Why Experience Matters to Privacy: How Context-Based Experience Moderates Consumer Privacy Expectations for Mobile Applications

Kirsten Martin
Strategic Management and Public Policy, School of Business, George Washington University, Washington, DC 20052. E-mail: martink@gwu.edu

Katie Shilton
College of Information Studies, University of Maryland, 4121H Hornbake Building, South Wing, College Park, MD 24971. E-mail: kshilton@umd.edu

Two dominant theoretical models for privacy—individual privacy preferences and context-dependent definitions of privacy—are often studied separately in information systems research. This paper unites these theories by examining how individual privacy preferences impact context-dependent privacy expectations. The paper theorizes that experience, provides a bridge between individuals' general privacy attitudes and nuanced contextual factors. This leads to the hypothesis that, when making judgments about privacy expectations, individuals with less experience in a context rely more on individual preferences such as their generalized privacy beliefs, whereas individuals with more experience in a context are influenced by contextual factors and norms. To test this hypothesis, 1,925 American users of mobile applications made judgments about whether varied real-world scenarios involving data collection and use met their privacy expectations. Analysis of the data suggests that experience using mobile applications did moderate the effect of individual preferences and contextual factors on privacy judgments. Experience changed the equation respondents used to assess whether data collection and use scenarios met their privacy expectations. Discovering the bridge between 2 dominant theoretical models enables future privacy research to consider both personal and contextual variables by taking differences in experience into account.

Received November 4, 2014; revised December 16, 2014; accepted December 17, 2014

 DOI: 10.1002/asi.23500

Introduction

Individuals use mobile applications to socialize, communicate, play, shop, bank, and search for information. As of January 2014, 90% of American adults have a cell phone (Pew Research Internet Project, 2013), 63% of adult cell owners use their phones to go online (Pew Research Internet Project, 2013), and 50% download mobile applications (Duggan, 2013). During these activities, mobile applications collect personal data to facilitate both services and advertising. According to the *Wall Street Journal*, mobile advertising grew 110% in 2013 to a $7.1 billion industry, which marks the third year of triple-digit growth (Perlberg, 2014). This industry relies on user data to track consumer behavior and target advertisements.

As the technical possibilities for data collection expand, and the prevalence of data sharing in the mobile industry grows, it is unclear what data uses consumers will tolerate. In the US, consumer privacy expectations are a common indicator of whether corporate data collection and use are considered acceptable to users (Strickland & Hunt, 2005). Consumers, organizations, and regulators struggle to address privacy expectations for these new forms of data collection across a diverse set of activities (Boyles, Smith, & Madden, 2012; Urban, Hoofnagle, & Li, 2012). Additionally, perceived privacy violations can cause consumer backlash against technology developers (Jackson, Gillespie, & Payette, 2014; Kang, 2013). Providing better guidance on consumers' privacy norms and expectations can support developers and firms as they work to create both innovative and fair technologies.

The information systems literature, and privacy scholarship in general, is divided as to how to examine users' privacy expectations: privacy expectations are understood

both as individual, highly variable preferences of a static definition of privacy (Buchanan, Paine, Joinson, & Reips, 2007; Chellappa & Sin, 2005; Westin, 1970; Yao, Rice, & Wallis, 2007) as well as social norms that depend on ~~the~~ contextual factors such as the type of information collected, who receives the information, the purpose of the information, and how the information is handled (Barth, Datta, Mitchell, & Nissenbaum, 2006; Nissenbaum, 2009; Vasalou, Joinson, & Houghton, forthcoming). These two bodies of privacy literature are often examined separately or even presented in opposition. Little work has been done to identify how these two theoretical approaches might intersect.

This paper examines the relationship between individual privacy dispositions and contextual privacy judgments in the mobile application space. ~~Inspired~~, by theory that explains how experience in an industry modifies consumers' trust judgments, we hypothesize that consumers' experience within a context impacts not only their trust judgments but also their privacy expectations. In other words, we hypothesize that experience provides the bridge between general privacy attitudes and nuanced contextual factors in privacy judgments. We surveyed American users of mobile applications and asked them to make judgments about whether context-specific scenarios met their privacy expectations. We chose to investigate the privacy expectations of mobile application users because the mobile application space presented both innovative forms of data collection and use, as well as difficult questions about whether or how that data collection and use should be controlled or regulated (Federal Trade Commission, 2012). The survey used a factorial vignette method to ask respondents to judge whether common and realistic forms of data collection (e.g., harvest of application usage, location, image, and contact data) and data use (tracking and targeting) by mobile applications met their privacy expectations.

Analysis of these data illustrated significant relationships between individual factors (e.g., consumer experience and general privacy attitudes) and contextual privacy expectations. While the reported frequency of app use did not impact individuals' privacy expectations directly or their general belief that privacy is important, more frequent users of mobile apps relied *less* on their general privacy disposition in making particular privacy judgments and gave greater weight to contextual factors. Frequent users of mobile applications showed greater sensitivity to contextual variables; less frequent users relied primarily on their general privacy attitude to make judgments. Experience impacted how vignettes met privacy expectations by changing the equation respondents used to assess whether vignettes met their expectations.

First, the Literature Review discusses related work on individual privacy preferences as well as contextual integrity, and discusses the role of experience in the related domain of consumer trust. Next, the Methods section describes the survey procedures. We explore the findings in the Results section, and the implications in the Discussion section. Finally, the Conclusion explores the impact of these findings for mobile application developers and policymakers.

Literature Review: Privacy and Trust in Mobile Application Use

Consumers and Privacy Judgments

In information systems research, two major approaches to understanding consumer privacy expectations focus on (a) individuals' privacy attitudes and (b) situational or contextual privacy expectations. The first approach explores privacy expectations as individual dispositions, beginning with Westin's (1970) surveys that categorized individuals on a spectrum between privacy fundamentalists and the privacy unconcerned. This approach continues in surveys and studies which ask consumers general questions about their privacy preferences (Boyles et al., 2012; Urban et al., 2012) or in privacy research focused on privacy concerns of consumers as an individual attribute (Buchanan et al., 2007; Smith, Milberg, & Burke, 1996; Yao et al., 2007). Surveys in this space range from asking about broad concerns, such as *How concerned are you about threats to your privacy . . .* (Nguyen, Bedford, Bretana, & Hayes, 2011), to more specific concerns about particular issues, such as surveys which measure aspects of concern for information privacy (CFIP) (Smith et al., 1996). In addition to concern for privacy, studies have asked about individuals' valuation of privacy, interpreted as the general *value that individuals assign to the protection of their personal data* (Acquisti & Varian, 2005; Acquisti, John, & Loewenstein, 2013; Chellappa & Sin, 2005). Importantly, privacy is defined independently of context in this research, and investigators measure the degree to which individuals value, or are concerned with, a fixed definition of privacy.

