

Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices

Kirsten Martin & Katie Shilton

To cite this article: Kirsten Martin & Katie Shilton (2016) Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices, *The Information Society*, 32:3, 200-216, DOI: [10.1080/01972243.2016.1153012](https://doi.org/10.1080/01972243.2016.1153012)

To link to this article: <http://dx.doi.org/10.1080/01972243.2016.1153012>



© 2016 Kirsten Martin and Katie Shilton
Published with license by Taylor & Francis©
Kirsten Martin and Katie Shilton



Published online: 13 Apr 2016.



Submit your article to this journal



View related articles



View Crossmark data

Putting mobile application privacy in context: An empirical study of user privacy expectations for mobile devices

Kirsten Martin^a and Katie Shilton^b

^aDepartment of Strategic Management and Public Policy, George Washington University, Washington, DC, USA; ^bCollege of Information Studies, University of Maryland, College Park, Maryland, USA

ABSTRACT

Users increasingly use mobile devices to engage in social activity and commerce, enabling new forms of data collection by firms and marketers. User privacy expectations for these new forms of data collection remain unclear. A particularly difficult challenge is meeting expectations for contextual integrity, as user privacy expectations vary depending upon data type collected and context of use. This article illustrates how fine-grained, contextual privacy expectations can be measured. It presents findings from a factorial vignette survey that measured the impact of diverse real-world contexts (e.g., medical, navigation, music), data types, and data uses on user privacy expectations. Results demonstrate that individuals' general privacy preferences are of limited significance for predicting their privacy judgments in specific scenarios. Instead, the results present a nuanced portrait of the relative importance of particular contextual factors and information uses, and demonstrate how those contextual factors can be found and measured. The results also suggest that current common activities of mobile application companies, such as harvesting and reusing location data, images, and contact lists, do not meet users' privacy expectations. Understanding how user privacy expectations vary according to context, data types, and data uses highlights areas requiring stricter privacy protections by governments and industry.

ARTICLE HISTORY

Received 6 November 2014

Accepted 6 September 2015

KEYWORDS

Applications; contextual integrity; Internet; mobile; privacy; social contract

Mobile devices raise a number of privacy-related issues (Johnson 2004; Curry, Phillips, and Regan 2004; Shilton 2009). They enable the gathering of a wealth of data on location, motion, communications content, in-application activities, and other traces of mobile use that can be pieced together to understand and influence users' behavior. Moreover, mobile applications enable new data collection actors, including application developers (e.g., Rovio Games), mobile providers (e.g., AT&T), operating system providers (e.g., Google), and device manufacturers (e.g., BlackBerry, Apple). These data may also be shared with third-party tracking or advertising companies.

In the United States, user privacy expectations are an ethical and legal standard by which the appropriateness of data collection or use can be judged (Waldo, Lin, and Millett 2007). Moreover, meeting of consumer privacy expectations has been shown to be linked with increased trust in firms, increased consumer likelihood to transact with firms (Eastlick, Lotz, and Warrington 2006), and increased consumer purchases (McCole, Ramsey, and Williams 2010; Hoffman, Novak, and Peralta 1999). Hence, helping mobile application companies meet

users' privacy expectations will not only make the mobile application marketplace more ethical, it will also enable these firms to retain existing customers and to grow their customer base.

However, how firms should best meet user privacy expectations in the mobile space is an unanswered question. This article addresses both the academic and practical challenges of understanding user privacy expectations in the mobile sector. It provides empirical support to the theory of privacy as contextual integrity (Nissenbaum 2009), while simultaneously providing nuanced practical guidance for firms struggling to address user privacy expectations. It presents findings from a factorial vignette survey designed to understand users' privacy expectations for mobile devices across diverse real-world contexts. This survey, which was conducted four times from May 2013 to February 2014 for a total of 1,915 respondents who rated 77,480 vignettes, was designed to examine (a) how users hold different privacy expectations based on the social context of their mobile activity and (b) how contextual factors such as who (the data collection actor, e.g., the application

CONTACT Kirsten Martin  martink@gwu.edu  Department of Strategic Management and Public Policy, George Washington University, 2201 G Street, Washington, DC 20052, USA.

This is an Open Access article. Non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly attributed, cited, and is not altered, transformed, or built upon in any way, is permitted. The moral rights of the named author(s) have been asserted. © 2016 Kirsten Martin and Katie Shilton



developer or mobile phone provider), what (the type of information received or tracked by the primary organization), why (application context, e.g., games, weather, social networking, navigation, music, banking, shopping, and productivity), and how (the use of data, e.g., the amount of time data is stored or how that data is reused) affect users' privacy expectations.

There is a common misconception that contextual definitions of privacy mean that privacy is difficult to understand or implement, because user expectations can vary based upon so many factors. The results of this survey illustrate, however, that privacy concerns can be thought of as predictably contextual. Within an industry (in this case, mobile applications), certain combinations of data types and data uses are largely acceptable to consumers, and other combinations are unacceptable. Finding the combinations of factors that are acceptable to consumers is the key for firms that hope to practice privacy by design.

The following sections review literature supporting the research approach and method, explain survey construction and deployment, and discuss the findings and their significance.

Literature review

Fair information principles (FIP), and in particular notice and choice, serve as one source of guidance for self-regulation within the industry (Culnan and Williams 2009; Bowie and Jamal 2006; Milne and Culnan 2002), and could be adapted to the mobile sector (Federal Trade Commission 2012). However, practical and philosophical problems persist. "Choice" is a problematic concept when individuals perceive that opting out of application usage has more costs than benefits (Cate 2010). And surveys and experiments have shown that individuals make judgments about privacy expectations and violations regardless of the content of privacy notices (Milne, Culnan, and Greene 2006; McDonald and Cranor 2008; Beales and Muris 2008; Martin 2013; Nissenbaum 2011).

Increasingly, an approach known as privacy by design is gaining popularity (Mayer and Narayanan 2013; Cavoukian 2012; Spiekermann and Cranor 2009). In privacy by design, privacy protection measures are built into technologies at the point of design. But privacy by design has been criticized for being vague in its proscriptions (Kroener and Wright 2014). In particular, Popescu and Baruh (2013) call for close examination of the affordances of mobile technology, including location tracking and browsing histories, as well as the difficult problem of legitimacy of consent secured from what many argue amounts to a captive audience (Popescu and Baruh 2013). Both firms and regulators relying upon privacy by

design need specific guidance around privacy expectations for mobile devices.

A promising new approach—privacy as contextual integrity—posits that privacy expectations about the transmission and uses of information are dependent upon social context (Nissenbaum 2009). Here, individuals are seen as providing information within a particular context with an understanding of the privacy rules that govern that context. Shopping online, talking in the break room, and divulging information to a doctor are all governed by different information norms. As Nissenbaum states, "The crucial issue is not whether the information is private or public, gathered from private or public settings, but whether the action breaches contextual integrity" (Nissenbaum 2004, 134).

Privacy as contextual integrity joins a growing body of theoretical scholarship examining privacy norms and expectations within a specific set of relationships, situations, or contexts (Nissenbaum 2009; Solove 2006; Martin 2012). Context-dependent definitions of privacy suggest that instead of measuring privacy concerns and expectations as static attributes of individuals, privacy concerns and expectations are best defined and measured as context-specific reactions (Xu et al. 2012; Lin et al. 2012). 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. Within a contextual model, individuals are able to exercise informational self-determination (Buitelaar 2014) by discriminately sharing information based on actual and hypothetical social contracts (Martin 2012; Martin forthcoming). Rules for information flow within a social context take into account the purpose of the information exchanged, as well as risks and harms associated with sharing information (Culnan and Bies 2003; Li, Sarathy, and Xu 2010; Martin 2012; Xu et al. 2009; Milne and Gordon 1993; Dunfee, Smith, and Ross Jr 1999; Culnan 1995; Phelps, Nowak, and Ferrell 2000; Heeney 2012). These rules might take into account:

- Who/recipients—people, organizations, technologies who are the senders, recipients, and subjects of information.
- What/information—the information types or data fields being transmitted.
- How/transmission principles—the constraints on the flow of information.
- Why—the purpose of the use of information (Nissenbaum 2009).

