I've Got 99 Problems, But Vibration Ain't One: A Survey of Smartphone Users' Concerns

Adrienne Porter Felt
Serge Egelman
David Wagner

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2012-70
http://www.eecs.berkeley.edu/Pubs/TechRpts/2012/EECS-2012-70.html

May 9, 2012

Acknowledgement

We would like to thank Coye Cheshire and Galen Panger for their insightful comments and discussion. This material is based on work supported by a Facebook Fellowship, NSF grant CCF-0424422, and the Intel Science and Technology Center for Secure Computing. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of Facebook, the National Science Foundation, or Intel.

I've Got 99 Problems, But Vibration Ain't One: A Survey of Smartphone Users' Concerns

Adrienne Porter Felt, Serge Egelman, and David Wagner
University of California, Berkeley
{apf, egelman, daw}@cs.berkeley.edu

ABSTRACT

Smartphone operating systems expose a wide range of functions and user data to third-party applications. However, past research on mobile privacy has focused exclusively on the risks pertaining to sharing location. To expand the scope of smartphone security and privacy research, we surveyed 3,115 smartphone users about 99 risks associated with 54 smartphone privileges. We asked participants to rate how upset they would be if a given risk occurred. Based on this data, we ranked risks by the number of users who stated that they would be "very upset" if the risks occurred. We then performed an open-ended survey of 41 smartphone users, which let respondents discuss the risks in their own words. The follow-up study confirmed that people find the lowest-ranked risks merely annoying but might seek legal or financial retribution for the highest-ranked risks. Our ranking could be used to guide the severity or selection of warnings on smartphone platforms. Notably, our results show that location is not a high-ranked user concern.

Categories and Subject Descriptors

D.4.6 [**Operating Systems**]: Security and Protection

General Terms

Human Factors, Security

Keywords

Smartphones, risks, warnings

1. INTRODUCTION

Smartphone operating systems like Android and iOS expose rich APIs to third-party applications. These APIs allow applications to use hardware (e.g., vibrator and camera), change phone settings, and read data (e.g., text messages and contacts). Unfortunately, malicious and unscrupulous application authors have taken advantage of these resources, to the displeasure of users [8, 10, 18].

Android warns users about dangerous application *permissions* during installation. Applications must request permissions in order to access sensitive APIs, and each permission corresponds to an install-time warning. Studies have shown that the permission warnings confuse users [9, 14]. One problem is that Android controls a large number of privileges with permissions, and the multitude of warnings is difficult for users to process and recall [9]. iOS takes the opposite approach: users are only asked to provide consent for two application permissions (reading location and sending notifications). However, iPhone users were outraged when they discovered that applications access other resources without their approval [4]. These experiences suggest that Android and iOS are not asking users about the right permissions.

In order to guide the future selection and design of smartphone warnings, we performed two surveys to rank the level of user concern about a wide range of smartphone resources. In our first survey, we asked 3,115 smartphone users to rate their level of concern about 99 risks corresponding to 54 smartphone permissions. We ranked the risks using the percentage of respondents who said they would be "very upset" if the risks occurred. Appendix B provides the full ranking. In our second survey, we asked 42 smartphone users to state their reactions to low-ranked, medium-ranked, and high-ranked risks in their own words. The open-ended responses validate the ranking: participants viewed low-ranked risks as manageable annoyances, whereas they viewed high-ranked risks as severe offenses that may require the help of their service provider, law enforcement, a lawyer, etc.

Notably, we find that warnings in Android and iOS do not correspond to users' concerns. Future permission system should consider user concerns when deciding which privileges are protected with warnings. We also find that improper location sharing is only a mid-ranked risk: users are more concerned about many other permissions. As such, mobile privacy research should be refocused to consider other resources and privileges beyond location.

Contributions. We contribute the following:

- We created a ranking of the risks of 54 smartphone application permissions based on user concerns.
- We find that respondents' concerns about data sharing depend on how the application is using the data.
- We find that location-related risks rank as the second-lowest of data types and in the bottom half of all risks.
- We collected descriptions of users' anticipated reactions if they were to encounter various risks.

