Reasoning about privacy in mobile application install decisions: Risk perception and framing

Siok Wah Tay [a], Pin Shen Teh [b], Stephen J. Payne [c]

[a] Department of Computer Science, The University of Manchester, Manchester, M13 9PL, United Kingdom, siokwah.tay@manchester.ac.uk

[b] Department of Operations, Technology, Events and Hospitality Management, Manchester Metropolitan University, Manchester, M15 6BH, United Kingdom, p.teh@mmu.ac.uk

[c] Department of Computer Science, University of Bath, Bath, BA2 7AY, United Kingdom, S.J.Payne@bath.ac.uk

Correspondence should be addressed to Siok Wah Tay; siokwah.tay@manchester.ac.uk

Abstract

Data sharing has become prevalent with the rapid growth of mobile technologies. A lack of awareness and understanding of privacy practices often results in the installation of privacy-invasive applications (apps) which could potentially put users' personal data at risk. This study aimed to explore how users' risk perception could be shifted towards more privacy-aware decisions through generation fluency and framing manipulations. It is an online study composed of three components, an experiment and two questionnaires. We manipulated the availability of privacy worries, by asking participants to generate either 2 or 10 privacy worries. Generating 10 worries was experienced as difficult, whereas generating 2 was easy. The difficult experience led to downgraded perception of risk, and consequently increased likelihood of installing a low privacy rated fictional app. Therefore, we suggest that improving generation fluency of privacy concerns could encourage users' adoption of a more conservative judgment strategy when installing an app, safeguarding them against privacy-invasive apps.

Keywords

Privacy; mobile application; generation fluency; risk perception; framing; decision-making

1. Introduction

Over the past few years, the demand for mobile applications (apps) has increased rapidly. Worldwide mobile app downloads had surpassed 90 billion by 2016 and the time spent in apps had reached nearly 900 billion hours (Thompson, 2017). The growing adoption of mobile apps in a wide range of areas has resulted in an enormous increase in the collection, use and sharing of personal data through mobile apps. Although greater data sharing promises benefits and conveniences, at the same time it raises privacy concerns.

Despite large-scale personal data collection through mobile apps, it appears that users are not always made aware of what happens with their data. A study conducted by the Future of Privacy Forum (FPF) revealed that a significant number of mobile apps did not provide users with basic privacy notes about how their personal data would be collected, used and shared (Future of Privacy Forum, 2016). In a study by Enck et al. (2010), it was found that half of the mobile apps in the study transferred private data, e.g. location, to third parties for the purposes of advertising or analytics, without user consent. There also exist apps which covertly collect users' data even when these are unnecessary for the apps' operation. Felt et al. (2011) revealed that about one-third of 940 apps from the Android marketplace were detected as overprivileged due to errors in API documentation and the lack of developer understanding.

Privacy policies or notices have been widely adopted by websites and apps. Ironically, users typically do not read privacy notices because they are often lengthy, complicated, and difficult to understand (Liccardi et al., 2014b; Schaub et al., 2015). The well-known Heuristic-Systematic Model (HSM) of information processing assumes that systematic processing demands greater cognitive capacity and effort than heuristic processing (Chaiken et al., 1989). Hence, we posit that in the context of

mobile app installation, especially inexperienced users who do not possess enough knowledge to understand and process the complexity of permission and privacy information of an app, are likely to make an app-install decision heuristically. Furthermore, due to the ubiquitous nature of mobile devices, users are often on the move and tend to make a quick decision. Providing privacy notices on mobile devices is further challenged by interface restrictions, such as limited screen space. Conventional privacy notices are thus likely to be insufficient in communicating privacy information to users of typical mobile devices. The inefficacy of prevailing privacy communication mechanisms in the mobile app ecosystem is a rising issue in need of attention. Poor privacy communications could result in a low level of risk awareness and perception of users towards mobile apps, leading to privacy-compromising app-install decisions.

2. App-install Decisions

Of course, judged risk to privacy is only one of several factors that will contribute to an individual app-install decision, and this very fact helps explain why privacy-related information might play a smaller role than it perhaps should. In general, before installing an app, we might expect a user to compute some kind of cost-benefit analysis, but we might also expect that, like most time-limited decisions under uncertainty, this process will be heuristic rather than exhaustive: certain aspects of the decision space will be more salient in some decisions than others, according to the user's state of mind, and their current processing resources, as well as how readily information is made available, either by the app itself or by relevant materials such as reviews.

On the benefits side of the decision, a user must judge to what extent the app will be useful to them. We know from Harris et al. (2016) and Shen (2015) that users are influenced by this and we know from Bonné et al. (2017) that user will derive this information from reviews and from perceived popularity of the app (see also (Gu et al., 2017)) as well as from more specific information about functionality.

On the costs side there is expense, available space on one's mobile, and risks to privacy. As with usefulness, the judgment of privacy risks will depend at least on direct and indirect information, but it will also depend on the user's risk-aversion and on the extent to which privacy-risks are activated during the decision, which will depend on salience of information and on the way the user's habitual level of risk-aversion informs the particular heuristic decision (we suggest, somewhat speculatively, that risk aversion might vary at different time scales, with local fluctuations around a more stable baseline).

In the light of this simple, orienting analysis, the aim of this research is to explore how risk perception and risk information interact in the context of (simulated) mobile app-install decisions. In particular, the following research questions (RQs) were formulated.

- **RQ1.** Can risk-aversion be activated or de-activated in an experimental situation, by a simple generation fluency manipulation?
- **RQ2.** Does positive versus negative framing of risk information affect users' perception of risks associated with a low privacy rated mobile app?
- **RQ3.** Do post hoc reports confirm that users attend to privacy information during an app-install decision-making process when the information is made available alongside other app attributes?

3. Related Work

3.1 Privacy and Risk Perception

According to Westin (1968), privacy is "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others." In the mobile app context, users should be given control over their personal data, that is, at the most basic level, users should have the right to exercise real choice over the collection, use, and sharing of their personal data. The Information Commissioner's Office (ICO) defines personal data as any data, by itself or in combination with other information, that can be used to identify an individual (Information Commissioner's Office, n.d.). For example, name and address of individuals are personal data. There are personal data, labelled as sensitive personal data, which should be securely protected. Such data include individuals' health and biometric data, etc.

Risk perception can be generally defined as the subjective evaluation of the likelihood of a misfortune and its accompanying consequences (Sjöberg et al., 2004). In the context of mobile app installation, we define risk perception as mobile users'

assessment of the probability of the occurrence of a privacy breach, e.g. data leak, intentional or unintentional data disclosure, caused by the app, and how concerned the users are about the negative consequences. The literature in risk perception has shown an influential impact of risk perceptions on people's judgment and decision making in a wide array of contexts. Particularly, in the mobile app environment, studies have shown that perceived risk moderates the influence of app reputation on users' attitude towards adopting an app (Shen, 2015) and that subjective risk perception reduces one's intention to install an app (Harris et al., 2016).

3.2 Judgment and Decision Making

Given our emphasis on heuristic decision making under uncertainty, we wish in particular, to explore two classical issues that would be predicted to influence any particular app-install decision. First, with respect to how active a role risk-perception might play, we explore a fluency manipulation designed to activate or deactivate privacy risk during the decision process. Second, following other investigators, we explore how risk information might be framed, and whether this influences the way it is processed.

3.2.1 Retrieval Fluency

Retrieval fluency, also known as accessibility experience, indicates the ease or difficulty with which an individual recalls information relevant to the judgment at hand (Schwarz, 2004). This subjective experience is formed when people generate the same thoughts with varying degrees of difficulty. In a classic ease-of-retrieval study, Schwarz et al. (1991) claimed that in the process of making judgments, people not only rely on what comes to mind, but also take into account how easy or difficult is the process of bringing the information to mind. In that study, participants were told to generate 6 or 12 instances of their assertive or unassertive behaviours, and subsequently rate their assertiveness. Retrieving 6 instances was experienced as easy, indicating high fluency, whereas retrieving 12 instances was relatively difficult, indicating low fluency. The results demonstrated that participants who generated 6 instances of their assertive (or unassertive) behaviours rated themselves as more assertive (or unassertive) on average than those who generated 12 instances.

