Peeling the Lemons Problem with Risk Communication for Mobile Apps

Behnood Momenzadeh
School of Informatics and Computing, Indiana University Bloomington

Jean Camp
School of Informatics and Computing, Indiana University Bloomington

Abstract

Information asymmetry is a common challenge in information security. This information asymmetry arguably exists in app markets, where people do not understand permissions and have little information on the security of apps. It is not feasible to compare apps based on security and privacy in current app stores. Solving this lemons market problem requires the creation of signals that allow users to differentiate between otherwise indistinguishable goods: more or less secure. In the case of mobile app selection, effective signals should distinguish apps from with lower or higher quality in terms of information security and privacy. To function, such signals should be meaningful, available at or before a decisions is made, and easy to understand. We used the lock icon as a cue, due to its connection with security, and implemented a ratings scale based on We developed an extended Play Store that embedded information security signals. We recruited sixty participants to test the interaction using tablets running Jelly Bean with the cues, as well as the standard user ratings, download count, and permissions interface. The result was that participants chose apps with higher security ratings, and accepted apps with lower ratings or lessor download counts to obtain apps with higher security ratings. We conclude with comparing our results to the users' behavior in Android Market and show how improving security in the Android ecosystem, can be an economic solution to a lemons market challenge.

1 Introduction

There is evidence that privacy and security are both subject to information asymmetry, in other words these are lemons markets [4]. In the smart phone domain, a permissions model is used to provide information to support informed choice and address this information asymmetry. These permissions allow access to phone functionality and user data. Yet, significant research has shown that people do not understand permissions. Others argue that permissions illustrate the privacy paradox or the control paradox. Yet, if the app market is a lemons market with respect to security and privacy, this can explain part of this paradox in that individuals cannot distinguish apps based on their permissions and the associated information risks. If the current interaction model is inadequate to support permissions-aware decision making, then the existence of signals can make a difference in selection of apps.

Choosing to download or use an app can be simple decision of evaluating its costs and benefits. The benefits can be said to include security and privacy. And information exfiltration is part of the cost, like through paying via personal info. With the current interaction design it is difficult for an individual to evaluate the security costs and benefits. It is possible to compare apps based on privacy and security. Yet this requires 1) an understanding of the permissions model, 2) an understanding of the risk associated with specific permissions, and 3) the cognitive work to compare the various options provided in the permissions manifests. One way to address this is to provide easy to understand indicators that distinguish between low risk and high risk apps, which we refer to here as a signal.

To explore the efficacy of such a signal, we developed an Android Play Store interaction that included permissions indicators which are visible in the listing of apps. We illustrated that this changed the choices about which apps to install with users handling a Android Nexus 7 tablets. With an interaction that provides information about the permissions aggregated into a single indicator, individuals choose apps having better privacy/security rating over those with more downloads or higher app ratings.

2 Related Work

In fact, evidence suggests that users do not understand permissions. A common statistic is that only 17% of people look at permissions [19]. In previous work, we have found less than 8% [32] view permissions, and in this work permissions were viewed in only 7% of apps chosen. When simply asked if they were comfortable in their understanding of permissions, 74% said they are [24]. One possibility is that people have become desensitized to over-privileging [24]; for example, a person indicated she previously avoided apps asking for location but abandoned that effort as futile [18].

Multiple approaches have been proposed to increase user comprehension of permissions to improve decision-making. [42, 7] leveraged crowd-sourcing to improve descriptions and understanding of permissions. Similarly, our ratings [29, 33] build on the tradition of peer production and crowd-sourcing in security [13, 14, 3]. Thus far, however, the predominate social effect on app downloads appears to be the importance of download counts as a source for decision support [26]. Such indicators may be quite misleading, as Morton found, in both focus groups and a large scale survey. Participants in both experiments reported by Morton, indicated concern about privacy, and faith that widespread adoption indicated acceptable privacy and security settings. [31]

Considering the cognitive work necessary to compare the various options provided in the permissions manifests as mentioned above. People are cognitive misers and this applies in security and privacy [2]. This is particularly important in mobile computing because individuals may respond differently to actual devices as opposed to a simulated device. Tasks on tablets and mobile devices are quite likely to have different cognitive responses than desktop computers. The simple proximity of an item to a hand increases the likelihood of response to a stimuli. Such proximity also reduces the speed of evaluation while increasing the comprehension [41]. Both distance from the face and the choice of single or double hands changes cognitive response [12], with two hands requiring a lessor interruption for the same focus. Since a significant component of our experiment embeds comprehensive, we choose the more labor-intensive method of building a functional app store and recruiting individuals for actual use of the app over the option of a larger sample size in MTurk.