Key to all contextual definitions of privacy is how the main components work together—who receives the information, what type of information, how is it used, and for what purpose—within a particular context. An

individual's accounting for these contextual rules, and how they work together, is also known as the *privacy calculus*, where privacy norms are developed with the costs and benefits of sharing information in mind (Martin 2013).

The privacy as contextual integrity approach alerts us to two considerations in the realm of mobile communication. First, tactics to address privacy expectations on mobile devices should depend on the context of the exchange. For example, location data may be required for contexts such as navigation but inappropriate for a flashlight application (Kang 2013); anonymity may be appropriate for contexts such as Internet search, but inappropriate in a context such as social networking. Second, data types cannot be deemed "private" or "public" across contexts. Tactics such as behavioral advertising, data collection and retention, and tracking may be appropriate and within the contextually defined privacy norms in one context while inappropriate in another. To enable privacy by design that accounts for contextual privacy norms, empirical work is needed to identify the data, practices, and recipients appropriate for different mobile application contexts.

Researchers such as Lin et al. (2012) have begun this effort by measuring sensitive data types and user reactions to the purpose of data collection in the mobile ecosystem. This research expands on such efforts by measuring the full range of contextual factors suggested by Nissenbaum's (2009) theoretical work, including variables such as application context, recipients of data, and transmission principles. Since addressing mobile privacy expectations is the goal of organizations and regulatory bodies, understanding how those expectations change in different contexts based on the contextually defined privacy norms would help managers and regulators identify which contexts require different privacy protections.

Methods

To investigate whether and how privacy expectations vary across mobile activity contexts, a survey was conducted using factorial vignette methodology (Wallander 2009). The factorial vignette survey methodology was developed to investigate human judgments (Rossi and Nock 1982; Jasso 2006; Wallander 2009) and asks respondents to rate a series of hypothetical vignettes. A set of vignettes is generated for each respondent, where the vignette factors are independent variables controlled by the researcher and randomly selected. The respondent is shown a screen displaying a hypothetical vignette, and asked a single question to evaluate whether the vignette meets their privacy expectations. Over the course of the survey, participants are shown a total of 40 vignettes.

Statistical techniques are then used to identify the relative importance of the factors driving the outcome for the respondents. The vignettes vary based on relevant factors and are controlled and presented by the investigator to ensure that intercorrelations among vignette characteristics are zero. This method allows the researchers to simultaneously examine multiple factors (e.g., changes in context and types of information sharing) by providing respondents with systematically varied vignettes.

This design mitigates several major concerns in empirical privacy research. As noted by a recent Federal Trade Commission report, traditional surveys are limited in their ability to measure privacy expectations of individuals (Federal Trade Commission 2010, fn 72). First, privacy surveys are fraught with respondent bias, as respondents tend to inflate their concern for privacy when asked directly about the topic (Hui, Teo, and Lee 2007). For example, despite reporting a general "concern for online privacy" (Buchanan et al. 2007), users seldom provide false information or alter their privacy settings in online applications (Gross and Acquisti 2005). In addition, individuals often have difficulty articulating the relative importance of factors that constitute their privacy expectations across different contexts online, such as shopping, seeking medical advice, search, and playing games. As noted by the recent FTC report, traditional surveys are limited in their ability to measure privacy expectations of individuals (Federal Trade Commission 2010, fn 72); indirect measurements are sometimes necessary (Braunstein, Granka, and Staddon 2011).

A set of vignettes was displayed on a screen for each respondent, chosen randomly with replacement from a vignette universe as the respondent took the survey. For each vignette, the associated rating, factor levels, and the vignette script were preserved, as well as the vignette sequence number. The vignette format is provided in the Appendix with a sample vignette and the vignette template. The vignette wording and format were pilot tested for clarity with students in several courses at a Mid-Atlantic public university.

General survey design

The independent variables tested drew upon factors important to privacy as contextual integrity (Nissenbaum 2009) and privacy as a social contract (Martin 2012; Martin forthcoming). Each survey respondent was shown a series of vignettes that varied based on:

- Who: the data collection actor—the primary organization collecting information, such as application developer or mobile phone provider.
- What: the type of information received or tracked by the primary organization.

- Why: the application context—for example, playing games, checking weather, participating in social networking, navigating using maps, listening to music, banking, shopping, and organizing personal productivity.
- How (used): transmission principles—for example, how the data are reused or stored.

The two surveys cover general online activities and are related (e.g., tracked data can also be used for targeted advertising). However, the number of factors grew too large for robust findings when the vignettes were combined. The vignettes focused on targeted advertising were therefore run separately from those focused on tracking users. A general design is illustrated in [Figure 1](#), with details for each factors explained in the following.

Independent variables

Defining “why”

Defining meaningful social contexts is one of the challenges inherent in understanding privacy in context. This study took an approach to defining “context” that replicates 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, and so on. Therefore, we used the function of the application as a stand-in for context. Contexts were chosen and named according to how they are identified by the two major application stores: the Apple iTunes store and Google Play. Next, we consulted industry data on the most 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](#)). Mobile application contexts were then chosen based on a combination of popularity and diversity. We chose the most popular application contexts, as well as those that are known to have sensitive data in face-to-face transactions, such as medical and banking. The final social contexts chosen to test were:

- Games.
- Weather.
- Social networking.
- Navigation.

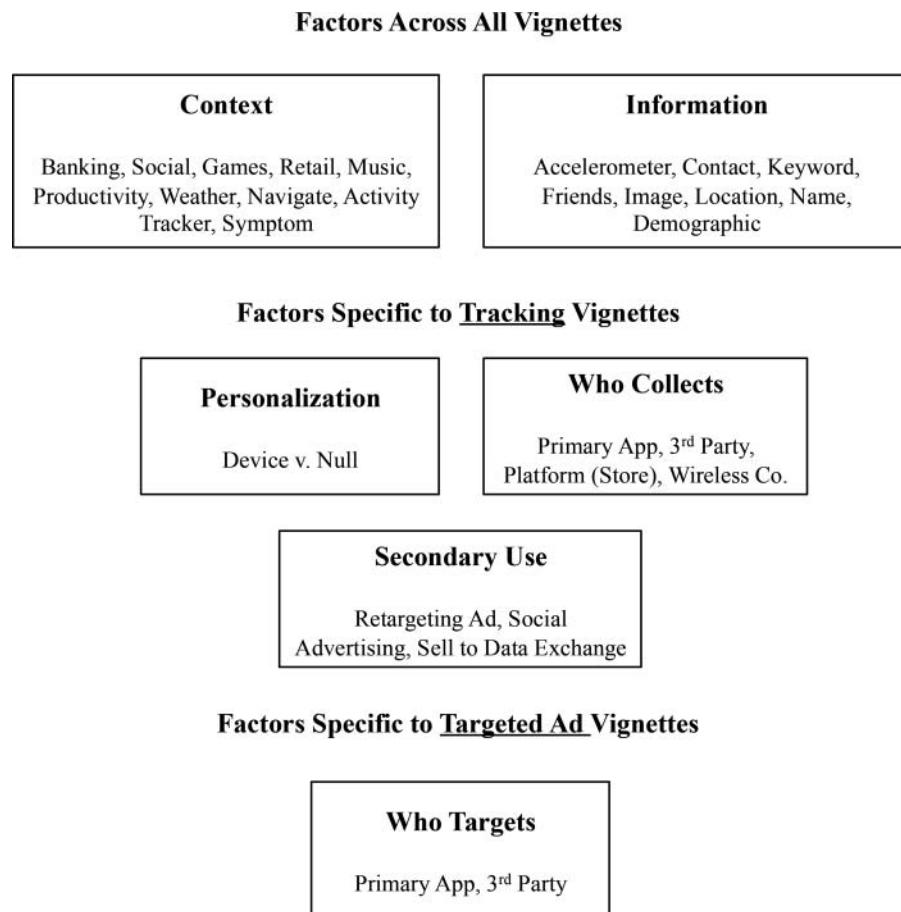


Figure 1. Factors for vignette surveys.

- Music.
- Banking/finance.
- Shopping/retail.
- Productivity.

