Short Paper: "Here I am, now pay me!": Privacy Concerns in Incentivised Location-sharing Systems

Luke Hutton
School of Computer Science
University of St Andrews
St Andrews, Fife, UK
lh49@st-andrews.ac.uk

Tristan Henderson
School of Computer Science
University of St Andrews
St Andrews, Fife, UK
tnhh@st-andrews.ac.uk

Apu Kapadia
School of Informatics and Computing
Indiana University
Bloomington, IN, USA
kapadia@indiana.edu

ABSTRACT

Social network sites, location-sharing services and, more recently, applications enabling the quantified self, mean that people are generating and sharing more data than ever before. It is important to understand the potential privacy impacts when such personal data are commercialised, to ensure that expectations of privacy are preserved.

This paper presents the first user study of incentivised location sharing, where people are given a direct monetary incentive to share their location with a business or their social network. We use Nissenbaum's framework of contextual integrity in a preliminary user study (n=22) to investigate potential privacy risks with such services. We find that monetisation changes why people share their data, but not the frequency of disclosures. Our results motivate further study and are useful for designers of location-sharing systems and researchers who wish to leverage the diverse range of personal data that are available in a privacy-sensitive manner.

Categories and Subject Descriptors

H.1.2 [**Information Systems**]: User/Machine Systems

Keywords

privacy; contextual integrity; location-sharing

1. INTRODUCTION

The rise of smartphones and mobile sensing smart devices is enabling the quantified self,[1] whereby vast amounts of personal data are collected or generated, and optionally shared with other people, services, and businesses. Such self-tracking is the logical extension of context-sharing applications such as Foursquare; as new sensors become available, new devices have exploited these and new applications have arisen to enable data collection and sharing. At the same time, operators of services using these data have aimed to develop sustainable business models, leading to increased commercialisation of people's data. For instance, advertisers now use location-sharing services to reach highly-targeted groups of people, who receive incentives such as cash or discounts, in exchange for promoting the business to their social networks, and releasing some personal and demographic information to the advertisers.

While previous work has identified the economics of privacy decisions in abstract scenarios [6], or suggested mechanisms for exchanging data based on people's values [2], it is also necessary to study what happens when incentives are explicitly introduced into a service's business model. It is well-known that people who use location-sharing applications are concerned about privacy [18], but how does the introduction of incentives affect these privacy concerns, or people's uses of such services? Do people's decisions change for the worse as a result of incentives, and how, or indeed should, we improve this? In this paper we look at what we term *incentivised location-sharing* (ILS) services, to determine whether they may constitute a risk to individual privacy. While the use of location-based services has been well studied, this is the first user study to examine the potential risks in embedding financial incentives in traditionally social-oriented online sharing.

To determine whether incentives may create new privacy violations, we use Nissenbaum's model of contextual integrity [16]. As the ILS context is not well-studied, we conduct a pilot user study in which 22 people use an ILS application for one week, and receive financial incentives to share their location with businesses and their social network. The user study allows us to better understand the expectations of people using such an application, their behaviour and motivations for disclosing their location for a financial incentive, and how the design of the application affects how people use it. Our early findings can help application developers to deliver incentives for disclosures in a way which preserves people's comfort and privacy, while delivering benefits to advertisers and developers. Our findings also motivate a set of research questions for further work.

2. BACKGROUND AND RELATED WORK

So-called "social-driven" location-sharing services, that allow people to share their location with their social network using a smart device, have existed for some time [11, 23]. More recently, systems have emerged which adopt the fundamental principles of these location-sharing services, but in addition deliver incentives to people for disclosing their location. We refer to such systems as *incentivised location-sharing* services, and these range from companies offering geo-fenced discounts to nearby users [15]; existing location-based services offering location-specific discounts for users who check in, such as Foursquare offering discounts to the

[1] http://quantifiedself.com/

WiSec'14, July 23–25, 2014, Oxford, UK.

http://dx.doi.org/10.1145/2627393.2627416.

"mayor" of Starbucks [22] and RadioShack [19]; to advertising-specific mobile applications such as Quidco.

To study these ILS services, we employ the contextual integrity framework. Contextual integrity has been widely used to study privacy in various services [3, 5, 9, 12, 21] but ours is the first study to use contextual integrity for location-sharing and advertising, and the first to conduct a user study to collect empirical data on these services. Privacy concerns and behaviour in location-sharing services have been extensively researched, with user studies [1, 17, 24] evincing the social norms and practices that govern the expectations of users, and their motivations for using such services.