2. RELATED WORK

Mobile privacy research has traditionally focused on location tracking and sharing. Numerous studies have examined users' privacy concerns about sharing mobile location data with friends, acquaintances, and advertisers (e.g., [1, 2, 3, 7, 12, 13, 15, 19]). This prior literature thoroughly establishes that users have privacy concerns about sharing location data with their social contacts and companies.

However, smartphone operating systems provide applications with the ability to access a number of resources beyond location data. Smartphone APIs let applications read many types of data (e.g., photographs) and make changes to the phone (e.g., delete data). Few studies have explored the space of smartphone privacy and security beyond location. Roesner et al. [17] studied smartphone users' privacy and security expectations for copy-and-paste, photography, and text messaging in addition to location. Muslukhov et al. [16] asked a number of smartphone users about the value and sensitivity of eleven types of data on their phones. We aim to further expand the scope of research on users' smartphone concerns by studying users' opinions about 99 risks associated with 54 smartphone permissions.

3. METHODOLOGY

We asked 3,115 smartphone users to rate how upset they would be if an application performed certain actions on their phones without user approval. We then performed a follow-up survey where 42 respondents provided open-ended explanations of their feelings about applications performing undesirable actions, which contextualizes the ratings.

3.1 Rating Survey

The purpose of the large-scale rating survey was to create an index that ranks the risks of allowing applications to access smartphone resources by degree of user concern.

3.1.1 Instrument Design and Validation

Our instrument was designed to elicit user concerns about different resources. We faced two design constraints. First, we aimed to measure opinions about risks rather than application features. (For example, a user might view an application that deletes files as useful or harmful depending on whether the deletion was intentional.) Second, we did not want to scare participants by mentioning malware or viruses. We suspected that participants would report high levels of concern for *any* action that they were told is associated with malware. As such, we needed to ensure that respondents were aware that we were asking about undesirable actions, without mentioning how or why those actions were initiated (e.g., by malware).

We performed two preliminary surveys that asked respondents about situations in which applications performed an action "without my knowledge" or "when you believed [the app] had no reason to do so." The results of these surveys were inconsistent. Subsequent interviews revealed that participants were unsure whether the listed actions were negative side-effects or positive features. We conducted one-on-one interviews and a focus group with Craigslist-recruited smartphone users to generate new wording.

We validated our final instrument by asking four smartphone users to take the survey and speak with an interviewer. These participants were selected from applicants on Craigslist to represent a diverse cross-section of smartphone users. We found that the respondents understood that the questions asked about risks rather than features, and they used the full range of the scale. When asked to describe how the scenarios in the questions could occur, all four participants listed both buggy and compromised applications. Some participants also mentioned viruses, bad UI design, or aggressive marketing. This indicated that all four participants had a firm grasp of the meaning of the questions without specifically focusing on malware.

3.1.2 Instrument

The survey began by asking respondents to think about negative side-effects of applications, rather than features:

> *Every once in a while, an app might do something on your phone without asking you first. Depending on what the app does to your phone, your feelings could range from indifference (you don't care) to being very upset.*

We then asked participants about various risks:

> *How would you feel if an app [insert risk], without asking you first?*

For example: "How would you feel if an app added new contacts, without asking you first?" Respondents answered using a horizontal five-point scale that ranged from "Indifferent" to "Very upset," with unlabeled intermediate points.

Each survey participant saw 12 questions on one page, selected at random from a set of 99 potential questions. Appendix B shows the risks that participants were asked about. We compiled the set of questions by assigning risks to Android, Windows Phone 7, and iOS permissions. The three platforms define a total of 191 permissions, but we grouped equivalent permissions (e.g., "power device on or off" and "force device reboot") and discarded irrelevant permissions (e.g., "enable application debugging") to arrive at 54 permissions. We then assigned risks to the actions using documentation and domain expertise. Some actions are associated with multiple risks, and we assigned least four risks to each type of smartphone data[1]: "publicly shared your [data type]," "shared your [data type] with your friends," "shared your [data type] with advertisers," "sent copies of [data type] to their servers (but didn't share them with anyone else)."