In a study on health risk judgment, Rothman and Schwarz (1998) found that the perceived self-relevance of a health issue could influence one's adoption of judgment strategy. The study showed that participants used heuristic processing strategy and based their judgments on retrieval fluency when the health issue was not considered self-relevant (i.e. no family history of the disease). On the other hand, when the health issue was considered self-relevant (i.e. with family history of the disease), participants turned to systematic processing strategy.

A host of studies have observed the effects of fluency experiences on decisions in various contexts, e.g. consumer choice behaviours (Novemsky et al., 2007) and on individuals' confidence level in judgments (Kelley and Lindsay, 1993), but to the best of our knowledge, no previous study has examined the impact of fluency on users' decision making about mobile technologies. Our study focused on investigating the impacts of generation fluency on individuals' subjective risk perception, which might subsequently influence their app-install decisions (our use of the name "generation fluency" rather than "retrieval fluency" reflects the fact that we asked participants to list concerns that were not necessarily retrieved directly from memory).

3.2.2 Framing Effects

Tversky and Kahneman (1981) demonstrated that presenting exactly the same choice in different ways can change people's responses to the choice. This phenomenon, known as the framing effect has been tested in a wide array of studies involving decision making in various domains. In particular, presenting an attribute or characteristic of an object in a positive or negative fashion is described as attribute framing (Levin et al., 1998). A classic study by Levin and Gaeth (1988) showed that framing an attribute of an object positively leads to more favourable evaluations of the object than framing the same attribute negatively. This claim has been supported by a number of later studies in a variety of contexts including funding allocation (Duchon et al., 1989), evaluation of organisational performance (Schoorman et al., 1994) and consumer judgment (Levin et al., 1996).

Framing effects have also been extended to the domain of privacy. Johnson et al. (2002) conducted research about permission marketing to observe how positive, i.e. opt-in option, and negative, i.e. opt-out option, framing affected people's decisions on their preference to be notified about future health-related online surveys. In the context of mobile privacy, Choe et al. (2013)

examined the impact of framing on mobile users' perception of an app by manipulating framing using visuals instead of text descriptions. Similarly, a later study by Gates et al. (2014) observed framing effects on apps' risk communication. These two studies are reviewed below.

3.3 Privacy Information Communication

In recent years, several studies have explored new ways to effectively communicate privacy or risk information to users, in an effort to assist them in shaping a more privacy-conscious mobile app adoption or usage behaviour.

Kelley et al. (2013), using a role-playing task in which users considered whether to download real apps, demonstrated the importance of making privacy part of the app-install decision-making process. These authors designed Privacy Facts, a checklist consisting of ten privacy facts, and incorporated this checklist into an app's download page. The privacy facts included the types of information which the app collects and third-party modules the app uses. The findings showed that when privacy information was presented clearly on an app's main screen, users shifted their preferences towards apps which requested less access to personal data.

To complement the privacy facts approach, Harbach et al. (2014) proposed a new app permission display, with the aim of drawing users' attention to permission risks by displaying the potential consequences of installing an app, using personalised examples. For instance, the approach randomly selects and displays an image file stored on the user's device, indicating that the image will be accessed by the app if the app is installed. The study found that highlighting personal relevance in an app's permission interface evoked a negative effect which drew users' attention to permission information before downloading the app.

Liccardi et al. (2014b) integrated a privacy sensitivity score into the Google Play store interface. The sensitivity score (Liccardi et al., 2014a) is measured by the number of sensitive permissions requested by an app when the network permission is also requested. More specifically, sensitive permissions enable an app to access users' personal data, and with the network permission, these personal data could be potentially transmitted to third parties. An online experiment, using fictional apps, found that in the context in which app-selection decision was entirely based on permissions, users preferred apps with fewer permission requests without considering the types of permission requested, implying a lack of awareness and understanding of permission information. On the other hand, when the sensitivity score was employed, users were guided to choose less risky apps.

In a study on users' app recommendation decisions, Kang et al. (2015) proposed Privacy Meter, a meter presented as a slider bar to indicate risks. The results of a role-playing laboratory study showed that the privacy meter is most effective in visualising risks in comparison with the Google Play store permissions displays (before and after version 4.8.20) and the privacy facts checklist of Kelley et al. (2013).

Zhang and Xu (2016) introduced two types of privacy nudges, i.e. frequency and social nudges. The frequency nudge shows how often an app accesses user data such as location and call log. The social nudge indicates the percentage of other users of the same app who allow the app to access their data. An online study using a fictional app showed that privacy nudges could influence users' perception of an app. The frequency nudge constructs a negative perception of the app, e.g. low perceived usefulness, mediated by the emotion of creepiness. The social nudge reflects social norms, making users feel more comfortable with sharing their personal data with the app. In a similar study, Almuhimedi et al. (2015) introduced privacy nudges that actively alert users to the data accessed by apps installed on their mobile phones. The study asserted that the privacy nudges complement a permission manager, empowering users to better control their privacy by revising their current permission settings.

To observe framing effects on users' perception of an app, Choe et al. (2013) represented an app's privacy rating using two semantically equivalent visuals framed either positively or negatively. For example, with a rating scale from 1 to 5, a positively-framed privacy rating of 1 (represented by one green plus sign) is equivalent to a negatively-framed privacy rating of 4 (represented by four red minus signs). Interestingly, unlike most previous studies of framing effects that demonstrated that negative framing yields less favourable evaluations compared with positive framing, this study revealed that a privacy rating in a positive frame is more influential than the one in a negative frame, in protecting users from low privacy rated apps. The authors suggested that people might associate "more signs" with "better" when interpreting the privacy rating, and thus

making a low privacy rated app in a positive frame, i.e. one plus sign, appear less favourable than the one in a negative frame, i.e. four minus signs.

A similar online study was conducted by Gates et al. (2014) to investigate the effects of framing in the app-install context. The summarised risk information of an app was displayed using visual symbols, i.e. circles in a scale of 4, in two conditions: risk and safety conditions. In the safety condition, more filled circles indicate less risk. Conversely, in the risk condition, more filled circles imply more risk. Participants were asked to install only the apps with "the most safety/the least risk" or "some safety/little risk" apps. The results showed that participants process an app's risk information faster in the safety condition rather than in the risk condition, indicating that they were more risk-averse when the positive frame, i.e. safety condition, was used. However, framing effects did not occur in app selection performed in a more realistic setting in which the concept of the "entire app" was applied. These observations suggested that framing effects are present under the conditions in which people make decisions quickly when automatic processing is involved.

Our study extended the work of Choe et al. (2013) in three respects. First, we observed the effects of manipulating users' generation fluency on privacy concerns and investigated how this priming could shape users' risk perception of a low privacy rated app. Additionally, we also examined how generation fluency manipulation might interact with framing effects to influence users' app-install decision. To the best of our knowledge, this idea is novel in the context of privacy in mobile app-install decisions. Second, we tested the effects of generation fluency and framing in a more realistic app-install environment in which all common app attributes including user ratings and reviews as well as download counts were presented, whereas the work of Choe et al. (2013) emphasised user rating and different levels of privacy rating. Third, to design the framing stimuli, inspired by the work of Kang et al. (2015) which highlighted the effectiveness of privacy meter in risk communication, we created a new version of the privacy-rating visualisation. We incorporated the notion of making privacy part of the decision-making process suggested by Kelley et al. (2013) into our privacy rating design. One aspect of the reviewed studies is the variety of experimental contexts that have been used, from laboratory to online studies and with fictional or real apps. The study we report is conducted online, and uses a fictional app and a hypothetical install decision.