Similarly, other work has examined online advertising [27] or location-based advertising [26]. In particular Kelley et al. study the sharing of location with advertisers [10] with a user study of 27 participants, although this is not the ILS form of advertising that we study here. Other work has also examined incentives and location-sharing, e.g., Cramer et al. identify emergent norms in the use of Foursquare, with location disclosures motivated by a desire to share interesting events as an impression management technique, and sometimes to endorse local businesses [4]. More recently, Patil et al. find increased rates of incentivised disclosures, with many people disclosing their location to receive rewards such as discounts or coupons [18].

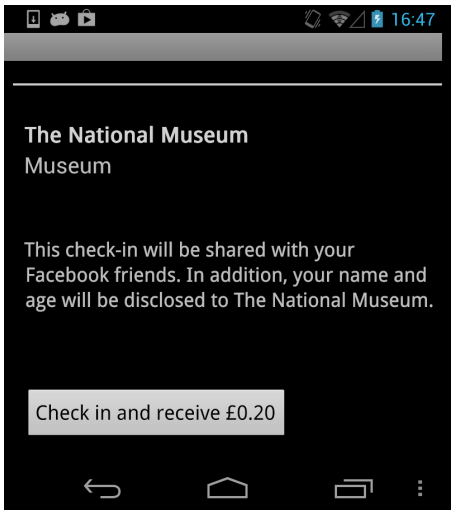

Figure 1: Screenshot of the incentivised location sharing application created for our user study. One group of participants are shown information about the flow of PII before confirming a check-in.

3. CONTEXTUAL INTEGRITY

Nissenbaum proposes contextual integrity as a theoretical framework for describing privacy in terms of information flows, arguing that information is not inherently public or private, but privacy is violated when norms governing the appropriate flow of information are not respected [16]. These context-specific "informational norms" govern all aspects of life in which information flows between actors. Individuals, groups and sections of society have expectations about how their information is collected, processed, and transmitted to other parties, and they perceive their privacy to be violated when these expectations are not met.

For example, while many would consider medical data to be "private", they might accept that it is appropriate to disclose information about medical conditions to a doctor to receive a diagnosis and treatment. If that doctor were to then gossip about their conditions to their friends, this would be a clear violation of privacy. It is the *appropriateness* of the flows of information in different contexts that determines privacy, not the information itself.

Our specific focus is on the new privacy concerns and violations created by the introduction of incentives to location-sharing and mobile advertising applications. These violations might arise because people may be incentivised to make decisions about sharing sensitive data that they might otherwise be unwilling to do so. If so, one way to improve systems is to aid this decision-making process, and previous work [20, 25] has indeed indicated that providing feedback to users can aid in decisions about sharing data.

We apply contextual integrity to ILS by conducting a user study to capture the expectations of users before and after using a prototype ILS service, and measure their motivations for disclosing their location when financially incentivised. Assessing these factors, and the role of feedback, allows us to conduct a preliminary analysis about the potential risks of such services, identifying questions for further research.

4. METHOD

To apply Nissenbaum's framework of contextual integrity, we conducted a week-long pilot user study with 22 participants, to identify the prevailing norms and expectations necessary to determine whether ILS constitutes a potential risk to privacy. We chose to run the study for seven days based on recommendations from the experience sampling literature [7] and from running such studies in the past.

For this user study, we developed an application for Android smartphones which closely resembled the interface and feature-set of existing commercial applications such as Quidco and Foursquare. The application consisted of a widget which used the Google Places API to periodically update and display the names of businesses close to the participant's current location. From the widget, the participant could select a nearby business and check in, thus creating a Facebook status update, in exchange for a small financial incentive. At the start of the study, participants chose six of their Facebook friends who would be able to view these stories, representing a cross-section of close friends, acquaintances, and colleagues. By choosing a small subset of people to share locations with, we mitigate potential adverse effects of the study, considering our interest in potentially inappropriate disclosures, while still making participants consider the social impact of their disclosures to a diverse audience. Participants could pause the application for a short period of time if they did not want location data to be collected.

4.1 Study design

Participants were randomly assigned to one of three conditions, which affected the feedback displayed to participants immediately before they checked in, and the value of the cash incentive:

- Low incentive, no feedback (**LowNo**): Participants received £0.10 for each check-in, and were not actively reminded that PII would be disclosed to the business.
- High incentive, no feedback (**HighNo**): Participants received £0.20 for each check-in, and were not actively reminded that PII would be disclosed to the business.
- High incentive, feedback (**HighYes**): Participants received £0.20 for each check-in, and were reminded that their name and age would be disclosed to the business.