The second approach to privacy scholarship defines privacy as a collective (rather than individual), contextually dependent phenomenon. This approach posits that privacy expectations are based on social norms within particular information contexts (Nissenbaum, 2009). Those privacy norms dictate what data are acceptable to collect, who can have access to it, whether ~~it~~ should be kept confidential, and how ~~it~~ can be shared and reused. Called "privacy as contextual integrity" by Nissenbaum (2009), this approach suggests that instead of measuring privacy concerns or expectations as static attributes of individuals, we measure privacy as responses to context-specific rules. When privacy expectations are context-specific, norms around what information should be disclosed and gathered and for what purpose are developed within a particular community or context. A context-specific definition of privacy is consistent with a social contract approach to privacy expectations (Culnan & Bies, 2003; Li, Sarathy, & Xu, 2010; Martin, 2012; Xu, Zhang, Shi, & Song, 2009), in which rules for information flow take into account the purpose of the information exchange as well as risks and harms associated with sharing information. The process of taking contextual variables into account is also known as the *privacy calculus*, where privacy norms are developed with the costs and benefits of sharing information in mind. However, because "contexts" can be difficult to precisely define (Vasalou et al.,

DOI: 10.1002/asi

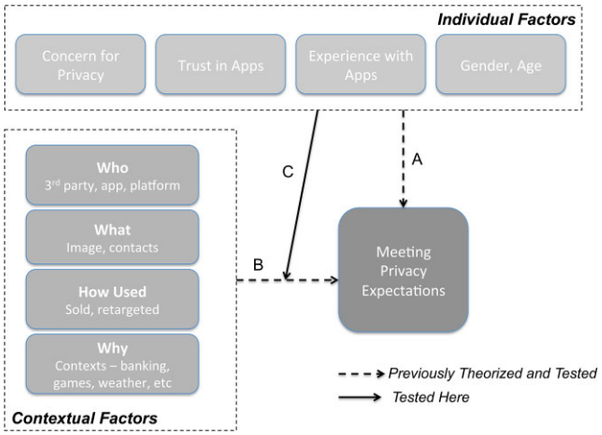

FIG. 1. Conceptual model of privacy judgments.

forthcoming), it is a challenge to measure contextual variables (Shilton, 2012).

Integrated Model of Privacy Judgments

These two approaches to understanding privacy expectations—one focused on individual beliefs or attributes and one focused on contextual definitions—are normally studied independently, as shown in Figure 1 (arrows A and B). Limited work has sought to identify how individual beliefs and contextual factors work in tandem. Scholars have examined how individuals' concerns for privacy can be shifted or modified by context (Malhotra, Kim, & Agarwal, 2004), yet work on context-dependent definitions of privacy have yet to consider how individual attributes may impact how contextual factors are considered. In fact, previous work suggests that contextual factors do not explain all of the variance in individuals' privacy judgments (Martin, 2013; Shilton & Martin, 2013), and individual differences in perceptions, beliefs, and attitudes remain important to privacy expectations. We model the possibly moderating relationship of individual factors on consideration of contextual factors in Figure 1 (arrow C).

Experience, Institutional Trust, and Privacy

To investigate how individual differences might moderate contextual considerations, we turned to a related branch of theory discussing consumer trust. Trust is defined as the willingness to accept vulnerability to the actions of an organization (Mayer, Davis, & Schoorman, 1995), and may be directed towards an individual, group, organization, or institution (Pirson, Martin, & Parmar, 2013). Trust has been found to be closely related to privacy (Pavlou, 2011). Trust may be a more salient predictor of behavior than privacy concerns (Eastlick, Lotz, & Warrington, 2006; Sultan & Rohm, 2004; Van Slyke, Shim, Johnson, & Jiang, 2006), leading to scholars to call for privacy research to include the effects of trust (Belanger, Hiller, & Smith, 2002; Pavlou, Liang, & Xue, 2007; Van Slyke et al., 2006).

Generalized trust is a nonreflective attitude of the public towards a social institution and is measured by surveys of public trust in congress or public trust in business (Stevenson & Wolfers, 2011). Institutional trust captures individuals' assessment of "favorable conditions" for participating in transactions through norms, procedures, and controlling mechanisms, and is specific to a context such as industry or type of business (e.g., banking, social networking) (Pirson, Martin, & Parmar, 2014). Finally, stakeholder trust is an individual's willingness to accept vulnerability to the actions of a particular organization. Stakeholder trust is reflective, informed, and influenced by the actions of the firm (e.g., with enough bad actions, firms can destroy stakeholder trust). Importantly for this research, experience facilitates individuals' transitions from generalized trust to stakeholder trust. Individuals rely on generalized and institutional trust to make decisions when they have limited experience with a firm (Pavlou & Gefen, 2004). Through familiarization and experience, individuals begin to form stakeholder trust by taking into account the firm's behavior.

We began to draw analogies between this literature and a progression from individual privacy preferences to reactions to the context-appropriateness of data use scenarios. Noting this similarity led to the hypotheses we will explore here: that experience is not only an important influence on trust, but also an influence on the progression from generalized privacy beliefs to contextual privacy judgments. Similar to how experience modifies trust judgments, this project hypothesizes that individuals with less experience in a context rely more on generalized privacy beliefs, and individuals with more experience in a context are increasingly influenced by contextual factors and norms. This hypothesis is explored in more detail in the following section.

Hypotheses on the Moderating Effect of Experience on Privacy Judgments

Trust develops over time. Initial trust is not based on experience, but instead on an individual's general disposition (McKnight, Cummings, & Chervany, 1998). Interactions between an individual and firm then increase the knowledge on which individuals base judgments (Pirson et al., 2014). In a similar manner, we hypothesize that as consumers become more experienced in a context, they learn more about and become accustomed to the context's informational norms. As they learn these norms through experience, they need not rely on their general privacy beliefs and will begin to take contextual factors into account while making privacy judgments. Therefore, we hypothesize:

> **H1:** Frequency of application use moderates the role of individual dispositions in making specific privacy judgments: more frequent users of mobile apps rely *less* on general privacy dispositions in making particular judgments about meeting privacy expectations.

DOI: 10.1002/asi

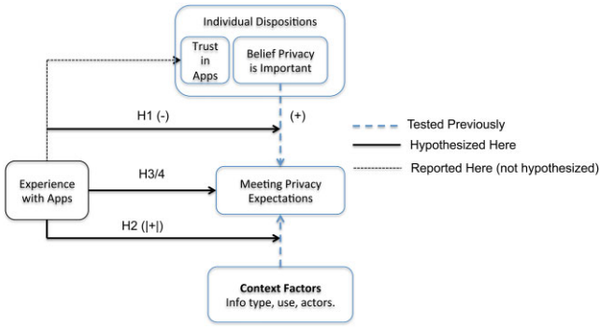

FIG. 2. Hypotheses of the role of mobile application experience in privacy judgments.

> **H2:** Frequency of application use moderates the relative importance of contextual factors in making specific privacy judgments: more frequent users of mobile apps rely *more* on contextual factors in making particular judgments about meeting privacy expectations.

An individual's decision to transact with a firm is based on their perceived trust in the firm or in a broader institution, as well as perceived risk (Gefen & Pavlou, 2012). Individuals are more willing to transact, and with greater frequency, if they perceive lower risk based on an understanding of the "rules of transaction conduct" (Gefen & Pavlou, 2012, p. 942). These hypotheses frame privacy norms and expectations as an important "rule of conduct" for mobile applications, and predict that individuals who transact with mobile applications more frequently will have a better understanding of, and agreement with, rules of engagement. Insiders have a better understanding of contextual privacy norms compared to outsiders (Martin, 2012). We therefore predict that high-frequency users of mobile applications have less variance in the privacy judgments across respondents (i.e., show more agreement with each other regarding privacy expectations), as well as more certainty in their own judgments, leading to the following hypotheses:

> **H3:** Users with more frequent use of mobile applications will show less variation in their contextual privacy judgments across respondents.

