Privacy exposed

Wottrich, V.M.


1

INTRODUCTION AND DISSERTATION OUTLINE

INTRODUCTION

Few devices know more about us than the smartphones and tablet PCs in our pockets. By downloading and using mobile applications ("apps"), we constantly and often unwittingly leave invisible traces that are collected, stored, used, and auctioned off by third parties, such as app developers, data brokers, analytics companies, and marketers (Eastin, Brinson, Doorey, & Wilcox, 2016; Sipior, Ward, & Volonino, 2014; Thurm & Kane, 2010; van Dijck, 2014). The tracking of sensitive information, such as location information (e.g., GPS), address book contacts, in-app behavior, or sensor data (e.g., cameras, microphones) has benefits for both consumers and businesses. By agreeing to the information requests of mobile apps, consumers can, for instance, get unprecedented, instant, and often free access to information, entertainment, and social interaction at any time and from any place (Eastin et al., 2016). For businesses, highly personalized, rich data generated through consumers' use of mobile apps create valuable insights about consumers' actual behavior, thoughts, and sentiments (Buck, Horbel, Kessler, & Christian, 2014; van Dijck, 2014), which can be used to improve business efficiency by narrowly tailoring marketing efforts and sales strategies to individual customers (Ashworth & Free, 2006; Nissenbaum, 2009; Vesanen, 2007).

However, the individual and economic opportunities associated with the data collection and usage practices of mobile apps do not come without risks. These practices also raise concerns about consumer privacy (Boyles, Smith, & Madden, 2012). A growing industry is assembling the gathered data into consumer profiles, which are often shared via various platforms and networks with third-party delivery services, marketers, and analytics companies (Sipior et al., 2014; Thurm & Kane, 2010). In fact, a recent report commissioned by the Norwegian Consumer Council revealed that the 21 most popular apps in Norway (e.g., Facebook, Snapchat, Tinder) transmitted potentially sensitive consumer data to approximately 600 different first- and third-party domains[1] (Pultier, Harrand, & Brandtzæg, 2016). The gathered information is regularly stored in large data assemblages, which are vulnerable to security breaches and routinely contain errors and inaccuracies (Nissenbaum, 2009). Moreover, the information is often used for discriminating between users in buying situations, social sorting, and (hidden) manipulation (Bauman & Lyon, 2013; Esposti, 2014; Nissenbaum, 2009; Turow, 2013).

[1] First-party domain refers to websites or apps users explicitly aim to interact with (e.g., Facebook). Third-party domain refers to parties outside the first-party domain (e.g., external advertising networks publishing ads on Facebook).

Currently, users barely have influence on the gathering of personal information via mobile apps. As they cannot selectively grant or decline certain permission requests or simply "opt out" of the tracking, the guiding principle is often "all-or-nothing": accept the information request or do not install the app (Egelman, Felt, & Wagner, 2013; Sipior et al., 2014; Thurm & Kane, 2010). So far, there is mixed evidence concerning how consumers respond to this situation. On the one hand, prior research has shown that app users will uninstall or decline to install apps due to concerns about privacy and that they are willing to pay premium prices for apps that access less personal information (Boyles et al., 2012; Egelman et al., 2013; Shklovski, Mainwaring, Skúladóttir, & Borgthorsson, 2014). On the other hand, there are also signs that app users do not consider or understand the (consequences of the) information requests of apps and engage in privacy-risky behavior in exchange for convenience, functionality, or small financial gains when they think the utility of the app is high enough (Felt et al., 2012; Good et al., 2005; Grossklags & Acquisti, 2007). As these mixed findings show, it is still unclear how consumers respond to data collection and usage practices of mobile apps. Therefore, this dissertation investigates (1) the status quo of privacy protection behavior, (2) the drivers of information disclosure, and (3) the consequences of information disclosure in the privacy-sensitive context of mobile apps.

What Does Privacy Actually Mean?