The Appendix contains the factors and the possible levels for each in a table, as well as the vignette templates and sample vignettes.

Defining "who" and "what"

The investigators chose prominent real-world actors in the mobile marketplace to test as data collection actors. These included mobile application development firms, platform providers, and telecommunications companies (ACT 2012). For the targeted advertising survey, vignettes included either the application development firm, or a third party placing an ad. For the tracking survey, vignettes included the primary app development firm, the wireless provider, the platform provider (the app store), or a third party as the collector of the information. The investigators also chose data types based on real-world data collection capabilities of mobile phones. Most smartphones can collect a range of data including location, accelerometer, demographic data, contacts, keywords, user name, and images (Boyles, Smith, and Madden 2012). Information type did not vary across the two surveys as shown in Figure 1.

Defining "how"

Tracking vignettes included age (the length of time data was stored, in months), personalization (whether data was tied to a unique identifier for your mobile device), and the secondary use (what the collecting organization does with the information, such as retargeting, data exchange, or social advertising).

Sample vignettes

The investigators constructed two types of vignettes to test norms around how mobile data is used. These included scenarios about using data to provide targeted advertisements and scenarios about using data to track users. This generated vignettes such as the following, where underlining highlights the factors (independent variables) that would systematically change (template and additional examples are included in the Appendix).

Targeting vignette sample.

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 vignette sample.

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.

Control variables

The respondents' age and gender were collected before participants began the vignettes, and were used in the regression analysis in addition to three control questions to gauge users' overall trust in applications, their generalized level of concern about application privacy, and their level of experience using mobile applications. Respondents were asked to judge the frequency of their own use of mobile applications, from rarely to multiple times per day. In addition to experience, age, and gender, the respondent was asked:

Tell us how much you agree with the statements below. On the sliding scale below, with a rating to the left being "strongly disagree" and to the right being "strongly agree."

The rating task stated: "In general, I trust mobile applications." Rating on this question was used as a measure for general level of trust in applications as an institution (Pavlou and Gefen 2004). Trust is important to understand online experiences and privacy in particular (Hoffman, Novak, and Peralta 1999). The second rating task stated: "In general, I believe privacy is important." Ratings on this question were used as a measure of general level of privacy concern, which has been shown to vary across individuals and impact privacy expectations (Sheehan 2002).

Dependent variable: Privacy rating tasks

For each vignette, respondents were given a single rating task with the constant prompt:

Tell us how much you agree with the statements below. Using a sliding scale from -100 to 100, with -100 indicating "strongly disagree" and 100 indicating "strongly agree."

Respondents were then given the prompt: "This application meets my privacy expectations." Respondents could adjust a sliding bar to indicate their response from -100 to 100.

Sample

The surveys were deployed four times every three months from May 2013 to February 2014 to a total of 1,925 United States-based respondents to rate more than 77,000 vignettes using Amazon's Mechanical Turk (see Table 1). Because mobile applications are an emerging industry, user perceptions may be highly time sensitive. In addition, exogenous events, such as the Snowden-NSA revelation in June 2013, could skew the results and render the findings of a single survey less generalizable. In order to identify stable, consistent attributes of user privacy expectations, the survey was conducted four times during 2013–2014.¹

Though use of Mechanical Turk (mTurk) for survey deployment has been criticized (Lease et al. 2013; Ross et al. 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 et al. 2011; Berinsky, Huber, and Lenz 2012).² General statistical information about the sample is provided in Table 1. The average age of the respondents was 31 years old for both tracking and targeting surveys, and the sample was 58% male for targeting vignettes and 55% male for tracking vignettes. On average, the respondents trusted applications generally, with an average score of +20.26 when taking the targeted advertising survey and +12.97 when taking the tracking survey. In addition, the respondents report privacy to be important to them, with an average score of +79.82 and +79.24, respectively, for targeting and tracking surveys.

Analysis

The factorial vignette approach allows the researcher to examine (a) the elements of information used to form judgments, (b) the weight of each of these factors, and (c) how different subgroups of the respondents agree on (a) and (b) (Nock and Guterbock 2010). These factors and their associated coefficients are the equations-inside-the-head (Jasso 2006) of respondents as to judgments of privacy expectations. In this case, they represent the privacy calculus of the individual respondents.

The data were analyzed on two levels: variation in privacy judgments attributable to the contextual factors, and variation in privacy judgments attributable to the respondent-level control variables. For the targeted advertising surveys, 976 respondents rated 40 vignettes each, resulting in 39,320 rated vignettes (e.g., 39,320 observations). For the tracking surveys, 949 respondents rated 38,160 vignettes as shown in Table 1. In order to show consistencies and inconsistencies in the relative importance in the privacy factors over time, the four subsamples, as well as the overall sample, are used in the analysis and illustrated in the tables that follow.

The resulting data set can be thought of as having two levels: Level 1 is the contextual factors, and Level 2 is the respondent control variables. If I is the number of the respondents with Level 2 individual variables, and K is the number of vignettes answered with Level 1 factor variables, the general equation is:

$$Y_{ij} = \beta_0 + \sum \beta_k V_{jk} + \sum \gamma_h R_{hi} + u_i + e_j. \quad (1)$$

Table 1. Sample characteristics for each sample 2013–2014.

Targeted ad vignettes	May 13		August 13		November 13		February 14		Combined
Users	247		243		244		242		976
Vignettes	10,000		9,880		9,760		9,680		39,320
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean
Age	32.21	9.74	30.58	9.62	30.57	8.78	32.86	9.59	31.56
Male	55.2%		57.9%		68.0%		52.9%		58.5%
Privacy is Important	82.11	22.97	78.77	27.43	81.36	24.70	76.97	28.10	79.82
I trust mobile apps	19.29	42.48	22.22	44.33	19.91	49.43	19.62	42.44	20.26
Mean (DV)	-22.24	27.84	-16.30	27.30	-18.42	28.13	-15.00	28.93	-18.01
R ² of Users	0.82		0.84		0.86		0.84		0.84
ICC	22.50%		20.40%		21.70%		23.00%		22.10%
Tracking users vignettes	May 13		August 13		November 13		February 14		Combined
Users	244		233		238		234		949
Vignettes	9,880		9,400		9,520		9,360		38,160
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean
Age	31.40	10.76	30.69	10.44	31.98	10.09	34.27	11.14	32.08
Male	44.1%		64.3%		63.0%		50.9%		55.5%
Privacy is important	79.27	30.00	78.74	27.06	79.95	26.98	79.00	27.16	79.24
I trust mobile apps	13.63	44.52	17.13	48.27	13.75	50.21	7.33	47.22	12.97
Mean (DV)	-43.53	32.29	-38.55	32.33	-43.36	29.62	-45.31	33.36	-42.70
R ² of users	0.78		0.83		0.81		0.79		0.80
ICC	35.40%		32.70%		30.30%		37.50%		34.10%

where Y_{ij} is the rating of vignette j by respondent i , 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, β_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_j is a vignette-level residual. The model conceptualizes the ratings as a function of the contextual factors described in the vignette (ΣV_k) and the characteristics of the respondent (ΣR_h) as already hypothesized. For the targeted advertising survey, the vignette factors above would be represented by Eq. (2):

$$\begin{aligned} \Sigma \beta V_{\text{targeting}} = & \beta_{1-8} \text{Information} + \beta_{1-10} \text{Context} \\ & + \beta_{1-2} \text{AdType} \end{aligned} \quad (2)$$

And the equation for the tracking users vignettes would be the following:

$$\begin{aligned} \Sigma \beta V_{\text{tracking}} = & \beta_{1-8} \text{Information} + \beta_{1-10} \text{Context} \\ & + \beta_{1-2} \text{CollectingActor} + \beta_1 \text{Personalization} \\ & + \beta_{1-3} \text{SecondaryUse} + \beta_1 \text{StorageMths} \end{aligned} \quad (3)$$

As the data can be modeled at two levels, multilevel modeling was used to control for and measure individual variation in privacy judgments. Multilevel modeling (xtmixed in STATA) accounts for the possibility that the error terms were not equal across individuals, and later that the intercepts and coefficients may vary across respondents with random intercept and random slope models.