The stronger cognitive response may be one cause of the high variance in user engagement with security in mobile devices, as observed by [20]. That work concluded that people need concise, precise, and simple to use security interactions for these to be effective because of low levels of user engagement. Yet the counter argument can be made, that perhaps people are reasonably unconcerned. For example, "When 24¢is too much", examined a willingness to pay to hide or to prevent exposure of information [23]. While there were some willing to pay a premium to hide information, most valued information concealment on the order of pennies. In later work in the mobile domain, a plurality of users stated a willingness to accept risk to obtain either the desired functionality, a free app, or a combination [24]. It is possible that people will not respond to risk cues. Certainly individuals ignore warnings on desktops [15, 9], so there is no certainty that they would engage with comparable icons on a mobile phone.

If we consider privacy and permissions comparable to end user licenses agreements rather than warnings, then this visibility may increase acceptability and decrease regrets. For example, when individuals were provided clear information about end user licensing agreements up front, there was little change in removal of software. In addition, individuals told up front were more accepting of the EULA conditions than those told afterwards [21, 22].

Alternatively, people may care about privacy but lack understanding. Providing aggregate information about permissions requirements increases users expressions of their awareness of permissions and associated risks [28]. These findings further support an argument that privacy is a lemons market [39, 30] where additional clarifying information can change decision-making.

There is an argument that people will change their behavior when given simple indicators of risk. In previous research, Tsai and colleagues found that clear information about privacy on a website resulted in consumers willingness to pay a premium for privacy [38]. They found a willingness to pay for items that are widely considered privacy-sensitive (a sex toy) and items that are not considered privacy-sensitive (batteries, ones that were not compatible or related to the toy). Changing a privacy interaction on the web changes willingness to share information [27, 1]. Similarly, using the socially-based icon of eyeballs, individuals were shown to change privacy behavior, albeit inconsistently [34].

In terms of the permissions model previous research such as [18] has consistently found that developers do not understand the permissions model. Systematic over-privileging resulted from confusion in permission naming and permission inheritance. Developers also bundle permissions, requesting entire classes when only one permission is needed. In the case of wifi permissions, code reuse and popular but confusing documentation seem to create confusion and risk. Apps even include systems permissions that are refused in practice by the OS. Given the complexity of permissions for developers, expecting users to understand them seems optimistic.

There are strong arguments that people do not un-

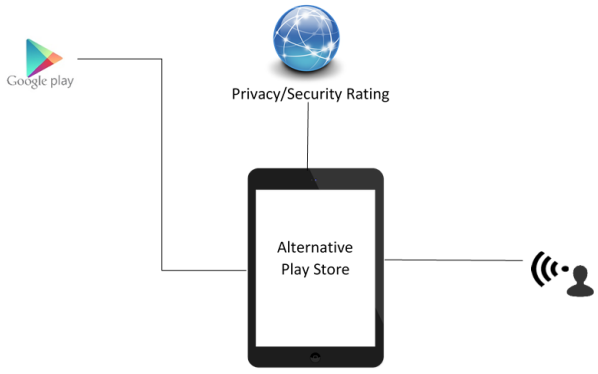

Figure 1: Our system overview

derstand permissions and may change their behavior if provided clear risk information. There is also research showing that people accept over-privileging, and that individual concerns about information risks do not result in changed behaviors. In this experiment we test these conflicting arguments by building on the underlying concept of information asymmetry and providing risk signals.

While we used a permission/privacy rating for our alternative Play Store, there has been numerous research projects to quantify the risk of an app. We categorize these projects into 4 main categories. 1-The ones which mainly use permissions [8, 36]: These projects focus on the permissions an app request with regard to the category of its app, for instance if a flashlight app asks for access to contacts, it looks as a malicious activity for these rating systems. 2- Static analysis [17, 35]: In this category projects focus on analyzing the source code of an app in order to find out if there is any sort of data leak or not. 3- Dynamic analysis [16, 11, 5]: These projects use the data flow in apps to check if the app is using user data maliciously. 4-User intent [40, 29, 33]: In these projects, researchers try to infer the intent of the user when he/she uses an app and grants a permission. Our system is designed to accept all of these approaches as long as they are consistent across all apps and could be interpreted into a rating. For the purpose of this research we used privacygrade, a project that falls into the user intent category.