There is a vast landscape of theoretical and empirical work on the concept of privacy, spanning disciplines from political science to information law, philosophy, psychology, communication, media and information studies, computer science, and marketing (Nissenbaum, 2009). It is probably due to these multidisciplinary perspectives that there is not one overarching definition of privacy. This dissertation approaches the concept from a communication science perspective, distinguishing four aspects of privacy (Burgoon, 1982): (a) *physical privacy*, which is concerned with an individual's physical (in)accessibility and control over spatial intrusions; (b) *social privacy*, which refers to the process of regulating distance and proximity in social encounters; (c) *psychological privacy*, which captures an individual's control over affective and cognitive inputs and outputs; and (d) *informational privacy*, which is concerned with an individual's level of control over the collection and dissemination of personal information. As privacy in the mobile app context is mainly related to the collection and usage of personal information, the focus of this dissertation lies on *informational privacy*.

Existing literature on informational privacy and law most commonly characterizes privacy as a form of control (Nissenbaum, 2009). In this regard, Westin's (1967) work is often mentioned. He defines (informational) privacy as "the claim of individuals,

groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others" (Westin, 1967, p. 7). Westin's conceptualization of privacy has been used widely in informational privacy literature, yet, during the past years, another approach has gained more and more attention: Nissenbaum's (2009) framework of contextual integrity.

According to Nissenbaum (2009), "privacy is neither a right to secrecy nor a right to control but a right to *appropriate* flow of personal information" (p. 127). The basic assumption of the contextual integrity framework is that we are living in different contexts, broadly defined as "abstract representations of social structures experienced in daily life" (Nissenbaum, 2009, p. 134). The term "context" can be interpreted in various ways. It can, for instance, refer to social domains, such as education or family, or to technological platforms, such as social network sites or mobile apps (Nissenbaum, 2015). Each of these contexts is governed by context-relative informational norms, that is, norms "that are specifically concerned with the flow of personal information—transmission, communication, transfer, distribution, and dissemination—from one part to another, or others" (Nissenbaum, 2009, p. 140). Whether a context-specific information flow is appropriate depends on the type of information, the actors involved in the transmission (e.g., sender, receiver), and the conditions under which the information transmission occurs (Nissenbaum, 2015). In the health care context, for instance, it would be appropriate if the general practitioner asks the patient to undress for the medical examination. When practices and actions (e.g., undressing) are in line with the context-relative informational norms (e.g., patients sometimes must undress in front of the doctor), contextual integrity is achieved. However, when the informational norms are violated, for instance, when the hairdresser asks the client to undress, the information flow is inappropriate, and our sense of privacy is disturbed. Thus, context-specific informational norms structure people's privacy expectations and their responses toward information requests. Nissenbaum's (2009) definition of privacy and her framework of contextual integrity form the overarching conceptual backbone of this dissertation. However, this dissertation also considers different aspects of privacy, such as the *control* aspect introduced by Westin (1967).

The Mobile App Context

Drawing on Nissenbaum's (2009) contextual integrity framework, this dissertation focuses on the context of mobile devices, in general, and on mobile apps, in particular. The term mobile devices here refers to smartphones and tablet PCs. Mobile apps are defined as "computer programs designed to run on mobile devices" (Sipior et al., 2014, p. 179). These contexts were chosen for the following reason: Mobile devices

and mobile apps are deeply ingrained into our everyday lives. In the U.S., for instance, Internet users spend, on average, 71% of their digital media time on mobile devices (ComScore, 2017), and a total of 87% of that mobile media time is spent on mobile apps (Smart Insights, 2017). According to the Federal Trade Commission (FTC), an independent agency of the U.S. government responsible for consumer protection, mobile devices "can facilitate data collection and sharing among many entities, including […] application developers, analytics companies, and advertisers to a degree unprecedented in the desktop environment" (FTC, 2013, p. 2). This might have serious consequences not only for our privacy and security but also for the society we live in.

The personal data generated through the use of digital systems, such as mobile apps, are often centrally stored in large data assemblages, which are vulnerable to security breaches, such as break-ins by hackers, data loss, or accidental disclosure to bad actors (Nissenbaum, 2009). Just recently, in November 2017, the popular taxi app Uber disclosed that, a year earlier, hackers had stolen 57 million rider and driver accounts, including their phone numbers, e-mail addresses and names (Isaac, Benner, & Frenkel, 2017). Information that has been stolen during data breaches, such as the one Uber experienced, could be used for identity theft (Reyns, 2013) and fraud on the Internet (Narayanan, Koo, & Cozzarin, 2012). To give an example, criminals apparently used the information stolen in September 2017 from the U.S. credit-reporting agency Equifax for applying for mortgages, credit cards, and student loans, as well as for filing insurance claims, thus racking up substantial debts (Harney, 2017). Apart from that, large data assemblages are vulnerable to routinely present errors and inaccuracies, which might negatively affect our life chances and choices. For instance, if our *data-double* is factually flawed, we might get higher insurance rates or experience problems with getting a mortgage (Bauman & Lyon, 2013; Nissenbaum, 2009).