On the last page of the survey, we collected demographic information about participants and their cell phones.

3.1.3 Deployment and Demographics

We deployed the survey on Mechanical Turk for 13 days. Participants were paid $1 each for completing the survey, and we limited the survey to respondents in the United States. We filtered responses for validity based on users' survey completion time, responses to short open-ended questions, and the self-reported type of phone. After filtering, we obtained 3,115 valid responses from smartphone users.

Participants' ages ranged from 18 to 80 (μ=29.7), while 47.9% were female and 51.9% were male. Although the population was younger than the U.S. population overall (65% of respondents were below the age of 30), it was only slightly younger than U.S. smartphone user demographics [6]. Participants reported many occupations: healthcare workers,

[1]Due to a survey programming error, we accidentally omitted one of these risks for three types of data.

software engineers, financial advisors, federal government employees, graphic designers, etc. However, the predominant occupations were students, stay-at-home parents, and the unemployed. Completed levels of education ranged from some high school to doctorates.

Participants reported owning the following smartphones: 49.5% Android phones, 39.7% iPhones, 7.7% Blackberries, and 1.7% Windows phones. The remainder stated that they owned Palm, Symbian, or multiple phones. There was no incentive to lie about phone ownership because we paid participants regardless of whether they owned smartphones.

3.2 Open-Ended Survey

The purpose of the open-ended survey was to associate participants' own words with the large-scale survey ratings.

3.2.1 Instrument

The open-ended survey asked participants the following short essay questions about risks:

1. How would you feel if an app [insert risk], without asking you first?
2. Why would you feel that way?
3. What would you do if this happened?

Each participant was asked about three of nine risks, which we selected based on the results of the large-scale study. We chose the three lowest-ranked risks, three mid-ranked risks, and three of the highest-ranked risks:

- *Lowest-ranked risks:*
  - vibrated your phone
  - connected to a Bluetooth device (like a headset)
  - turned your flash on
- *Mid-ranked risks:*
  - added new contacts
  - took screenshots when you're using other apps
  - un-muted a phone call
- *Highest-ranked risks:*
  - deleted all of your contacts
  - sent premium text messages from your phone (they cost money)
  - made phone calls to 1-900 numbers (they cost money)

We showed each respondent one of the lowest-ranked risks, one of the mid-ranked risks, and one of the highest-ranked risks. We displayed one page for each risk, and participants could not move backwards in the survey. The last page of the survey collected demographic data.

3.2.2 Deployment and Demographics

We deployed the survey on Mechanical Turk for two days. Participants were paid $8 each. The survey was advertised as being worth $3, with an additional $5 reward for complete sentences and correct grammar. We ran the survey until we had 42 valid responses (with a target of 40) from people in the United States. The participants were evenly split by gender, with an average age of 30.3. All of the respondents said that they have used smartphone applications, with an average of 29 applications installed on their phones.

Risk	VUR Rate
permanently disabled (broke) your phone	98.21%
made phone calls to 1-900 numbers (they cost money)	97.41%
sent premium text messages from your phone (they cost money)	96.39%
deleted all of your contacts	95.89%
used your phone's radio to read your credit card in your wallet	95.15%
publicly shared your text messages	94.48%
deleted all of the information, apps, and settings on your phone	94.39%
publicly shared your e-mails	93.37%
deleted all of your other apps	93.14%
shared your text messages with your friends	92.49%
...	...
inserted extra letters into what you're typing	45.48%
read files that belong to other apps	44.33%
sent your phone's unique ID to their servers (but didn't share it with anyone else)	42.16%
added new browser bookmarks	39.22%
sent the list of apps you have installed to their servers (but didn't share it with anyone else)	34.92%
turned the sound on your phone down really low	36.96%
sent your location to their servers (but didn't share it with anyone else)	29.88%
turned your flash on	29.67%
connected to a Bluetooth device (like a headset)	27.47%
vibrated your phone	15.62%

Table 1: The highest- and lowest-ranked risks.

