decision making. Our research aims at applying and extending lessons from the nascent field of soft paternalism to the field of privacy decision making. This paper presents an overview of our research agenda in this area. First, we introduce the research exploring cognitive and behavioral biases in privacy decision making. Then, we examine current academic studies and industry products that focus on influencing privacy (and security) decision making, and that therefore may be compared to nudging interventions. Finally, we discuss how we are integrating soft paternalistic mechanisms in our research on privacy in location sharing applications and social networks.

FROM HURDLES IN PRIVACY DECISION MAKING TO SOFT PATERNALISM

Findings from behavioral economics and behavioral decision research have highlighted hurdles in human decision making that lead, sometimes, to undesirable outcomes. The hurdles are often due to lack of information or insight, cognitive limitations and biases, or lack of self-control [16]. Because of those hurdles, individuals may end up making decisions that they later regret. Those decisions may include (not) saving for retirement, (not) eating well, or smoking cigarettes [11]. They may also include decisions about protecting too much, or not enough, personal information [3]. Privacy decisions are complex and often taken in conditions of information asymmetry (that is, individuals may not have full knowledge of how much of their personal information is being gathered, and how it is being used). Furthermore, privacy decision making may be overwhelming: the cognitive costs associated with considering all the ramifications of a disclosure may hamper decision making [3]. Finally, cognitive biases may affect one's propensity to reveal personal

CHI 2011, May 7 - 12, 2011, Vancouver, BC, Canada.
ACM 978-1-4503-0268-5/11/05.

information: for instance, heightened control of one's personal information may, paradoxically, make the user overconfident about sharing information [5].

Paternalistic policies try to solve decision-making hurdles by mandating decisions for individuals. Such policies are often heavy-handed and generate externalities [11]. Soft paternalism, on the other hand, avoids coercion; it seeks to steer users in a direction (believed to be more desirable based on the user's own prior judgement, or on external empirical validation), without impinging on her autonomy. A soft paternalistic solution, for instance, would consist of making an individual aware of the biases, lack of information, or cognitive overload that may affect her decision.

Nudges are tools of soft-paternalism, and may be used to ameliorate privacy (as well as security) decision making [2]. Their application to scenarios involving mobile devices is particularly appealing. In the case of insecure communication channels, or covert data collection through a mobile device, a nudge may take the form of an alert that informs the user of the risk. In the case of mobile devices that store sensitive information (which could be accessed by strangers if the phone was misplaced), a nudge might discourage users from storing private data on mobile phones. When information is being disclosed through a smart-phone, nudges may provide alerts about the recipients, contexts, or type of data being shared.

Many different types of nudging interventions are possible. Some simply consist of informing the user — in which case they relate to privacy research on informed consent. Some focus on making systems simpler to use — in which case, privacy nudges fall into the realm of research on privacy usability. However, other nudges aim at countering specific cognitive and behavioral biases, such as neutralizing the detrimental effects of immediate gratification biases in privacy decision making [1] by altering the individual's perception of the sequence of costs and benefits associated with revealing sensitive information.

The literature on soft paternalism applied to privacy decision making is in its infancy, and therefore extremely scarce. However, a number of recent studies and products focus on mechanisms that may be categorized as nudges. We present a brief overview of them in the following sections.

PRIVACY NUDGES IN THE LITERATURE

Previous research on the drivers of privacy concerns has demonstrated that users' attitudes towards security and privacy are influenced by numerous factors, including information available, personal beliefs, economic valuations, moral reasoning, social values, cognitive biases, and so on. Therefore, providing adequate information, making privacy tools more evident, or rewarding and punishing users as they make safer or riskier decisions are all ways of nudging or influencing privacy behavior. The privacy literature offers some examples of these approaches.

For example, recent experimental research has shown that users are interested in protecting their privacy and may even pay for it, *if* appropriate tools and salient, simple, and compact privacy information are offered. Specifically, one series of studies explored the impact of making information about privacy practices on web sites more accessible to buyers. The results showed that online customers are more likely to shop online from websites that exhibit more protective privacy policies. Additionally, those customers are willing to pay a premium for privacy. Furthermore, privacy indicators displayed at the moment an individual is shopping online may have an impact on consumer decisions. In particular, they increase the willingness to pay for privacy; however, if the indicator is provided only after the shopper has already chosen the website from which to buy, the user will not change their already-made decisions. The authors find that timing is essential when trying to help people to protect their privacy [17], [6]. Similarly, another study found that merely priming Facebook users with questions about their online disclosure behavior and the visibility of their Facebook profiles was sufficient to trigger changes in their disclosure behavior [13]. Application interface design is also important, and should help users notice when changes in context generate changes in information flows and then help them to maintain their privacy [7].

In the context of location sharing applications, providing feedback to users whose location has been requested by others has been shown to have both positive and negative implications [8]. It can prevent excessive requests and hence protect people's privacy. However, unless appropriate notifications are used, feedback receivers could also be annoyed. In addition, notifications may inhibit users from requesting others locations and hence affect system usage.