> **H4:** Users with more frequent use of mobile applications will show greater certainty in their contextual privacy judgments.

Our hypotheses are summarized in Figure 2.

Methods

To measure the effects of contextual variables in addition to personal variables, we conducted a survey using factorial vignette methodology (Wallander, 2009). ~~Factorial vignette survey methodology~~ was developed to investigate human judgments (Jasso, 2006; Rossi & Nock, 1982; Wallander, 2009). In a factorial vignette survey, a set of vignettes is generated for each respondent, where a set of vignette factors (the independent variables) controlled by the researcher are randomly selected, and respondents are asked to evaluate the resulting hypothetical situation. The evaluation is accomplished through a rating task designed to capture the level of an outcome corresponding to the survey's unit of analysis (Jasso, 2006). The factorial vignette method~~ology~~ enables researchers to simultaneously examine multiple factors (Ganong & Coleman, 2006). Using vignettes enables researchers to capture respondents' *equations-inside-the-head* (Jasso, 2006) about judgments of complex constructs such as privacy expectations. In this study, the rating task included factors such as types of mobile applications, data type, and data uses, and asked respondents to indicate the degree to which an application in the vignette met their privacy expectations. Statistical techniques were used to identify the relative importance of each vignette factor in driving the respondents' outcome.

We constructed vignettes by varying several contextual factors comprising our independent variables, as will be described here. A software interface created a deck of vignettes for each respondent randomly with replacement as the respondent took the survey. Vignette rating, factor levels, the vignette script, and the vignette sequence number were preserved. We tested vignette wording for clarity with graduate students in several courses at the University of Maryland, College Park.

Control Variables

The factorial vignette survey method~~ology~~ captures the relative importance of contextual factors for respondents. By capturing individual attributes separately, we were able to examine the relative impact of both individual and contextual privacy factors. Before beginning the vignettes, the system supplied respondents with several questions to measure individual attributes, used as control variables to compare respondents. We collected respondents' age and gender, as well as responses to two rating tasks important to the hypotheses. The first asked participants to rate on a scale from *strongly disagree* to *strongly agree* their agreement with the statement: "In general, I trust mobile applications." This rating captured respondents' institutional trust in the mobile application industry. The second rating task asked for their agreement with the statement: "In general, I believe privacy is important." This rating captured respondents' general privacy belief.[1] Finally, respondents reported their frequency of use of mobile applications (Table 1), given a scale of options ranging from never to frequently.

[1]Although individual privacy belief has traditionally been measured with a multi-question indicator (Buchanan et al., 2007) such as the Internet Users' Information Privacy Concern (IUIPC) scale, recent research has shown that responses to these questions are highly correlated, and a single question can capture generalized privacy belief. See (Garg, Benton, & Camp, 2014) for details.

DOI: 10.1002/asi

TABLE 1. Respondent self-reported frequency of application use distribution.

		Targeting Survey		Tracking Survey		Combined	
		Respondents in Sample		Respondents in Sample		Respondents in Sample	
Frequency	AppUse	#	%	#	%	#	%
Never	1	13	1.3%	18	1.9%	31	1.6%
Rarely	2	80	8.2%	76	8.0%	156	8.1%
1–2 Weekly	3	28	2.9%	29	3.1%	57	3.0%
1–2 Daily	4	256	26.2%	238	25.1%	494	25.7%
All the time	5	592	60.7%	572	60.3%	1,164	60.5%

Independent Variables

Defining meaningful social contexts is a challenge for research into privacy as contextual integrity. We chose to define contexts according to industry, which replicates the structure of how mobile phone software is built and delivered. Applications are usually developed and marketed for a single purpose: communicating with your bank, playing a game, keeping your calendar, etc. We first classified industry contexts according to their identification by the two major application stores: the Apple iTunes store and Google Play. We also sought industry data on dominant uses of mobile applications. According to an industry survey, e-mail and calendaring, instant messaging (IM), office and personal productivity, web conferencing, and e-commerce are the most popular uses of mobile applications (Columbus, 2013).

Using these data, we chose mobile application industry contexts based on a combination of popularity and diversity. We chose the most popular application contexts, as well as those known to have sensitive data in face-to-face transactions, such as medical and banking contexts.

- Games
- Weather
- Social networking
- Navigation
- Music
- Banking/Finance
- Shopping/Retail
- Productivity

After choosing industry contexts to test, we chose a series of other independent variables based on the factors that Nissenbaum (2009) identifies as relevant to privacy as contextual integrity: who the data collection actors are, the type of data collected, how the data are used, and why the data are collected (see also Martin, 2012). Each survey respondent was shown a series of vignettes that varied based on:

- *Who*: The primary organization collecting information, such as application developer, mobile phone provider, or a third party;
- *What*: The type of information received or tracked by the primary organization, including location, accelerometer, demographic, images, keywords, name, friends, or contact lists;
- *Why*: The application purpose—playing games, checking weather, participating in social networking, navigating using maps, listening to music, banking, shopping, and organizing personal productivity.[2]
- *How (used):* The transmission principles (Nissenbaum, 2009)—for example, how the data are reused or stored, and whether it is used to target ads or track users over time.

We developed realistic scenarios by choosing organizations, types of information, application purposes, and transmission principles known to be employed in today's mobile sector (Urban et al., 2012). Figure 3 contains an example of the vignette survey given to respondents. For example, type of information included location, accelerometer, demographic, contacts, keywords, name, images, and friends, all of which are currently collected by mobile applications in the marketplace. Targeting vignettes contained ad type (what the organization does with the information, either using it to target their own ads or selling it to a third party). Scenarios about tracking mobile data also included retention (the length of time data were stored, in months), personalization (whether data were tied to a unique identifier for your mobile device), collection (who collects the information, such as the primary organization, your wireless provider, your platform provider, third-party tracking), and secondary use (what the collecting organization does with the information, such as retargeting, data exchange, or social advertising). This list of factors generated vignettes like the following (italics highlight factors that systematically changed) in Figure 4:

Dependent Variable: Vignette Ratings

For each vignette, respondents were instructed: "Tell us how much you agree with the statement below. Using a sliding scale from −100 to 100, with −100 indicating *strongly disagree* and 100 indicating *strongly agree*." For each vignette, respondents rated their agreement with the prompt,

[2]This survey includes two distinct measures of frequency of application use. Before rating vignettes, respondents provided a self-reported frequency of application use. A hypothetical measure of frequency and tenure was used and was also included as an independent variable in vignettes. Neither hypothetical frequency nor tenure was found to be a significant factor in meeting privacy expectations.