In addition, the aggregation and analysis of personal information can lead to discrimination, social sorting, and (hidden) manipulation. Currently, companies often perform a "highly controversial form of social profiling and discrimination by customizing our media content on the basis of marketing reputations we don't even know we have" (Turow, 2013, p. 2). Some people are classified as "targets", while others are considered as "waste". For example, some online travel agencies show Apple MacBook users costlier travel options than Windows users because it has been found that Mac users spend, on average, more on hotels than other users (Mattioli, 2012). As becomes clear from this example, the treatment we receive in the commercial marketplace, the access to certain kinds of consumer offers, news, and information, as well as the price of goods and services (e.g., flight tickets) often depends on what companies already know about us (Bauman & Lyon, 2013; Esposti, 2014; Nissenbaum,

2009; Turow, 2013). As a result, our society is more and more segmented into different *filter bubbles* in which everyone of us gets to see a different picture of the world (Pariser, 2012; Turow, 2013). This might have consequences for the world we live in.

After Brexit and the election of Donald Trump as 45th president of the United States of America, there have been various speculations about the role data-driven marketing and campaigning practices, such as those of the micotargeting firm Cambridge Analytica (CA), played in these events (e.g., Grassegger & Krogerus, 2016; Persily, 2017). While CA claims that their approach to data-driven communications played an integral part in Trump's electoral win (Cambridge Analytica, 2016), there is, to date, no (empirical) evidence to prove this claim (Taggart, 2017). Nevertheless, one thing is clear: The data collection and usage practices that came with the technological advancements during the past years are transforming our society. The boundaries between what is public and what is private have been blurred, and the smallest details of our lives are monitored more closely than ever before (Bauman & Lyon, 2013)—especially by mobile devices and apps. Our **privacy is exposed**.

How Do Consumers Respond to the Data Collection and Usage Practices of Mobile Apps?

To better protect consumer privacy in the rapidly changing online environment, the European Union (EU) has recently introduced a comprehensive reform of data protection rules (EU General Data Protection Regulation (GDPR), 2016), which entered into force in May 2016 and which shall apply from May 2018. The GDPR does not only regulate how businesses collect and protect consumer information, but it also aims to give citizens back control over their personal data (European Commission, 2015). Firms that wish to track (mobile) Internet users need to provide them with clear and complete information on the purpose of the tracking, and they need to obtain users' consent before storing and accessing personal data (EU Data Protection Directive, 1995; EU General Data Protection Regulation, 2016). Speaking in terms of Nissenbaum's (2009) contextual integrity framework, in the mobile app context, the transmission principle (i.e., "a constraint on the flow (distribution, dissemination, transmission) of information from party to party in a context" (p. 145)) requires the app users' knowledge ("notice") of the data collection and usage practices of apps and the users' permission ("consent").

To comply with these requirements, mobile apps typically provide app users with an overview of the information they want to access before the download. Users of the market-leading operating system Android, for instance, get to see a permissions page containing a list of all the data collected by the app, and they have to accept the app's permissions request before they can download it (Kelley, Cranor, & Sadeh,

2013; Kesswani & Lin, 2016). In theory, this should help in reducing the information asymmetry between companies and consumers (McDonald & Cranor, 2008) so that consumers are empowered enough to protect their privacy. However, consumers have difficulties understanding the language of privacy statements (Milne, Culnan, & Greene, 2006) and they barely read the privacy policies of apps (Felt et al., 2012; Liu, 2014; Shklovski et al., 2014). Against this backdrop, the question arises:

RQ 1: What is the status quo of privacy protection behavior in the mobile app context?