Factorial vignette surveys can be perceived as long and respondents could suffer from fatigue. While the

survey took, on average, approximately 10 minutes to complete, respondent fatigue was checked by controlling for later vignettes in the respondents' sequence. The sequence number for each rating (numbers 1–40) was stored, and a dummy variable was created for low sequence (numbers 1 and 2) and high sequence (numbers 35–40 and numbers 38–40) vignettes. If the associated dummy variable was significant when included in the regression, the respondent either had respondent fatigue or a learning-curve effect. Controlling for later vignettes revealed that respondent fatigue was not a factor. However, we did find a respondent learning curve: Respondents took one or two vignettes to become acclimated to the methodology. The analysis was run minus the first two vignettes for each respondent and the results remained the same. The results in the following include all respondent ratings. In addition, the rating task was designed to capture the respondents' normative judgments about mobile applications and not a respondent's intention to transact or actual use of an application.

Results

Overall privacy judgments

Based on the overall results presented in Table 1, tracking scenarios met privacy expectations to a lesser extent than targeting scenarios (mean = -42.70 and -18.01, respectively), a result consistent over the four samples as depicted in Figure 2. Both types of vignettes did not meet privacy expectations on average, as both means are negative. In addition, the respondent-level R^2 , developed by regressing the rating task on to the contextual factors for each respondent ($N = 40$), was larger for targeting ($\beta = 0.842$) as compared to tracking ($\beta = 0.801$;

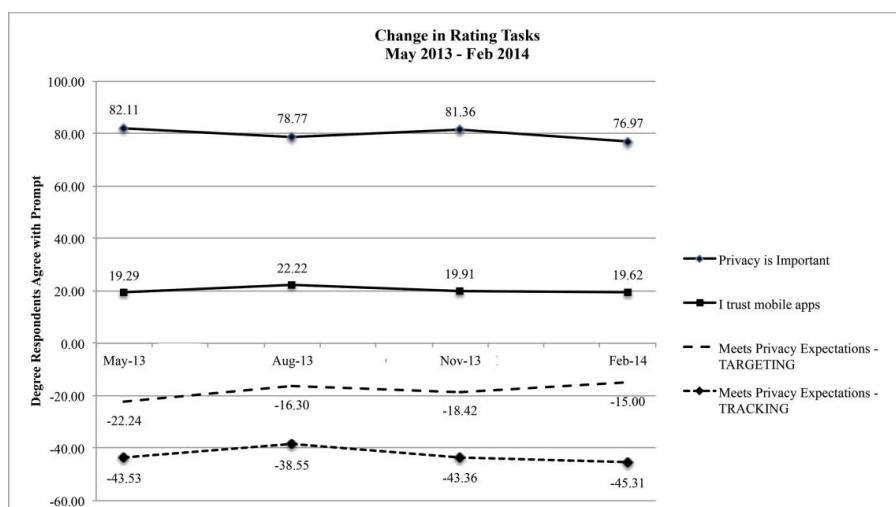


Figure 2. Charts of main control factors and mean vignette rating across samples.

$t = -6.185$, $p = .00$). This illustrates that individual respondents were slightly more certain of their judgments for targeted advertising as compared to tracking, indicating that participants felt more familiar with, and have formed a more certain opinion of, targeted advertising. The intraclass correlation coefficient, which measures the variance in the privacy judgment attributable to the grouping variable (the individual) using the random intercept model with no fixed effects (no variates included), was 34.1% for tracking versus 22.1% for targeted advertising. While tracking evoked a larger negative response, it also was a subject of less certainty to individual respondents and greater variance (less agreement) across respondents.

Individual attributes

Age had a significant impact in users' privacy expectations: Older respondents rated scenarios lower (meeting privacy expectations to a lesser extent) even when controlling for respondents' general belief that privacy is important and institutional trust in applications. Respondents who trust applications more also rated scenarios as meeting privacy expectations to a greater degree for both targeted advertising ($\beta = 0.206$, $p = .00$) and tracking ($\beta = 0.219$, $p = .00$). In addition, respondents who reported a greater belief that privacy is

important judged scenarios to meet privacy expectations less for both targeted advertising and tracking. However, the explained R^2 for adding individual control variables to the null random intercept model in Table 2 is only 4.8% for targeted advertising and 7.4% for tracking; while significant, the low explained R^2 suggests additional contextual factors are significant to explain the rating task.

Relative importance of contextual privacy factors

Responses from both targeting and tracking vignettes were used to run regressions of the rating task (responses to "This application met my privacy expectations") onto vignette-level and respondent-level factors using hierarchical analysis. By adding blocks of factors corresponding to the major theoretical drivers of privacy expectations, the analysis captures not only the significance of each block of factors but also the explained R^2 .

Both tracking and targeting vignettes were run as two separate regression equations. Important distinctions occurred in response to information types, secondary use, actors, and context. Figure 3 illustrate the relative importance of the contextual factors in driving privacy expectations described in detail next. Tables 3–6 include the multi-level regression results for each block of vignette factors.

Table 2. Importance of individual attributes in meeting privacy expectations.

Targeted ad vignettes	May 13		August 13		November 13		February 14		Combined		
	Parameter	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>	Estimate	<i>p</i>
Control variables											
Male	2.054	.53		3.109	.33	1.223	.71	-6.170	.06	-1.211	.52
Age**	-0.125	.63		-0.415	.12	-0.969	.00	-0.602	.02	-0.632	.00
TrustApps**	0.250	.00		0.160	.00	0.182	.00	0.288	.00	0.206	.00
PrivacyImportant**	-0.202	.00		-0.275	.00	-0.376	.00	-0.136	.02	-0.256	.00
_cons	8.668	.37		41.552	.00	69.950	.00	35.300	.00	48.031	.00
Average rating	-22.24		-12.65		-25.69		0.00		-16.30		
<i>N</i>	10,000		9,880		9,760		9,680		39,320		
ICC	22.5%		20.4%		21.7%		23.0%		22.1%		
sd(_cons)	26.66		26.00		26.89		27.74		26.96		
									explained R^2	4.80%	
									log ratio χ^2	218.98	
									<i>p</i>	.00	
Tracking users vignettes											
Male	-0.345	.93		2.353	.54	4.306	.25	4.093	.27	3.822	.08
Age*	-0.416	.14		-0.199	.45	-0.468	.10	-0.338	.19	-0.353	.02
TrustApps**	0.174	.00		0.264	.00	0.144	.00	0.255	.00	0.219	.00
PrivacyImportant**	-0.299	.00		-0.254	.00	-0.237	.00	-0.355	.00	-0.277	.00
_cons	31.389	.00		24.129	.01	27.688	.01	29.520	.00	27.375	.00
Average Rating	-43.54		-38.56		-43.36		-45.31		-42.70		
<i>N</i>	9,880		9,400		9,520		9,360		28,280		
ICC	35.4%		32.7%		30.3%		37.5%		34.1%		
sd(_cons)	28.66		27.05		26.92		27.91		27.51		
									explained R^2	7.37%	
									log ratio χ^2	220.12	
									<i>p</i>	.00	

* $p < .05$ for combined sample.

** $p < .01$ for combined sample.

Table 3. Importance of information type in meeting privacy expectations.