4. RESULTS

We used the results of the large-scale survey to rank the severity of the risks (Section 4.1). The open-ended survey provides supplementary qualitative data to add context to the large-scale survey results (Section 4.2). We then discuss how these results can be interpreted (Section 4.3).

4.1 Risk Rankings

Our goal is to rank the severity of potential risks based on users' concerns. In our large-scale survey, respondents rated how upset they would be if certain risks occurred. Our resulting metric for the severity of a risk is the percentage of respondents who indicated that they would be "very upset" if the given risk occurred. We refer to this metric as the *VUR rate* (the "very upset" respondent rate). We obtained an average of 376.7 ratings per risk.

The highest-ranked risk is "permanently disabled (broke) your phone," with a 98.2% VUR rate. The lowest-ranked risk is "vibrated your phone," with a 15.6% VUR rate. The risks that evinced the highest levels of concern involve permanent data loss or financial loss (e.g., sending premium text messages or spying on credit card numbers), and the lowest-ranked risks pertain to phone settings or sending data to servers. Table 1 shows the ten highest-ranked and ten lowest-ranked risks, and Appendix B provides the VUR rates for all of the 99 risks in our survey.

We surveyed respondents about four types of data sharing: public sharing, sharing with friends, sharing with advertisers, and removing the data from the phone without sharing it with another party. For all of the data types, publicly sharing the data is approximately twenty percentage points more concerning than sending the data to a server. Shar-

	Low-Ranked Risks				Mid-Ranked Risks				High-Ranked Risks			
	Avg	vibrate	Blue-tooth	flash	**Avg**	added contacts	screen-shots	un-muted	**Avg**	deleted contacts	$ SMS	$ calls
nothing	**21%**	29%	20%	15%	**12%**	18%	0%	15%	**5%**	0%	5%	10%
tinker with app	**33%**	50%	27%	23%	**12%**	6%	33%	0%	**0%**	0%	0%	0%
uninstall the app	**62%**	42%	73%	69%	**74%**	71%	67%	85%	**76%**	67%	80%	80%
contact developer	**12%**	7%	7%	23%	**17%**	18%	8%	23%	**40%**	25%	45%	5%
write a review	**5%**	0%	7%	8%	**14%**	12%	8%	23%	**21%**	25%	20%	20%
contact press	**5%**	0%	13%	0%	**5%**	6%	0%	8%	**5%**	0%	10%	0%
call service provider	**0%**	0%	0%	0%	**2%**	0%	8%	0%	**17%**	17%	15%	20%
replace/wipe phone	**0%**	0%	0%	0%	**5%**	0%	8%	0%	**5%**	8%	0%	10%
contact authorities	**0%**	0%	0%	0%	**2%**	6%	8%	0%	**19%**	25%	15%	20%
pursue legal action	**0%**	0%	0%	0%	**0%**	0%	0%	0%	**12%**	17%	5%	20%

Table 2: Forty-two survey respondents told us how they would react if certain risks occurred. We categorized their responses; some responses fall into multiple categories.

ing with friends and advertisers rank in the middle, between public sharing and sending the data to a server. Notably, illicit location sharing has the lowest or second-lowest VUR rates of the eleven smartphone data types in our survey.

We consider the percentage of "very upset" respondents instead of medians because the responses were not normally distributed. (Despite this, ordering the risks by medians returns a very similar ranking.) We observed different amounts of diversity of opinion between risks. Eighteen risks had a standard deviation greater than 1, whereas six had a standard deviation less than .37. In general, the risks with high VUR rates have low standard deviations, whereas the risks with low VUR rates have larger standard deviations. One interpretation is that there is user consensus about what *is* very upsetting, but not about what is *not* very upsetting. It may also be an artifact of our five-point scale: some users might have selected something stronger than "very upset" if such an option were available, which would have resulted in greater variance among high-ranked risks.