PRIVACY NUDGES IN INDUSTRY

Examples of industry products or solutions that influence decision making in regards to privacy (either to better protect the user, *or instead* to influence her to reveal more information) take various forms, and some have been applied to mobile devices. Some of these solutions may be interpreted as soft paternalistic for privacy protection, in the sense that they nudge towards privacy. They include privacy/security usability solutions, simplifications of privacy settings, or tests and delays before one can post information. More frequent, however, are the examples of products and solutions that nudge individuals to *give up* even more of their privacy, surrendering sensitive information. These include privacy defaults that are open, lack of usability in privacy settings interfaces, poorly designed warnings, and other rewards for sharing data or encouraging friends to share data.

Connections in social applications

Some applications provide information about who can see your data, who has seen your data, or how many people can see your data. For instance, Flickr.com, a video and image sharing website, provides information on each user-owned picture stating who can see it, followed by a link to edit the privacy settings for that picture. This may be a nudge towards privacy, as users may decide to share certain photos with friends, and share other photos with everyone.

Social networking sites often show the number of connections a user has. These connections may be called followers, friends, or ties. In some cases, connections can have access to all the user's information that is on the application. Twitter and Google Buzz are examples of sites that prominently show the number of connections. In the case of LinkedIn.com, a job searching social network, the user may prefer to add additional connections, even with people they don't know well, in order to grow their job-searching network. However, by opening their information to more connections, they may be compromising their privacy. These applications may nudge users towards increasing their connections and revealing *more* information. Indeed, several online social networks such as Facebook.com and LinkedIn.com periodically encourage users to add new connections by searching the user's email accounts for email contacts.

Connections such as friends in Facebook and followers in Twitter do not set the boundaries for information flow. One's connections may be able to share information with other unintended recipients, or even make it available to the public. In Twitter, for example, re-tweets allow connections to pass on information without the original sender's control. In Facebook, default privacy settings usually allow sharing of individual's information with friends of friends. Therefore, the information provided about the number of connections may mislead the user about the privacy of their data and decrease the likelihood that the user will take an information-protective stance.

Privacy Settings

The privacy settings allowed in an application impact the user's ability to control how their information is shared. Both the default settings and the usability of the settings user interface create nudges towards and away from privacy [10, 12, 13].

Some websites make privacy options very simple. For example, Pandora.com, an online music station, explicitly gives users two options regarding their profile page: make private or keep public. These clear options allow a user to choose without understanding complex details or settings. Conversely, the lack of granularity may encourage users to make everything public.

Several tools provide simple ratings of privacy settings. PrivacyCheck,[1] and ProfileWatch,[2] give Facebook settings a privacy score. Other services provide a user-friendly layer on the Facebook privacy settings, allowing the user to change the settings. For example, Privacy Defender[3] provides a sliding color scale that allows the user to set their Facebook options as more or less private. These software services actively encourage stricter privacy settings.

Reduction of Information Disclosure

If an individual expects she may be likely to post information she may later regret, software exists to discourage her from doing so. Sophisticated users may choose to employ software tools to prevent excess disclosure. For example, the Social Media Sobriety Test, `socialmediasobrietytest.com`, and Mail Goggles on Gmail `googlelabs.com` both allow the user to set certain hours of the week when they may typically embarrass themselves, such as weekend evenings after trips to the bar. During these hours, social network sites or Gmail may be blocked until the user can complete a dexterity or cognitive test. The user has the option to bypass the test. Alternatively, a user may set up a warning system if a message is likely to be poorly interpreted. ToneCheck `tonecheck.com` scans emails written in Outlook to discover whether the tone is off-putting, and will ask the user to confirm before sending it. This may help discourage users from sending or posting regrettable information.

Other tools may discourage users from posting information by reminding the user who can see it. NetNanny is a tool that parents can user to protect their children online. It will show a message every time a child posts on a social network. This message reminds the child that her parents will see the post as well `netnanny.com`.

ONGOING WORK WITH MOBILE APPLICATIONS

By studying and understanding the specific biases and user actions in regards to mobile applications, we hope to suggest and test nudges that will help users make decisions that improve their satisfaction and well being. We are moving towards that goal by first understanding users' needs, preferences, biases, and limitations about privacy, and second by using that information to evaluate the efficacy of techniques that exploit biases to improve decision making. As an example, we are currently pursuing foundational studies with two applications developed at Carnegie Mellon: a location sharing application called Locaccino [15] and a privacy agent for Twitter.

Locaccino is a unique location sharing application that allows users to control the conditions under which they make their location visible to others. This includes controlling the times and days of the week when different groups of people can see the user's location as well as the specific locations where the user is willing to be visible. For instance, a user can specify rules such as "I'm willing to let my colleagues see my location but only when I am on company premises and only 9am-5pm on weekdays." Research conducted by our group has shown that this level of expressiveness is critical to capturing the location sharing preferences many people have when it comes to disclosing their locations to others across a broad range of scenarios [4]—in contrast to the much narrower set of scenarios supported by location sharing applications such as Foursquare.