3 Experiment Design

Do people change their app selections when provided information that distinguishes apps based on their permission requests? If there is significant change in behavior this supports an argument that individuals will act on permissions if the information is clearly and simply presented. Such an argument is the essence of information asymmetry. Whether the change in behavior is a result of improved usability or risk communication, the availability of information in comprehensible, usable, and transparent form is essential to a functioning market. If there is no change, this would support the contention that individuals want free apps, and the benefits outweigh the risks. As our interest is in market impact and ordering of selections, we developed an app store with the goal of being cognitively identical to the current Play Store with the only difference being the presence of signals about the risk (or safety).

If there is no change in behavior then our distribution of apps selected should be the same as the distribution of apps selected in the Android Play Store. That is, we should be sampling from a well known distribution and our sample should be representative.

Every app in the current Android market place has a rating out of five and apps are presented in order of popularity and match to the search. The apps were presented in the same order in our work. To calculate the privacy and security of each app, we use ratings from privacygrade. Privacygrade assigns ratings of A+ to D for different apps which we convert to 5 to 2 locks. If the app is too new and therefore it has not been processed by privacygrade yet, the app will be assigned 1 lock.

Permissions are consistent across apps and readily available to consumers and developers. As such they provide a consistent risk metric that is already designed to support decision-making. Our prototype pulls information from privacygrade.org [29, 33]. However, any arbitrary risk rating could be implemented as long as it is consistent across categories and the apps within these categories. We are explicitly not making the argument that our risk ratings are optimal. The underlying ratings and the interaction are separate functionality. In fact, any organization could promulgate a version of our interaction and provide their own risk ratings; for example, a consumer protection agency may place higher priority on privacy while a security company may provide ratings as a service to encourage employees to consider risk in a BYOD environment.

Previous work has tested a number of different cues or signals using simulations and MTurk. [32] uses different mental models and padlocks were found the most effective and most representative of mental models for privacy and security and it verified the result of another research project which had shown communicating privacy using padlocks is the best representation [6].

To test competing arguments about the efficacy of risk information in decision making we constructed an alternative Play Store, as shown in Figure 1. We used Android Nexus 7 tablets with Jelly Bean.

We used normalized aggregate permissions ratings for each app in each category and displayed those in a scale of one to five. We added this rating to an otherwise iden-

tical Play Store and implemented this augmented interaction on an Android Nexus 7 tablets. We recruited a diverse participant population through outreach at the public library and the Farmers Market. Our sixty participants chose apps from our alternative Play Store and are compared against Android Play Store users. With the users of Android Play Store selecting apps based on user ratings and downloads, while the experimental group's choices reflected permissions/privacy on top of what Play Store users use.

The purpose of building on permissions as a basis for ratings was three fold. First, we wanted to use information that is already provided in the marketplace as opposed to embedding a new standard or ratings mechanism for risk. Second, permissions provide information about security and privacy, whereas a privacy-only or security-only rating would have to be generated by the researchers. It is reasonable to assume that the mobile team at Android has a better understanding of risk concerns on their own platform than we could develop for this experiment. Third, and most importantly, we seek to make a contribution about supporting decision-making in the marketplace. It is not our intention to define the best possible rating system. Therefore, informed by the literature on risk communication and human decision-making, we offered an aggregate rating for each app based on category and specific permissions requested.

We used a play store that was developed as part of μg [**?**] project. Although it is not maintained anymore, we found it the most suitable for our needs. It uses the Google play store APIs to get information like user ratings, number of downloads, descriptions, list of permissions and etc. We added our privacy rating to it, automated the user login, changed the GUI to include our privacy rating, added permissions button, changed the way installation worked (which we will talk about later in experiment design section). We made sure that we send minimal requests to fetch privacy ratings both not to flood the server with duplicate requests and to optimize the performance of our alternative play store. We save the rating fetched for a version of a particular app so we do not have to resend the request unless the user connected to the play store logs out which is a way to introduce a new method of fetching privacy ratings and clearing the past fetched ratings.

