

How Short is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices

Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib,
Norman Sadeh, Lorrie Faith Cranor, Yuvraj Agarwal
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA

}@andrew.cmu.edu

ABSTRACT

Privacy policies are often too long and difficult to understand, and are therefore ignored by users. Shorter privacy notices with clearer wording may increase users' privacy awareness, particularly for emerging mobile and wearable devices with small screens. In this paper, we examine the potential of (1) shortening privacy notices, by removing privacy practices that a large majority of users are already aware of, and (2) highlighting the implications of described privacy practices with positive or negative framing. We conducted three online user studies focused on privacy notice design for fitness wearables. Our results indicate that short-form privacy notices can inform users about privacy practices. However, we found no effect from including positive or negative framing in our notices. Finally, we found that removing expected privacy practices from notices sometimes led to less awareness of those practices, without improving awareness of the practices that remained in the shorter notices. Given that shorter notices are typically expected to be more effective, we find the lack of increased awareness of the practices remaining in the notice surprising. Our results suggest that the length of an effective privacy notice may be bounded. We provide an analysis of factors influencing our participants' awareness of privacy practices and discuss the implications of our findings on the design of privacy notices.

1. INTRODUCTION

The purpose of a privacy policy is to make users aware of a system's or company's practices related to collection, sharing, use, and storage of personal information. In theory, a company's privacy policy contains all the information that users need to be aware of a company's privacy practices and to make informed decisions about which companies to entrust with their personal information. In practice, privacy policies are too long, leading to user fatigue and users ignoring privacy policies [12, 33, 40]. Recognizing this problem, the Federal Trade Commission (FTC) has called for clearer and shorter privacy notices [16].

Prior research has examined short-form privacy notices, which are condensed versions of privacy policies that include the main practices, but may remove some degree of nuance or detail. Research studies have found that standardized short-form privacy notices can increase user awareness of privacy practices [15, 28, 29]. Other research and reports have suggested that focusing privacy notices on unexpected practices may increase awareness and effective transparency, reducing the potential for user surprise, and reducing the burden on users [6, 17, 41]. Prior work has also shown that presenting information with a positive or negative framing can also change users' perceptions and awareness of privacy practices [1, 2, 3, 22]. Our research builds upon prior work, examining three important questions.

Our first research question is whether removing from notices those privacy practices that most participants already expect to occur, would lead to greater overall awareness of an organization's privacy practices. We hypothesize participants will have higher awareness of privacy practices remaining in notices, since the notices will be shorter and more focused. In addition, participants should have similar awareness of practices that were removed, as these would be practices most participants would already expect without a notice.

Our second research question examines the effect of notice framing on user awareness about privacy practices. We compare positively and negatively framed notices against a neutral baseline.

Our third research question examines the effectiveness of short-form privacy notices in the context of fitness wearables. The effectiveness of short-form notices on increasing user awareness has been shown in several contexts [29, 31]. However, while the fitness wearable companies we surveyed (Fitbit, Misfit, Jawbone) have made some attempt to use clear language in their privacy policies, none utilized short-form privacy notices at the time of our study [19, 26, 36]. Fitbit had a plain-language illustrated version, but it was still fairly long when fully expanded. We picked fitness wearables for this study given their increasing popularity [25] and the fact that they typically collect a number of privacy-sensitive data items for their functionality (e.g., detailed physical activity of the user) leading to security and privacy concerns [24].

We conducted three online user studies to analyze notice design format, participants' baseline knowledge, and notice length and framing for the Fitbit Surge watch (shown in Figure 1). We conducted the design format study to compare the effectiveness of four candidate short-form notice designs. The baseline knowledge study served to determine which privacy practices a large majority of users would already be aware of. The notice framing and length

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22–24, 2016, Denver, Colorado.



Figure 1: A Fitbit Surge Watch, which we used as a representative Fitness wearable.

study was a 3 (lengths) x 3 (framing) study, with a control condition, to answer the research questions outlined above. All studies were approved by Carnegie Mellon University’s Institutional Review Board.

The results from our design format study showed our four short-format notice designs resulted in similar awareness of privacy practices, so we chose a format loosely based on the format of Fitbit’s existing online privacy policy. The results from our second study showed a wide range of awareness rates about Fitbit’s individual privacy practices and allowed us to identify 6 practices expected by at least 85% participants to remove from the medium and short version of the policy, and an additional 7 practices expected by at least 70% of participants to also remove from the short version of the policy.