Participants in the feedback condition were shown the information depicted in Fig. 1, while other participants were only informed that their check-ins would be disclosed to their Facebook friends before checking in. These incentive levels were chosen based on the distribution of incentives we find in commercial applications such as Quidco. The high incentive level was set at £0.20 because we were interested in seeing whether differences would manifest even between marginally different levels of micro-incentives, and to avoid compelling lower-income participants to check in out of financial need, which we are not investigating in this study. Participants were not aware that there were other conditions, nor how the incentives were chosen.

Before joining the study, all participants were asked to read the application's privacy policy, which specified that the business indicated would receive some PII in return for the financial incentive. Before beginning to use the application, participants completed a pre-briefing questionnaire, consisting of 15 questions drawn from the 'collection', 'control', 'awareness', and 'secondary use' dimensions of the Internet Users' Information Privacy Concerns (IUIPC) scale [13]. To these we added a question identifying expectations in the ILS context. Immediately after completing the study, participants were asked to complete the same questionnaire, allowing us to identify a relationship between different feedback conditions and a change in privacy attitudes.

In addition to recording the participant's activity during their seven days of participation, participants received an automatically-generated end-of-day questionnaire each night, based on their activity during the preceding day. This allowed us to capture qualitative data about the motivations for activity within the application, and to clean anomalous outliers (such as accidental interface taps) within hours of the activity occurring.

4.2 Recruitment

Participants were recruited through advertisements on Facebook and mailing lists aimed at university students and staff. Participants were not screened, with the only requirement being possession of an Android smartphone and a Facebook account. 39 participants were recruited in total, of whom 22 completed, and 17 prematurely left the study. The majority of those who did not complete the study installed the application to their mobile device but did not complete the registration and consent process. 9 participants were in the LowNo condition, 6 in the HighNo condition, and 7 in the HighYes condition. To prevent the study being affected by cultural differences in privacy expectations, we recruited all participants from the United Kingdom. Recruitment for the experiment positioned our system as a new commercial application, to closely align participant expectations with that of existing commercial applications.

4.3 Remuneration

Participants were told they would earn money for sharing their locations. Rather than provide the exact amount promised by the application, all participants were given an Amazon voucher of equal value at the end of the study, surpassing the value any participant accrued during normal use of the application. This strategy was employed due to ethical concerns about financially rewarding some participants more than others.

4.4 Ethical considerations

Data collection used our framework for privacy-sensitive handling of social network data [8]. This ensured that the proportionate set of data necessary to execute the study was collected, and that personal Facebook data were sanitised and deleted at appropriate points during the lifecycle of the study.

All participants were informed of the study's deception after the experiment closed in an email providing their remuneration and explaining the motivation for the study.

The experimental design and all recruitment materials were approved by the relevant ethics committee (our institution's equivalent of an Institutional Review Board).

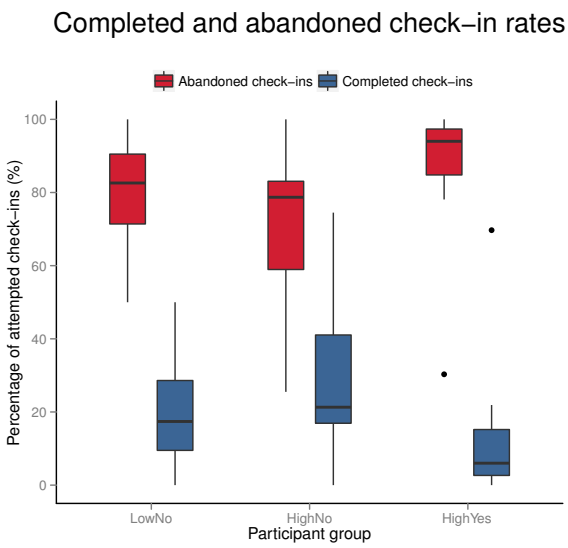

Figure 2: Participants who receive higher incentives (the HighNo and HighYes groups) check in more often, unless more feedback about PII flows is provided.