4. Method

To address the research questions, we designed an online study in which participants considered a fictional mobile app and reported their willingness to install it. The study consisted of two phases. The first phase comprised an experiment followed by a main questionnaire. In the experiment, participants were given a thought generation task and an app-install decision task. In the main questionnaire, participants answered questions about their mobile app usage and privacy experiences. In the second phase, we conducted a follow-up questionnaire with a small subsample of participants. In the follow-up questionnaire, participants were asked about the reasons for their responses to the tasks performed earlier, in the first phase; this enabled qualitative insights into individuals' attitudes and decision making with regard to mobile app installation.

4.1 Participants

We recruited volunteers to participate in our study through mailing lists, social media and research participants recruitment websites, i.e. Call for Participants (www.callforparticipants.com) and Survey Circle (www.surveycircle.com).We recruited 135 participants for the experiment. 37 of them did not complete the study and were therefore removed from the analyses. A further 7 participants were excluded on the basis of findings reported in the first section of our Results. To be eligible for the study, participants were required to fulfil the following criteria: (1) Aged 18 or over (2) Own a mobile device including a smartphone or tablet (to reduce influence due to unfamiliarity with mobile devices and apps) (3) Currently live in the UK and (4) Fluent in English (as disfluency in English might lead to difficulty in interpreting and performing the tasks). Sixty-one participants were females, 26 were males, and 4 chose not to disclose their gender. Participants' age ranged from 18 to 54, with a median age group of 25 to 34. The majority of participants (74 of 91) were university students. Most of the participants (56 of 91) identified themselves as native or bilingual speakers of English. 21 and 14 participants felt they had full professional proficiency and professional working proficiency, respectively. Our power calculations suggested we needed ~125 participants to detect a medium size main effect with 80% power.

4.2 Experimental Design

The experiment had two independent variables (fully crossed in a 2x2 between-subjects factorial design), one dependent variable, and one covariate, as shown in Table 1. Generation Fluency (FLUENCY) and Privacy Framing (FRAMING) were the independent variables. FLUENCY refers to the subjective experience accompanying an individual's thought generation process. To manipulate FLUENCY, we aimed to generate differential difficulties of experience of thought generation. To engender high FLUENCY, we asked half of the participants to list 2 privacy worries/concerns when using apps with a low privacy rating. To engender low FLUENCY, the remaining participants were asked to produce 10 mobile privacy worries. The number of worries to be generated in the experiment was determined based on a pilot study (Tay, 2017) which showed that generating 2 privacy worries was an easy experience indicating high fluency, whereas generating 10 privacy worries was experienced as difficult, resulting in low fluency. The second independent variable, FRAMING manipulated the presentation of privacy information of the app. We designed a fictional mobile app with a low privacy rating and framed its privacy information either positively or negatively. The design of the fictional mobile app interfaces will be further discussed in Section 4.3. The dependent variable in the experiment was the Subjective App Installation Likelihood (INSTALL) reported by the participants after presentation of the fictional app following the thought generation task. Participants' self-reported INSTALL was measured by a rating on a 10-point scale (1 = definitely would not install; 10 = definitely would install). The covariate was Perceived App Usefulness (USEFUL) which indicates how participants perceived the usefulness of our app. We expected that USEFUL may influence users' app-install decision because users are likely to adopt an app if they perceive the app as useful or beneficial, as suggested by Harris et al. (2016) and Shen (2015).

To address the research questions, we defined two hypotheses as stated below.

H1. Users are more likely to install a low privacy rated app when generation fluency for privacy concerns is low than when it is high.

H2. Users are less likely to install a low privacy rated app when the app's privacy is framed positively than when it is framed negatively.

To test the hypotheses, we used a 2 (FLUENCY: high versus low) x 2 (FRAMING: positive versus negative) factorial between-subjects design, yielding four conditions: (1) high FLUENCY and positive FRAMING (condition denoted as High-FLUENCY & Pos-FRAMING), (2) high FLUENCY and negative FRAMING (condition denoted as High-FLUENCY & Neg-FRAMING), (3) low FLUENCY and positive FRAMING (condition denoted as Low-FLUENCY & Pos-FRAMING) and (4) low FLUENCY and negative FRAMING (condition denoted as Low-FLUENCY & Neg-FRAMING).

Table 1. A list of independent and dependent variables, and covariate used in the experiment.

Types of Variable	Variables	Descriptions
Independent Variables	Generation Fluency (FLUENCY)	Subjective experience of ease or difficulty in generating thoughts. It consists of two levels: high (easy) and low (difficult).
	Privacy Framing (FRAMING)	Privacy information presented in opposite ways (two levels): positive and negative.
Dependent Variable	Subjective App Installation Likelihood (INSTALL)	Participants' subjective likelihood to install the app measured on a 10-point scale: 1 = definitely would not install; 10 = definitely would install.
Covariate	Perceived App Usefulness (USEFUL)	Participants' rating of how useful an app would be to them, measured on a 10-point scale: 1 = not useful at all; 10 = very useful.

4.3 App Interface Design

As shown in Figure 1, we introduced a fictional mobile app named My Fun Language Buddy (MFLB), a language learning app, and created its interfaces for the experiment. To reduce biases in participants' decision making, the following factors were considered in the design of the app interface.

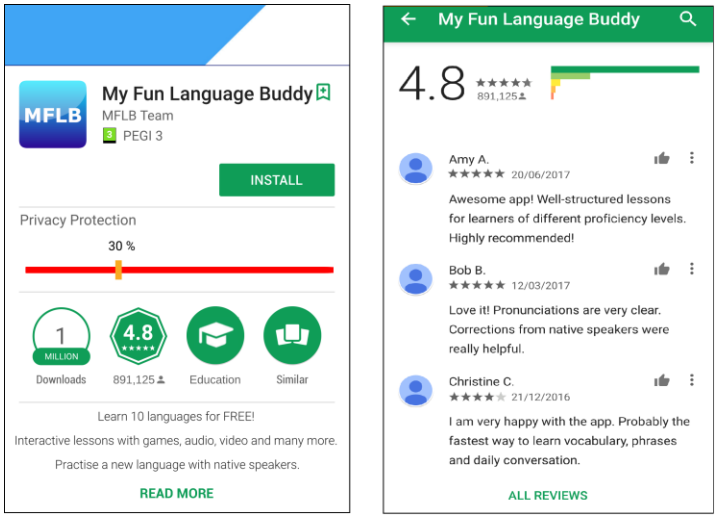

Figure 1. App Description page (left) of MFLB as it appeared in the positive frame and User Review page (right).

- *App brand.* People generally prefer familiar brand names over unfamiliar ones when making choices (Coates et al., 2006). To avoid app-install decisions made based on brand familiarity, we created a fictional app, i.e. MFLB with a non-existent developer name, MFLB Team, rather than using an established app.
- *App type.* In terms of the type of app, we intended to avoid apps favouring only certain user groups which could potentially influence users' app-install decision. An example of such apps is a gardening app which is highly likely to attract only gardening enthusiasts. As a result, language, a common app type was selected based on its suitability for users of all ages and genders. In addition, we created a scenario to engage users in the app-install decision-making process in a more practical and realistic way. The scenario will be presented in Section 4.5.1 (Experimental Task 2).
- *App platform.* Given the open-source nature and wide adoption of Android, we decided to design an app display which is analogous to one of the standard Android app interfaces based on the Google Play store version 8.0.26 run on Android 7.0.