Our final study, examining the effects of short-form notice length and framing on privacy awareness, provided a number of interesting results. We found that participants in the medium short-form notice conditions were similarly aware of privacy practices as those in the long short-form notice conditions. Removing expected practices from the medium notices did not impact awareness significantly of either the removed or remaining practices. However, participants in the shortest short-form notice conditions were less aware of the practices removed only from the shortest notices, with no significant change in awareness of the practices also removed from the medium notice or those that remained. We also found no significant difference in awareness from positive or negative framing in the notices. While not finding an effect does not prove that such an effect does not exist, it does suggest that the effect, at least in this context, is likely to be small. We discuss the implications of our results at the end of this paper.

2. RELATED WORK

Here we discuss prior work on privacy notice design in three areas: short-form privacy notices, framing, and delivery methods.

2.1 Short-form Privacy Notices

It is fairly rare for individuals to read a privacy policy in its entirety. Prior work has shown two key reasons for this: the complexity of privacy policies, and their length. Privacy policies are generally written in complex legalese or are purposefully vague, making it hard for readers to understand them [12, 27]. In fact, research has shown that not only do users struggle to make sense of privacy policies, but that even experts can disagree on the meaning of certain statements [42]. In addition, prior work has suggested that an individual would have to spend 244 hours each year to read the privacy policies of websites they visit [33]. As a result, the FTC and others have called for privacy notices to be made both clearer and shorter, in order to increase comprehension [16, 17].

Prior work has shown that short-form notices summarizing the

key privacy practices of an organization can provide significant benefits to user awareness over a traditional privacy policy [28, 29]. However, including all of the relevant information in a privacy notice, even in a compact form, may still result in overly long notices, and leaving out unexpected privacy practices can hide information and impair transparency [34].

Others have suggested that focusing on unexpected practices is important for user understanding. A recent FTC staff report suggested that when “data uses are generally consistent with consumers’ reasonable expectations, the cost to consumers and business of providing notice and choice likely outweighs the benefits” [17]. Rao et al. studied mismatches between user privacy expectations and practices disclosed in privacy policies. They found that mismatches (e.g. unexpected practices) comprise a relatively small set of practices described in privacy policies, and that creating privacy notices focusing on these practices could reduce user burden [41]. Ayres and Schwartz proposed warning labels to highlight unexpected terms in contracts [6]. Ben-Sahar and Chilton found that a warning label focusing on unexpected privacy practices benefited user comprehension, although they did not find any behavioral change associated with this increase in user comprehension [9].

Layered notices, short notices that link to a full policy containing more information, may allow for the benefits of a short-form notice, as well as avoiding the appearance of hiding unexpected practices [13, 35, 37]. However, users may consent to the first layer of the notice they encounter, without delving into the following layers [34, 43].

Other work has examined the potential of using machine learning and natural language processing to extract answers to specific questions from privacy policies and display it using a web browser plugin [47, 49]. Similarly, browser plugins have been developed to display summaries of computer-readable privacy policies [14].

We seek to reach a compromise between length and inclusion of relevant information in a short-form privacy notice. Our approach is to determine the privacy practices that are unexpected by most participants, and ensure that those are included in even the shortest privacy notice, while removing practices that users generally expect. We hypothesize that doing so will provide the benefits of a shorter notice without the downsides of leaving out unexpected privacy practices or relegating them to a secondary layer.

2.2 Framing

In addition to the content of a privacy notice, the way in which privacy practices are explained can also have a major effect on users’ perception and retention of those practices. Perception of the relative importance, or sensitivity, of certain types of information can strongly affect a users’ willingness to share it. Prior work has shown that providing reasons for privacy practices [44, 45], or communicating risks and implications [20], can grab users’ attention, change their level of concern over practices, and cause them to reflect on privacy practices more deeply. Research has shown that including personal examples, such as the number of data accesses associated with mobile permissions, can lead to even greater concern, and therefore reflection [5, 7, 22].

Framing can also ease users’ concerns over privacy. Studies have found that framing notices with more positive, misleading, or misdirecting statements can direct users’ attention away from the implications of privacy practices, and thus decrease their awareness of these practices [1, 2, 3].