Targeted ad vignettes	May 13		August 13		November 13		February 14		Combined	
Parameter	Estimate	p	Estimate	p	Estimate	p	Estimate	p	Estimate	p
Information										
AccellInfo**	1.676	.30	-16.113	.00	-20.370	.00	-8.469	.00	-15.124	.00
ContactInfo**	-59.656	.00	-71.138	.00	-72.571	.00	-66.511	.00	-70.099	.00
KeywordInfo**	18.478	.00	8.508	.00	11.799	.00	14.399	.00	11.490	.00
FriendsInfo**	-8.976	.00	-16.712	.00	-21.121	.00	-22.782	.00	-20.347	.00
ImageInfo**	-65.768	.00	-78.153	.00	-81.713	.00	-76.390	.00	-78.843	.00
LocationInfo**	-2.195	.17	-9.706	.00	-20.145	.00	-10.142	.00	-13.328	.00
NameInfo**	-17.760	.00	-23.104	.00	-20.143	.00	-15.057	.00	-19.512	.00
(null = Demo)									explained R^2	25.79%
									log ratio χ^2	15933.08
									p	.00
Tracking users vignettes										
AccellInfo**	-0.621	.68	-6.184	.00	-3.020	.05	-0.709	.64	-3.261	.00
ContactInfo**	-20.692	.00	-28.803	.00	-18.964	.00	-17.676	.00	-21.811	.00
KeywordInfo**	-2.492	.10	-4.226	.01	-2.105	.17	-3.455	.02	-3.245	.00
FriendsInfo**	-8.226	.00	-10.408	.00	-6.386	.00	-6.403	.00	-7.770	.00
ImageInfo**	-25.106	.00	-36.423	.00	-26.203	.00	-24.257	.00	-29.044	.00
LocationInfo**	-4.651	.00	-9.855	.00	-5.171	.00	-5.917	.00	-7.016	.00
NameInfo**	-4.992	.00	-12.946	.00	-6.043	.00	-6.718	.00	-8.601	.00
(null = Demo)									explained R^2	2.93%
									log ratio χ^2	1760.54
									p	.00

* $p < .05$ for combined sample.** $p < .01$ for combined sample.**Information type (what)**

Generally, the type of information mattered to respondents' privacy judgments. The contact information ($\beta = -70.10$) and image information (-78.84) were the most (negatively) influential types of information, followed by the individual's name (-19.51), friend information (-20.35), accelerometer (-15.17), and location (-13.33). All of these data types negatively impacted meeting privacy expectations for targeted advertising compared to using demographic information. However, using keywords (11.49) positively impacted meeting privacy expectations for targeted advertising compared to using demographic information. In other words, individuals prefer that applications use keywords rather than

demographic data to serve advertisements, but did not expect the use of contact information or image information. The same pattern emerged for tracking scenarios: All data types negatively impacted meeting privacy expectations compared to demographic information, with the contact information (-21.81) and image information (-29.04) being consistently the least expected two types of information to track. In addition, the explained variance for the type of information was large and significant for targeting and tracking surveys.³

Interestingly, the use of accelerometer information became significant only after the May 2013 sample. This may indicate that respondents became aware of this type of information in the summer of 2013, or

Table 4. Importance of second use in meeting privacy expectations.

Tracking users vignettes	May 13		August 13		November 13		February 14		Combined	
Parameter	Estimate	p	Estimate	p	Estimate	p	Estimate	p	Estimate	p
Personalization										
devicePersonal*	-1.316	.08	-1.227	.14	-1.278	.10	-0.661	.38	-1.081	.02
(null = Null)										
Second Use										
DataExchange2nd**	-44.858	.00	-43.666	.00	-50.113	.00	-47.972	.00	-47.171	.00
SocialAd2nd**	-21.343	.00	-19.349	.00	-24.386	.00	-21.102	.00	-21.613	.00
(null = Retarget)										
Storage Months**	-0.499	.00	-0.728	.00	-0.602	.00	-0.783	.00	-0.702	.00
(null = Demo)									explained R^2	12.57%
									log ratio χ^2	8113.44
									p	.00

* $p < .05$ for combined sample.** $p < .01$ for combined sample.

Table 5. Importance of actor in meeting privacy expectations.

Targeted ad vignettes	May 13		August 13		November 13		February 14		Combined	
	Parameter	Estimate	p	Estimate	p	Estimate	p	Estimate	p	Estimate
AdType										
ThirdPartyAd*	-1.013	.22	-1.300	.13	-1.038	.21	-0.910	.28	-1.066	.03
(null = Primary)									explained R^2	0.01%
									log ratio χ^2	4.39
									p	.04
Tracking users vignettes										
Collecting Actor										
ThirdPartyCollect**	-8.927	.00	-6.246	.00	-4.720	.00	-5.711	.00	-5.591	.00
PlatformCollect	2.120	.05	0.214	.85	-1.024	.35	-0.301	.78	-0.416	.52
WirelessCollect*	0.164	.88	-1.940	.10	-1.753	.11	-1.223	.25	-1.687	.01
(null = Primary)									explained R^2	0.27%
									log ratio χ^2	145.91
									p	.00

* $p < .05$ for combined sample.** $p < .01$ for combined sample.

became more sensitive to its use. In addition, judgments about the use of friend and location information for targeted advertising became increasingly negative after May 2013.

Secondary use (how)

For tracking scenarios, the secondary use of information was the most important factor impacting privacy expectations. Selling to a data exchange ($\beta = -47.17$) and using tracked information for social advertising to contacts and friends (-21.61) both negatively impacted meeting privacy expectations. In addition, the amount of time the tracked information was stored negatively impacted meeting privacy expectations across all contexts aside from games and navigation. This may indicate user recognition that storing information can have benefits in gaming (e.g., saved games) and navigation (e.g., saved routes) contexts.

Actor (who)

The presence of a third-party collector in tracking scenarios meant vignettes met privacy expectations slightly less ($\beta = -5.59$). Similarly, the presence of a third-party advertiser in targeted advertising vignettes had a small but significant negative impact on whether the vignettes met privacy expectations ($\beta = -1.07$, $p = .03$). The explained variance for the block of factors varying the actors was minimal for both targeted advertising vignettes (0.01%) and tracking users vignettes (0.27%). This means who is involved was less important to users overall than the type of information collected and how the information was used.

Context

The impact of context on privacy expectations was complicated. Generally, as a factor in a multilevel linear regression, context did not significantly and consistently

Table 6. Importance of website context in meeting privacy expectations.

Targeted ad vignettes	May 13		August 13		November 13		February 14		Combined	
	Parameter	Estimate	p	Estimate	p	Estimate	p	Estimate	p	Estimate
Context										
BankingCxt**	-13.022	.00	-9.010	.00	-8.366	.00	-10.976	.00	-9.400	.00
SocialCxt	0.440	.79	-0.409	.83	1.834	.33	-3.562	.06	-0.655	.55
GamesCxt	1.328	.42	3.715	.05	-1.737	.36	-1.042	.58	0.439	.69
MusicCxt	-0.418	.80	0.207	.91	2.622	.16	-2.061	.27	0.250	.82
ProductivityCxt	-1.469	.37	-0.803	.67	-3.244	.08	-1.691	.36	-1.984	.07
WeatherCxt	-0.635	.70	2.492	.19	-0.001	1.00	-2.841	.13	-0.057	.96
NavigateCxt	0.370	.82	2.570	.17	1.193	.53	-1.608	.39	0.729	.50
ActivityCxt	n/a	.00	0.681	.72	0.534	.77	-1.770	.34	-0.160	.88
SymptomCxt	n/a	.00	-4.883	.01	-5.295	.00	-3.841	.04	-4.633	.00
(null = Retail)									explained R^2	0.35%
									log ratio χ^2	118.76
									p	.00

* $p < .05$ for combined sample.** $p < .01$ for combined sample.