4 Methodology

A core design goal was to make the experimental interaction as close as possible to the actual experience of a user interacting with the Android Play Store marketplace. So the user could search, choose, download and install the apps he/she chose to do. But after running our first pilot experiment we saw that downloading apps took considerable time, resulting in fewer people willing to complete the entire experiment. In our pilot running an experiment could take up to 60 minutes (much longer than 15 minutes we had estimated for each participant). As a result, we decided not to download the apps and only show a "successful installation" pop up as soon as the participant clicked on the download/install button.

The essences of our experiment is asking users to select apps and evaluating the resulting decision in terms of permissions (our risk measure), downloads, and community rating (i.e., stars). As shown in Figure 4 the user could have read the permissions before installing the app using the button below the screen shots of the app.

We asked each participant to select four apps from four different categories for a total of sixteen apps. We provided the tablet, and asked participants to use specific search terms for each category. The categories were Flashlight, Photos, Games and Weather. To make sure all the participants saw the same results we ensured that each participant used the exact same term (i.e., the name of the categories mentioned above). The experiment was subject to IRB review and approval. We did not describe the purpose of the experiment as being grounded in security. We did not bring the participant attention span to the indicators. Our goal was not to inform participants about risk but to observe their changes.

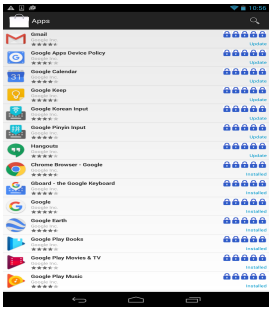

Our Play Store-my apps

Google Play Store-my apps

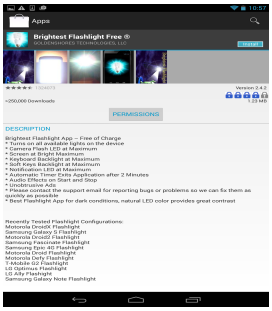

Our Play Store-app details

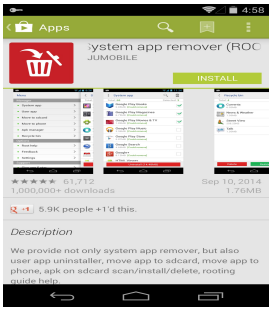

Google Play Store-app details

After the participants completed the selection of apps, we asked them to fill out two surveys to describe their experience with the marketplace. One of the surveys was a workload task survey, the NASA task load index [25].

The second survey consisted of demographic ques-

tions and also the participants' self-reported habits while installing applications from app store.

When the keyword for each category was searched during the experiment, participants were presented with many apps to choose from. We compare top 5 apps with most installs during experiment with top 5 apps with most downloads from app store which were shown in the search results.

5 Results

If our participants are representative samples of Android participants then the distribution of selection of apps should match the distribution of selection of apps in the Play Store. In fact, only those apps with high security ratings (showing low risk) were selected in the top five chosen by our participants. That is, participants mostly chose apps with high security ratings. This is particularly clear in the weather category, where the most popular app was the same for both cases but those rated as more secure were chosen next in our experiment over more popular but less secure apps. Similarly, in Photos category participants chose more secure apps over other more popular apps with more downloads, more familiarity and more popular design. Specially, *PicsArt Photo Studio and Collage* was just selected by 3 of our participants in our experiment for the Photos category. Similarly the over-privileged *Widget Forecast Radar* was systematically rejected by our participants.

Comparing results show that in 2 (Photos and Weather) out of 4 categories, participants have decided to choose apps with higher security rating over apps with more downloads while in the other 2 categories there is no evidence that shows our permission-based rating has affected participants' decisions as the most downloaded apps had high security ratings themselves.

We had 58% male and 42% female participants. Around 22% of the our participants said they will "almost every time" or "always" check permissions while installing an app. However, only 7% of our installations were preceded with a check of application's permissions. A majority (79.6%) of the participants declared they have refused to continue with the installation of an app because of its permissions before. In our experiment there was one instance in which a user did not continue with the installation after viewing the permissions. When asked about participants' priorities when installing apps, around 48% of the participants prioritized application's features over other criteria for choosing apps. The criteria participants could choose from were: ads, permissions, rank, reviews, friends' suggestions, popularity, features, and design. In another survey we asked the participants about the Play Store itself and around 78% of the participants found it easy to work with our Play Store.

