Mobile Privacy Expectations in Context

Katie Shilton, University of Maryland College Park
Kirsten Martin, George Washington University

An increasing amount of social activity and commerce is performed using applications running on mobile devices such as phones and tablets. During these activities, mobile applications collect increasing amounts of personal data. Consumers, organizations, and regulators struggle to address privacy expectations for these new forms of data collection across a diverse set of activities. This paper describes findings from empirical research employing a context-based survey to understand consumers' privacy expectations for mobile devices across diverse real-world contexts (e.g., consulting medical applications, navigating using map applications, or playing music or games, etc). The project tests the hypotheses that (a) individuals hold different privacy expectations based on the social context of their mobile activity, and (b) contextual factors such as *who* (the data collection actor, e.g. the application developer or mobile phone provider), *what* (data attributes, e.g. the type of information received or tracked by the primary organization), *why* (application purpose, e.g. playing games, checking the weather, participating in social networking, navigating using maps, listening to music, banking, shopping, and organizing personal productivity) and *how* (the use of data by the collector, e.g. the amount of time data is stored or how that data is reused) affect individuals' privacy expectations.

This paper reports on survey findings that identify contextual factors of importance in the mobile data ecosystem. Our survey demonstrated that overall, very common activities of mobile application companies such as harvesting and reusing location data, accelerometer readings, demographic data, contacts, keywords, name, images and friends do not meet users' privacy expectations. But these differences are modulated by both information type and social context. Addressing privacy expectations for mobile devices is an explicit goal of US regulatory bodies and firms rely on understanding privacy expectations to gain the trust of users. Understanding how consumer privacy expectations change in different data use and business contexts can help regulators identify contexts that may require stricter privacy protections and help firms and managers better meet privacy expectations of users. Study results help us understand one aspect of mobile privacy: the expectations of consumers as they vary by context. These expectations have direct implications for researchers, business leaders, policy experts, and consumers.

Introduction

Individuals increasingly do their banking, play games, find doctors, pay bills, catch up with friends, and find and rate restaurants using mobile devices. As of May 2012, 10% of Internet traffic stems from mobile devices – up from just 1% 18 months earlier. While an increasing amount of social activity and commerce is performed using dedicated software loaded onto mobile devices – or *mobile applications* – consumers, organizations, and regulators struggle to understand the privacy expectations across a diverse set of activities. The White House has begun to encourage federal agencies to deal with the privacy consequences of mobile applications (Federal Trade Commission, 2012; Strickling, 2012).

Notice and consent policies have been the dominant mechanism used by organizations and regulators to address privacy expectations online and they will likely be the dominant mechanism for mobile application regulation. Yet even proponents agree that notice and consent does not sufficiently

address privacy issues in certain contexts. In contrast, *contextual privacy norms* allow privacy expectations to vary across contexts. Contextual privacy norms assume that privacy expectations about the type and transmission of information are dependent upon the context (Nissenbaum, 2009). This promising direction of privacy scholarship provides the theoretical backbone for empirical work to discover consumers privacy expectations *in context*. Contextual privacy norms posit that expectations about use and transmission of information are dependent upon the context. Individuals exchange information with particular people for a specific purpose, and expectations around what constitutes appropriate flows of information vary across such contexts. The theory suggests that tactics to address privacy expectations with mobile devices should depend on the context of the exchange.

This paper describes findings from empirical research employing a context-based survey to understand consumers' privacy expectations for mobile devices across diverse real-world contexts. The study asks:

1) How do individuals' privacy expectations change between mobile application contexts?

2) What factors change these expectations?

The project uses a factorial vignette survey to test the hypotheses that (a) individuals hold different privacy expectations based on the social context of their mobile activity, and (b) contextual factors such as *who* (the data collection actor, e.g. the application developer or mobile phone provider), *what* (data attributes, e.g. 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 individuals' privacy expectations. Two surveys were designed to investigate judgments about privacy expectations around targeted advertising and tracking users in the mobile space. Amazon Mechanical Turk was used to gather 979 respondents to judge over 39,000 vignettes describing different mobile application scenarios. The data was gathered at 3-month intervals (May and August 2013) as part of an ongoing longitudinal study to identify if privacy expectations are changing.

This design mitigates several concerns in privacy research. First, privacy surveys are also fraught with respondent bias where respondents inflate their concern for privacy that may not reflect their true attitude (Hui, Teo, and Lee, 2007). For example, despite a reported general 'concern for online privacy', 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, researching, and playing games. As noted by the recent FTC report, traditional surveys are limited in their ability to measure privacy expectations of individuals (FTC, 2010, fn 72). In fact, people appear unconcerned about privacy until violated or breached, and they do not act according to privacy preference they claim to have in traditional surveys (Spiekermann & Cranor, 2009, p. 71). For example, respondents express a concern for privacy at a macro level,[1] yet their general concern does not translate to privacy protecting behavior within specific contexts (Gross and Acquisti, 2005).

Background

Privacy is a persistent government and societal concern, and the introduction of new technology has consistently challenged commonly-held expectations of privacy. While US Internet users saw flat

[1] The general concern for privacy is captured in statements such as "how concerned are you, if at all, about Internet privacy?" See also Tufekci (2008) and Young and Quan-Haase (2009).

growth in 2011 overall, mobile Internet users experienced 31% year-over-year growth in 4Q2011 (Meeker, 2012). 55% of all cell phone users access the Internet through their device, which doubled since 2009 (Boyles, Smith, & Madden, 2012). 17% of cell phone users now do *most* of their browsing on their phone (Smith, 2012). While mobile information technology enables novel possibilities for human interaction, these new possibilities change the privacy landscape in important ways (Johnson, 2004). For example, mobile technologies enable information to move quickly and to a wide audience with many-to-many communication while also allowing information to be easily captured, indexed, saved, and searched (Gelman, 2009; Tufekci, 2008). But this growing market currently lacks standards of practice for addressing privacy, and most uses of this data are currently unregulated (Federal Trade Commission, 2012; Shilton, 2009).

Mobile devices change the manner in which information can move, the people who can see it, the actions that can be taken, and the amount of time it is retained. First, new information is available through smartphones, tablets, and e-readers. Mobile data such as location information, activities, motion information, text, and sound are more easily gathered through mobile devices. Second, new data collection methods and actors are present with mobile applications. Mobile information may be collected by application developers (e.g. Rovio Games), mobile providers (e.g. AT&T), or operating system providers (e.g. Google), device manufacturers, and smartphone or tablet brands (e.g., Blackberry, Apple) in addition to the telecommunication provider all have access and collect mobile data in addition to the 'standard' 3rd party tracking companies online. Finally, the smartphone or tablet is more personal and connected to the individual. Many carry their mobile devices with them at all times and the device is 'owned' by them in a way that differs from the shared personal computer or laptop in a family, at a university, or in a business. Individuals even personalize their mobile devices and use them in a way that suggests a different degree of privacy.