To simulate a real-word app-install environment, we designed the interfaces of our app to mimic the standard Google Play Store Listing (Google Play Store version 8.0.26 run on Android 7.0). Every Android app has a dedicated Store Listing page to showcase its features. A typical Store Listing includes major attributes such as app icon, title, developer, category, e.g. Education, Social, Transport, etc., feature image, description, content rating, e.g. PEGI 3, 12+, etc., varied by country or region, download counts, and user ratings and reviews. All these app attributes are made available on the primary screen without requiring users to open another screen in order to view them. However, no privacy information about the app is displayed on the primary screen. To integrate privacy information into the primary screen of Store Listing, we created a privacy rating, and to make it salient, we placed it directly below the "install" button, preceding other common app attributes including total downloads and average user rating. We designed two interfaces which are the App Description and User Review pages to present Store Listing. To maintain the look and feel of the standard Google Play Store Listing, the only modification that we made was the inclusion of the privacy rating in the App Description page, replacing the app's screenshots. The User Review page contained the app's overall user rating, and individual user rating and review; it was designed to provide participants with more information about the app before installation. We varied the privacy rating on the App Description page in either a positive or negative way while controlling other app attributes to be identical in all experimental conditions.

Figure 2. Privacy Protection Rating in positive frame (left) and Privacy Risk Rating in negative frame (right) used in the experiment. The two ratings are semantically equivalent.

Adopting the concept of privacy meter suggested by Kang et al. (2015), we used a slider bar to represent our privacy rating. As shown in Figure 2, our rating slider bar presents a unipolar scale depicting increasing degrees of privacy (or risks) in a positive (or negative) frame. Our privacy rating exhibits privacy as a percentage. In a positive frame, higher percentage represents more privacy and lower percentage represents less privacy. In a negative frame, higher percentage represents less privacy and lower percentage represents more privacy. MFLB was designed to be a low privacy rated app. To represent the low level of privacy of MFLB, we used a 30% Privacy Protection Rating for the positive frame and a 70% Privacy Risk Rating for the negative frame. Both the positive and negative ratings are semantically equivalent. In computer interface design, effective use of colour can help users form a good mental model (Wright et al., 1997) and the red colour is associated with warnings (Braun and Silver, 1995; Wang et al., 2011). Therefore, we created our privacy rating slider bar in red to alert users to the low level of privacy of the app. In the experiment, we assumed that MFLB had been assessed by privacy experts based on the number of dangerous permissions requested by MFLB, and the assessment results were reflected by the privacy rating. This assumption was communicated to the participants.

4.4 Apparatus

The experiment and main questionnaire were hosted on PsyToolkit (www.psytoolkit.org), a toolkit commonly used in cognitive and psychological research for developing and running experiments and questionnaires online (Stoet, 2017, 2010). We programmed our online study using PsyToolkit scripts and tested the study before publishing it for pilot and real data collection. Each experimental condition has its own designated PsyToolkit online study page. To ensure all conditions have equal sample size, we created a script to sequentially allocate participants to one of the four conditions in the order in which they accessed our main study website. The follow-up questionnaire was conducted using Google Forms, a popular online tool for creating and analysing surveys.

4.5 Procedure

The study was conducted online. All participants who voluntarily took part in the study went through two experimental tasks and a main questionnaire. First, participants recruited received a link to our main study website. After clicking on the link, they were automatically directed to a specific study page hosted on PsyToolkit, depending on which of the four experimental conditions they were assigned to. Each experimental condition had its own study page, resulting in four specific study pages. On the introduction screen of each page, participants were informed that the study was concerned with mobile privacy. Other related information such as requirements, tasks, instructions, approximate completion time and researcher's contact information, was also presented. Before starting the study, participants were provided with the following reminders. Such emphases were to prevent effects from external factors irrelevant to the context of mobile privacy, which may influence the subjective generation experience in the experiment:

- *Close other browser tabs on the computer or switch off the phone or anything else distracting, e.g. music or chat software, television, etc.* (This is to avoid attributing the thought generation experience to distraction.)
- *Avoid using mobile devices to complete the study. We expected the difficulty of entering text input using a smartphone or tablet and thus did not recommend the use of mobile devices in this study*. (This is to avoid attributing the experience to input difficulty.)
- *Complete the study in one sitting*. (This is to avoid breaks between the experimental tasks.)

In addition, our experiment was set up in such a way that it only recorded the answer on the first attempt. Furthermore, participants were made aware that their participation was entirely voluntary and they had the option to withdraw from the study at any time or omit any question during the process. Detailed information about data confidentiality and anonymity was also explained. After giving informed consent to take part in the study, participants were directed to the experiment.

4.5.1 Experiment

The experiment consisted of two consecutive tasks. In Task 1, participants in all four conditions were presented with the identical general statements below about privacy in mobile apps.

> *"There are many mobile apps now which collect information about users such as location, contacts, etc. Therefore, there is a risk that the data will be shared with other parties."*

Next, based on the assigned condition, each participant was required to generate either 2 (high FLUENCY) or 10 (low FLUENCY) privacy worries when using a mobile app with a low privacy rating. The participants in High-FLUENCY & Pos-FRAMING and High-FLUENCY & Neg-FRAMING conditions were asked to provide 2 privacy worries whereas the participants in Low-FLUENCY & Pos-FRAMING and Low-FLUENCY & Neg-FRAMING conditions were asked to list 10 worries. The following instruction was given to the participants who were required to list 2 worries concerning mobile privacy:

> *"Please list TWO (2) things that you would worry about concerning using a mobile app with low privacy rating."*

Participants entered their responses into the text boxes provided. Given the possibility of requiring help from other sources which may affect the accessibility experience, participants were reminded to provide answers based on their own knowledge without looking up the information online. After all the text boxes were filled, participants proceeded to Task 2.

In Task 2, participants were presented with our fictional app MFLB with a privacy rating framed either positively or negatively and an accompanying scenario. MFLB consisted of two pages, the App Description page (Figure 1 (left) for positive frame) and User Review page (Figure 1 (right)). For the App Description page, the participants in High-FLUENCY & Pos-FRAMING and Low-FLUENCY & Pos-FRAMING conditions saw the 30% Privacy Protection Rating whereas the participants in High-FLUENCY & Neg-FRAMING and Low-FLUENCY & Neg-FRAMING conditions saw the 70% Privacy Risk Rating. In all conditions, participants were presented with the same User Review page. Below the app interface, participants were given the scenario which read as follows:

> *"Suppose you are learning a new language and will be sitting for a speaking test soon. You need a language learning app which can enable you to practise your speaking skills. You come across an app called My Fun Language Buddy (MFLB) as shown above in a mobile app market place.*
>
> *Figure 1 shows the app's description. Figure 2 shows the app's user ratings and reviews. MFLB teaches ten languages including the one that you are learning and offers users the ability to connect with native language speakers.*
>
> *An assessment of the app's privacy was conducted by privacy experts. A privacy protection rating (as shown in Figure 1) was given based on the number of dangerous permissions (e.g. calendar, contacts, location, etc.) requested by the app that could potentially share your data with other third parties."*

In the above scenario, Figure 1 indicated the App Description page and Figure 2 indicated the User Review page as they appeared in the online study. Subsequently, participants rated the likelihood of installing the app on a 10-point scale, numbered from 1 (definitely would not install) to 10 (definitely would install). This rating was the dependent variable in our study. We also recorded the time each participant spent on this task to observe any abnormally long delay in response, indicating the possible occurrence of any side effects between the thought generation and app-install decision-making tasks that may render generation fluency uninformative for the judgment.

After completing the two tasks, participants were asked to answer three questions relating to the tasks that they had performed. The questions required participants to assess the usefulness of the app, their interest in the app, and the difficulty of generating the requested number of privacy-related worries along 10-point scales, with 1 representing not at all useful, not at all interested, or not at all difficult and 10 representing very useful, very interested, or very difficult. The questions provided direct measures of perceived difficulty in generating privacy worries, interest in app, and perceived app usefulness (USEFUL), respectively.

4.5.2 Main Questionnaire

Following the experiment, participants completed a questionnaire comprising two parts: (1) experiences with mobile devices, apps and privacy, and (2) basic demographics.