An often-made assumption in prior literature is that consumers respond to businesses' data requests in an informed, rational manner, weighing the costs of the information trade against its benefits (Aguirre, Mahr, Grewal, de Ruyter, & Wetzels, 2015; Ashworth & Free, 2006; Fife & Orjuela, 2012; Keith, Thompson, Hale, Lowry, & Greer, 2013; H. Li, Sarathy, & Xu, 2010). Supporting this assumption, research has shown that consumers who are concerned about their privacy engage in risk-reducing behavior, such as withholding or falsifying personal information, using privacy-enhancing techniques (e.g., encryption), rejecting unnecessary cookies, or refraining from using certain websites (Milne & Culnan, 2004; Sheehan & Hoy, 1999; Wirtz, Lwin, & Williams, 2007; Youn, 2009). In the mobile app context, app users reported they would uninstall or decline to install apps due to concerns about privacy and that they are willing to pay premium prices for apps that access less personal information (Boyles et al., 2012; Egelman et al., 2013; Shklovski et al., 2014).

However, research has also shown that many Internet and app users tend to accept almost all information requests, or they simply ignore them (Felt et al., 2012; Zuiderveen Borgesius, 2015). Moreover, it has been demonstrated that the immediate benefits from information disclosure (e.g., app use) may trump the delayed benefits (e.g., privacy protection), even among privacy-conscious users (Acquisti, Brandimarte, & Loewenstein, 2015; Acquisti & Grossklags, 2005; John, Acquisti, & Loewenstein, 2011). In the mobile app context, a higher perceived benefit of the information disclosure made consumers select riskier privacy settings in apps (Keith et al., 2013) and positively influenced their intention to download an app (Eling, Widjaja, Krasnova, & Buxmann, 2013). This evidence raises doubts about the assumption of rationality in privacy decision-making. Supporting these doubts, there is a stream of research stating that even if individuals have access to comprehensive information on privacy risks and protection possibilities, they might not be able to process this information to formulate rational privacy-sensitive decisions (e.g., Acquisti & Grossklags, 2005).

Human rationality is bounded, which limits our ability to acquire and apply information (Simon, 1982). Given these mixed findings on the drivers of information disclosure in the mobile app context, the following research question is formulated:

RQ 2: Which factors drive information disclosure in the mobile app context?

Highly personalized, rich data generated through consumers' use of mobile apps are often considered as "the oil of the 21st century" (BBC, 2013) because the data provide valuable insights about consumers (Buck et al., 2014; van Dijck, 2014), which can be used to narrowly target individuals with various forms of marketing communications (Ashworth & Free, 2006; Nissenbaum, 2009; Vesanen, 2007). It is also due to the tracking possibilities of apps that marketers are increasingly interested in employing mobile apps to interact with potential and existing customers (Eastin et al., 2016; Steel, 2013; Thurm & Kane, 2010; Wang, Kim, & Malthouse, 2016). These kinds of apps are often referred to as "branded apps", here defined as mobile apps that (1) prominently display a brand identity or logo, (2) originate from brands offering products and/or services in exchange for money, and (3) are used by companies as an additional marketing tool. So far, existing research on branded apps has mainly focused on their advertising effectiveness (Bellman, Potter, Treleaven-Hassard, Robinson, & Varan, 2011; Kim, Wang, & Malthouse, 2015). What has been neglected is that the data collection and usage practices of branded apps can be considered as intrusive by consumers (Boyles et al., 2012), which might have a damaging effect on consumers' app and brand perceptions (Pan & Zinkhan, 2006). Consequently, the following research question was formulated:

RQ 3: What are the consequences of information disclosure in branded apps on consumers' app and brand perceptions?

DISSERTATION OUTLINE AND FINDINGS

The aim of this dissertation is to investigate how consumers currently respond to data collection and usage practices of mobile apps. The following research questions will be answered:

1. What is the status quo of privacy protection behavior in the mobile app context?
2. Which factors drive information disclosure in the mobile app context?
3. What are the consequences of information disclosure in branded apps on consumers' app and brand perceptions?

This dissertation consists of four empirical studies that are based on seven different datasets gathered among more than 4,000 participants. Each study is presented in a separate chapter and has its own abstract, theoretical background, method, results, and conclusions. All studies are published or submitted for publication in peer-reviewed academic journals. Figure 1.1 provides a graphical overview of the outline of this dissertation.