Privacy as Contextual Integrity. A promising direction of privacy scholarship – privacy as contextual integrity – posits that privacy expectations about the type and transmission of information are dependent upon the social context (Nissenbaum, 2009). Individuals exchange information with particular people for a specific purpose—people vote, go to the doctor, do taxes, hang out with friends – and expectations around what constitutes appropriate flows of information vary across such contexts. According to privacy as contextual integrity, individuals do not hold universally-applicable definitions of privacy; instead, individuals give access to 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 each 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, p. 134). Maintaining privacy norms entails the "information gathering and dissemination be appropriate to that context and obey the governing norms of distribution within it" (Nissenbaum, 2004, p. 101).

Within a given context – e.g. education, healthcare, employment, banking, gaming, social networking, shopping – norms of contextual integrity govern the information flow between individuals and organizations (Nissenbaum, 2009). Communities develop privacy norms around:

- Actors – people, organizations, technologies who are the senders, recipients, and subjects of information.
- Attributes – the information types or data fields being transmitted.
- Transmission principles – the constraints on the flow of information.

These facets can be seen as working in concert. For example, for a medical professional, attributes such as medical history and overall medical concerns are expected to be transmitted to the doctor and her staff. However, in a different context, such as the workplace, the same information would be considered inappropriate. Even a question about medical history is deemed inappropriate and a violation of privacy expectations when at work. Within a context, for every given set of actors and attribute, there exists a transmission principle. Similarly, for every given set of actors and transmission principle, there exists a set of attributes. Key to Nissenbaum's privacy as contextual integrity is how the main components work together – actor, attributes, and transmission principles – within a particular context.

Privacy Expectations and Mobile Apps. Privacy as contextual integrity suggests that tactics to address privacy expectations on mobile devices should depend on the context of the exchange. For example, notice and consent may be required for some contexts but not for all contexts; anonymity may be appropriate for some contexts, such as internet search, but inappropriate in a context such as social networking. In addition, information cannot be deemed 'private' or 'public' across contexts. Such a designation is contextually defined. Therefore, current attempts to tag information online as having a universal transmission principle, such as public, private, or sensitive, could unnecessarily restrict the flow of information in some circumstances and inappropriately pass on information in other contexts (see also FTC, 2010, p. 31). Finally, tactics such as behavioral advertising, data collection and retention, or tracking, may be appropriate and within the contextually-defined privacy norms in one context while inappropriate in another.

Empirical work is needed to identify the data, practices, and recipients for different mobile application contexts. Thus far, empirical studies have tested for the presence of a preconceived definition of privacy within a group of people and have assumed individuals have a persistent understanding of privacy expectations across contexts. 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 contexts in mobile activity, the researchers conducted a survey using factorial vignette methodology (Wallander, 2009), in which respondents answer questions based on a series of hypothetical vignettes. This method allows the researchers to simultaneously examine multiple factors – e.g. changes in context and types of information sharing – by providing respondents with rich vignettes which are systematically varied. This study answers the research questions:

1. How do individuals' privacy expectations vary across mobile application contexts?
2. How do privacy judgments about mobile applications change over time (if at all)?

Toward this end, the factorial vignette survey methodology, developed to investigate human judgments (Rossi and Nock 1982; Jasso 2006; Wallander 2009), was employed. In a factorial vignette survey, a set of vignettes is generated for each respondent, where the vignette factors or independent variables are controlled by the researcher and randomly selected, and respondents are asked to evaluate these hypothetical situations. In the factorial survey approach, each respondent rates the level of an outcome (in this case, the degree to which the mobile application in the vignette meets the privacy expectations of the respondent) corresponding to the unit of analysis described in the vignette (Jasso

2006). The respondent is presented a large set of hypothetical vignettes and statistical techniques are used to identify the implicit factors and their relative importance driving the outcome variable for the respondents.

The vignettes vary based on relevant factors and are controlled and presented by the investigator to ensures intercorrelations among vignette characteristics are zero. Factorial survey methodology allows for the simultaneous experimental manipulation of a large number of factors through the use of a contextualized vignette (Ganong and Coleman 2006).[2] 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 Gutterbock 2010). These factors and their associated coefficients are the *equations-inside-the-head* (Jasso 2006) of respondents as to judgments of privacy expectations.

The vignettes were constructed by varying several factors along dimensions or levels. A deck of vignettes for each respondent was randomly created with replacement as the respondent was taking the survey from a vignette universe. For each rated vignette, the associated rating, factor levels, and the vignette script was preserved as well as the vignette sequence number. The vignette format is also provided in the appendix below with a sample vignette and the vignette template. We pilot-tested the vignettes for clarity with students in several courses at the University of Maryland.

Independent Variables

Each survey respondent was shown a series of vignettes which 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 purpose – playing games, checking weather, participating in social networking, navigating using maps, listening to music, banking, shopping, and organizing personal productivity – as well as how frequently and for how long an application has been used;
- How (used): Transmission principles – e.g., how the data is reused or stored.

Social contexts of mobile use. Defining meaningful social contexts is one of the challenges inherent in understanding privacy in context. We have chosen an approach 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 social contexts according to how they are identified by the two major application stores: the Apple iTunes store, and Google Play. See Appendix for the list of possible contexts and categories from these sources. We also sought industry data on the most dominant uses of mobile applications. According to an industry survey, email and calendaring, Instant Messaging (IM), office & personal productivity, web conferencing, and e-commerce are the most popular uses of mobile applications (Columbus, 2013).

We chose mobile application contexts 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

[2] In comparison, in experiments, factors are designed orthogonal to each other but manipulated one at a time; however, in a traditional survey, many factors are examined but are not necessarily orthogonal to each other (Appelbaum, Lennon, and Aber, 2006).

to face transactions, such as medical and banking. The final social contexts chosen to test in the first set of surveys (May 2013) were:

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

During the second round (August 2013), we added two additional sensitive application contexts to see if these would have a greater impact on results. To the previous list, we added:

- Activity Monitor
- Symptom Checker

Application Privacy Factors. Factors such as the overall purpose of the application (context), the amount of time spent using the application (tenure and frequency), and the type of information gathered and used (such as location, accelerometer, or demographic data) were included across both targeting and tracking situations. The appendix contains the factors and the possible levels for each vignette type in a table as well as the vignette templates and sample vignettes.

We formed two types of vignettes: stories about targeted advertisements and stories tracking mobile data use. In addition to social context, tenure, and frequency, the **targeting** vignettes contained statements about *information* (the type of information received or tracked by the primary organization, such as Location, Accelerometer, Demographic, Contacts, Keywords, Name, Images, and Friends), and *ad type* (what the organization does with the information, either using it to target their own ads or selling it to a third party). This generated vignettes like the following – underlining highlights the factors that would systematically change.