In the first part of the questionnaire, participants answered a series of seven closed-ended questions pertaining to mobile device and app usage which included the type of mobile operating systems used, the average number of mobile apps used per day, and the frequency of installing a new mobile app. Next, using 10-point scales, participants rated their expertise in mobile privacy and how concerned they were about the permissions an app has access to when downloading the app. After that, each participant was required to state whether or not they had experienced a violation of privacy, e.g. disclosure of personal data without their consent. This question was designed to help us to gain insights into participants' perceived self-relevance of privacy risks. We assume that a privacy violation experience would render privacy risks more personally relevant. Specifically, we are interested in finding out whether such a perception of self-relevance will elicit systematic processing in mobile app privacy risk judgments.

In the second part of the main questionnaire, participants provided basic demographic information including gender, age, occupation, educational background. In addition, participants were asked to assess their English proficiency. Previous work (Schwarz et al., 1991) suggested that attributing the accessibility experience to irrelevant sources influences the impact of subjective experience on forming judgments. Considering participants may attribute the difficulty of generation to English proficiency instead of app privacy risks, the questions pertaining to English proficiency were created for data screening purpose to ensure that participants were proficient at performing the thought generation task. After completing the questionnaire, participants were debriefed about the objective of the study and thanked. As a token of appreciation for completing the study, participants could choose to be entered into a prize draw to win a £10 Amazon Gift Voucher.

4.5.3 Follow-up Questionnaire

This questionnaire, containing four open-ended questions, was designed to gain a deeper understanding of the salience of our proposed privacy rating and to investigate the reasons behind the app-install decision making at the individual level. Participants who were willing to provide feedback on their responses to the experiment received a link to our online follow-up questionnaire. The participation was voluntary. First, a screenshot of the privacy rating the participants saw in the experiment was presented as a reminder. In the experiment, participants rated whether they would install the MFLB app, found the app useful, and were interested in the app. Based on their responses earlier, each participant was asked to give reasons for their ratings. Finally, they indicated whether the privacy rating shown in MFLB was helpful and provided reasons for their answers.

5. Results

5.1 Experiment and Main Questionnaire

5.1.1 App Experience and Privacy Concerns

We performed a data screening process to filter out participants who did not meet the study requirements as well as inconsistent responses concerning generation fluency and perceived difficulty. In such responses, participants indicated that listing 10 privacy worries as "not difficult at all" but failed to complete the list. After data screening, there were 91 participants (High-FLUENCY & Pos-FRAMING: 21, High-FLUENCY & Neg-FRAMING: 23, Low-FLUENCY & Pos-FRAMING: 25, Low-FLUENCY & Neg-FRAMING: 22).

The majority of participants (55 of 91) used 1 to 5 apps per day. In terms of how often participants install a new app on their mobile device, it was found that 46 of them installed a new app once per month, 26 installed a new app once per 3 months, 13 installed a new app once per 6 months or less often, and only 6 installed a new app at least one per week or more often. Nineteen participants reported that they had experienced a violation of privacy. Further, participants rated their IT knowledge (Mean = 6.36), expertise concerning mobile privacy (Mean = 4.81) and privacy concern ('*In general, when deciding to install an app, how concerned are you about the permissions the app has access to?*', Mean = 6.47) along 10-point scales.

5.1.2 Manipulation Check

We used participants' self-reported perceived difficulty for the generation fluency (FLUENCY) manipulation check. A two-way Analysis of Variance (ANOVA) examined the effectiveness of the FLUENCY manipulation. Consistent with our pilot study (Tay, 2017), this study found a significant difference in the perception of difficulty in generating privacy worries, $F(1, 87) = 84.55$, $p < .001$, partial $\eta^2 = .493$. As expected, the ANOVA revealed that generating 10 worries ($M = 7.663$, $SE = .310$) was experienced as more difficult than generating 2 worries ($M = 3.562$, $SE = .320$), indicating successful manipulation of FLUENCY. Table 2 shows the mean perception of difficulty in generating privacy worries in each experimental condition, demonstrating that across all conditions, participants found it more difficult to generate 10 (low FLUENCY) rather than 2 (high FLUENCY) privacy worries.

Table 2. Mean perception of difficulty in generating privacy worries concerning mobile apps in all experimental conditions.

Number of Privacy Worries Generated	**Privacy Framing (FRAMING)**	
	Positive	*Negative*
2	3.429	3.696
10	7.280	8.045

***Note**: Perceived difficulty was measured on a 10-point scale (1 = not at all difficult; 10 = very difficult).*

5.1.3 App-install Decision

Figure 3 shows the mean INSTALL of all the four conditions. The results of the two-way Analysis of Covariance (ANCOVA) after controlling for the effects of USEFUL are shown in Table 3. The 2 (FLUENCY: high, low) x 2 (FRAMING: positive, negative) ANCOVA used an alpha level of 0.05.

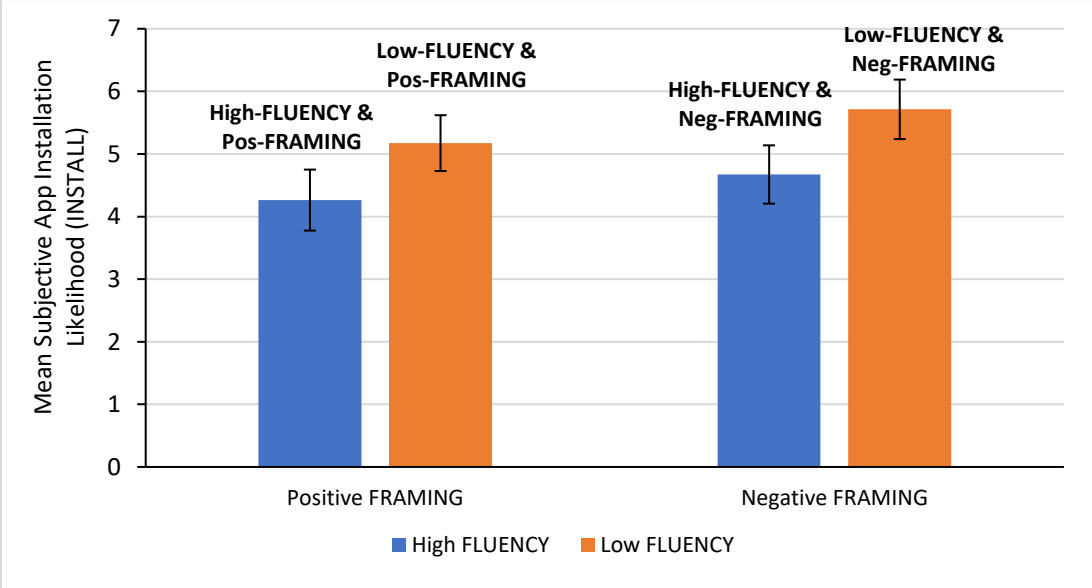

Figure 3. Mean Subjective App Installation Likelihood (INSTALL) per condition, with Perceived App Usefulness (USEFUL) as the covariate. INSTALL was measured on a 10-point scale (*1 = definitely would not install to 10 = definitely would install*). Error bars represent standard errors.

Table 3. ANCOVA summary table.

	SS	*df*	*MS*	*F*	*p*	η^2
USEFUL	130.805	1	130.805	26.326	<.001	.234
FLUENCY	21.565	1	21.565	4.340	.040	.048
FRAMING	5.069	1	5.069	1.020	.315	.012
FLUENCY x FRAMING	.097	1	.097	.020	.889	<.001

***Note**: SS = Type III Sum of Squares, MS = Mean Square, η^2 = Partial Eta Squared.*

Test of H1: Generation Fluency (FLUENCY) Effect

As shown in Table 3, ANCOVA results revealed a statistically significant difference in INSTALL between the high and low FLUENCY groups while adjusting for USEFUL, F (1, 86) = 4.340, p = .040, partial η^2 = .048. Figure 3 shows that, participants who had to generate 10 privacy worries (Low FLUENCY) rated a higher INSTALL (M = 5.173, SE = .446 for Low-FLUENCY & Pos-FRAMING and M = 5.713, SE = .475 for Low-FLUENCY & Neg-FRAMING) than those who generated 2 privacy worries (M = 4.263, SE = .487 for High-FLUENCY & Pos-FRAMING and M = 4.672, SE = .466 for High-FLUENCY & Neg-FRAMING). This FLUENCY main effect reflected a general trend for both positive and negative framing. Hence, H1 was supported.