6 Discussion

In general, it is necessary for purchasers to have clear and correct information at the time of purchase for a market to function [4]. In our work we applied these long-established principles to the design of decision support for privacy awareness in the selection of apps.

It was not entirely certain that additional risk information will actually reduce risk-taking. In fact, perceptions of risk mitigation sometimes increases risk taking. Studies in risk communication have sometimes shown that individuals find risk more acceptable if the exposure to the risk is voluntary; and the individual exposed is capable of mitigating the impact of risk. That is, shifting the nexus of control may increase aggregate risk-taking. In privacy, this response is called the 'control dilemma' [10]; that is the perception of control increases data sharing. Yet unlike some risks, longitudinal exposure may increase mitigating activity [37, 10].

The goal of the NASA TLX was to evaluate if there was significant cognitive load in using our experimental marketplace. We were unable to make significant conclusions about task load with our participants. Only one person reported the workload of this app store as being "quite demanding". A longer term comparison with a larger group may allow us to make stronger assertions than simply that a highly usable Play Store is not burdensome. In retrospect, this is not surprising.

7 Conclusions

In the mobile market permissions control access to user data and phone functionality. When forced to make a trade-off between risk and benefits, individuals with simplified indicators may choose lower benefit and higher security options, as shown in the choices for weather and photos. However, when there is one choice with far more downloads, that will still be users' top choice.

From the results of this and previous work [32] we can conclude that providing security ratings wont affect the participants decision about a dominant app as we see that privacy ratings has not affected participants' first choices in choosing *Weather - The Weather Channel*, which still is participants' top choice despite having less security rating.

Nevertheless, the average app ratings for top 5 choices of our participants and the average app ratings for the top 5 choices of the Play Store are quite similar. When app ratings are similar participants choose apps with better security ratings over apps with more downloads which might mean that with the same quality, participants will choose security over popularity.

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Super-Bright LED Flashlight	38	5	4.6	>500 millions
Tiny Flashlight + LED	26	5	4.4	100-500 millions
Color Flashlight	34	5	4.2	50-100 millions
Brightest Flashlight Free	20	4	4.7	50-100 millions
Brightest LED Flashlight	15	5	4.5	50-100 millions
Weighted Average	—	4.85	4.46	188 millions

Table 1: Flashlight - Most Downloads

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Super-Bright LED Flashlight	38	5	4.6	>500 millions
Color Flashlight	34	5	4.2	50-100 millions
Tiny Flashlight + LED	26	5	4.4	100-500 millions
Brightest Flashlight Free	20	4	4.7	50-100 millions
Flashlight Galaxy S7	16	5	4.8	1-5 millions
Weighted Average	—	4.85	4.56	181 millions

Table 2: Flashlight - Participants' Choices

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Google Photos	39	5	4.4	>500 millions
PicsArt Photo Studio and Collage	3	3	4.4	100-500 millions
Photo Editor Pro	20	5	4.3	100-500 millions
Photo Grid and Photo Collage Maker	15	5	4.5	100-500 millions
Photo Lab Picture Editor FX	24	5	4.4	50-100 millions
Weighted Average	—		4.56	181 millions

Table 3: Photos - Most Downloads

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Google Photos	39	5	4.4	>500 millions
PhotoDirector Photo Editor App	25	5	4.6	10-50 millions
Photo Lab Picture Editor FX	24	5	4.4	50-100 millions
Gallery	23	5	4.3	5-10 millions
Photo Editor Pro	20	5	4.3	100-500 millions

Table 4: Photos - Participants' Choices

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Subway Surfers	23	5	4.5	>500 millions
Fruit Ninja Free	39	5	4.3	100-500 millions
Piano Tiles 2	15	5	4.7	100-500 millions
slither.io	12	5	4.3	100-500 millions
PAC-MAN	20	5	4.0	50-100 millions

Table 5: Games - Most Downloads

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Fruit Ninja Free	39	5	4.3	100-500 millions
Subway Surfers	23	5	4.5	>500 millions
Super Smash Jungle World	22	5	4.2	10-50 millions
PAC-MAN	20	5	4.0	50-100 millions
Wheel of Fortune Free Play	16	5	4.5	5-10 millions

Table 6: Games - Participants' Choices

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Weather - The Weather Channel	40	4	4.3	50-100 millions
AccuWeather	31	5	4.3	50-100 millions
Weather and Clock Widget Android	14	4	4.4	50-100 millions
Widget Forecast Radar	3	4	4.5	10-50 millions
Weather Underground	19	5	4.5	5-10 millions