5. RESULTS

In total, our 22 participants completed 212 check-ins, and abandoned 471, and the most active user checked in 15 times in one day. We first examine the overall differences between our groups of participants to understand the effects of our feedback and incentive conditions. Our sample size is too small to make confident statistical inferences, but our findings motivate further work to investigate the specific research questions our results identify.

5.1 Less feedback induces greater sharing

Fig. 2 shows the proportion of completed check-ins and abandoned check-ins. We found that participants in the HighNo condition exhibited the most variable behaviour. A number of participants performed more check-ins than lower-paid participants, but this was not consistent across the group. We did, however, note a reduction in disclosure rates for participants in the HighYes condition, where most users only completed less than 10% of check-ins. Behaviour was the most consistent within this group, and the higher variance among non-feedback groups indicates that the absence of such feedback generally induced more sharing. Those who received more feedback abandoned more relative to the number of completed check-ins. We note that higher incentives without feedback leads to greater variance in check-in rates, indicating an influence on behaviour. Some participants in all conditions did not complete any check-ins throughout the study, despite continuing to interact with the application.

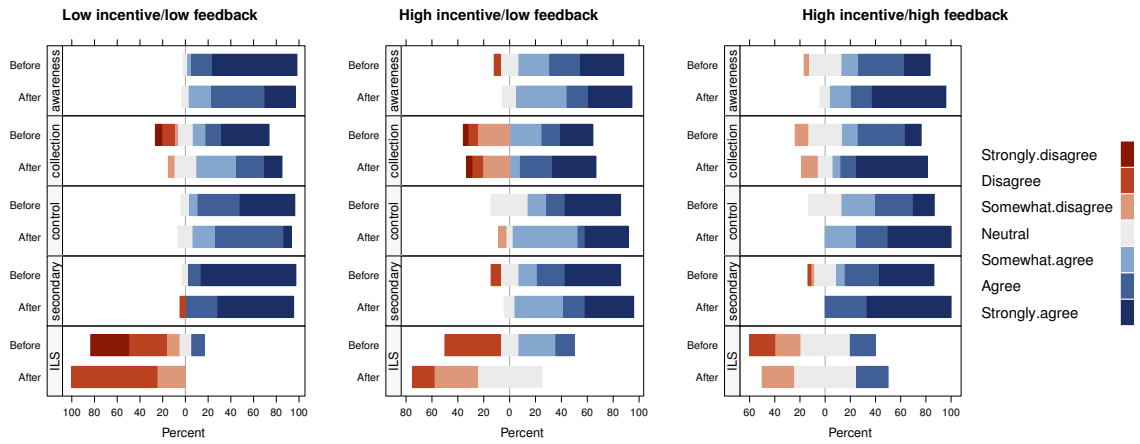

Figure 3: Diverging stacked bar chart showing how participants responded to a questionnaire about online privacy concerns before and after participating in the study. For each dimension tested, positive answers indicate increased concern or awareness about relevant privacy practices. Participants reported high concern and awareness about online privacy issues generally, and these were reinforced when more feedback about PII flows was provided by the application. Many believed, however, that companies are "entitled" to personal information in exchange for money, and comfort with this practice increased with use of a mobile advertising check-in application, as shown by the positive tendency on the "ILS" dimension.

5.2 Feedback engenders support for ILS

Fig. 3 shows the results of the IUIPC survey before and after the study. Participants exposed to more feedback about the flow of PII during the study reported greater agreement on most dimensions in our post-brief survey, particularly 'secondary use' (concern about information being used for reasons not originally sanctioned), 'control' (loss of control leads to privacy violation), and 'awareness' (of privacy practices). Participants in low feedback conditions did not significantly alter their responses in the debrief questionnaire.

We found that participants who received less feedback were less comfortable with ILS at the end of the study, although we saw increased comfort with ILS services for those in the higher feedback conditions, represented by the greater positive tendency for this question in Fig. 3.

Participants in both high incentive conditions reported greater concern about secondary use of their PII, a fundamental aspect of ILS, while lower-paid participants' concern did not change as much. This suggests neither incentives nor feedback significantly impact people's concern about the use of their information.

Our examination of attitudes before and after the study reveals that concern and awareness of privacy issues generally increases through participating in the study. Interestingly, only participants who were provided feedback about the flows of their information believed that ILS was a more legitimate practice than at the start of the study, while confidence in ILS dropped for our other participants. This mirrors our finding that high feedback participants are much more comfortable with the disclosure of their PII. We believe this can be attributed to a combination of such participants feeling more empowered by the transparent explanation of how their information is used, while other participants, without the same *in situ* assurances, may have exhibited a priming effect from the study being bookended by our IUIPC questionnaire. Participants did not frequently report such concerns in our end-of-day questionnaires, lending further support to this theory.