Test of H2: Privacy Framing (FRAMING) Effect

ANCOVA results in Table 3 revealed that there was no statistically significant difference in INSTALL between positive and negative framing, F (1, 86) = 1.020, p = .315. Figure 3 shows that when FLUENCY is low, participants who saw a positively-framed privacy rating rated a lower INSTALL (M = 5.173 for Low-FLUENCY & Pos-FRAMING) than those who saw a negatively-framed privacy rating (M = 5.713 for Low-FLUENCY & Neg-FRAMING). The same trend can be found in high FLUENCY conditions: (M = 4.263 for High-FLUENCY & Pos-FRAMING) and (M = 4.672 for High-FLUENCY & Neg-FRAMING). Despite this trend, no significant main effect of FRAMING was detected, and therefore H2 was rejected.

As illustrated in Figure 3, the INSTALL rating was the highest in the Low-FLUENCY & Neg-FRAMING condition in which generation was experienced as difficult and negative framing was employed. However, the interaction effect between FLUENCY and FRAMING on INSTALL was not significant, F (1, 86) = .020, p = .889, as shown in Table 3.

Perceived App Usefulness (USEFUL)

Participants' self-reported perceived usefulness was used as the covariate, USEFUL. The two-way ANCOVA showed that the covariate of USEFUL was significantly correlated with the dependent variable of INSTALL, F (1, 86) = 26.326, p < .001, partial η^2 = .234 (see Table 3). This was further supported by a Pearson correlation coefficient which revealed a positive correlation between the two variables, r(89) = .48, p < .001, indicating that as USEFUL increased, INSTALL also increased.

Other Factors

In addition to the main effects of FLUENCY and USEFUL, we found that the 72 participants who did not have a privacy violation experience (M = 5.29, SD = 2.4) rated that they were more likely to install a low privacy rated app than the 19 with previous privacy violation experience (M = 3.74, SD = 2.786), as assessed by an independent samples t-test, M = 1.555, 95% CI [0.283, 2.827], t(89) = 2.428, p = .017.

To further examine if the effect of FLUENCY on INSTALL depended on participants' prior experience of privacy violation, we conducted a two-way ANCOVA, with FLUENCY and previous experience of privacy violation as the independent variables and USEFUL as the covariate. The results showed that there was no significant interaction effect between FLUENCY and privacy violation experience, F (1, 86) = 2.147, p = .147. Consistent with the general fluency findings, participants who had experienced a privacy violation rated that they would more likely install a low privacy rated app after generating 10 (M = 4.910, SE = .651, N = 11) rather than 2 worries (M = 2.575, SE = .757, N = 8).

Surprisingly, we found that there was only a slight difference in the reported privacy concern between participants who had prior privacy violation experience ($M = 7.11$, $SD = 1.883$, $N = 19$) and those who did not have such experience ($M = 6.31$, $SD = 2.147$, $N = 72$), and this difference was not significant, as assessed by an independent samples t-test, $M = .800$, 95% CI [-0.275, 1.874], $t(89) = 1.479$, $p = .143$. In addition, it was found that INSTALL was positively related to users' levels of interest in the app, $r(89) = .52$, $p < .001$, as evaluated by a Pearson correlation coefficient.

5.2 Follow-up Questionnaire

In the experiment and main questionnaire, participants rated the likelihood of installing MFLB, the usefulness of MFLB, and how interested they were in MFLB. The follow-up questionnaire required participants to give reasons for these ratings and express their perception of effectiveness of our app privacy rating. The responses collected from 10 participants (High-FLUENCY & Pos-FRAMING: 3, High-FLUENCY & Neg-FRAMING: 2, Low-FLUENCY & Pos-FRAMING: 3, Low-FLUENCY & Neg-FRAMING: 2) who voluntarily partook in this questionnaire were analysed.

5.2.1 Demographics

Of the 91 qualified participants in the experiment, 10 voluntarily completed the follow-up questionnaire. These participants' ages ranged from 25 to 54 (Median: 25 to 34). 5 of them were females, 3 were males and 2 chose not to disclose their gender. Four participants reported that they had experienced a privacy violation. Seven participants were postgraduates studying at the universities across the UK whereas the remaining participants were employed or self-employed individuals. They were from diverse educational and/or professional backgrounds, including Computer Science, Engineering, Biology, Environmental Science, Digital Marketing, Management, Education, and Museum Studies. We found that participants had a mean of 7.3, 6.0, and 6.9 for IT knowledge, mobile privacy expertise, and privacy concern, respectively. The three items were measured using 10-point scales, with higher values representing higher degrees or magnitudes.

5.2.2 Qualitative Analysis

Subjective App Installation Likelihood

The mean of 4.0 revealed that the ten participants did not show high intention to install the app. The analysis of the collected responses showed that app privacy was evidently the major consideration for the majority (eight participants) when making the app-install decision. This indicates that users are likely to attend to privacy information during an app-install decision-making process if the information is made available. Other considerations included app security, user reviews, personal interests and comparisons with similar apps.

App Usefulness

The average perceived app usefulness was 7.3. The attributes that contributed to perceived app usefulness were identified and could be broadly classified into four categories: (1) user reviews and recommendations (2) app features, functionalities, and capabilities (3) convenience and (4) cost.

Interest in App

The reported level of interest had a mean of 6.2, implying that generally most of the participants were not highly interested in the app. After analysing the answers of participants, we found that the interest level was closely related to the purpose of having the app, the magnitude of the need for the app, and personal interests.

Perceived Effectiveness of App Privacy Rating

We queried participants if they found our app privacy rating, i.e. 30% Privacy Protection and 70% Privacy Risk ratings presented in MFLB, helpful and to provide the reasons for their answers. One of the participants answered the question from the viewpoint of an app developer rather than an app user, and thus the response was excluded from this discussion. All the remaining participants perceived the privacy rating as helpful. Four participants said that the privacy rating conveys the privacy level of the app in a simpler and clearer way. For example, P7 stated "Yes, this is the easiest way for users to identify the level of privacy.", and P1 who saw a 30% Privacy Protection Rating, explained that 30% of app's privacy protection was comprehensible. Similarly, P6 who was presented with the 30% Privacy Protection Rating, stated "Yes, simplified way to understand privacy level 30%, the red colour indicates a threat level, most people understand the traffic light representation."

P10 mentioned that the privacy rating was very helpful because it reflected the results of the privacy assessment conducted by privacy experts as stated in our online study. There were also participants who provided reasons relating to their own privacy behaviours and attitudes. For example, P8 mentioned that the privacy rating was helpful as it would remind him or her to be more cautious when installing an app. On the other hand, P9 stated "Yes, although I haven't thought a lot about the risk of my information being shared in the past." These findings suggested that an effective privacy communication tool can enable users to attend to and interpret privacy information during an app-install decision-making process.

6. Discussion and Conclusions

Our study generally confirms our view of the app-installation decision as a heuristic process, recruiting knowledge to weigh costs against benefits. Our quantitative and qualitative findings support the suggestion that an app's functionalities, user reviews and recommendations, and the cost of the app, are all factors that might influence this process. In this section, we summarise and discuss the key findings with respect to our research questions.

RQ1. Can risk-aversion be activated or de-activated in an experimental situation, by a simple generation fluency manipulation?

The major novel finding of this study is that users' risk perception can be influenced by the subjective experience of ease or difficulty of bringing privacy worries or potential risks to mind. This finding is interesting, not only because it suggests that fluency per se might usefully be manipulated during the application-installation process, with practical benefits, but because it shows, more generally, how risk perception is volatile, in that it might be dependent on subjective states that are transient and particular to the moment of the installation process.