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.

Stories about **tracking** mobile data also featured *information* types, as above. In addition, 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), *collection* (who collects the information, such as the primary organization, your wireless provider, your platform provider, 3rd party tracking), *secondary use* (what the collecting organization does with the information, such as retargeting, data exchange, or social advertising). This generated vignettes like the following:

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 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.

Control Variables

The respondents' age and gender are used in the regression analysis in addition to two control questions to gauge their overall trust in application and concern about application privacy. In addition to 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' to the right being 'strongly agree.' The rating task stated 'In general, I trust mobile applications.' Their rating on this question was used as their general level of trust in applications. The second rating task stated, 'In general, I believe privacy is important.' Finally, respondents were asked to judge the frequency of their own use of mobile applications – from rarely to multiple times per day.

Dependent Variable: Privacy Rating Tasks

For each vignette, respondents were given a rating task depending on the survey type 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'. For the privacy expectations survey, the respondents were given the prompt, 'This application meets my privacy expectations.'

Sample

The mobile privacy surveys were deployed two times (May and August 2013) to 979 respondents to rate over 39,000 vignettes using Amazon's Mechanical Turk – see Table 1. Though use of Mechanical Turk for survey deployment is controversial (Lease et al., n.d.; Ross, Irani, Silberman, Zaldivar, & Tomlinson, 2010), studies have shown that mTurk workers are more representative of the US 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 delivering this survey. [3]

Analysis

The data in this study was analyzed on two levels: the vignette level factors including privacy factors, and the respondent-level control variables. For the May survey, 497 respondents rated 40 vignettes each, resulting in 19,880 rated vignettes (e.g. 19,880 observations) for the privacy expectations surveys (May -- targeting: 215/8,600 respondents/vignettes; tracking: 98/3,920). For August, 482 respondents rated 19,280 vignettes as shown in Table 1.

If N is the number of the respondents with level 2 demographic variables and K is the number of vignettes answered with level 1 factor variables, the general equation is:

(1) $$Y_{nk} = \beta_0 + \Sigma\beta_j V_{jk} + \Sigma\gamma_h R_{hn} + u_n + e_k$$

[3] Amazon Mechanical Turk (MTurk) is an online labor market where requestors, such as academics, post jobs and the workers, such as the resondents, choose jobs to complete. For a full description, see Mason & Suri (2011), for how MTurk samples are more representative of the U.S. population than in-person convenience samples, see Berinksy, Huber, and Lenz (2012), and for the external and internal validity of MTurk, see Horon, 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. Horton, Rand, and Zeckhauser (2011) illustrate how behavioral economics experiments are successfully replicated on MTurk.

where Y_{nk} is the rating of vignette k by respondent n, V_{jk} is the j^{th} factor of vignette k, R_{hn} is the h^{th} characteristic of respondent n, β_0 is a constant term, β_j and γ_h are regression coefficients, u_n is a respondent-level residual (random effect), and e_{ik} is a vignette-level residual. The model conceptualizes the ratings as a function of the factors of the situation described in the vignette and the characteristics of the respondent as hypothesized above.[4] As the data can be modeled at two levels – the vignettes and the individual respondents – multi-level 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.

Results

Tracking versus Targeting

Based on the broad results in Table 1, tracking scenarios met privacy expectations to a lesser extent than targeting scenarios both May (β = -43.53 v. -22.24 respectively) and August (-38.55 v. -16.30). Both types of vignettes did not meet privacy expectations on average, as both means are negative. However, both tracking and targeting vignettes trend positive between May and August, as shown in Figure 4, suggesting the same mobile application scenarios met privacy expectations of users to a greater extent by August. Finally, the respondent-level R2, created by running regressions for each respondent based on their 40 vignettes, was larger for targeting (β = 0.820) as compared to tracking (0.776) in May, and this gap narrowed in August. This illustrates that individual respondents were more slightly more certain of their judgments for targeted advertising as compared to tracking users. Figure 4 also illustrates the certainty across respondents in judging tracking and targeting vignettes. So while tracking evoked a larger negative response, it also was a subject of less certainty within and across individual respondents.

Relative Importance of Contextual Privacy Factors

In addition, the overall sample of both targeting and tracking vignettes was used to run regressions of the rating task ("This application met my privacy expectations") onto vignette-level and respondent-level factors as shown in Table 2 and 3. Both tracking and targeting vignettes for May and August were run as four separate regression equations. Important distinctions occurred in response to *information types, secondary use, actors,* and *context*.

Information (What). Generally, the type of information mattered to respondents. The use of contact information (β = -59.66), image information (-65.77), and the individual's name (-17.76) negatively impacted meeting privacy expectations for targeted advertising compared to using demographic information. However, using keywords (18.48) positively impacted meeting privacy expectations for targeted advertising compared to using demographic information. This trend held for both August and May. In other words, respondents expected keyword collection in targeted advertising, but did not expect use of contact information, image information, or the individual's name. However,

[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 where the respondents take 1-2 vignettes to get acclimated to the methodology. The analysis was run minus the first 2 vignettes for each respondent and the results remained the same.

both location and accelerometer information had negative impact on meeting privacy expectation for targeted advertising in August after being insignificant in May. For tracking users, the use of contact information (-20.69), image information (-25.11), location information (-4.65), and friend information (-8.23) all negatively impacted meeting privacy expectations. These trends also held from May to August. However, the impact of each of these factors varies across contexts.

Tables 4 and 5 include the multi-level regression equations *for each context* to identify how the relative importance of information type varies across contexts. For example, for targeting vignettes, accelerometer information has no effect on privacy expectations in May, with the exception of navigation and weather contexts, where the information has a positive impact on privacy expectations ($\beta = 11.52$ and 8.67 respectively). Respondents expected accelerometer information to be used for targeting in both navigation and weather applications in May. Yet by August, accelerometer information had a consistent and significant negative impact on meeting privacy expectations. In both May and August, location information had a positive impact for weather applications, but a negative impact for music, social, banking, and retail applications. In May, collecting information about friends had a negative impact only on social, banking, and retail applications. By August, using friend information had a consistently negative impact across contexts. Interestingly, using both contact and image information was consistently found to have a negative impact on privacy expectations across contexts – e.g. using this information never met respondents' privacy expectations.

Secondary Use (Why). For tracking users, the secondary use of information was the most important factor in meeting privacy expectations. Selling to a data exchange ($\beta = -44.86$) and using tracked information for social advertising to contacts and friends (-21.34) both negatively impacted meeting privacy expectations. This trend held from May to August (Table 3) and across all contexts (Table 5). In addition, the amount of time the tracked information was stored was not significant in May, but negatively impacts meeting privacy expectations for retail, social networking, music, productivity, and weather applications in August.