Individual respondents' scores are not directly comparable to each other because they received different questions, but we can compare groups of respondents. Women rank risks higher than men do ($\mu_M = 4.47, \mu_W = 4.55$; $p < 0.0005$, $z = -4.269$, Wilcoxon-Mann-Whitney test), although the effect size is very small ($d = 0.18$). People above the age of 50 rank risks higher than people below the age of 30 do ($\mu_{<30} = 4.46, \mu_{>50} = 4.67$; $p < 0.0005$, $z = 5.943$, Wilcoxon-Mann-Whitney test), with a medium effect size ($d = 0.51$). We do not find a significant difference between types of phones ($\chi^2 = 4.487$, $p = 0.6110$, Kruskal-Wallis test).

4.2 Open-Ended Responses

The VUR rate is a relative metric that allows us to compare risks against each other. However, the metric does not provide us with any context for how users interpret "very upset." Our open-ended survey assigns user-supplied meaning to the metric. It also serves as a second measure to evaluate whether there are differences in users' concerns across risks.

We asked participants what they would do if certain risks occurred. Table 2 displays the frequency with which participants mentioned certain reactions. Respondents' stated reactions fell in the following categories:

- **Nothing.** Some participants stated that they would ignore the risk or simply reverse the undesirable action.
- **Tinker.** Participants said that they would try to change the application's settings or determine what UI element was causing the undesirable behavior. This was often the first of multiple proposed steps. For example, *"I would first try to change the settings so that it doesn't connect. If I can't find the settings to turn such a feature off, I would immediately delete the app."*
- **Uninstallation.** The most common recourse was to uninstall the application.
- **Contact the developer.** Many people said that they would try to contact the developer of the application to complain or request a refund.
- **Reviews.** Some participants said that they would try to make others aware of the application's problems by writing negative reviews. For example, *"...for the first time ever, I would probably review [the] app. I would type (probably even in all caps!) about what it does."*
- **Contact the press.** Participants sometimes said that they would warn other users by contacting blogs or *"watchdog news groups."*
- **Contact the phone company.** Several participants said that they would contact their service provider to reverse charges or restore data. Surprisingly, many participants in this category said that they would blame their service providers for negative application behavior. For example, one respondent wrote, *"If this happened I would consult my service provider to try and retrieve my contacts, and probably cancel my service."* Another said, *"I would simply switch to another phone company if I could not uninstall the defective app."*
- **Replace or wipe the phone.** Although none of the risks in the survey were permanent side-effects, some participants said that they would get a new phone or wipe their existing phone so that it would be like having a new phone. One participant wrote, *"[If] this happened and I could not turn off this feature in the settings, then I would not continue using the phone and I would try to either get a refund or to sell it."* Another person said that he would *"smash [his] phone to bits."*
- **Contact authorities.** Some participants said that they would notify authorities about the application's misbehavior so that the application would be punished or removed from the store. For example, *"I would call up the FBI or other organizations to look into how my information might have been mishandled."*

- **Legal action.** In some cases, participants wrote that they would seek legal action against the application developer. For example, *"...I may seek legal counsel to solve the issue and perhaps receive compensation for the inconvenience and trouble that the application developer put me through."*

As Table 2 shows, participants' reactions increased in severity from the lowest-ranked risks to the highest-ranked risks. Participants' responses to risks with similar rankings are fairly similar. This supports the validity of the ranking from the large-scale rating survey. For low-ranked risks, participants would attempt to resolve the situation themselves or complain. Responses to mid-range risks contain a greater emphasis on complaining in reviews or to the developer. For high-ranked risks, many participants would seek help from external parties like service providers, police, or lawyers.

4.3 Limitations

The ratings and open-ended questions are not absolute measures of user concern because our surveys explicitly asked respondents about privacy and security. Surveys that directly ask questions about privacy suffer from inflated user concerns about privacy [5], and therefore are not a reliable measure of the absolute level of concern. We expect this applies to our study as well. We intentionally primed respondents to think about the negative side-effects of applications because we did not want users to mix risks and features. Instead, our surveys provide a basis for comparing risks against each other. The same set of priming biases are applied equally to all of the risks presented in the surveys.