DOI: 10.1002/asi

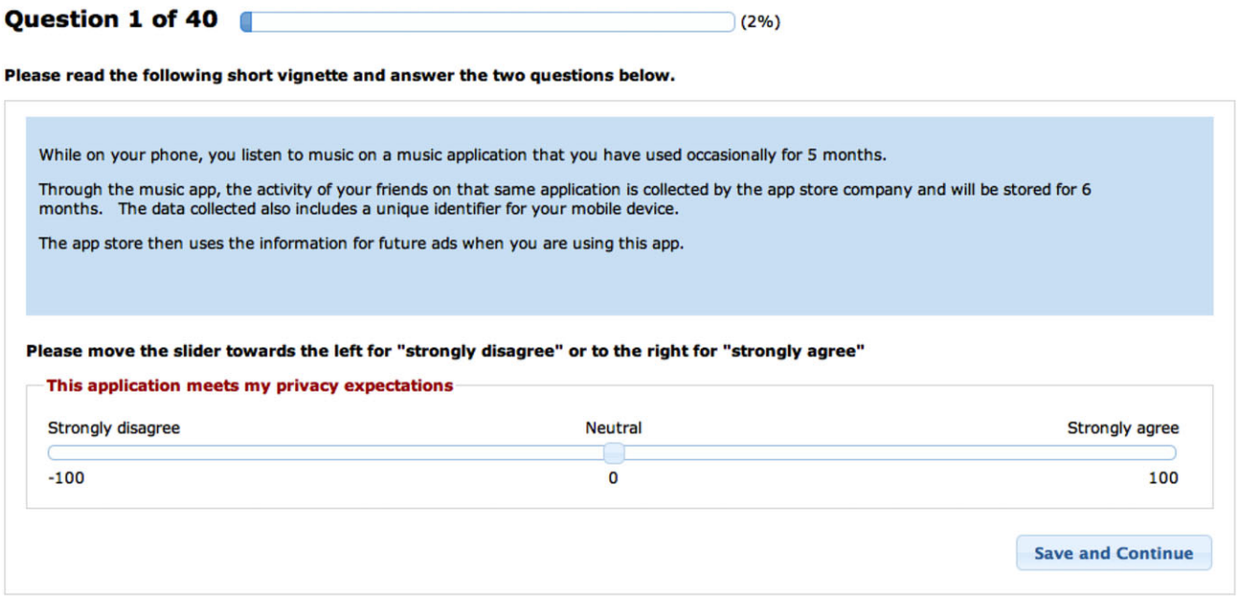

FIG. 3. Screenshot of a vignette survey question. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

SAMPLE VIGNETTES (additional in Appendix):
Targeting Sample Vignette: While using your phone, you check updates on a social networking application that you have used occasionally for less than a month. The social networking app shows you an advertisement for another application they sell based on your phone contact list.
Tracking Sample Vignette: While on your phone, you update your to-do list on a scheduling app that you have used infrequently for 3 months. Through the scheduling app, your phone contact list is collected by the app store company and will be stored for less than a week. The app store company then uses the information to show future ads to your friends and contacts.
Rating Task: This website meets my privacy expectations (strongly disagree….strongly agree)

FIG. 4. Sample vignettes.

"This application meets my privacy expectations." Ratings were used as the dependent variable to analyze what independent or control variables made a difference to their ratings.

Sample

The surveys were deployed four times (May, August, and November of 2013 and February of 2014) to a total of 1,925 respondents who rated over 77,000 vignettes through Amazon's Mechanical Turk (see Table 1).[3] Although use of Mechanical Turk for survey deployment is controversial (Lease et al., 2013; Ross, Irani, Silberman, Zaldivar, & Tomlinson, 2010), studies have shown that mTurk workers are more representative of the U.S. population than the samples often used in social science research (Behrend, Sharek, Meade, & Wiebe, 2011; Berinsky, Huber, & Lenz, 2012). Given resource constraints and the need to use a web-based platform for survey delivery, Mechanical Turk provides the best available platform for this survey. In a separate survey on privacy expectations for websites using a similar method, we compared Mechanical Turk results with results from a nationally representative sample from KnowledgeNetworks (GfK). The mTurk sample produces the same

[3]Amazon Mechanical Turk (mTurk) is an online labor market where requestors post jobs and workers choose jobs to complete. For a full description, see Mason and Suri (2012). For details on how mTurk samples are more representative of the U.S. population than in-person convenience samples, see Berinsky et al. (2012). For the external and internal validity of mTurk, see Horton, Rand, and Zeckhauser (2011). In sum, respondent samples on mTurk are found to be representative of the general population with high internal and external validity.

DOI: 10.1002/asi

TABLE 2. Respondent statistics.

	Survey	
	Targeting	Tracking
Users	976	949
Vignettes	39,320	38,160
Age	31.6	32.1
Male	58.5%	55.5%
Privacy is Important	79.82	79.24
I trust mobile apps	20.26	12.97
Mean (DV)	−18.01	−42.70
R2 of Users	0.84	0.80
ICC	22.1%	34.1%

theoretical generalizations as the GfK survey, illustrating the ability to build generalizable theory from Mechanical Turk samples in online privacy studies (Martin, 2013).

Sample demographics are provided in Table 2. The average age of the respondents was 31.4 years ~~old~~ and the sample was 56.3% male. On average, respondents trusted applications generally with an average score of 20.34 when taking the targeted advertising survey and 15.58 when taking the tracking survey. In addition, the respondents found privacy to be important with an average score of 80.27 and 78.95, respectively, for targeting and tracking surveys.[4]

Analysis

We analyzed data on two levels: at the vignette level to test independent variables (contextual factors: level 1), and at the respondent level to test control variables (individual factors: level 2). If *I* is the number of the respondents with level 2 individual variables and *J* is the number of vignettes answered with level 1 factor variables, the general equation is:

$$Y_{ij} = e_0 + e_k V_{jk} + \Sigma\gamma_h R_{hi} + u_i + e_j \quad (1)$$

where Y_{ij} is the rating of vignette k by respondent in, V_{jk} is the k^{th} factor of vignette j, R_{hi} is the h^{th} characteristic of respondent i, β_0 is a constant term, s_k and γ_h are regression coefficients for *k* vignette factors and *h* respondent factors, u_i is a respondent-level residual (random effect), and e_k is a vignette-level residual. The model conceptualized the ratings as a function of the contextual factors described in the vignette (ΣV_k) and the characteristics of the respondent (ΣR_h) as hypothesized earlier.

As the data can be modeled at two levels—the vignettes and the individual respondents—multilevel modeling was used to control for and measure individual variation in privacy judgments. Both OLS regressions as well as hierarchical regressions (xtmixed in STATA) were used to analyze the data to account for the possibility that the error terms were not equal across individuals. Finally, a respondent-specific equation (Jasso, 2006) was developed by regressing the rating task on to the contextual factors for each respondent ($N = 40$). A new data set was formed with 976 rows for the targeting vignettes and 949 rows for the tracking vignettes with a privacy equation for each respondent. The respondent-specific equation includes the respondent's intercept, the relative weight for each contextual factor, and a respondent-specific R^2 as in equation (2) below and used to test H4.

$$Y_i = \beta_i + \Sigma\beta_k V_k + e_i \quad (2)$$

The hypotheses and findings are summarized in Table 3.

TABLE 3. Hypotheses and results.