Actor (Who). For tracking users, only third-party collectors impacted whether a scenario met privacy expectations. The platform, wireless provider, and primary application were treated statistically identically. The presence of a third party collector meant the vignette met privacy expectations slightly less ($\beta = -8.93$) whereas the presence of the platform as a collector of information had a slightly positive impact on privacy expectations (+2.12) in May and no impact in August. In particular, third party collection had a negative impact in retail, navigation, and activity monitoring in August.

Context. The impact of context on privacy expectations was complicated. Generally, as a factor in a linear regression, context did not significantly impact the rating task directly, with the exception of banking, where all behaviors met the privacy expectations of users to a lesser extent than in other contexts ($\beta = -13.02$ for targeting and -10.30 for tracking). The addition in August of two other categories we assumed to be sensitive – activity monitoring and medical symptom checking – appeared not to impact privacy expectations. Instead, data type remains the most important variable in meeting privacy expectations, and varies across contexts. Tables 4 and 5 illustrate how the relative importance of data type as a factor in privacy expectations changed across contexts.

Discussion

The results indicate that overall, what are 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 these differences are modulated by both information type and social context. For example, consumers seem to expect navigation and weather applications to use location and accelerometer data, and consumers seem to expect a link between harvesting keywords and targeted advertising.

Some data types, however, did appear to be particularly sensitive – or at least, particularly surprising to respondents. Harvest of both image and contact data failed to meet consumer privacy expectations in any context. This may be because it is not widespread knowledge that these data *can* be harvested by mobile applications, or it may be that these data types are particularly sensitive.

In our initial survey, banking was the only social context in which respondents had clear negative reactions to being both targeted and tracked. To examine this result, we added two more presumably sensitive contexts to the second survey: activity monitoring and medical symptom checking. However, banking remained the only context in which respondents had clear negative reactions to all factors. Ongoing work will determine whether other contexts may be more sensitive, or whether data type will remain the most decisive variable.

Finally, we included a free-text response section at the end of the survey where participants could leave reactions. Comments like "Is this really happening with our privacy?" are evocative of the overall findings – that current practices in the mobile application space do not meet user privacy expectations. Whether this will change over time as users come to understand what their apps can do remains to be seen.

Future Work

We plan to deploy the survey every 3-6 months for at least the next two years to test for any longitudinal changes. Though this is a short time period, the mobile application space is changing quickly. While Internet use growth was flat in 2011, mobile Internet use is up 31% over the previous year as of 4Q2011. Understanding if and how privacy expectations change during this interesting and evolving period of time with illustrate the temporary conceptual muddles (Moor, 1997) created by novel technology.

Conclusions

Surveys deployed in May and August of 2013 to measure consumer privacy expectations for mobile application demonstrate how complicated the space of privacy expectations can be. Consumers expect particular data types, such as location and accelerometer data, to be used to improve services in the contexts of navigation and weather applications, but they do not expect this data to be used for targeting. Consumers do, however, expect keyword harvesting to improve targeting. And they do not expect contact and image information to be harvested in any context.

Longitudinal data comparison also suggests some interesting trends. Between May and August 2013, the amount of time data was stored came to matter more to consumers, and both location and

accelerometer information had a more consistently negative impact on consumers' privacy expectations. We can speculate on causes for these shifts, ranging from growing consumer knowledge of mobile application data markets, to broad public attention to privacy in the US press due to revelations about the NSA's domestic and international data collection programs. Ongoing work is needed to monitor how these trends change over time, conduct analysis into *why* these changes are occurring, and inform both policymakers and application developers of the privacy expectations of consumers.

Figure 1: Mean Rating Tasks for Targeting/Tracking and May/Aug.