Our study relies on self-reported data. Users might act differently when confronted with actual problems on their phones. However, this bias is equally present for all risks, so it does not affect the validity of our ranking.

We relied on Mechanical Turk workers for survey data. As discussed in Section 3.1.3, the workers who completed our study did not proportionately represent the smartphone population in terms of occupation. Our survey did not reach many highly-paid professionals, who may have different concerns. However, our survey was taken by a large number of participants with varying ages and socio-economic statuses. Secondary studies may be needed to target specific groups that could plausibly have their own privacy and security concerns, such as doctors (patient data), lawyers (client data), or executives (corporate data).

5. IMPLICATIONS AND DISCUSSION

Warnings. Our results could be used to guide warning design: warnings for higher-ranked risks should be more emphatic than warnings for lower-ranked risks, in order to convey the appropriate level of danger. Android currently divides permissions into three levels of severity, but their categorization differs from our survey respondents' concerns in many cases. For example, Android places the `SET_TIME` permission in the highest-severity category, yet it falls in the bottom third in our ranking. Conversely, access to photos ranks in the top quartile in our study, yet Android does not restrict access to photos with any permission at all.

iOS only prompts users for consent for location data and pop-up notifications. However, we find that both rank low in comparison to other privileges; this may indicate that iOS does not ask users about the correct privileges. Although iOS applications cannot perform all of the actions in our ranking, they can perform many of the actions that rank higher than location and notifications without a consent dialog. Given our ranking, iPhone users' complaints about the lack of a consent dialog for contacts [4] is not surprising.

Location. Most mobile privacy and security research has focused on location. However, we find that improper location sharing is not viewed as dangerous in comparison to the other risks of using applications. All of the location-related risks rank in the bottom half of risks, and location is the second-lowest data type. Consequently, we believe that the privacy community should refocus their efforts on other types of smartphone data that evoke higher levels of user concern, such as text messages, photos, and contacts.

Sharing Data. Respondents' concerns about illicit data sharing depend on who the data is being shared with. Based on this finding, we suspect that warnings about data access that do not specify where the data is being sent to do not provide users with enough information to gauge the risk of sharing the data. This motivates further work on tools like AppFence [11] that tell users whether data is being sent to advertisers or other known third parties. Additionally, developers could provide annotations that reflect their privacy policies, and this information could be incorporated into warnings or data access requests.

Service Providers. In our qualitative study, 19% of participants said that they expect their service provider to remedy any data loss or data theft. Several specifically stated that they would consider switching service providers if an application misbehaved. In practice, service providers do not offer default backup services, nor do they control what applications are listed in application markets. This suggests that some users may not understand the security or liability implications of installing applications. In contrast, some participants said that they would contact their service providers to refund fraudulent SMS or phone charges; this expectation is likely well-founded, depending on the service provider.

Direct Observation. In future work, we plan to explore ethical ways of measuring users' reactions to negative experiences with applications. This would allow us to collect data on how users respond to application misbehavior without priming participants or relying on self-reported data. However, this poses numerous ethical challenges: we cannot actually delete user data on phones or perform similar harmful actions, and a study that relies on deception could still prove emotionally difficult for participants.

6. CONCLUSION

We surveyed 3,115 smartphone users on Mechanical Turk about potential risks of smartphone applications. Participants rated how upset they would be if the risks occurred. From this data, we developed a ranking of risks by user concern. A follow-up, open-ended survey of 41 smartphone users found that users view the lowest-ranked risks are annoyances that they can resolve, whereas the highest-ranked risks are serious offenses that may require external parties. Our ranking could be used to guide warning design, and our results show that location is not a high-ranked user concern. We also found that some users hold service providers responsible for abusive applications.

Acknowledgments

We would like to thank Coye Cheshire and Galen Panger for their insightful comments and discussion. This material is based on work supported by a Facebook Fellowship, NSF grant CCF-0424422, and the Intel Science and Technology Center for Secure Computing. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of Facebook, the National Science Foundation, or Intel.