Previous study (Rothman and Schwarz, 1998) suggested that one's adoption of judgment strategy may depend on the perceived self-relevance of a judgment task. According to this idea, people adopt heuristic processing strategy and base their judgments on retrieval fluency when the task is not considered self-relevant. In contrast, for people who consider the task as personally relevant, they will opt for systematic processing strategy. However, our study did not discover a relationship between perceived self-relevance and the adoption of judgment strategy. We found that participants who had experienced a privacy violation expressed higher intentions of app installation after generating 10 rather than 2 privacy worries, indicating that the app-install decision was based on the generation fluency experience despite the presence of a privacy violation experience. Hence, we suspect that the degree of perceived self-relevance of privacy risks varies in the extent to which an individual has been affected by the privacy violation. These subjective factors of how users perceive self-relevance of privacy risks may include the severity and impact of the violation experienced, time elapsed since the violation, type of privacy violation experienced, etc.

As a result, we conjecture that in our study, participants who had experienced a privacy violation exhibited a lower degree of perceived self-relevance, insufficient to eliminate the effects of generation fluency. This assumption is supported by our study results which showed that participants who had experienced a privacy violation did not appear more concerned about the permissions requested by an app than those without a violation experience, suggesting that such experience may have not significantly affected their privacy attitude. Given that our study primarily focused on examining the effects of generation fluency on individuals' app-install decision, a more direct test would be needed to assess if the degree of perceived self-relevance affects users' adoption of judgment strategy, i.e. heuristic or systematic strategy, and consequently shifts their app-install decision in the opposite direction.

RQ2. Does positive versus negative framing of risk information affect users' perception of risks associated with a low privacy rated mobile app?

Despite the fact that positive framing has previously been found to be more effective in safeguarding users against low privacy rated apps (Choe et al., 2013), we did not find a significant framing effect. As shown in Figure 3, the findings of our study discovered that generally, participants did not express a high intention to install the low privacy rated app presented. This suggests that both our positively and negatively framed privacy ratings are effective in conveying privacy information to participants. This is also supported by the qualitative analysis, in which participants reported that the privacy ratings in both frames facilitated easy interpretation of the app's privacy protection or risk level. We would not like to draw many conclusions from our non-replication of a framing effect. Our study, after all, used only a single fictional app, with a single privacy-rating level.

RQ3. Do post hoc reports confirm that users attend to privacy information during an app-install decision-making process when the information is made available alongside other app attributes?

To address RQ3, we conducted a small-scale qualitative exploration of the motivations behind the app-install decisions made by participants using a follow-up questionnaire. The post hoc analysis found that most participants paid attention to privacy if this information was brought to them when making an app-install decision. This suggests that presenting privacy information in a simpler and salient way enables users to consider privacy as part of the decision-making process, and hence protecting them from privacy-invasive apps. It is worth stressing that our main interest was not to design an approach to replace the current permission dialogue and privacy policies employed by mobile app marketplaces. In fact, we sought to complement the existing permission system by adding a privacy rating to the app's primary screen to assist users in forming a privacy perception at a glance. Because most users have problems understanding privacy policies and permissions, and tend to make judgments heuristically when it comes to installing a new mobile app, we believe that a simple privacy rating is helpful in reducing users' cognitive load in understanding the complexity of apps' privacy practices, and hence preventing superficial decision making.

Given the nature of online studies, the findings obtained might be subject to external factors such as participants' concentration, emotions, and feelings while they completed the experiment and questionnaires. To control for external effects, we made our best efforts to provide participants with reminders, e.g. turning off devices that could cause distractions, prior to performing the study. In addition, we required participants to not seek assistance from other sources, e.g. search engines, in the task of generating privacy worries. Nevertheless, future studies could be carried out in more controlled settings to optimise the study of generation fluency. Furthermore, considering that the sample of our study was primarily comprised of university students and young people with moderate app usage and installation frequency, our study could be developed in terms of sampling method to yield a larger and more representative sample.

Our study contributes to the existing literature of privacy risk perceptions in mobile app-install decision-making by demonstrating that one's risk perception can be changed by the experienced ease or difficulty of generating concerns during an app-install decision-making process. We note that users are more likely to install a low privacy rating app if the generation of privacy concerns is experienced as difficult. This finding provides directions for future research on approaches to improving risk perceptions among mobile users during an app-install decision-making process by decreasing their perceived difficulty of generating privacy concerns. For example, future studies could design optimal methods to make privacy concerns vivid and easy to bring to mind (e.g. through frequent exposure), so that such information could be easily accessible at the time of app-install decision-making. We believe that this idea could complement privacy information communication tools proposed in earlier studies in helping users to make more privacy-aware decisions.

Declaration of Competing Interest

None.

Funding

Not applicable.

Acknowledgments

This article is based on an MSc thesis submitted to the University of Bath by the first author, under the supervision of the third author.

References

Almuhimedi, H., Schaub, F., Sadeh, N., Adjerid, I., Acquisti, A., Gluck, J., Cranor, L.F., Agarwal, Y., 2015. Your Location Has Been Shared 5,398 Times!: A Field Study on Mobile App Privacy Nudging, in: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15. ACM, New York, NY, USA, pp. 787–796. https://doi.org/10.1145/2702123.2702210

Bonné, B., Peddinti, S.T., Bilogrevic, I., Taft, N., 2017. Exploring decision making with Android's runtime permission dialogs using in-context surveys. Presented at the Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017), pp. 195–210.

Braun, C.C., Silver, N.C., 1995. Interaction of signal word and colour on warning labels: differences in perceived hazard and behavioural compliance. Ergonomics 38, 2207–2220. https://doi.org/10.1080/00140139508925263

Chaiken, S., Liberman, A., Eagly, A.H., 1989. Heuristic and systematic information processing within and beyond the persuasion context, in: Unintended Thought. Guilford Press, New York, NY, US, pp. 212–252.

Choe, E.K., Jung, J., Lee, B., Fisher, K., 2013. Nudging People Away from Privacy-Invasive Mobile Apps through Visual Framing, in: Kotzé, P., Marsden, G., Lindgaard, G., Wesson, J., Winckler, M. (Eds.), Human-Computer Interaction – INTERACT 2013, Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 74–91.

Coates, S.L., Butler, L.T., Berry, D.C., 2006. Implicit memory and consumer choice: the mediating role of brand familiarity. Applied Cognitive Psychology 20, 1101–1116. https://doi.org/10.1002/acp.1262

Duchon, D., Dunegan, K.J., Barton, S.L., 1989. Framing the problem and making decisions: the facts are not enough. IEEE Transactions on Engineering Management 36, 25–27. https://doi.org/10.1109/17.19979

Enck, W., Gilbert, P., Chun, B.-G., Cox, L.P., Jung, J., McDaniel, P., Sheth, A.N., 2010. TaintDroid: An Information-flow Tracking System for Realtime Privacy Monitoring on Smartphones, in: Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation, OSDI'10. USENIX Association, Berkeley, CA, USA, pp. 393–407.

Felt, A.P., Chin, E., Hanna, S., Song, D., Wagner, D., 2011. Android Permissions Demystified, in: Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS '11. ACM, New York, NY, USA, pp. 627–638. https://doi.org/10.1145/2046707.2046779

Future of Privacy Forum, 2016. FPF Mobile Apps Study [WWW Document]. URL https://fpf.org/wp-content/uploads/2016/08/2016-FPF-Mobile-Apps-Study_final.pdf (accessed 12.1.16).