The first empirical study reported in **chapter 2** explores the status quo of privacy protection behavior in the context of mobile apps (RQ 1). By means of an online survey conducted among 1,593 Dutch app users, it is demonstrated that app users' current knowledge about the data collection and usage practices of mobile apps is very limited. Moreover, app users only report moderate levels of privacy concern, perceived vulnerability, self-efficacy, and protection motivation, and rather low levels of protection behavior. In addition, the study shows that app users are more likely to engage in privacy protection, if they feel vulnerable, concerned, and think that they can protect themselves from the data collection and usage practices of apps. Surprisingly, higher levels of knowledge about the data collection and usage practices of apps were not associated with more, but with less, protection motivation and behavior. These findings raise doubts concerning the assumption of informed privacy decision-making in the context of mobile apps.

Chapter 3 and chapter 4 investigate the drivers of information disclosure in the mobile app context (RQ 2). **Chapter 3** examines the privacy trade-off for mobile app downloads. Two separate online experiments (n_1 = 183; n_2 = 687) examine the effects of app value and app intrusiveness on app users' intention to accept permission requests in the app download stage and the role app users' privacy concerns play in this regard. The results show that app users do seem to engage in a privacy trade-off when downloading mobile apps, in which they weigh the costs of the information

trade (i.e., app intrusiveness, privacy concerns) against the benefits (i.e., app value). However, in this trade-off, the immediate value gained from information disclosure via mobile apps trumps the cost related to the information trade. **Chapter 4** focuses on a different kind of app, which has been often used before mobile apps became popular. The chapter investigates how online games originating from brands (i.e., advergames) "seduce" users to actively share personal information. A lab experiment (n = 181) examines the effects of advergame customization features and trust in the brand advertised in the game on personal information disclosure and brand attitude. Moreover, it examines to what extent privacy concerns moderate these effects. The results show that customization features and brand trust have a positive influence on brand attitude and information disclosure, but this influence is strongly conditioned by consumers' privacy concerns.

Chapter 5 focuses on the consequences of the data collection and usage practices of branded mobile apps (RQ 3). Three separate online experiments (n_1 = 190; n_2 = 201; n_3 = 1,004) explore the effect of branded app intrusiveness on consumers' app and brand perceptions. The results show that app intrusiveness has a negative effect on app and brand perceptions for fictitious branded apps, but not for real branded apps.

After these four empirical chapters, **chapter 6** presents the five main conclusions of this dissertation. Based on these conclusions, practical implications for policy makers, mobile app users, and marketers are provided. Moreover, suggestions for future research are discussed.

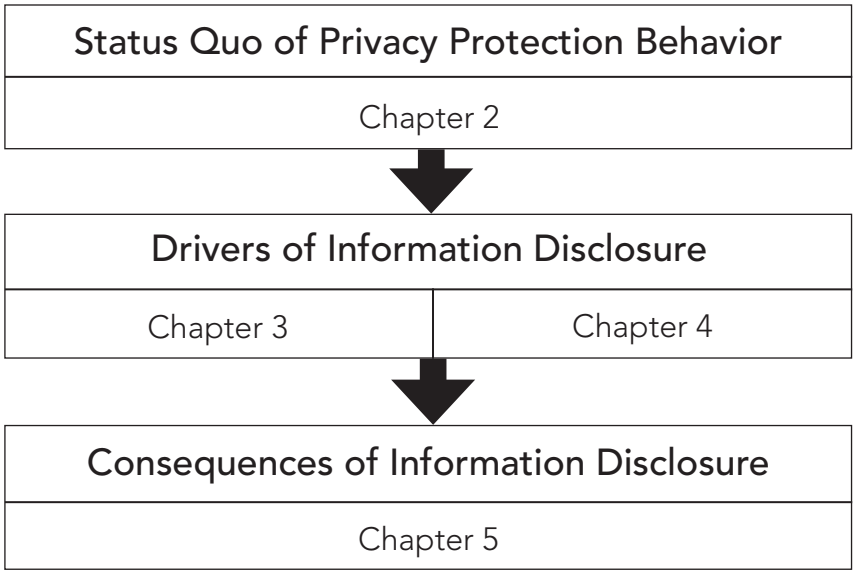

Figure 1.1 Dissertation outline.