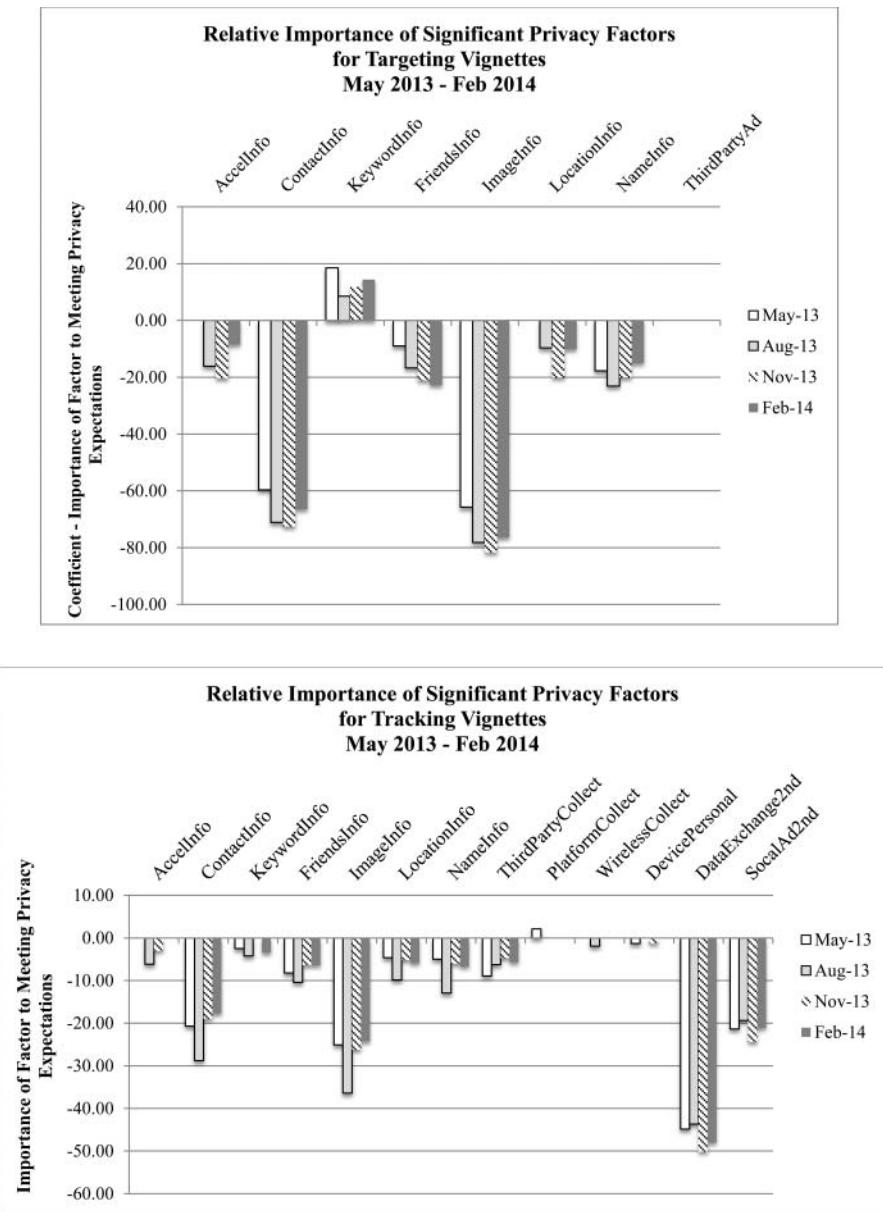


Figure 3. Relative importance of privacy factors for tracking vignettes.

impact the rating task directly, with the exception of banking ($\beta = -9.40$ for targeted advertising and -8.69 for tracking users) and symptom checking applications ($\beta = -4.63$ for targeted advertising and -5.43 for tracking users).

Figures 4 and 5 illustrate how application context can impact the relative importance of information are judged to meet privacy expectations. For example, use of location information to target ads had a positive impact on meeting privacy expectations within weather and navigation applications ($p = .00$), but a negative impact within social ($p = .05$) and gaming ($p = .00$) applications. Collecting information about friends to target ads had a negative impact on meeting privacy expectations within banking and symptom applications ($p = .00$), but

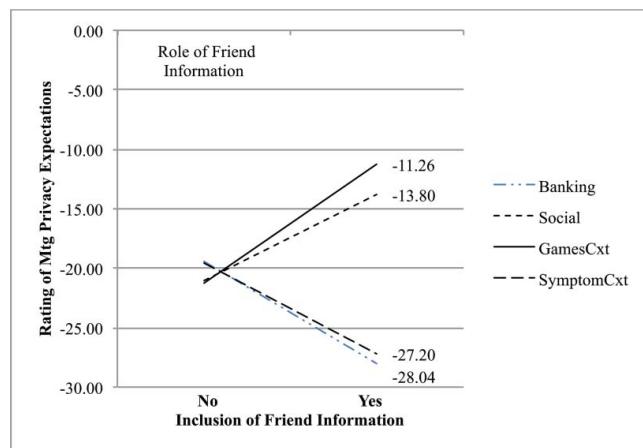


Figure 4. Interaction between application context and type of information for targeted advertising (friend information).

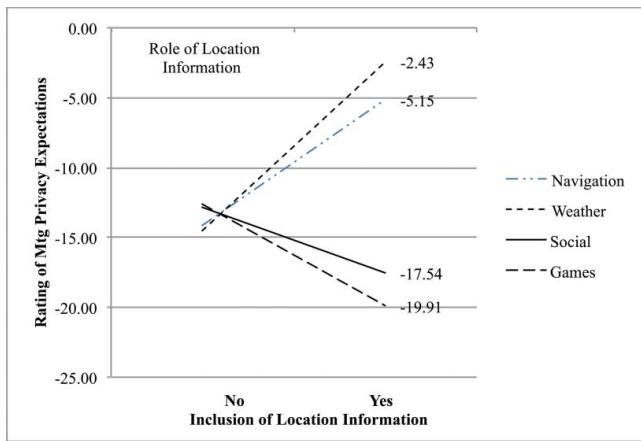


Figure 5. Interaction between application context and type of information for targeted advertising (location information).

positive impact within social and gaming applications ($p = .00$). Figures 4 and 5 illustrate the important role context can have on the relative impotence of using information in targeted advertising.

Discussion

Users displayed nuanced judgments about the ways that data type, context, and use impact their privacy expectations. Using factorial vignette method, we can measure these nuances and provide guidance to developers and businesses interested in practicing privacy by design. The findings have practical implications for application companies struggling to user-friendly privacy policies. The findings also have scholarly implications for researchers studying information privacy. We discuss each of these in detail next.

Implications to practice

The results indicate that overall, very common activities of mobile application companies (harvesting and using data such as location, accelerometer readings, demographic data, contacts, keywords, name, images, and friends) do not meet users' privacy expectations. But users are not monolithic in their privacy expectations. For example, users expect navigation and weather applications to use location and accelerometer data, and users expect a link between harvesting keywords and targeted advertising. The results show that nuanced, contextual privacy concerns can be measured within an industry. The resulting conclusions about which data types are sensitive when put to what uses in which contexts can help firms engage in privacy by design to meet the expectations of users. For example, navigation applications should feel confident collecting users' location data,

but should not collect image data. Navigation applications can make design changes to avoid data besides location, supporting privacy by design. This article contributes evidence for contextually defined privacy rules based on who asks, what information is collected, in what context the information is collected, and whether it is used for targeting ads and tracking users.

Some data types were particularly sensitive regardless of context—or at least particularly surprising to respondents. Harvest of both images and contact information failed to meet user privacy expectations. This may be because it is not widely known that these data can be harvested by mobile applications, or it may be that these data types are particularly sensitive. Application developers engaged in privacy by design may wish to avoid collecting, using, or selling images or contact information from phones.

Symptom-checking and banking were both sensitive contexts for users, and scenarios using these types of applications were judged more harshly. Developers of applications in the medical and financial sectors may need to be particularly protective of user privacy to retain public trust. In addition, gaming apps had much in common with social apps: Users' overall privacy judgments and the relative importance of privacy factors such as location and friend information were statistically equal. Gaming and social applications may be considered similar contexts. Data on the nuances of consumer privacy expectations give application developers, as well as other industry stakeholders such as platform providers and application stores, more power to self-regulate to protect user privacy and build consumer trust.

Implications for scholarship

The findings also provide strong empirical support for rule utilitarian approaches such as privacy as contextual integrity (Nissenbaum 2004; Nissenbaum 2009), privacy as a social contract (Martin 2012), and some versions of the privacy calculus (Culnan and Bies 2003; Martin 2013; Xu et al. 2009). Factors such as the type of information collected, the secondary usage of the information, and the industry context significantly impact users' privacy judgments. Individual dispositions about a concern for privacy or belief that privacy is important were of limited importance in privacy judgments about a particular scenario. Respondents were quite nuanced in their consideration of the contextual factors: The use of contact list information in banking and medical applications was considered a greater violation than when used in social or gaming applications. Similarly, the use of location information in social and gaming applications was a greater violation than when used for navigation and weather applications.