Gates, C.S., Chen, J., Li, N., Proctor, R.W., 2014. Effective Risk Communication for Android Apps. IEEE Transactions on Dependable and Secure Computing 11, 252–265. https://doi.org/10.1109/TDSC.2013.58

Gu, J., Xu, Y. (Calvin), Xu, H., Zhang, C., Ling, H., 2017. Privacy concerns for mobile app download: An elaboration likelihood model perspective. Decision Support Systems 94, 19–28. https://doi.org/10.1016/j.dss.2016.10.002

Harbach, M., Hettig, M., Weber, S., Smith, M., 2014. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '14. ACM, New York, NY, USA, pp. 2647–2656. https://doi.org/10.1145/2556288.2556978

Harris, M.A., Brookshire, R., Chin, A.G., 2016. Identifying factors influencing consumers' intent to install mobile applications. International Journal of Information Management 36, 441–450. https://doi.org/10.1016/j.ijinfomgt.2016.02.004

Information Commissioner's Office, n.d. What is personal data? [WWW Document]. URL https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/key-definitions/what-is-personal-data/ (accessed 11.27.18).

Johnson, E.J., Bellman, S., Lohse, G.L., 2002. Defaults, Framing and Privacy: Why Opting In-Opting Out1. Marketing Letters 13, 5–15. https://doi.org/10.1023/A:1015044207315

Kang, J., Kim, H., Cheong, Y.G., Huh, J.H., 2015. Visualizing Privacy Risks of Mobile Applications through a Privacy Meter, in: Lopez, J., Wu, Y. (Eds.), Information Security Practice and Experience, Lecture Notes in Computer Science. Springer, Cham, pp. 548–558. https://doi.org/10.1007/978-3-319-17533-1_37

Kelley, C.M.., Lindsay, D.S., 1993. Remembering Mistaken for Knowing: Ease of Retrieval as a Basis for Confidence in Answers to General Knowledge Questions. Journal of Memory and Language 32, 1–24. https://doi.org/10.1006/jmla.1993.1001

Kelley, P.G., Cranor, L.F., Sadeh, N., 2013. Privacy As Part of the App Decision-making Process, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13. ACM, New York, NY, USA, pp. 3393–3402. https://doi.org/10.1145/2470654.2466466

Levin, I.P., Gaeth, G.J., 1988. How consumers are affected by the framing of attribute information before and after consuming the product. Journal of Consumer Research 15, 374–378. https://doi.org/10.1086/209174

Levin, I.P., Jasper, J.D., Gaeth, G.J., 1996. Measuring the Effects of Framing Country-Of-Origin Information: a Process Tracing Approach. NA - Advances in Consumer Research 23, 385–389.

Levin, I.P., Schneider, S.L., Gaeth, G.J., 1998. All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects. Organizational Behavior and Human Decision Processes 76, 149–188. https://doi.org/10.1006/obhd.1998.2804

Liccardi, I., Pato, J., Weitzner, D.J., 2014a. Improving User Choice Through Better Mobile Apps Transparency and Permissions Analysis. Journal of Privacy and Confidentiality 5. https://doi.org/10.29012/jpc.v5i2.630

Liccardi, I., Pato, J., Weitzner, D.J., Abelson, H., De Roure, D., 2014b. No Technical Understanding Required: Helping Users Make Informed Choices About Access to Their Personal Data, in: Proceedings of the 11th International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services, MOBIQUITOUS '14. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, pp. 140–150. https://doi.org/10.4108/icst.mobiquitous.2014.258066

Novemsky, N., Dhar, R., Schwarz, N., Simonson, I., 2007. Preference Fluency in Choice. Journal of Marketing Research 44, 347–356. https://doi.org/10.1509/jmkr.44.3.347

Rothman, A.J., Schwarz, N., 1998. Constructing Perceptions of Vulnerability: Personal Relevance and the Use of Experiential Information in Health Judgments. Pers Soc Psychol Bull 24, 1053–1064. https://doi.org/10.1177/01461672982410003

Schaub, F., Balebako, R., Durity, A.L., Cranor, L.F., 2015. A Design Space for Effective Privacy Notices. Presented at the Eleventh Symposium On Usable Privacy and Security ({SOUPS} 2015), {USENIX} Association, pp. 1–17.

Schoorman, F.D., Mayer, R.C., Douglas, C.A., Hetrick, C.T., 1994. Escalation of Commitment and the Framing Effect: An Empirical Investigation. Journal of Applied Social Psychology 24, 509–528. https://doi.org/10.1111/j.1559-1816.1994.tb00596.x

Schwarz, N., 2004. Metacognitive Experiences in Consumer Judgment and Decision Making. Journal of Consumer Psychology 14, 332–348. https://doi.org/10.1207/s15327663jcp1404_2

Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., Simons, A., 1991. Ease of retrieval as information: Another look at the availability heuristic. Journal of Personality and Social Psychology 61, 195–202. https://doi.org/10.1037/0022-3514.61.2.195

Shen, G.C.-C., 2015. Users' adoption of mobile applications: Product type and message framing's moderating effect. Journal of Business Research 68, 2317–2321. https://doi.org/10.1016/j.jbusres.2015.06.018

Sjöberg, L., Moen, B.-E., Rundmo, T., 2004. Explaining risk perception. An evaluation of the psychometric paradigm in risk perception research. Norwegian University of Science and Technology.

Stoet, G., 2017. PsyToolkit: A Novel Web-Based Method for Running Online Questionnaires and Reaction-Time Experiments. Teaching of Psychology 44, 24–31. https://doi.org/10.1177/0098628316677643

Stoet, G., 2010. PsyToolkit: A software package for programming psychological experiments using Linux. Behavior Research Methods 42, 1096–1104. https://doi.org/10.3758/BRM.42.4.1096

Tay, S.W., 2017. The Role of Risk Perception and Framing on the Use of Privacy Information in Mobile Application Download Decisions. The University of Bath.

Thompson, E., 2017. App Annie 2016 Retrospective — Mobile's Continued Momentum [WWW Document]. App Annie Content. URL https://www.appannie.com/en/insights/market-data/app-annie-2016-retrospective/ (accessed 3.31.17).

Tversky, A., Kahneman, D., 1981. The framing of decisions and the psychology of choice. Science 211, 453–458. https://doi.org/10.1126/science.7455683

Wang, N., Xu, H., Grossklags, J., 2011. Third-party Apps on Facebook: Privacy and the Illusion of Control, in: Proceedings of the 5th ACM Symposium on Computer Human Interaction for Management of Information Technology, CHIMIT '11. ACM, New York, NY, USA, pp. 4:1–4:10. https://doi.org/10.1145/2076444.2076448

Westin, A.F., 1968. Privacy and freedom. Washington and Lee Law Review 25, 166.

Wright, P., Mosser-Wooley, D., Wooley, B., 1997. Techniques & Tools for Using Color in Computer Interface Design. XRDS 3, 3–6. https://doi.org/10.1145/270974.270976

Zhang, B., Xu, H., 2016. Privacy Nudges for Mobile Applications: Effects on the Creepiness Emotion and Privacy Attitudes, in: Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing, CSCW '16. ACM, New York, NY, USA, pp. 1676–1690. https://doi.org/10.1145/2818048.2820073

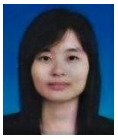

Siok Wah Tay received the B.Sc. degree in Security Technology from Multimedia University, Malaysia, and the M.Sc. degree in Human-Computer Interaction from the University of Bath, UK. She is currently pursuing a Ph.D. degree in Computer Science at the University of Manchester, UK. Her research interests include security, the Internet of Things and human-computer interaction.

Pin Shen Teh received the Ph.D. degree in Computer Science from University of Manchester, UK in 2018. He is currently with the Department of Operations, Technology, Events and Hospitality Management, Manchester Metropolitan University, UK. His research interests include security, biometrics, pattern recognition, machine learning and mobile authentication.

Stephen Payne is professor of human-centric systems in the Department of Computer Science at the University of Bath. Before moving to Bath, Payne was a Professor of Psychology in Cardiff University and (briefly) a Professor in Manchester Business School. Previously he worked at IBM T.J. Watson Research Center. Payne has worked on cognitive approaches to Human-Computer Interaction since his PhD (1985), on Task-Action Grammars.