Hypothesis	Findings
H1 Frequency of application use moderates the role of individual dispositions in making specific privacy judgments	**Supported.** Frequency of app use is associated with less reliance on general privacy belief while making privacy judgments as shown in Tables 4a and 4b and Figure 5—particularly for scenarios around targeted advertising.
H2 Frequency of application use moderates the relative importance of contextual factors in making specific privacy judgments.	**Supported.** Frequency of app use is associated with greater reliance on contextual factors while making privacy judgments as shown in Tables 4a and 4b. The more frequently a respondent uses applications, the greater weight they place on contextual factors in making privacy judgments.
H3 Users with more frequent application use will show less variation in their contextual privacy judgments across respondents	**Supported.** Individuals with more frequent application use had *less variance* across respondents in their privacy judgments as shown in Tables 4a and 4b. In other words, high-frequency users have greater agreement across respondents.
H4 Users with more frequent use of mobile applications will show greater certainty in their contextual privacy judgments.	**Supported.** Individuals with more frequent application use had greater certainty in their privacy judgments. In other words, respondents with more frequent application use showed greater internal consistency in their privacy judgments as shown in Figure 6.

Results

On average, the scenarios presented did not meet user privacy expectations: Average ratings of both the

[4]Respondent fatigue was checked by controlling for later vignettes in the respondents' sequence (the sequence number of the vignette was captured and ranged from 1–40). While respondent fatigue was not a factor, we found a respondent learning curve to be important to check; respondents appear to take 1–2 vignettes to acclimate to the format. The analysis was run minus the first two vignettes for each respondent and the results remained the same.

DOI: 10.1002/asi

targeting and tracking vignettes were negative, indicating that respondents did not agree with the statement "This application meets my privacy expectations." Although the range of scenarios presented was representative of data collection and use in the mobile sector, respondents found the scenarios to be surprising. This finding bolsters research that indicates American consumers are concerned about common forms of data collection and use in the mobile sector (Urban et al., 2012) and the broader data marketplace (Madden, 2014).

The dependent variable—the degree to which a vignette met the respondent's privacy expectations—was regressed on the vignette contextual factors as well as the individual control attributes. This multilevel analysis allows the respondents' intercept and error terms to vary: 22.1% of the rating task variance for targeting scenarios, and 34% for tracking scenarios, were attributable to individuals.

In order to understand the relationship between general trust and application use, the respondent's trust in applications rating (the degree to which the respondent agreed with the statement "In general, I trust mobile applications") was regressed on the frequency of applications use (1 = *never*, . . . 5 = *all the time*). For each increase in the frequency of application use, the respondents' trust in apps increased +9.23 ($p = .00$) for the targeting survey and +9.44 ($p = .00$) for the tracking survey, even while controlling for individual attributes such as age, gender, and the individual's general privacy belief. The more frequently respondents reported using applications, the higher their general trust in applications.

Next, to examine how respondents' use of mobile applications related to their specific privacy judgments, the reported frequency of application use was included as a control variable in a general regression of the rating task contextual and individual factors; reported frequency of application use was *not* found to be significant. In addition, the average rating of vignettes was regressed on reported frequency of application use for targeting ($\beta = -0.702$, $p = .52$) and tracking ($\beta = 0.680$, $p = .57$), further reinforcing the finding that reported frequency of application use *is not* significant for privacy judgments. When measured over all respondents, greater frequency of application use was correlated with meeting privacy expectations to a greater extent, but when individual attributes were included—such as age, gender, general trust in applications, and general belief that privacy is important—the relationship between frequency of application use and privacy judgments about specific scenarios was not significant.

Hypotheses 1 and 2 suggested the reported frequency of application use moderated respondents' balance between relying on their general beliefs about privacy versus contextual privacy factors when making privacy expectation judgments. Specifically, H1 suggested that reported frequency of application use moderated the role of individual dispositions in making specific privacy judgments: the more experience users reported with mobile applications, the less impact their individual general privacy belief had on their specific judgments. To test H1, the sample was first split into low- (1–3), medium- (4), and high-frequency (5) users. The privacy judgment rating task was regressed on the vignette factors and individual variables as in Table 4a,b for each of the subsamples. The coefficient for the individual general privacy belief rating was the relative importance of this indicator in making a particular judgment. After comparing the coefficients across samples, that is, a Chow test (Chow, 1960), the results in Table 4a,b supported H1: the relative importance of individual general privacy belief in making particular judgments *decreased* (the coefficient trends towards zero) for targeted advertising vignettes ($\beta_{low} = -0.37$; $\beta_{med} = -0.32$; $\beta_{high} = -0.19$; $prob(\chi^2) = 0.00$) and *decreased somewhat* for tracking vignettes ($\beta_{low/med} = -0.37$; $\beta_{high} = -0.24$; $prob(\chi^2) = 0.00$). Second, this moderating impact of reported frequency of use was also tested with an interaction term (AppUseLow*PrivacyImport) and graphed in Figure 5. For targeting vignettes, Figure 5 shows the steeper slope for

TABLE 4A. Regressions for low-, medium-, and high-frequency users for targeted adverting vignettes.*

Targeting vignettes			
	Low use	Med use	High use
Context			
BankingCxt	**–6.68**	**–7.00**	**–10.77**
SocialCxt			
GamesCxt			
MusicCxt	**6.04**		
ProductivityCxt			–2.60
WeatherCxt			
NavigateCxt			
ActivityCxt	n/a		
SymptomCxt	n/a		**–6.28**
(null = Retail			
Information			
AccelInfo	**–10.92**	**–10.10**	**–18.32**
ContactInfo	**–60.04**	**–61.02**	**–76.19**
KeywordInfo	5.01	**11.26**	**12.65**
FriendsInfo	**–27.80**	**–15.91**	**–21.16**
ImageInfo	**–65.03**	**–73.14**	**–84.28**
LocationInfo	**–8.03**	**–11.90**	**–14.93**
NameInfo	**–6.44**	**–14.05**	**–24.78**
(null = Demo)			
AdType			
ThirdPartyAd	**–2.82**		
(null = Primary			
Control Variables			
Male			
Age		**–1.33**	**–0.50**
AgeOver30		**12.95**	
TrustApps	**0.20**	**0.19**	**0.21**
PrivacyImportant	**–0.37**	**–0.32**	**–0.19**
_cons	**49.44**	**60.20**	**43.61**
Average Rating	–20.25	–17.79	–15.71
N	3,600	7,120	18,160
ICC	24.2%	25.7%	20.0%

Note. *Key for Table 4a,b: **bold $p < .05$;** gray $p < .10$; blank $p > .10$ and not significant.

DOI: 10.1002/asi

TABLE 4B. Regressions for low-, medium-, and high-frequency users for tracking user vignettes.