7. REFERENCES

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

[2] L. Barkhuus. Privacy in location-based services, concern vs. coolness. In *Workshop on Location System Privacy and Control at MobileHCI*, 2004.

[3] L. Barkhuus and A. Dey. Location-based services for mobile telephony: a study of users' privacy concerns. In *INTERACT*, pages 702–712, 2003.

[4] C. Bonnington. Apple Says Grabbing Address Book Data Is an iOS Policy Violation. *Wired: Gadget Lab*, February 15 2012. `http://www.wired.com/gadgetlab/2012/02/apple-responds-to-path/`.

[5] A. Braunstein, L. Granka, and J. Staddon. Indirect Content Privacy Surveys: Measuring Privacy Without Asking About It. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2011.

[6] comScore Data Mine. US Smartphone Owners by Age. `http://www.comscoredatamine.com/2011/06/us-smartphone-owners-by-age`, June 2011.

[7] S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2005.

[8] A. P. Felt, M. Finifter, E. Chin, S. Hanna, and D. Wagner. A survey of mobile malware in the wild. In *Proceedings of the 1st ACM workshop on Security and Privacy in Smartphones and Mobile Devices (SPSM)*, 2011.

[9] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: User attention, comprehension, and behavior. Technical Report UCB/EECS-2012-26, February 2012.

[10] D. Goodin. Backdoor in top iPhone games stole user data, suit claims. *The Register*, November 2009.

[11] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall. These Aren't the Droids You're Looking For: Retrofiting Android to Protect Data from Imperious Applications. In *ACM Conference on Computer and Communications Security (CCS)*, 2011.

[12] G. Iachello, I. Smith, S. Consolvo, M. Chen, and G. D. Abowd. Developing privacy guidelines for social location disclosure applications and services. In *Proceedings of the 2005 Symposium on Usable Privacy and Security (SOUPS)*, 2005.

[13] P. G. Kelley, M. Benisch, L. F. Cranor, and N. Sadeh. When are users comfortable sharing locations with advertisers? In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2011.

[14] P. G. Kelley, S. Consolovo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall. A Conundrum of Permissions: Installing Applications on an Android Smartphone. In *Proceedings of the 2012 Workshop on Usable Security (USEC)*, 2012.

[15] S. Lederer, J. Mankoff, and A. K. Dey. Who wants to know what when? Privacy preference determinants in ubiquitous computing. In *CHI Extended Abstracts on Human Factors in Computing Systems*, 2003.

[16] I. Muslukhov, Y. Boshmaf, C. Kuo, J. Lester, and K. Beznosov. Understanding Users' Requirements for Data Protection in Smartphones. In *ICDE Workshop on Secure Data Management on Smartphones and Mobiles*, 2012.

[17] F. Roesner, T. Kohno, A. Moshchuk, B. Parno, H. Wang, and C. Cowan. User-Driven Access Control: Rethinking Permission Granting in Modern Operating Systems. In *Proceedings of the IEEE Security & Privacy Symposium*, 2012.

[18] S. Thurm and Y. I. Kane. Your apps are watching you. *The Wall Street Journal*, December 2010.

[19] J. Wiese, P. G. Kelley, L. F. Cranor, L. Dabbish, J. I. Hong, and J. Zimmerman. Are you close with me? are you nearby?: investigating social groups, closeness, and willingness to share. In *Proceedings of the 13th International Conference on Ubiquitous Computing (UbiComp)*, 2011.

APPENDIX

A. RESEARCH ETHICS

We received advance approval from the Institutional Review Board at the University of California, Berkeley to perform this work. Survey data was collected anonymously. The interviews that we conducted to test the large-scale survey instrument were not anonymous, but we did not ask interviewees to provide any confidential or sensitive information.