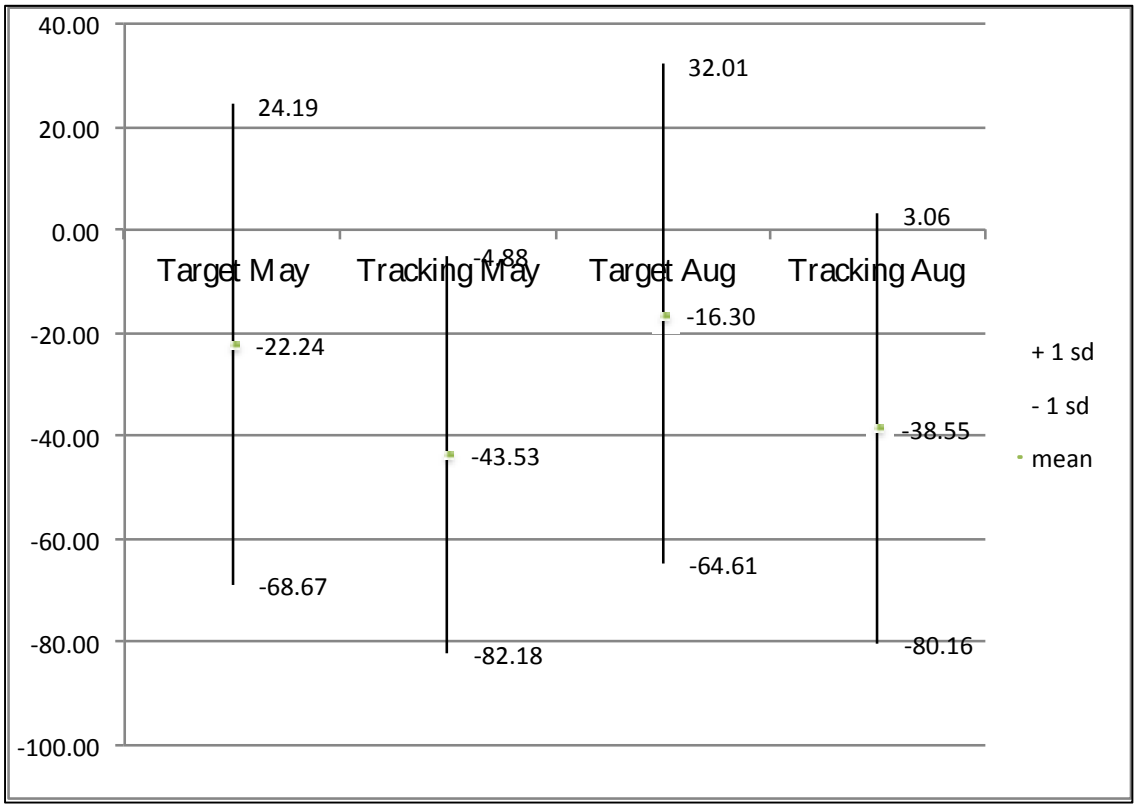

FIGURE 2: Charts of Main Regression comparing coefficients –

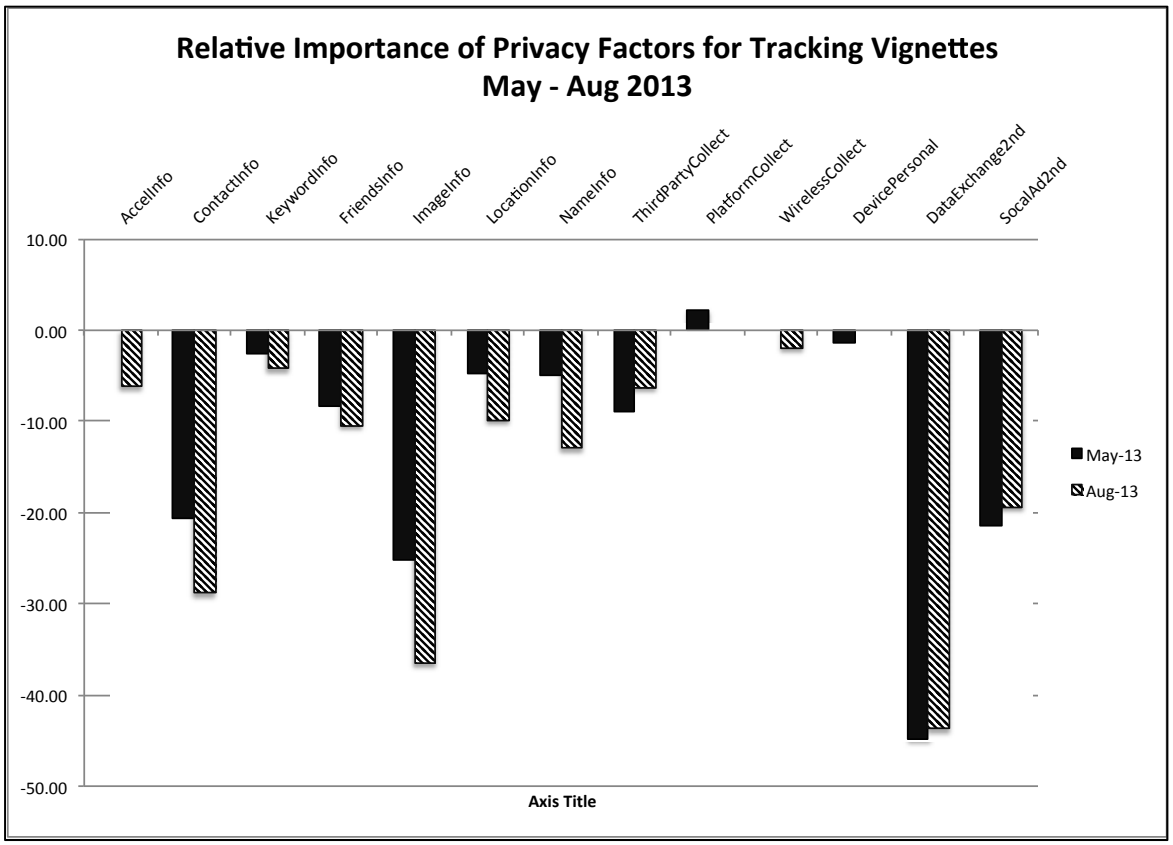

FIGURE 3: Relative Importance of Privacy Factors for Targeted Advertising

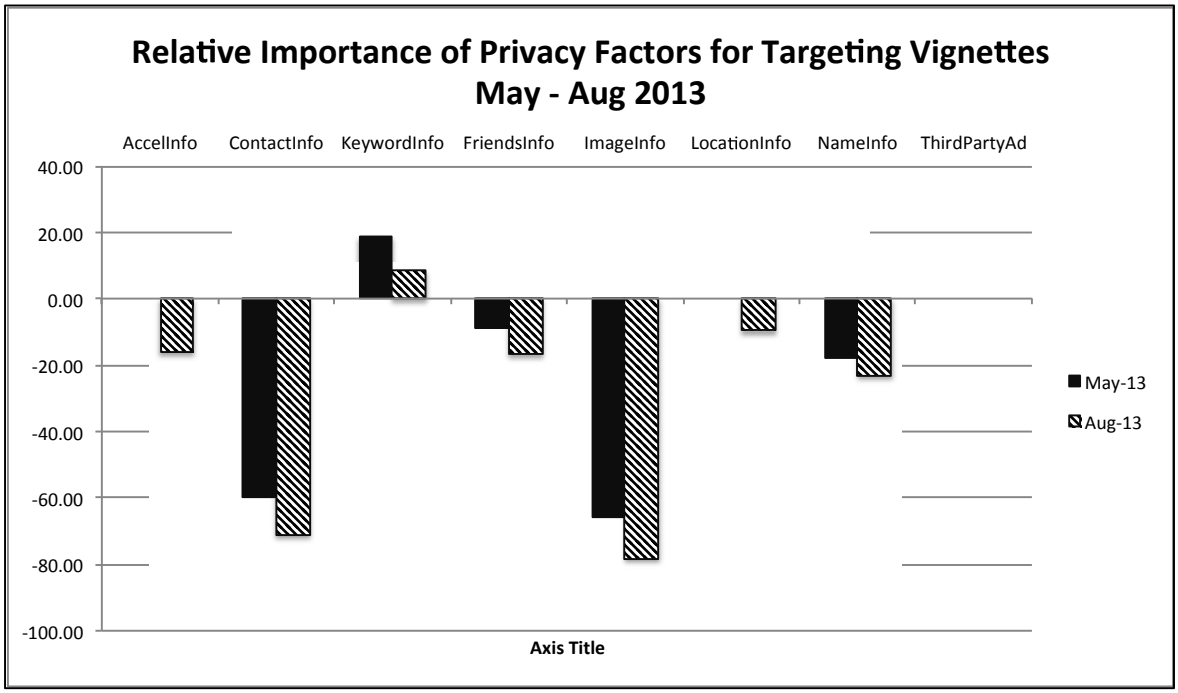

FIGURE 4: Mean Privacy Important and Trust in Apps Control Variables Plus Tracking and Targeting Rating Task.