Tracking vignettes			
	Low use	Med use	High use
Context			
BankingCxt	**–8.43**	**–8.40**	**–8.90**
SocialCxt			
GamesCxt			
MusicCxt			
ProductivityCxt			
WeatherCxt			
NavigateCxt			
ActivityCxt			
SymptomCxt	**–8.66**	–3.14	**–5.48**
(null = Retail)			
Information			
AccelInfo		**–6.76**	**–2.48**
ContactInfo	**–18.66**	**–21.69**	**–22.63**
KeywordInfo	**–4.30**		**–3.26**
FriendsInfo	**–5.74**	**–7.94**	**–8.05**
ImageInfo	**–23.43**	**–28.19**	**–30.64**
LocationInfo	**–4.37**	**–7.20**	**–7.63**
NameInfo	–3.61	**–8.42**	**–9.84**
(null = Demo)			
Collecting Actor			
ThirdPartyCollect	**–3.99**	**–8.42**	**–4.75**
PlatformCollect			
WirelessCollect		–2.05	**–2.21**
(null = Primary			
Personalization			
DevicePersonal	–2.08	–1.40	
(null = Null)			
Second Use			
DataExchange2nd	**–37.51**	**–47.16**	**–49.57**
SocalAd2nd	**–18.98**	**–18.90**	**–23.41**
(null = Retarget			
Storage Months	**–0.71**	**–0.93**	**–0.59**
Control Variables			
Male	**15.98**	**10.41**	
Age			**–0.55**
AgeOver30			
TrustApps	0.16	0.21	**0.23**
PrivacyImportant	–0.29	–0.37	–0.24
_cons		**36.74**	**34.70**
Average Rating	–49.55	–43.05	–40.66
N	4120	7320	16840
ICC	45.2%	34.7%	31.5%

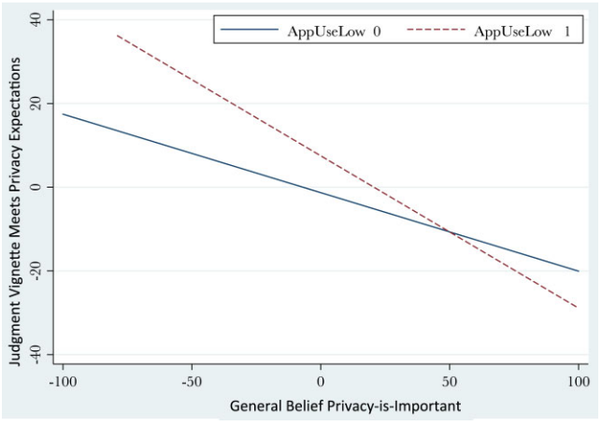

FIG. 5. Interaction between low frequency of app use and role of general privacy-is-important disposition. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

low-frequency users (Use <4), which illustrated a greater reliance on individual general privacy beliefs in making privacy judgments ($p = .02$). The results suggested that the impact of individual general privacy beliefs on privacy judgments varied based on the respondent's experience with applications. The less frequently a respondent reported using applications (and so the less experienced they were), the greater weight they placed on their general beliefs about the importance of privacy while making privacy judgments.

Hypothesis 2 suggested that reported frequency of application use also moderates the relative importance of contextual factors in making privacy judgments. To test H2, we conducted a Chow test comparing the relative importance of contextual factors across low-, medium-, and high-frequency respondents in Table 4a,b. The graphs below illustrate the impact of reported frequency of app use on the relative importance of contextual factors. The results support H2: contextual factors had greater importance in particular privacy judgments for high-frequency users compared to low-frequency users. Significant factors such as contact information ($\Delta\beta = |\ \beta_{high} — \beta_{low}| = 16.15$), image ($\Delta\beta = 19.25$), and name information ($\Delta\beta = 18.34$) for targeted advertising vignettes, as well as selling to a data exchange for tracking user vignettes ($\Delta\beta = 12.06$), increased their relative importance in driving privacy judgments as frequency of app use increased (for each $\Delta\beta$, $\text{prob}(\chi^2) = 0.00$).

Taken together, H1 and H2 suggested that individuals who use applications frequently placed a greater emphasis on contextual factors such as data type and data use in judging specific scenarios and placed less emphasis on their general belief about privacy. Respondents with low-frequency application use (less than daily) placed less importance on contextual factors and greater emphasis on their general privacy belief while making particular privacy judgments.

Hypothesis 3 suggested that users who reported greater frequency of mobile application use would show less variation in their contextual privacy judgments across respondents (i.e., would show more agreement with other respondents about privacy expectations). Hypothesis 3 was examined by comparing the intraclass correlation coefficient (ICC) produced in multilevel regressions, which measured the percent of variation in the dependent variable (meeting privacy expectations) attributable to the grouping variable (the individual). Table 4a,b includes the ICC for both targeted advertising vignettes and tracking user vignettes for low-, medium-, and high-frequency user samples. The

DOI: 10.1002/asi

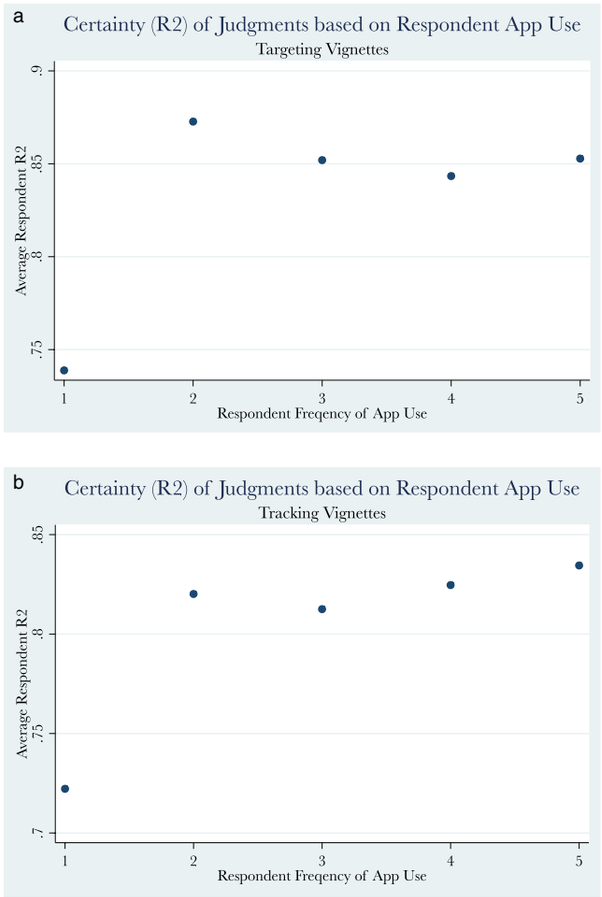

FIG. 6. **a,b**: Respondent level R2 by frequency of application use. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

results suggested that the variance in privacy judgments diverged *less* across respondents who reported greater frequency of mobile application use (24% for low-frequency and 20% for high-frequency users for targeting vignettes, and 45% for low- and 31% for high-frequency users for tracking vignettes). High-frequency mobile application users had less variation attributable to individual differences compared to low-frequency users, supporting H3. In other words, high-frequency users had greater agreement with each other about privacy expectations.

Hypothesis 4 suggested that users with more frequent use of mobile applications will show greater certainty in their contextual privacy judgments. To test H4, a respondent-level privacy equation was created by regressing the privacy judgment rating task on contextual factors (where $N = 40$ = number of vignettes rated by each respondent). A new data set was created, comprised of one row for each respondent with their own intercept, coefficients for the vignette factors, and respondent-level R^2. The results are graphed in Figure 6a,b showing a nonlinear relationship between reported frequency of application use and respondent certainty in privacy judgments. Specifically, nonusers rated contextual factors inconsistently, demonstrating that they were less certain of their privacy judgments (1 = *never* use mobile applications) than users of mobile applications.