B. FULL RESULTS

Table 3 shows the VUR rates for all of the 99 risks that were in our survey, ordered by rank. We have grouped questions about sharing the same data type into the same row. The tightest confidence interval for a VUR rate is ±1.4% at a 95% confidence level, and the widest confidence interval for a VUR rate is ±5.0% at a 95% confidence level. (The confidence interval depends on the number of ratings – $\mu = 376.7$ – and how close the VUR rate is to 50%.)

Risk	**Very Upset**	*Publicly*	*Friends*	*Advertisers*	*Servers*
permanently disabled (broke) your phone	98.21%				
made phone calls to 1-900 numbers (they cost money)	97.41%				
sent premium text messages from your phone (they cost money)	96.39%				
deleted all of your contacts	95.89%				
used your phone's radio to read your credit card in your wallet	95.15%				
shared your text messages		94.48%	92.49%	88.63%	75.48%
deleted all of the information, apps, and settings on your phone	94.39%				
shared your e-mails		93.37%	86.87%	83.96%	75.51%
deleted all of your other apps	93.14%				
recorded your credit card # when you entered it into a different app	92.35%				
shared your photos		90.60%	81.28%	87.26%	60.95%
changed your keylock/pattern/PIN	90.46%				
sent text messages from your phone	90.42%				
shared your contact list		-	73.59%	90.19%	69.29%
spammed your contacts with event invitations	89.73%				
made phone calls	89.62%				
sent spam to people on your contact list	87.95%				
shared your call history		87.77%	82.04%	85.8%	68.41%
shared your browsing history and bookmarks		-	85.68%	70.74%	54.59%
shared your calendar		83.68%	63.59%	73.47%	50.14%
recorded the passwords that you enter into other apps and websites	83.38%				
sent spam from your e-mail account	82.76%				
shared your e-mail address		80.70%	52.00%	82.31%	46.87%
deleted or changed files used by other apps on your phone	82.14%				
inserted spam messages at the end of a text message you sent	81.15%				
hung up your phone when you're talking	81.00%				
shared your phone's unique ID		78.92%	-	76.61%	42.16%
recorded you speaking with your phone's microphone	78.86%				
installed other apps onto your phone	78.46%				
took a photo with your front-facing camera	77.30%				
muted a phone call when you're talking	77.27%				
deleted all of the events on your calendar	76.89%				
posted to your Facebook wall	76.30%				
used your data plan to download data when you were roaming	75.57%				
turned your keylock/pattern/PIN off	74.72%				
deleted other apps' saved passwords	73.96%				
took a photo with your rear-facing camera	73.54%				
shared your location		71.57%	58.10%	62.80%	29.88%
un-muted a phone call	70.73%				
took screenshots when you're using other apps	70.23%				
deleted all of your browser bookmarks and RSS feeds	69.87%				
added new contacts	69.57%				
force quit all your other apps	69.29%				
used your data plan to download data	67.83%				
turned your Internet connection off while you were using the Internet	64.84%				
logged in to your Facebook account	64.34%				
prevented other apps from running	63.99%				
shared the list of apps you have installed		54.77%	52.36%	61.85%	34.92%
prevented your phone from being backed up to your computer	61.39%				
used your phone's unique ID to track you across apps	60.33%				
changed the time on your phone	60.00%				
logged in to your saved Google account	59.38%				
restarted your phone	57.56%				
drained your battery	55.61%				
set alarms on your phone	54.05%				
changed your phone's wallpaper	52.51%				
turned the sound on your phone up really high	52.12%				
showed you lots of pop-up notifications	51.90%				
prevented your phone from being backed up to the cloud	50.52%				
slowed down your phone	50.00%				
disconnected you from a Bluetooth device (like a headset) while you were using the Bluetooth device	48.63%				
turned your WiFi back on when you were on a plane	45.52%				
inserted extra letters into what you're typing	45.48%				
read files that belong to other apps	44.33%				
added new browser bookmarks	39.22%				
turned the sound on your phone down really low	36.96%				
turned your flash on	29.67%				
connected to a Bluetooth device (like a headset)	27.47%				
vibrated your phone	15.62%				

Table 3: The number of respondents who indicated they would be "Very upset" if a risk occurred.