These results give us some insight into the expectations of people before adopting an ILS service, and their values after having experienced such a system. We observe paradoxical results, with participants' general privacy concerns hardened, but their attitudes towards ILS more relaxed. Contextual integrity suggests if a new process perturbs the values of an existing context, there is a risk of privacy breaches. This appears to manifest in our results, as people are reporting greater concern, yet appear to be placated by the introduction of a financial incentive. Failure to reconcile these behaviours risks people feeling compelled to make disclosures they might otherwise consider inappropriate. Further work is needed to assess the significance of this effect.

5.3 Higher incentives change check-in motivations

Disclosures in traditional location-sharing applications are often motivated by attempts to build social capital, and this was a recurring theme in our study too. When asked in end-of-day questionnaires to explain why they checked in to certain locations, participants in all conditions frequently reported "wanting my friends to see that I was there". This motivation did not often coincide with those directly pertaining to the introduction of incentives, such as "I don't mind sharing my personal information in exchange for cash" and "wanting to promote the business", which were commonly reported independently. Interestingly, in 69.2% of cases where social capital was a motivation, it did not coincide with any other motivations, suggesting people may be motivated on two largely independent fronts — appealing to their social network, and promoting businesses for money.

Participants who were offered a higher incentive were more likely to intentionally promote local businesses to their social network. We also find that when participants cited promotion as a motivation, it coincided with no other motivation in 57.1% of cases, suggesting that participants treated these two use cases somewhat independently. Lower-paid participants did not often exhibit this motivation, suggesting they may have not considered the value of the incentive sufficient to deliberately act as an advertising agent for the business, even if the actual exposure of their check-ins was the same.

Participants who received less feedback cited social capital-building motivations most often when disclosing their location, suggesting

they treated the service in the same manner as other location-sharing services. Similarly, the lowest-paid participants were unlikely to cite financial or promotional motivations, suggesting they also considered the application to be much like any other location-sharing system, despite the additional PII flows our system introduces.

6. IMPLICATIONS

6.1 People value incentives over privacy

Despite high levels of online privacy concern, our results suggest many people may be comfortable with the notion of disclosing their location and personally identifiable information for a cash incentive. When our participants were exposed to feedback about the flows of their personal information, their overall privacy concern increases, but they become even more comfortable with incentivised location sharing information flows, believing companies are "entitled" to their personal information if they are paid. This result is consistent with previous findings that people generally value money over their PII [6], and has implications for further study of ILS. We find that location disclosures were instigated by a mix of social and financial motivations which often did not coincide, which may suggest people are attempting to reconcile the context of a traditional location-sharing service with ILS. This is cause for concern as, in our system, all location disclosures reached the same audience, leading to potentially inappropriate disclosures to one's social network. Application designers can address this issue by avoiding the conflation of social and incentivised disclosures through distinct interfaces for each, and allowing people to choose different audiences for alternative types of disclosures.

6.2 Feedback does not discourage sharing

Participants who received more feedback about the flow of their personal information were more comfortable with the practice of ILS at the end of the study, and only made slightly fewer location disclosures than other participants. This is an important finding for application designers to note, as it contradicts any intuition that full disclosure about how people's information is used might dissuade users. Rather, our results suggest people may be empowered by understanding how their information is used. While they are more discerning about when to share their information, they are also the most confident that they are aware of the privacy implications of using such services, and that the disclosure of PII is an appropriate outcome. Furthermore, the changes in motivations among our participants should be noted by designers, who ought to design services which satisfy both the social and financial reasons for their use. Where insufficient feedback was provided, we were concerned by the conflated motivations for sharing one's location, as participants struggled to reconcile the distinct use cases. Among participants who understood how their personal information was used, however, distinct social and financial thought processes were observed. Designers should provide *in situ* disclosures of how personal information is used, as our results suggests this satisfies people's privacy concerns, without severely affecting their willingness to use such services. In addition, incentivised and non-incentivised disclosures should be represented distinctly, to ensure people's motivations for making disclosures are aligned with the exposure of their information.

6.3 Potential privacy risks

Based on this preliminary study, there is evidence of potential risks to users of ILS services.