Discussion

Experience using mobile applications moderated the effect of individual preferences and contextual factors on privacy judgments. While reported frequency of app use did not impact either individuals' overall privacy expectations or their general belief that privacy is important, more frequent users of mobile apps relied *less* on their general privacy disposition in making particular privacy judgments, and gave greater weight to contextual factors. Experience impacted how vignettes met privacy expectations by changing the equation respondents used to assess whether vignettes met their expectations.

These results have important implications for practice. The results suggest that frequent users (a) trust applications more and (b) consider the application context while making privacy judgments. This suggests that a firm's changes in privacy practices or deviations from industry data use norms may be less meaningful to novices but important to high-frequency users. Firms may need to respect contextual privacy norms to avoid alienating high-frequency (and presumably high-value) users.

Results from this study have important implications for research on mobile applications in particular and privacy more generally. Since both the importance and certainty of contextual privacy judgments may vary according to individual attributes like experience, sampling strategies in privacy research may be critically important to studies of contextual privacy. Privacy studies which rely on samples of "insiders" (e.g., frequent mobile application users or even age-similar cohorts such as students) could be limited in their generalizability. The results further validate a social contract approach to privacy (Culnan & Bies, 2003; Martin, 2012) where privacy norms are better understood by those who contract within the community. For example, experience impacts the degree to which individuals agree with each other about privacy judgments, as well as the degree of individual certainty in judgments. In the language of a social contract approach to privacy, high-frequency users are contractors within the community with a better understanding of the internal privacy norms of that context.

Finally, the findings suggest that empirical privacy research should take into account both contextual variables and individual-level attitudes, concerns, and beliefs. These research trajectories are related, and privacy researchers should pursue them in tandem, as neither individual factors nor contextual factors can explain all of the variation in users' privacy expectations. While both individual attributes and contextual factors matter; how *much* each matter to privacy judgments may depend on the individual's experience in the context and should be the subject of future research.

DOI: 10.1002/asi

A limitation of this study is that we surveyed expressed preferences (and the factors which impact those preferences), rather than revealed preferences (Garg, Benton, & Camp, 2014). Although consumers may express particular privacy preferences, they do not always act in a way that is aligned with these expressions (Acquisti & Grossklags, 2008). We believe that expressed preferences are relevant and important to the mobile sector in particular, and privacy research in general, because they reveal factors that are likely to upset consumers. Comments in the free-text portions of our surveys recorded participants' general dismay at the (realistic) data use scenarios included in our vignettes. Although individuals may not stop using mobile applications because of perceived privacy violations, consumers may lose trust in firms who repeatedly violate expressed privacy preferences. We believe it is critical, therefore, to continue to measure and report on the factors that impact expressed privacy preferences.

This work has also not fully defined or measured experience in technology use. Future work should explore what factors contribute to users' increasing reliance on contextual norms. Do consumers learn contextual norms through feedback from applications, through positive or negative experiences with platforms or applications, or learning ways of navigating privacy settings? Probing the components of experience will help us better understand the relationship between experience, trust, and contextual privacy expectations.

Conclusion

Understanding the ways that individual experience within a context moderates the relationship between individual privacy preferences and reliance on contextual norms provides a new lens for empirical privacy scholarship. This paper suggests that context-dependent investigations of privacy could be strengthened by taking into account individual differences in experience, usage frequency, and general trust in the context under investigation. The results suggest that individual privacy preferences and contextual integrity are not completely separate theories, but instead are two important factors impacting people's privacy judgments. Finally, this difference can be measured with rich vignette methods and reported to developers and firms concerned with keeping consumers' trust while collecting data from mobile applications. Understanding the nuances behind variations in users' privacy preferences can enable fair and innovative technology development.

Acknowledgments

This work benefited from feedback from participants in the Technology Policy Research Conference, the Privacy Law Scholars' Conference, and the Future of Privacy Forum. We also thank Helen Nissenbaum, Mary Culnan, and Mary Madden for constructive feedback and continued support to develop the ideas explored here. This project was funded by the National Science Foundation (Grant # SES-1311823) and the University of Maryland ADVANCE Project seed grant program.

References

Acquisti, A., & Grossklags, J. (2008). What can behavioral economics teach us about privacy? In A. Acquisti, S. Gritzalis, C. Lambrinoudakis, & S. di Vimercati (Eds.), Digital privacy: Theory, technologies, and practices (pp. 363–377). New York and London: Auerbach Publications.

Acquisti, A., & Varian, H.R. (2005). Conditioning prices on purchase history. Marketing Science, 24(3), 367–381.

Acquisti, A., John, L.K., & Loewenstein, G. (2013). What is privacy worth? The Journal of Legal Studies, 42(2), 249–274.

Barth, A., Datta, A., Mitchell, J. C., & Nissenbaum, H. (2006). Privacy and contextual integrity: Framework and applications (pp. 184–198). Presented at the SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy (S&P'06), Washington, DC: IEEE Computer Society.

Behrend, T.S., Sharek, D.J., Meade, A.W., & Wiebe, E.N. (2011). The viability of crowdsourcing for survey research. Behavior Research Methods, 43(3), 800–813.

Belanger, F., Hiller, J.S., & Smith, W.J. (2002). Trustworthiness in electronic commerce: The role of privacy, security, and site attributes. The Journal of Strategic Information Systems, 11(3), 245–270.

Berinsky, A.J., Huber, G.A., & Lenz, G.S. (2012). Evaluating online labor markets for experimental research: Amazon.com's mechanical turk. Political Analysis, 20(3), 351–368.

Boyles, J.L., Smith, A., & Madden, M. (2012). Privacy and data management on Mobile Devices. Washington, D.C.: Pew Internet & American Life Project. Retrieved from http://www.pewinternet.org/Reports/2012/Mobile-Privacy.aspx

Buchanan, T., Paine, C., Joinson, A.N., & Reips, U.-D. (2007). Development of measures of online privacy concern and protection for use on the internet. Journal of the American Society for Information Science and Technology, 58(2), 157–165.

Chellappa, R.K., & Sin, R.G. (2005). Personalization versus privacy: An empirical examination of the online consumer's dilemma. Information Technology and Management, 6(2–3), 181–202.

Chow, G. (1960). Tests of Equality Between Sets of Coefficients in Two Linear Regressions. Econometrica, 28(3), 591–605. doi:10.2307/1910133.

Columbus, L. (2013). Roundup of mobile apps & app store forecasts, 2013. Forbes. Retrieved from http://www.forbes.com/sites/louiscolumbus/2013/06/09/roundup-of-mobile-apps-app-store-forecasts-2013/; June 9.

Culnan, M.J., & Bies, R.J. (2003). Consumer privacy: Balancing economic and justice considerations. Journal of Social Issues, 59(2), 323–342.

Duggan, M. (2013). Cell phone activities 2013. Washington, DC: Pew Internet & American Life Project.

Eastlick, M.A., Lotz, S.L., & Warrington, P. (2006). Understanding online B-to-C relationships: An integrated model of privacy concerns, trust, and commitment. Journal of Business Research, 59(8), 877–886.

Federal Trade Commission. (2012). Protecting consumer privacy in an era of rapid change: Recommendations for businesses and policymakers. Washington, DC: Federal Trade Commission.

Ganong, L.H., & Coleman, M. (2006). Multiple Segment factorial vignette designs. Journal of Marriage and Family, 68(2), 455–468.