The study suggests that users form quite nuanced judgments of data collection and use scenarios. This suggests that studies that focus on privacy expectations would be improved by including context-dependent factors to better illustrate individuals' privacy expectations. In addition, the individual-level measures were only minimally useful in explaining variances in respondents' privacy judgments, suggesting that general surveys about privacy concerns, privacy valuations, or privacy attitudes should be generalized with caution, as they may not apply across all contexts. Researchers should be aware of the potential nuances of respondents and the impact of exogenous events. Three factors became important to respondents in the summer of 2013 (concurrent with the Snowden revelations): the collection of accelerometer data, and the use of friend and location information for targeted advertising. The potential impact of current events is an important consideration for survey researchers.

Finally, we included a free-text response section at the end of the survey, where participants could leave reactions. Respondent comments like "Is this really happening with our privacy?" are evocative of the overall findings: Current practices in the mobile application space do not meet user privacy expectations. This provides openings for scholars in policy, ethics, and business to suggest new models for just and equitable data collection in the mobile sector based upon contextual norms. While our emphasis here is on enabling mobile industry self-regulation through privacy by design, future work should consider whether forms of data collection that fail to meet consumer expectations in any context should be subject to state or national regulation.

Conclusion

Surveys asking American mobile application users to make judgments about whether data collection and use scenarios met their expectations demonstrated how complicated the space of user privacy expectations can be. Users expect particular data types, such as location and accelerometer data, to be used in the contexts of navigation and weather applications, but they do not expect this data to be used for targeting of advertisements. Users do, however, expect keyword harvesting to improve targeting of advertising. And they do not expect contact and image information to be harvested in any context. These data illustrate the applicability of contextual integrity and privacy calculus approaches to understanding information privacy, and suggest that the mobile application industry must respond to nuanced data collection and use expectations to retain consumers' trust.

Notes

1. A limitation of this method of survey deployment is that the researchers could not prevent repeat respondents, as Mechanical Turk workers are pseudonymous.
2. Amazon Mechanical Turk (MTurk) is an online labor market where requestors post jobs and workers choose jobs to complete. In a parallel study of privacy expectations for websites, results from Amazon Mechanical Turk were compared to a nationally representative sample purchased from GfK/Knowledge Networks. The results of the two samples were statistically comparable. The results are available from the second author.
3. The coefficients of contact and image information in the targeted advertising regressions are 2–3 times larger than in the tracking regressions; the tracking vignettes on average are rated lower on the scale of meeting privacy expectations due to the large negative impact of secondary use of information described in the following.

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.

Funding

This project was funded by the National Science Foundation (grant SES-1311823) and the University of Maryland ADVANCE Project seed grant program.

References

ACT. 2012. *Apps across America: The economics and ecosystem of the mobile app market*. Washington, DC: ACT.

Android Apps on Google Play. 2013. https://play.google.com/store/apps?feature=corpus_selector (accessed July 14, 2013).

App Store Downloads on iTunes. 2013. <https://itunes.apple.com/us/genre/ios/id36?mt=8> (accessed July 14, 2013).

Beales, J. H., and T. J. Muris. 2008. Choice or consequences: Protecting privacy in commercial information. *University of Chicago Law Review* 75 (1):109–135.

Behrend, T. S., D. J. Sharek, A. W. Meade, and E. N. Wiebe. 2011. The viability of crowdsourcing for survey research. *Behavior Research Methods* 43 (3):800–813. <http://dx.doi.org/10.3758/s13428-011-0081-0>.

Berinsky, A. J., G. A. Huber, and G. S. Lenz. 2012. Evaluating Online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis* 20 (3):351–368. <http://dx.doi.org/10.1093/pan/mpr057>.

Bowie, N. E., and K. Jamal. 2006. Privacy rights on the Internet: Self-regulation or government regulation? *Business Ethics Quarterly* 16 (3):323–342.

Boyles, J. L., A. Smith, and M. Madden. 2012. *Privacy and data management on mobile devices*. Washington, DC: Pew

Internet & American Life Project. <http://www.pewinternet.org/Reports/2012/Mobile-Privacy.aspx> (accessed October 8, 2012).

Braunstein, A., L. Granka, and J. Staddon. 2011. Indirect content privacy surveys: Measuring privacy without asking about it. In *SOUPS 2011 Proceedings of the Seventh Symposium on Usable Privacy and Security* 15:1–15:14. New York, NY: ACM.

Buchanan, T., C. Paine, A. N. Joinson, and U.-D. Reips. 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.

Buitelaar, J. C. 2014. Privacy and narrativity in the Internet era. *The Information Society* 30 (4):266–281.

Cate, F. H. 2010. The limits of notice and choice. *IEEE Security Privacy* 8 (2):59–62. <http://dx.doi.org/10.1109/MSP.2010.84>.

Cavoukian, A. 2012. *Operationalizing privacy by design: A guide to implementing strong privacy practices*. Toronto, ON, Canada: Office of the Privacy Commissioner, Ontario, Canada. <http://www.privacybydesign.ca/index.php/paper/operationalizing-privacy-by-design-a-guide-to-implementing-strong-privacy-practices> (accessed July 17, 2013).

Columbus, L. 2013. Roundup of mobile apps & App Store forecasts, 2013. *Forbes*, June 9, sec. tech. <http://www.forbes.com/sites/louiscolumbus/2013/06/09/roundup-of-mobile-apps-app-store-forecasts-2013> (accessed July 14, 2013).

Culnan, M. J. 1995. Consumer awareness of name removal procedures: Implications for direct marketing. *Journal of Direct Marketing* 9 (2):10–19.

Culnan, M. J., and R. J. Bies. 2003. Consumer privacy: Balancing economic and justice considerations. *Journal of Social Issues* 59 (2):323–342. <http://dx.doi.org/10.1111/1540-4560.00067>.

Culnan, M. J., and C. C. Williams. 2009. How ethics can enhance organizational privacy: Lessons from the Choice-Point and TJX data breaches. *Management Information Systems Quarterly* 33 (4):673–687.

Curry, M. R., D. J. Phillips, and P. M. Regan. 2004. Emergency response systems and the creeping legibility of people and places. *The Information Society* 20 (5):357–69.

Dunfee, T. W., N. C. Smith, and W. T. Ross, Jr. 1999. Social contracts and marketing ethics. *Journal of Marketing* 63 (3):14–32.

Eastlick, M. A., S. L. Lotz, and P. Warrington. 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. 2010. *Protecting consumer privacy in an era of rapid change*. Washington, DC: Federal Trade Commission. <http://www.ftc.gov/sites/default/files/documents/reports/federal-trade-commission-bureau-consumer-protection-preliminary-ftc-staff-report-protecting-consumer/101201privacyreport.pdf> (accessed May 18, 2012).

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

Gross, R., and A. Acquisti. 2005. Information revelation and privacy in online social networks. In *WPES '05 Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*, 71–80. New York, NY: ACM.

Heeney, C. 2012. Breaching the contract? Privacy and the UK Census. *The Information Society* 28 (5):316–328.

Hoffman, D. L., T. P. Novak, and M. A. Peralta. 1999. Information privacy in the marketspace: Implications for the commercial uses of anonymity on the Web. *The Information Society* 15 (2):129–39.

Hui, K.-L., H. H. Teo, and S.-Y. Tom Lee. 2007. The value of privacy assurance: An exploratory field experiment. *MIS Quarterly* 31 (1):19–33.

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

Johnson, D. G. 2004. Is the global information infrastructure a democratic technology? In *Readings in cyberethics*, ed. R. A. Spinello and H. T. Tavani, 2nd ed., 121–133. Sudbury, MA: Jones and Bartlett.

Kang, C. 2013. Flashlight app kept users in the dark about sharing location data: FTC. *The Washington Post*, December 6, sec. Tech. 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 (accessed May 14, 2014).

Kroener, I., and D. Wright. 2014. A strategy for operationalizing privacy by design. *The Information Society* 30 (5):355–65.

Lease, M., J. Hullman, J. Bigham, M. Bernstein, J. Kim, W. Lasecki, S. Bakhshi, T. Mitra, and R. Miller. 2013. *Mechanical Turk is not anonymous* (SSRN Scholarly Paper ID 2228728). Rochester, NY: Social Science Research Network. <http://papers.ssrn.com/abstract=2228728> (accessed March 8 2013).