Our concern is that the prominence of 'traditional' location-sharing motivations such as social capital building and impression management, particularly among participants who received less feedback, suggests people may treat such applications in the same manner as any other social-driven location-sharing system, despite sensitive information being disclosed to advertisers, and the possibility that their social network will perceive incentivised check-ins to be of lower value [4]. When participants received clear feedback about the use of their personal information at the point of disclosure, they make slightly fewer disclosures, and while disclosures are still often socially motivated, they often constitute a deliberate effort to advertise that business to their social network. Current commercial applications often do not deliver this level of feedback, burying information about the flows of personal information within unread privacy policies [14], and we argue that if such feedback is provided, people are able to make more informed decisions about when and why their location will be disclosed. Participants who received more feedback also report slightly better awareness of online privacy issues.

We do not suggest that the introduction of new motivations for disclosure themselves constitute a breach of contextual integrity. In our user study, many participants managed the social and promotional aspects of the ILS context independent of each other. The relationship between the feedback within the application, and people's behaviour and wider attitudes towards online privacy, suggests the design of such an application can have a significant impact on people's relationship with technology. Participants who received more feedback about the flow of PII were the most comfortable with the practice of ILS.

Our results highlight interesting differences in how feedback and incentives affect behaviours and motivations in our application, but this pilot study has some limitations. Our study is limited to a small number of participants and levels of incentives, and indeed demographic information has not been studied. Future work should examine in more detail the effects of these variables over a longer period with a greater sample size. Using these results, we propose the following research questions, to be addressed in future work:

1. Do people who use location-sharing services have different expectations of privacy when their disclosures are financially motivated?
2. Do incentives perturb privacy norms in a location-sharing service to the extent that contextual integrity is violated?
3. Does greater transparency about the flow of personal information when people share their location for money affect behaviour, and reduce the risk of privacy violations?

7. CONCLUSION

This paper has presented the first study of what we term incentivised location-sharing — the introduction of financial incentives to traditional location-sharing and mobile advertising systems. To understand privacy violations in these systems, we employ Nissenbaum's contextual integrity framework and conduct a user study with 22 smartphone users.

Our results show that while monetisation does not change the frequency of location disclosures, people's motivations for sharing their location are altered. Also, people's concern and awareness of privacy issues increased during the study. We show that application designers can build applications to show people how their information is used in such a way to not dissuade people from using such services, while increasing their confidence in the practice. In future work we plan additional user studies to better understand the wider implications of these findings.

8. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under grants CNS-1016603 and CNS-1252697, and an Engineering and Physical Sciences Research Council (EPSRC) Doctoral Training Grant.

9. REFERENCES

[1] D. Anthony, T. Henderson, and D. Kotz. Privacy in location-aware computing environments. *IEEE Pervasive Computing*, 6(4):64–72, Oct. 2007. doi:10.1109/mprv.2007.83.

[2] C. Aperjis and B. A. Huberman. A market for unbiased private data: Paying individuals according to their privacy attitudes. *First Monday*, 17(5-7), May 2012. doi:10.5210/fm.v17i5.4013.

[3] L. Barkhuus. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proc. CHI*, pages 367–376, May 2012. doi:10.1145/2207676.2207727.

[4] H. Cramer, M. Rost, and L. E. Holmquist. Performing a check-in: emerging practices, norms and 'conflicts' in location-sharing using foursquare. In *Proc. MobileHCI*, pages 57–66, Aug. 2011. doi:10.1145/2037373.2037384.

[5] F. S. Grodzinsky and H. T. Tavani. Privacy in "the cloud": applying Nissenbaum's theory of contextual integrity. *ACM SIGCAS Computers and Society*, 41(1):38–47, Oct. 2011. doi:10.1145/2095266.2095270.

[6] J. Grossklags and A. Acquisti. When 25 Cents is too much: An Experiment on Willingness-To-Sell and Willingness-To-Protect Personal Information. In *Proc. WEIS*, 2007. Online at http://weis2007.econinfosec.org/papers/66.pdf.

[7] J. M. Hektner, J. A. Schmidt, and M. Csikszentmihalyi. *Experience sampling method: measuring the quality of everyday life*. SAGE Publications, Thousand Oaks, CA, USA, 2007.

[8] L. Hutton and T. Henderson. An architecture for ethical and privacy-sensitive social network experiments. *SIGMETRICS Perform. Eval. Rev.*, 40(4):90–95, Apr. 2013. doi:10.1145/2479942.2479954.