Garg, V., Benton, K., & Camp, L. J. (2014). The privacy paradox: a Facebook case study. In Proceedings of the 2014 Telecommunications Policy Research Conference (TPRC). Arlington, VA: SSRN. Retrieved from http://papers.ssrn.com.proxy-um.researchport.umd.edu/sol3/papers.cfm?abstract_id=2411672

Gefen, D., & Pavlou, P.A. (2012). The boundaries of trust and risk: The quadratic moderating role of institutional structures. Information Systems Research, 23(3 Pt -2), 940–959.

Horton, J. J., Rand, D. G., & Zeckhauser, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. Experimental Economics, 14(3), 399–425.

Jackson, S.J., Gillespie, T., & Payette, S. (2014). The policy knot: Re-integrating policy, practice and design in Cscw studies of social computing. In Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (pp. 588–602). New York: ACM.

Jasso, G. (2006). Factorial survey methods for studying beliefs and judgments. Sociological Methods & Research, 34(3), 334–423.

Kang, C. (2013, December 6). Flashlight app kept users in the dark about sharing location data: FTC. The Washington Post. Retrieved from http://www.washingtonpost.com/business/technology/flashlight-app-kept-users-in-the-dark-about-sharing-location-data-ftc/2013/12/05/1be26fa6-5dc7-11e3-be07-006c776266ed_story.html

Lease, M., Hullman, J., Bigham, J., Bernstein, M., Kim, J., Lasecki, W., . . . Miller, R. (2013). *Mechanical Turk is Not Anonymous* (SSRN Scholarly Paper No. ID 2228728). Rochester, NY: Social Science Research Network. Retrieved from http://papers.ssrn.com/abstract=2228728

Li, H., Sarathy, R., & Xu, H. (2010). Understanding situational online information disclosure as a privacy calculus. Journal of Computer Information Systems, 51(1), 62–71.

Madden, M. (2014). *Public perceptions of privacy and security in the post-Snowden Era*. Pew Research Center. Retrieved from http://www.pewinternet.org/files/2014/11/PI_PublicPerceptionsofPrivacy_111214.pdf

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.

Martin, K.E. (2012). Diminished or just different? A factorial vignette study of privacy as a social contract. Journal of Business Ethics, 111(4), 519–539.

Martin, K.E. (2013). Transaction costs, privacy, and trust: The laudable goals and ultimate failure of notice and choice to respect privacy online. First Monday, 18(12). Retrieved from http://firstmonday.org/ojs/index.php/fm/article/view/4838

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

Mayer, R.C., Davis, J.H., & Schoorman, F.D. (1995). An integrative model of organizational trust. Academy of Management Review, 20(3), 709–734.

McKnight, D.H., Cummings, L.L., & Chervany, N.L. (1998). Initial trust formation in new organizational relationships. Academy of Management Review, 23(3), 473–490.

Nguyen, D.H., Bedford, A., Bretana, A.G., & Hayes, G.R. (2011). Situating the concern for information privacy through an empirical study of responses to video recording. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 3207–3216). New York: ACM.

Nissenbaum, H. (2009). Privacy in context: Technology, policy, and the integrity of social life. Stanford, CA: Stanford Law Books.

Pavlou, P.A. (2011). State of the information privacy literature: Where are we now and where should we go. MIS Quarterly, 35(4), 977–988.

Pavlou, P.A., & Gefen, D. (2004). Building effective online marketplaces with institution-based trust. Information Systems Research, 15(1), 37–59.

Pavlou, P.A., Liang, H., & Xue, Y. (2007). Understanding and mitigating uncertainty in online exchange relationships: A principal-agent perspective. MIS Quarterly, 31(1), 105–136.

Perlberg, S. (2014). Nielsen and comScore duel over mobile ad-tracking. The Wall Street Journal. Retrieved from http://blogs.wsj.com/digits/2014/04/17/nielsen-comscore-mobile-ad-tracking/; April 17, New York.

Pew Research Internet Project. (2013, December 27). Mobile Technology Fact Sheet. Retrieved April 7, 2014, Retrieved from http://www.pewinternet.org/fact-sheets/mobile-technology-fact-sheet/

Pirson, M., Martin, K. E., & Parmar, B. L. (2013). Formation of Stakeholder Trust in Business and the Role of Personal Values (SSRN Scholarly Paper No. ID 2260527). Rochester, NY: Social Science Research Network. Retrieved from http://papers.ssrn.com/abstract=2260527

Pirson, M., Martin, K.E., & Parmar, B.L. (2014). Public trust in business. In B. Moriarty (Ed.), Public trust in business (pp. 116–152). Cambridge, UK: Cambridge University Press.

Ross, J., Irani, L., Silberman, M.S., Zaldivar, A., & Tomlinson, B. (2010). Who are the crowdworkers?: Shifting demographics in mechanical turk. In CHI '10 Extended Abstracts on Human Factors in Computing Systems (pp. 2863–2872). New York: ACM.

Rossi, P.H., & Nock, S.L. (1982). Measuring social judgments: The factorial survey approach. Beverly Hills, CA: Sage Publications.

Shilton, K. (2012). Participatory personal data: An emerging research challenge for the information sciences. Journal for the American Society of Information Science, 63(10), 1905–1915.

Shilton, K., & Martin, K.E. (2013). Mobile privacy expectations in context. In Proceedings of the 41th Research Conference on Communication, Information and Internet Policy (TPRC 2013). Arlington, VA: SSRN.

Smith, H.J., Milberg, S.J., & Burke, S.J. (1996). Information privacy: Measuring individuals' concerns about organizational practices. MIS Quarterly, 20(2), 167–196.

Stevenson, B., & Wolfers, J. (2011). *Trust in Public Institutions over the Business Cycle* (Working Paper No. 16891). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w16891

Strickland, L.S., & Hunt, L.E. (2005). Technology, security, and individual privacy: New tools, new threats, and new public perceptions. Journal of the American Society for Information Science and Technology, 56(3), 221–234.

Sultan, F., & Rohm, A.J. (2004). The evolving role of the internet in marketing strategy: An exploratory study. Journal of Interactive Marketing, 18(2), 6–19.

Urban, J.M., Hoofnagle, C.J., & Li, S. (2012). *Mobile Phones and Privacy*. Berkeley, CA: University of California at Berkeley—Center for the Study of Law and Society. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2103405

Van Slyke, C., Shim, J., Johnson, R., & Jiang, J.J. (2006). Concern for information privacy and online consumer purchasing. Journal of the Association for Information Systems, 7(1), 415–444.

Vasalou, A., Joinson, A., & Houghton, D. (forthcoming). Privacy as a fuzzy concept: A new conceptualization of privacy for practitioners. Journal of the Association for Information Science and Technology.

Wallander, L. (2009). 25 years of factorial surveys in sociology: A review. Social Science Research, 38(3), 505–520.

Westin, A.F. (1970). Privacy and freedom. New York: Atheneum.

Xu, H., Zhang, C., Shi, P., & Song, P. (2009). Exploring the role of overt vs. covert personalization strategy in privacy calculus. Academy of Management Proceedings, 2009(1), 1–6.

Yao, M.Z., Rice, R.E., & Wallis, K. (2007). Predicting user concerns about online privacy. Journal of the American Society for Information Science and Technology, 58(5), 710–722.