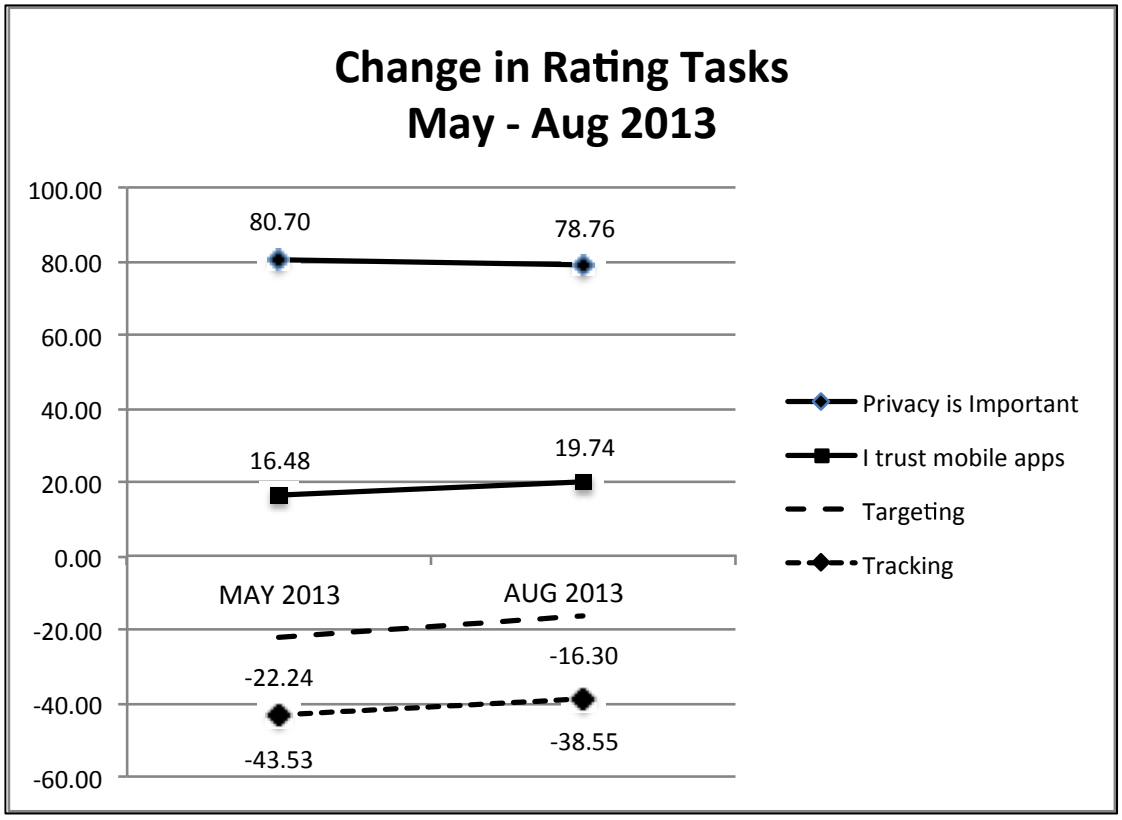

TABLE 1:

Date	**MAY 2013**		**May-13**		**May-13**	**AUG 2013**		**Aug-13**		**Aug-13**
Survey	**Targeting**		**Tracking**		**ALL RESP**	**Targeting**		**Tracking**		**ALL RESP**
Users	250		247		497	247		**235**		482
Vignettes	10,000		9880		19880	9880		9400		19280
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Mean	Std. Dev.	Mean	Std. Dev.	Mean
Age	32.21	9.74	31.40	10.76	31.81	30.58	9.62	30.69	10.44	30.64
Male	0.55		0.44		0.50	57.9%		64.3%		0.61
Privacy is Important	82.11	22.97	79.27	30.00	80.70	78.77	27.43	78.74	27.06	78.76
I trust mobile apps	19.29	42.48	13.63	44.52	16.48	22.22	44.33	17.13	48.27	19.74
Mean (DV)	-22.24	27.84	-43.53	32.29	-32.82	-16.30	27.30	-38.55	32.33	-27.15
R2 of Users	0.820	0.171	0.776	0.147	0.80	0.84		0.83		0.83
R2 of Sample	0.290		0.220		0.26	0.31		0.25		0.28

TABLE 2: Targeting Vignettes **KEY: Bold**: p < 0.05; Grey: p<0.10; Blank: p > 0.10

TARGETING VIGNETTES		
	May-13	**Aug-13**
tenure	**0.46**	
frequency	**0.58**	
Context		
BankingCxt	**-13.02**	**-9.01**
SocialCxt		
GamesCxt		3.71
MusicCxt		
ProductivityCxt		
WeatherCxt		
NavigateCxt		
ActivityCxt	n/a	
SymptomCxt	n/a	**-4.88**
(null = Retail		
Information		
AccelInfo		**-16.11**
ContactInfo	**-59.66**	**-71.14**
KeywordInfo	**18.48**	**8.51**
FriendsInfo	**-8.98**	**-16.71**
ImageInfo	**-65.77**	**-78.15**
LocationInfo		**-9.71**
NameInfo	**-17.76**	**-23.10**
(null = Demo)		
AdType		
ThirdPartyAd		
(null = Primary		
Control Variables		
Male		
Age		
AgeOver30		
TrustApps	**0.25**	**0.16**
PrivacyImporta	**-0.20**	**-0.28**
_cons	8.67	**41.55**
N	10000	9880
OLS R2	0.29	0.30
ICC	26.7%	24.5%

TABLE 3: MAYAUG 2013 Tracking Vignettes **KEY**: **Bold**: p < 0.05; Grey: p<0.10; Blank: p > 0.10

TRACKING VIGNETTES	May-13	Aug-13
Familiarity with site		
tenure		
frequency	**0.80**	
Context		
BankingCx	**-10.30**	**-7.50**
SocialCxt		
GamesCxt		
MusicCxt		**3.71**
Productivity(		
WeatherCxt		
NavigateCxt		
ActivityCxt	n/a	
SymptomC	n/a	**-4.03**
(null = Retail)		
Information		
AccelInfo		**-6.18**
ContactInf	**-20.69**	**-28.80**
KeywordIn	-2.49	**-4.23**
FriendsInf	**-8.23**	**-10.41**
ImageInfo	**-25.11**	**-36.42**
LocationIn	**-4.65**	**-9.86**
NameInfo	**-4.99**	**-12.95**
(null = Demo)		
Collecting Actor		
ThirdParty	**-8.93**	**-6.25**
PlatformC	**2.12**	
WirelessColl		-1.94
(null = Primary		
Personalization		
DevicePerso	-1.32	
(null = Null		
Second Use		
DataExcha	**-44.86**	**-43.67**
SocalAd2n	**-21.34**	**-19.35**
(null = Retarget		
Stoarge Months		
	-0.50	**-0.73**
Control Variables		
Male		
Age		
AgeOver30		
TrustApps	0.17	0.26
PrivacyIm	-0.30	-0.25
_cons	**26.57**	**24.13**
N	9880	9400
OLS R2	0.2197	0.2454
ICC	37.5%	32.0%

TABLE 4:
MAY & AUG 2013 REGRESSION OF PRIVACY FACTORS FOR **TARGETING** VIGNETTES BY CONTEXT
KEY: **Bold**: p < 0.05; Grey: p<0.10; Blank: p > 0.10