[9] E. A. Jones and J. W. Janes. Anonymity in a world of digital books: Google Books, privacy, and the freedom to read. *Policy & Internet*, 2(4):42–74, Jan. 2010. doi:10.2202/1944-2866.1072.

[10] P. G. Kelley, M. Benisch, L. F. Cranor, and N. Sadeh. When are users comfortable sharing locations with advertisers? In *Proc. CHI*, pages 2449–2452, May 2011. doi:10.1145/1978942.1979299.

[11] J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman. I'm the mayor of my house: examining why people use foursquare - a social-driven location sharing application. In *Proc. CHI*, pages 2409–2418, May 2011. doi:10.1145/1978942.1979295.

[12] H. R. Lipford, G. Hull, C. Latulipe, A. Besmer, and J. Watson. Visible flows: Contextual integrity and the design of privacy mechanisms on social network sites. In *Proc. CSE*, volume 4, pages 985–989, Aug. 2009. doi:10.1109/cse.2009.241.

[13] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, Dec. 2004. doi:10.1287/isre.1040.0032.

[14] A. M. McDonald and L. F. Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3):540–565, 2008. Online at http://www.is-journal.org/files/2012/02/Cranor_Formatted_Final.pdf.

[15] C. C. Miller. Take a step closer for an invitation to shop. *New York Times*, page B4, Feb. 23 2010. Online at http://www.nytimes.com/2010/02/23/business/media/23adco.html.

[16] H. Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books, Stanford, CA, USA, 2009.

[17] X. Page, B. P. Knijnenburg, and A. Kobsa. What a tangled web we weave: lying backfires in location-sharing social media. In *Proc. CSCW*, pages 273–284, 2013. doi:10.1145/2441776.2441808.

[18] S. Patil, G. Norcie, A. Kapadia, and A. J. Lee. Reasons, rewards, regrets: privacy considerations in location sharing as an interactive practice. In *Proc. Eighth Symposium on Usable Privacy and Security*, July 2012. doi:10.1145/2335356.2335363.

[19] M. Rost, L. Barkhuus, H. Cramer, and B. Brown. Representation and Communication: Challenges in Interpreting Large Social Media Datasets. In *Proc. CSCW*, pages 357–362, 2013. doi:10.1145/2441776.2441817.

[20] R. Schlegel, A. Kapadia, and A. J. Lee. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proc. SOUPS*, 2011. doi:10.1145/2078827.2078846.

[21] P. Shi, H. Xu, and Y. Chen. Using contextual integrity to examine interpersonal information boundary on social network sites. In *Proc. CHI*, pages 35–38, 2013. doi:10.1145/2470654.2470660.

[22] E. D. Spiegler, C. Hildebrand, and F. Michahelles. Social networks in pervasive advertising and shopping. In J. Müller, F. Alt, and D. Michelis, editors, *Pervasive Advertising*, chapter 10, pages 207–225. Springer, London, UK, 2011. doi:10.1007/978-0-85729-352-7_10.

[23] K. P. Tang, J. Lin, J. I. Hong, D. P. Siewiorek, and N. Sadeh. Rethinking location sharing: exploring the implications of social-driven vs. purpose-driven location sharing. In *Proc. Ubicomp*, pages 85–94, Sept. 2010. doi:10.1145/1864349.1864363.

[24] E. Toch, J. Cranshaw, P. H. Drielsma, J. Y. Tsai, P. G. Kelley, J. Springfield, L. Cranor, J. Hong, and N. Sadeh. Empirical models of privacy in location sharing. In *Proc. 12th ACM international conference on Ubiquitous computing*, pages 129–138, Sept. 2010. doi:10.1145/1864349.1864364.

[25] J. Y. Tsai, P. Kelley, P. Drielsma, L. F. Cranor, J. Hong, and N. Sadeh. Who's viewed you?: The impact of feedback in a mobile location-sharing application. In *Proc. CHI*, pages 2003–2012, Apr. 2009. doi:10.1145/1518701.1519005.

[26] R. Unni and R. Harmon. Perceived effectiveness of push vs. pull mobile location based advertising. *Journal of Interactive Advertising*, 7(2):28–40, 2007. doi:10.1080/15252019.2007.10722129.

[27] B. Ur, P. G. Leon, L. F. Cranor, R. Shay, and Y. Wang. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proc. SOUPS*, July 2012. doi:10.1145/2335356.2335362.