As part of our ongoing research, we are interested in better understanding how different elements of Locaccino functionality effectively nudge people in different directions. This includes experimenting with new interface designs as well as new ways of leveraging some of the machine learning techniques we have been developing, from exposing different sets of default privacy personas to users [14] to helping

[1] `http://rabidgremlin.com/fbprivacy`

[2] `http://atherionsecurity.com/idpro.html`

[3] `http://privacydefender.net`

them refine their privacy preferences [9]. We are looking at the preferences of like-minded users who have been using the system for a while and trying to use their preferences to guide new users. This would have the potential of reducing regret by giving new users the benefit of the experience acquired over time by others. We plan to explore to what extent such an approach can be made to work and to what extent it seems beneficial.

The Twitter privacy agent is an application we are building to help Twitter users behave in a more privacy protective way. We plan to build tools that will provide nudges that guide users to restrict their tweets to smaller groups of followers or discourage them from sending tweets from mobile devices that they may later regret. We plan to empirically test the impact of these nudges on user behavior. We will also examine whether fine-grained privacy controls result in more or less data sharing.

We expect our work on nudges in behavioral advertising, social networks, and location sharing to be effective for improving privacy decisions on mobile devices. We further hope our soft-paternalistic approach to have a broader impact, guiding the development of tools and methods that assist users in privacy and security decision making.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant CNS-1012763 (Nudging Users Towards Privacy), and by Google under a Focused Research Award on Privacy Nudges. This work has also been supported by NSF grants CNS-0627513, CNS-0905562, and by CyLab at Carnegie Mellon under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the Army Research Office. Additional support has been provided by Nokia, France Telecom, and the CMU/Portugal Information and Communication Technologies Institute

REFERENCES

1. A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Conference on Electronic Commerce (EC '04)*, pages 21–29, 2004.
2. A. Acquisti. Nudging privacy: The behavioral economics of personal information. *Security & Privacy, IEEE*, 7(6):82–85, 2009.
3. A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *Security & Privacy, IEEE*, 3(1):26–33, 2005.
4. M. Benisch, P. Kelley, N. Sadeh, and L. Cranor. Capturing Location-Privacy Preferences: Quantifying Accuracy and User-Burden Tradeoffs. *Journal of Personal and Ubiquitous Computing*, 2011.
5. L. Brandimarte, A. Acquisti, and G. Loewenstein. Misplaced Confidences: Privacy and the Control Paradox. Technical report, Mimeo, Carnegie Mellon University, 2010.
6. S. Egelman, J. Tsai, L. F. Cranor, and A. Acquisti. Timing is everything?: the effects of timing and placement of online privacy indicators. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, pages 319–328, New York, NY, USA, 2009. ACM.
7. G. Hull, H. R. Lipford, and C. Latulipe. Contextual gaps: Privacy issues on facebook. *Ethics and Information Technology*, pages 1–14, 2010.
8. L. Jedrzejczyk, B. A. Price, A. K. Bandara, and B. Nuseibeh. On the impact of real-time feedback on users' behaviour in mobile location-sharing applications. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, SOUPS '10, pages 14:1–12, New York, NY, USA, 2010. ACM.
9. P. Kelley, P. Hankes Drielsma, N. Sadeh, and L. Cranor. User-controllable learning of security and privacy policies. In *Proceedings of the 1st ACM workshop on Workshop on AISec*, pages 11–18. ACM, 2008.
10. Y.-L. Lai and K. L. Hui. Internet opt-in and opt-out: investigating the roles of frames, defaults and privacy concerns. In *Proceedings of the 2006 ACM SIGMIS CPR conference on computer personnel research*, pages 253–263, 2006.
11. G. F. Loewenstein and E. C. Haisley. *The Foundations of Positive and Normative Economics*, chapter 9: The Economist as Therapist: Methodological Ramifications of 'Light' Paternalism. Oxford University Press, 2008.
12. W. Mackay. Triggers and barriers to customizing software. In *Proceedings of the SIGCHI conference on Human factors in computing systems: Reaching through technology*, pages 153–160. ACM, 1991.
13. Ralph Gross and Alessandro Acquisti. Information Revelation and Privacy in Online Social Networks. In *Workshop on Privacy in the Electronic Society (WPES)*, pages 71–80, 2005.
14. R. Ravichandran, M. Benisch, P. Kelley, and N. Sadeh. Capturing social networking privacy preferences: can default policies help alleviate tradeoffs between expressiveness and user burden? In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 1. ACM, 2009.
15. N. Sadeh, J. Hong, L. Cranor, I. Fette, P. Kelley, M. Prabaker, and J. Rao. Understanding and capturing peoples privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, 2009.
16. R. Thaler and C. Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Yale Univ Pr, 2008.
17. J. Y. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, In press, 2010.