Table 7: Weather - Most Downloads

Name	Experiment Downloads	No. of locks	App Rating	No. of Downloads
Weather - The Weather Channel	40	4	4.3	50-100 millions
AccuWeather	31	5	4.3	50-100 millions
Yahoo Weather	27	5	4.4	10-50 millions
MyRadar Weather Radar	27	5	4.5	5-10 millions
Weather Underground	19	5	4.5	5-10 millions

Table 8: Weather - Participants' Choices

References

[1] M. S. Ackerman and L. Cranor. Privacy critics: Ui components to safeguard users' privacy. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '99, pages 258–259, New York, NY, USA, 1999. ACM.

[2] A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

[3] Y. Agarwal and M. Hall. Protectmyprivacy: detecting and mitigating privacy leaks on ios devices using crowdsourcing. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 97–110. ACM, 2013.

[4] G. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. In *Essential Readings in Economics*, pages 175–188. Springer, 1995.

[5] S. Arzt, S. Rasthofer, C. Fritz, E. Bodden, A. Bartel, J. Klein, Y. Le Traon, D. Octeau, and P. McDaniel. Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. *Acm Sigplan Notices*, 49(6):259–269, 2014.

[6] F. Asgharpour, D. Liu, and L. J. Camp. Mental models of security risks. In *International Conference on Financial Cryptography and Data Security*, pages 367–377. Springer, 2007.

[7] L. Bao, D. Lo, X. Xia, and S. Li. What permissions should this android app request? In *Software Analysis, Testing and Evolution (SATE), International Conference on*, pages 36–41. IEEE, 2016.

[8] D. Barrera, H. G. Kayacik, P. C. van Oorschot, and A. Somayaji. A methodology for empirical analysis of permission-based security models and its application to android. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 73–84. ACM, 2010.

[9] R. Böhme and J. Grossklags. The security cost of cheap user interaction. In *Proceedings of the 2011 workshop on New security paradigms workshop*, pages 67–82. ACM, 2011.

[10] L. Brandimarte, A. Acquisti, and G. Loewenstein. Misplaced confidences privacy and the control paradox. *Social Psychological and Personality Science*, 4(3):340–347, 2013.

[11] I. Burguera, U. Zurutuza, and S. Nadjm-Tehrani. Crowdroid: behavior-based malware detection system for android. In *Proceedings of the 1st ACM workshop on Security and privacy in smartphones and mobile devices*, pages 15–26. ACM, 2011.

[12] W. S. Bush and S. P. Vecera. Differential effect of one versus two hands on visual processing. *Cognition*, 133(1):232–237, 2014.

[13] L. J. Camp and A. Friedman. Peer production of privacy and security information.

[14] Z. Dong and L. J. Camp. Peersec: Towards peer production and crowdsourcing for enhanced security. In *HotSec*, 2012.

[15] S. Egelman, L. F. Cranor, and J. Hong. You've been warned: an empirical study of the effectiveness of

web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1065–1074. ACM, 2008.

[16] W. Enck, P. Gilbert, S. Han, V. Tendulkar, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)*, 32(2):5, 2014.

[17] W. Enck, D. Octeau, P. McDaniel, and S. Chaudhuri. A study of android application security. In *USENIX security symposium*, volume 2, page 2, 2011.

[18] A. P. Felt, E. Chin, S. Hanna, D. Song, and D. Wagner. Android permissions demystified. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 627–638. ACM, 2011.

[19] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the Eighth Symposium on Usable Privacy and Security*, page 3. ACM, 2012.

[20] A. Forget, S. Pearman, J. Thomas, A. Acquisti, N. Christin, L. F. Cranor, S. Egelman, M. Harbach, and R. Telang. Do or do not, there is no try: user engagement may not improve security outcomes. In *Symposium on Usable Privacy and Security (SOUPS)*, 2016.

[21] N. Good, R. Dhamija, J. Grossklags, D. Thaw, S. Aronowitz, D. Mulligan, and J. Konstan. Stopping spyware at the gate: a user study of privacy, notice and spyware. In *Proceedings of the 2005 symposium on Usable privacy and security*, pages 43–52. ACM, 2005.

[22] N. S. Good, J. Grossklags, D. K. Mulligan, and J. A. Konstan. Noticing notice: a large-scale experiment on the timing of software license agreements. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 607–616. ACM, 2007.