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

Lin, J., S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. 2012. Expectation and purpose: Understanding users' mental models of mobile app privacy through crowdsourcing. In *UbiComp '12 Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 501–510. New York, NY: ACM. <http://dx.doi.org/10.1145/2370216.2370290>.

Martin, K. 2013. Transaction costs, privacy, and trust: The laudable goals and ultimate failure of notice and choice to respect privacy online. *First Monday* 18 (12):online.

Martin, K. 2015. Understanding privacy online: Development of a social contract approach to privacy. *Journal of Business Ethics*: 1–19. <https://link.springer.com/article/10.1007/s10551-015-2565-9>.

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. <http://dx.doi.org/10.1007/s10551-012-1215-8>.

Mayer, J., and A. Narayanan. 2013. Privacy substitutes. *Stanford Law Review Online* 66:89–96.

McCole, P., E. Ramsey, and J. Williams. 2010. Trust considerations on attitudes towards online purchasing: The moderating effect of privacy and security concerns. *Journal of Business Research* 63 (9):1018–1024.

McDonald, A. M., and L. F. Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4:540–565.

Milne, G. R., and M. J. Culnan. 2002. Using the content of online privacy notices to inform public policy: A longitudinal analysis of the 1998-2001 US Web surveys. *The Information Society* 18 (5):345–59.

Milne, G. R., M. J. Culnan, and H. Greene. 2006. A longitudinal assessment of online privacy notice readability. *Journal of Public Policy & Marketing* 25 (2):238–49.

Milne, G. R., and M. E. Gordon. 1993. Direct mail privacy-efficiency trade-offs within an implied social contract framework. *Journal of Public Policy & Marketing* 12 (2):206–15.

Nissenbaum, H. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (1):119–58.

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

Nissenbaum, H. 2011. A contextual approach to privacy online. *Daedalus* 140 (4):32–48.

Nock, S., and T. Guterbock. 2010. Survey experiments. In *Handbook of survey research*, eds. P. V. Marsden and J. D. Wright, 837–864. Bingley, UK: Emerald Group Publishing.

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

Phelps, J., G. Nowak, and E. Ferrell. 2000. Privacy concerns and consumer willingness to provide personal information. *Journal of Public Policy & Marketing* 19 (1):27–41.

Popescu, M., and L. Baruh. 2013. Captive but mobile: Privacy concerns and remedies for the mobile environment. *The Information Society* 29 (5):272–86.

Ross, J., L. Irani, M. S. Silberman, A. Zaldivar, and B. Tomlinson. 2010. Who are the crowdworkers?: Shifting demographics in Mechanical Turk. In *CHI EA '10: CHI '10 extended abstracts on human factors in computing systems*, 2863–2872. New York, NY: ACM. <http://dx.doi.org/10.1145/1753846.1753873>.

Rossi, P. H., and S. L. Nock. 1982. *Measuring social judgments: The factorial survey approach*. Beverly Hills, CA: Sage.

Sheehan, K. B. 2002. Toward a typology of internet users and online privacy concerns. *The Information Society* 18 (1):21–32.

Shilton, K. 2009. Four billion little brothers?: Privacy, mobile phones, and ubiquitous data collection. *Communications of the ACM* 52 (11):48–53.

Solove, D. J. 2006. A taxonomy of privacy. *University of Pennsylvania Law Review* 154 (3):477–564.

Spiekermann, S., and L. F. Cranor. 2009. Engineering privacy. *IEEE Transactions on Software Engineering* 35 (1):67–82.

Waldo, J., H. S. Lin, and L. I. Millett. 2007. *Engaging privacy and information technology in a digital age*. Washington, DC: The National Academies Press.

Wallander, L. 2009. 25 Years of factorial surveys in sociology: A review. *Social Science Research* 38 (3):505–20.

Xu, H., H.-H. Teo, B. C. Y. Tan, and R. Agarwal. 2012. Research note Effects of individual self-protection, industry self-regulation, and government regulation on privacy concerns: a study of location-based services. *Information Systems Research* 23 (4):1342–1363. <http://dx.doi.org/10.1287/isre.1120.0416>.

Xu, H., C. Zhang, P. Shi, and P. Song. 2009. Exploring the role of overt vs. covert personalization strategy in privacy calculus. *Academy of Management Proceedings* 2009 (1):1–6. <http://dx.doi.org/10.5465/AMBPP.2009.44249857>.

Appendix

Factors common to all vignettes

Factor	Dimensions	In vignette
Context. The business of the primary organization. The underlying activity or purpose surrounding the exchange.	Games Weather Social networking Navigation Search (reference) Music Banking/finance Shopping/retail Productivity	play ... a game ... Game Look up the forecast with ... a weather ... weather Check updates on...a social networking...social networking Get direction on ... a map map Listen to music on ... a music... music Check your balance on ... a banking ... banking Shop on ... a retail ... Retail Update your to-do list on ... a productivity ... productivity
Tenure. Time since downloaded.	Months/years (continuous)	a week ago ... less than a month ago ... 2 ... 3 ... 4 ... 5 ... 6 ... 7 months ago ...
Frequency. Frequency of use. Information. The type of information received or tracked by the primary organization.	Hours per week (continuous) Location Accelerometer Demographic Contacts Keywords Name Images Friends	Very frequently ... frequently ... occasionally ... infrequently... rarely ... your location when you accessed the application how quickly you are moving your age and gender your phone contact list what you did in the current application your name pictures taken with your phone activity of your friends on that same application
Rating #1 This app has met my privacy expectations.	Strongly Disagree	Strongly Agree

Context chosen based on mobile app categories provided by the two major app stores, iTunes app store and Google Play. As of July 2013, iTunes (“App Store Downloads on iTunes” 2013) identifies application categories as:

- Books
- Business
- Catalogs
- Education
- Entertainment
- Finance
- Food & Drink
- Games
- Health & Fitness
- Lifestyle
- Medical
- Music
- Navigation
- News
- Newsstand
- Photo & Video
- Productivity
- Reference
- Social Networking
- Sports
- Travel
- Utilities
- Weather

Google Play (“Android Apps on Google Play” 2013) identifies application categories as:

- Games
- Books & Reference
- Business
- Comics
- Communication
- Education
- Entertainment
- Finance
- Health and Fitness
- Libraries & Demo
- Lifestyle
- Live Wallpaper
- Media & Video
- Medical
- Music & Audio
- News & Magazines
- Personalization
- Photography
- Productivity
- Shopping
- Social
- Sports
- Tools
- Transportation
- Travel & Local
- Weather
- Widgets

I. Pilot I—Targeting advertisements

Factor	Dimensions	In vignette
AdType.	Primary Org Ad 3rd Party Ad	Another application they sell Another company's mobile app

Vignette template

While using your phone, you {Context_alt} {Context_alt2} application that you have used {Frequency_alt} for {Tenure_alt}.

The {Context_alt3} app shows you an advertisement for {AdType_alt} based on {Information_alt}.

Sample 1

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.

II. Pilot II—Tracking data

Factor	Dimensions	In vignette
Age. Time stored	Continuous months	Less than a week, a month, 2 months, 4 months, 6 months ... 12 months
Personalization	NULL Device ID	a unique identifier for your mobile device
Collection. Who collects the information	Primary organization Wireless provider Platform provider 3rd party tracking	the mobile application ... app your wireless provider ... phone company the app store company ... app store an outside company's invisible tracking program ... tracking company
Second Use. What the collecting organization does with the information	Retargeting Data exchange Social advertising	uses the information for future ads when you are using this app sells the data in an online auction uses the information for future ads targeting your friends and contacts.

Vignette template

You are {Context_alt} {Context} application on your phone that you have used {Frequency} for about {Tenure}. On the {Context_alt3} app, {Information} {Information_alt} collected by {Collection} and will be stored for {Age}. The data collected also includes {Personalization}.

The {Collection_alt} then {Second Use}.

Sample 1

While on your phone, you update your to-do list a scheduling app application that you have used infrequently for 3 months.

Through the scheduling app, your phone contact list are 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.