	TARGETING VIGNETTES by CONTEXT				May-13			
	Retail	**Banking**	**Social**	**Games**	**Music**	**Productivity**	**Weather**	**Navigation**
tenure				1.06				
NewFrequency								**1.45**
Information								
AccelInfo							**8.67**	**11.52**
ContactInfo	**-67.78**	**-58.30**	**-58.22**	**-62.31**	**-65.16**	**-51.20**	**-56.60**	**-51.02**
KeywordInfo	**15.19**		**17.95**	**25.93**	**20.75**	21.02	32.99	19.37
FriendsInfo	**-19.24**	**-13.86**	**-15.06**					
ImageInfo	**-74.03**	**-55.26**	**-67.68**	**-62.39**	**-76.32**	-61.47	**-56.50**	-63.48
LocationInfo	-8.61	**-9.43**	**-11.61**		**-10.69**		**18.78**	
NameInfo	**-19.33**	**-17.31**	**-15.32**	**-18.77**	**-22.18**	-12.07	-16.04	-13.81
(null = Demo)								
AdType								
ThirdPartyAd			**-4.81**					-4.24
(null = Primary								
Control Variables								
Male		8.32						8.03
Age								
AgeOver30								
TrustApps	**0.28**	**0.23**	**0.25**	**0.28**	**0.23**	**0.28**	**0.23**	0.27
PrivacyImportant	**-0.23**	**-0.22**	**-0.24**	**-0.22**	-0.16	-0.22	-0.23	
_cons	16.15	-11.63	10.06	7.28	16.34	**16.65**	**-1.35**	**-5.58**
N =	1247	1276	1181	1234	1262	1235	1292	1273
R2 =	0.32	0.21	0.30	0.33	0.32	0.30	0.33	0.30
ICC	27%	36%	32%	30%	25%	32%	28%	29%
sd(_cons)	23.90	29.43	26.36	25.01	23.37	26.26	24.95	25.74

	TARGETING VIGNETTES by CONTEXT AUGUST 2013									
	Retail	**Banking**	**Social**	**Games**	**Music**	**Productivity**	**Weather**	**Activity**	**Navigation**	**Symptom**
tenure									-1.21	
NewFrequency		**3.29**								
Information										
AccelInfo	**-18.96**	**-11.15**	**-18.79**	**-23.47**	**-16.99**	**-13.86**	**-14.86**		**-16.78**	**-17.83**
ContactInfo	**-79.81**	**-54.53**	**-76.38**	**-72.52**	**-66.64**	-63.76	**-76.10**	**-76.25**	-75.91	**-73.61**
KeywordInfo			10.58		**14.37**		**16.06**	10.11	**9.11**	
FriendsInfo	**-20.19**	**-25.16**	**-13.78**	-10.34	**-12.55**	**-11.05**	**-14.18**	**-18.61**	**-12.63**	**-28.15**
ImageInfo	**-85.17**	**-72.90**	**-78.64**	**-84.55**	**-71.96**	-75.37	-77.84	**-75.65**	-78.14	**-77.88**
LocationInfo	**-15.86**		**-13.22**	**-16.66**	**-14.40**			-9.11		**-13.96**
NameInfo	**-24.26**	**-11.72**	**-21.29**	**-20.51**	**-24.90**	-25.35	-21.58	**-27.08**	-27.96	**-26.56**
(null = Demo)										
AdType										
ThirdPartyAd		-5.04								
(null = Primary										
Control Variables										
Male									**8.59**	
Age		-0.68							-0.99	
AgeOver30	**-13.25**									
TrustApps	**0.11**	**0.18**	**0.17**	**0.21**	**0.12**	**0.16**	**0.16**	**0.12**	**0.19**	**0.22**
PrivacyImportant	**-0.24**	**-0.29**	**-0.30**	**-0.22**	**-0.29**	**-0.22**	**-0.29**	**-0.29**	**-0.20**	**-0.28**
_cons	**30.80**	**31.24**	**40.90**	**36.64**	**35.87**	**35.75**	**36.73**	**33.76**	**48.72**	**37.66**
N =	988	975	985	929	977	1031	1016	990	1004	985
R2 =	0.35	0.26	0.33	0.33	0.29	0.28	0.34	0.35	0.36	0.31
ICC	22%	37%	28%	29%	28%	25%	31%	27%	24%	27%
sd(_cons)	21.86	30.02	24.78	25.29	25.10	23.84	26.49	23.91	23.25	24.77

TABLE 5:
MAY & AUG 2013 REGRESSION OF PRIVACY FACTORS FOR **TRACKING** VIGNETTES BY CONTEXT
KEY: **Bold**: p < 0.05; Grey: p<0.10; Blank: p > 0.10

	TRACKING VIGNETTES BY CONTEXT					May-13		
	Retail	**Banking**	**Social**	**Games**	**Music**	**Productivity**	**Weather**	**Navigation**
Familiarity with site								
tenure						**1.43**		
NewFrequen	**1.76**					-1.82		**1.77**
Information								
AccelInfo								
ContactInfo	**-20.03**	**-12.05**	**-19.85**	**-19.54**	**-23.64**	**-20.06**	**-27.05**	**-21.98**
KeywordInfo								
FriendsInfo					**-19.28**	-9.62	**-10.82**	
ImageInfo	**-22.43**	**-13.20**	**-21.37**	**-31.08**	**-32.52**	**-26.83**	**-28.73**	**-24.06**
LocationInfo					**-11.16**			
NameInfo			-7.33		-7.72		**-11.14**	
(null = Demo)								
Collecting Actor								
ThirdPartyC	**-9.78**		**-14.83**	**-9.49**	**-16.32**	**-11.51**	**-9.97**	**-5.77**
PlatformCol								5.18
WirelessColl					-5.93			5.62
(null = Primary								
Personalization								
DevicePerso								-4.03
(null = Null								
Second Use								
DataExchan	**-46.69**	**-41.36**	**-48.11**	**-46.57**	**-44.22**	**-46.13**	**-41.78**	**-44.70**
SocalAd2nd	**-21.81**	**-17.25**	**-24.36**	**-22.15**	**-23.88**	**-22.15**	**-19.51**	**-23.33**
(null = Retarget								
Stoarge Months								
	-0.63		-0.63				**-0.86**	
Control Variables								
Male								
Age			**-0.62**			**-0.62**	-0.64	
AgeOver30								
TrustApps	**0.19**	**0.13**	**0.13**	**0.21**	**0.15**	**0.19**	**0.15**	**0.21**
PrivacyImpo	**-0.33**	**-0.25**	**-0.21**	**-0.27**	**-0.33**	**-0.28**	**-0.29**	**-0.34**
_cons	24.37	2.24	**25.71**	**24.58**	**38.12**	**27.95**	**40.62**	17.77
N	1091	1097	1074	1054	1125	1096	1117	1124
R2	0.2558	0.153	0.24	0.2221	0.2603	0.2606	0.2181	0.2533
ICC	39%	43%	39%	34%	35%	37%	38%	40%
sd(_cons)	29.51	30.38	27.41	27.69	27.11	28.01	29.62	29.43