[23] J. Grossklags and A. Acquisti. When 25 cents is too much: An experiment on willingness-to-sell and willingness-to-protect personal information. In *WEIS*, 2007.

[24] M. A. Harris, R. Brookshire, K. Patten, and B. Regan. Mobile application installation influences: have mobile device users become desensitized to excessive permission requests? In *Proceedings of the Twentieth Americas Conference on Information Systems (AMCIS 2015)*, pages 13–15, 2015.

[25] S. G. Hart and L. E. Stavenland. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In P. A. Hancock and N. Meshkati, editors, *Human Mental Workload*, chapter 7, pages 139–183. Elsevier, 1988.

[26] P. G. Kelley, L. F. Cranor, and N. Sadeh. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3393–3402. ACM, 2013.

[27] S. Kolimi, F. Zhu, and S. Carpenter. Contexts and sharing/not sharing private information. In *Proceedings of the 50th Annual Southeast Regional Conference*, pages 292–297. ACM, 2012.

[28] L. Kraus, I. Wechsung, and S. Möller. Using statistical information to communicate android permission risks to users. In *Socio-Technical Aspects in Security and Trust (STAST), 2014 Workshop on*, pages 48–55. IEEE, 2014.

[29] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang. Expectation and purpose: understanding users' mental models of mobile app privacy through crowdsourcing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pages 501–510. ACM, 2012.

[30] G. R. Milne and M. J. Culnan. Strategies for reducing online privacy risks: Why consumers read (or don't read) online privacy notices. *Journal of Interactive Marketing*, 18(3):15–29, 2004.

[31] A. Morton. all my mates have got it, so it must be okay: Constructing a richer understanding of privacy concernsan exploratory focus group study. In *Reloading Data Protection*, pages 259–298. Springer, 2014.

[32] P. Rajivan and J. Camp. Influence of privacy attitude and privacy cue framing on android app choices. In *Authentication Workshop of the 12th Symposium on Usable Privacy and Security*. USENIX Association, 2016.

[33] J. Sadeh and J. I. Hong. Modeling users mobile app privacy preferences: Restoring usability in a sea of permission settings. In *Symposium on Usable Privacy and Security (SOUPS)*, volume 40, 2014.

[34] R. Schlegel, A. Kapadia, and A. J. Lee. Eyeing your exposure: quantifying and controlling information sharing for improved privacy. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 14. ACM, 2011.

[35] A.-D. Schmidt, R. Bye, H.-G. Schmidt, J. Clausen, O. Kiraz, K. A. Yuksel, S. A. Camtepe, and S. Albayrak. Static analysis of executables for collaborative malware detection on android. In *Communications, 2009. ICC'09. IEEE International Conference on*, pages 1–5. IEEE, 2009.

[36] W. Shin, S. Kiyomoto, K. Fukushima, and T. Tanaka. Towards formal analysis of the permission-based security model for android. In *Wireless and Mobile Communications, 2009. ICWMC'09. Fifth International Conference on*, pages 87–92. IEEE, 2009.

[37] F. Stutzman, R. Gross, and A. Acquisti. Silent listeners: The evolution of privacy and disclosure on facebook. *Journal of privacy and confidentiality*, 4(2):2, 2013.

[38] J. Y. Tsai, S. Egelman, L. Cranor, and A. Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268, 2011.

[39] T. Vila, R. Greenstadt, and D. Molnar. Why we can't be bothered to read privacy policies models of privacy economics as a lemons market. In *Proceedings of the 5th international conference on Electronic commerce*, pages 403–407. ACM, 2003.

[40] H. Wang, J. Hong, and Y. Guo. Using text mining to infer the purpose of permission use in mobile apps. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 1107–1118. ACM, 2015.

[41] X. Wang, F. Du, X. He, and K. Zhang. Enhanced spatial stimulus–response mapping near the hands: The simon effect is modulated by hand-stimulus proximity. *Journal of Experimental Psychology*, 40(6):2252, 2014.

[42] L. Yang, N. Boushehrinejadmoradi, P. Roy, V. Ganapathy, and L. Iftode. Short paper: enhancing users' comprehension of android permissions. In *Proceedings of the second ACM workshop on Security and privacy in smartphones and mobile devices*, pages 21–26. ACM, 2012.