TRACKING VIGNETTES BY CONTEXT AUGUST 2013

	Retail	Banking	Social	Games	Music	Productivity	Weather	Navigation	Activity	Symptoms
Familiarity with site										
tenure		**1.38**			**1.21**					
NewFrequency						**2.23**				
Information										
AccelInfo	**-15.85**		**-11.19**		**-17.75**		-8.39	16.31	-14.48	
ContactInfo	**-27.24**	**-22.02**	**-28.48**	**-34.55**	**-39.23**	**-32.40**	**-35.38**	**-22.22**	**-29.12**	**-21.41**
KeywordInfo		**-11.31**								
FriendsInfo	-9.97	**-18.43**			**-14.84**		**-20.47**		-9.63	
ImageInfo	**-49.47**	**-40.03**	**-41.17**	**-36.36**	**-45.44**	**-36.31**	**-43.37**	**-21.31**	**-32.05**	**-19.54**
LocationInfo	**-15.77**		**-13.17**	-10.24	**-15.52**		**-12.48**		**-20.53**	
NameInfo	**-28.56**	**-17.95**	**-15.60**	**-18.28**	**-19.37**	-10.37	**-11.98**		**-18.73**	
(null = Demo)										
Collecting Actor										
ThirdPartyCollect	**-11.42**							**-8.16**	**-8.23**	
PlatformCollect										
WirelessCollect				**-11.68**					-6.80	
(null = Primary										
Personalization										
DevicePersonal					-4.80					4.53
(null = Null										
Second Use										
DataExchange2nd	**-39.41**	**-39.48**	**-44.55**	**-47.21**	**-43.60**	**-45.94**	**-43.92**	**-47.80**	**-44.33**	**-42.02**
SocalAd2nd	**-16.93**	**-18.03**	**-17.69**	**-18.85**	**-17.85**	**-21.19**	**-17.51**	**-25.93**	**-22.68**	**-21.63**
(null = Retarget										
Stoarge Months										
	-0.95		**-0.88**	-0.81	**-0.99**		**-0.84**			
Control Variables										
Male										
Age		-0.59								
AgeOver30										
TrustApps	**0.22**	**0.23**	**0.28**	**0.28**	**0.31**	**0.28**	**0.33**	**0.27**	**0.22**	**0.28**
PrivacyImportant	**-0.33**	**-0.33**	**-0.35**		**-0.27**		**-0.26**	**-0.19**	**-0.33**	**-0.19**
_cons	**36.68**	**28.66**	**25.68**	21.17	22.93	17.80	**30.52**	18.12	**26.92**	0.96
N	951	937	951	893	1022	864	987	928	932	935
R2	0.2695	0.2635	0.2888	0.2319	0.2775	0.262	0.2973	0.2765	0.2236	0.2479
ICC	33%	41%	31%	35%	36%	33%	35%	31%	34%	33%
sd(_cons)	26.54	29.42	25.03	29.42	29.33	27.20	27.88	25.94	28.25	27.12

Bibliography

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

App Store Downloads on iTunes. (n.d.). Retrieved July 14, 2013, from https://itunes.apple.com/us/genre/ios/id36?mt=8

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. doi:10.3758/s13428-011-0081-0

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. doi:10.1093/pan/mpr057

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

Columbus, L. (2013, June 9). 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/

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

Lease, M., Hullman, J., Bigham, J., Bernstein, M., Kim, J., Lasecki, W., … Miller, R. (n.d.). *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

Meeker, M. (2012). Internet Trends. Presented at the D10 Conference, Kleiner Perkins Caufield Byers. Retrieved from http://kpcb.com/insights/2012-internet-trends

Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, *79*(1), 119–158.

Nissenbaum, H. (2009). Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford, CA: Stanford Law Books.

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, NY, USA: ACM. doi:10.1145/1753846.1753873

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

Smith, A. (2012). *Cell Internet Use 2012*. Washington, D.C.: Pew Internet & American Life Project. Retrieved from http://www.pewinternet.org/Reports/2012/Cell-Internet-Use-2012.aspx

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

Strickling, L. (2012, June 15). Putting the Consumer Privacy Bill of Rights into Practice. *National Telecommunications & Information Administration Blog*. Retrieved from http://www.ntia.doc.gov/blog/2012/putting-consumer-privacy-bill-rights-practice

Wallander, L. (2009). 25 years of factorial surveys in sociology: A review. *Social Science Research*, *38*(3), 505–520. doi:10.1016/j.ssresearch.2009.03.004

Appendix

As of July 2013, iTunes ("App Store Downloads on iTunes," n.d.) 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," n.d.) 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

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	play ….a game …. Game
	Weather	Look up the forecast with …..a weather….weather
	Social networking	Check updates on…a social networking…social networking
	Navigation	Get direction on…a map …. map
	Search (Reference)	
	Music	Listen to music on…a music…music
	Banking/Finance	Check your balance on…a banking…banking
	Shopping/Retail	Shop on ….. a retail …. Retail
	Productivity	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.	Hours per week *(continuous)*	Very frequently…frequently…occasionally…infrequently …rarely…
Information. The type of information received or tracked by the primary organization.	Location	your location when you accessed the application
	Accelerometer	how quickly you are moving
	Demographic	your age and gender
	Contacts	your phone contact list
	Keywords	what you did in the current application
	Name	your name
	Images	pictures taken with your phone
	Friends	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.

I. Pilot I – Targeting Advertisements

Factor		**Dimensions**		**In Vignette**
AdType. What the organization does with the information.		Primary Org Ad		Another application they sell
		3rd Party Ad		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.

Sample 2:

While using your phone, you update your to-do list on a productivity application that you have used rarely for 3 months.

The productivity app shows you an advertisement for another company's mobile app based on what you did in the application.

Sample 3:

While using your phone, you check updates on a social networking application that you have used occasionally for 6 months.

The social networking app shows you an advertisement for another company's mobile app based on your phone contact list.

Sample 4:

While using your phone, you get directions on a map application that you have used rarely for 6 months.

The map app shows you an advertisement for another company's mobile app based on your age and gender.

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	the mobile application… app
	Wireless provider	your wireless provider…phone company
	Platform provider	the app store company….app store
	3rd party tracking	an outside company's invisible tracking program …tracking company
Second Use. What the collecting organization does with the information	Retargeting	uses the information for future ads when you are using this app
	Data exchange	sells the data in an online auction
	Social advertising	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 scheduing 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.

Sample 2:

While on your phone, you play a game application that you have used occasionally for less than a month.

Through the game app, what you did in the current application is collected by The mobile application and will be stored for less than a week. The data collected also includes a unique identifier for your mobile device.

The mobile application then uses the information to show future ads to your friends and contacts.

Sample 3:

While on your phone, you play a game application that you have used occasionally for 7 months.

Through the game app, the pictures taken with your phone is collected by The app store company and will be stored for 15 months.

The app store company then sells the data in an online auction.