2.3 Delivery Methods

There has been substantial prior work examining the way in which privacy notices are delivered, including the timing [8], channel, and

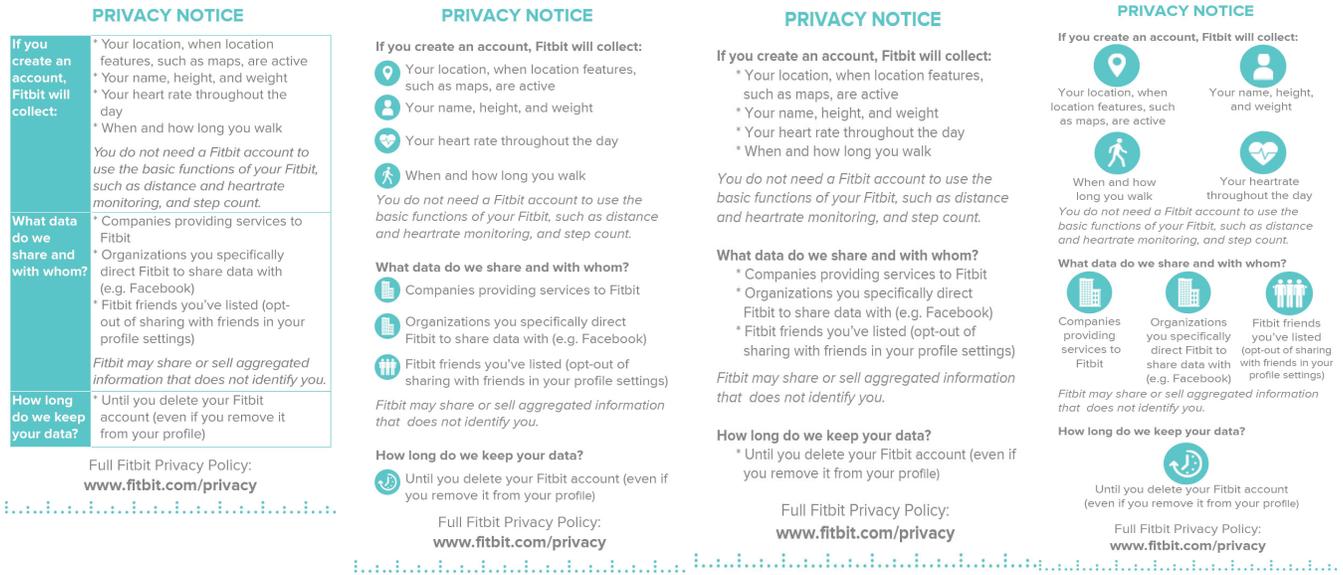


Figure 2: Privacy notice design formats tested in the first survey (left to right): Table format, bulleted icon format, bulleted format, and icon format. The privacy notices all show the same practices based on Fitbit’s privacy policy. The bulleted icon format was used in the second and third study.

modality of privacy notices [43]. Rather than displaying a single privacy notice when a device is first purchased or activated, prior work has examined the potential for showing privacy notices at regular frequencies, or in the form of ‘just-in-time’ notices that are sent just before a privacy sensitive activity is about to occur [4, 5, 7, 38, 39]. Other research has focused on making privacy notices integral to the function of the device, for example playing sounds when a photo is taken or data is sent to the cloud [10].

Finally, there has been significant research into formats for privacy policies and other notices [34]. Research on standardization of privacy policies [15, 18, 31], and privacy ‘nutrition labels’ [28, 29] has found that standardized tabular formats are beneficial. Good et al. found that users were more likely to notice short versions of end user license agreements (EULAs), but the notice format did not impact installation rates significantly [21]. Waddel et al. found that paraphrasing EULA content and splitting it into multiple pages increased comprehension [46]. However, it is not clear whether the change can be attributed to the paraphrasing or the multiple pages. In our studies, we isolate specific aspects to reduce confounding factors in order to gain deeper insights into notice effectiveness.

3. PRIVACY NOTICE DEVELOPMENT

We focused our research on the Fitbit Surge watch due to Fitbit’s leading market share in fitness wearables (22%) [25]. The Surge was the newest Fitbit device at the time we began our study. The content of the privacy notices we developed and tested are based on an analysis of Fitbit’s privacy policy from Dec. 9, 2014 [19], which was still Fitbit’s current privacy policy at the time of this writing. We included Fitbit’s collection, sharing, selling, and storage practices in our privacy notice designs. We did not include any practices relating to online tracking for individuals who visit Fitbit’s website, as these practices did not relate directly to the Fitbit device. Note that while our research was focused on a single fitness wearable’s privacy policy, we examined the privacy policies of other fitness wearable vendors (namely, Jawbone [26] and Misfit [36]) and found them to describe similar practices.

In the following sections, we describe our privacy notice development process. Our first step was to determine an effective privacy notice design format for the Fitbit device. Our second step was to determine which practices participants expected, even without a privacy notice. This informed our decisions about which practices to remove from the shorter versions of our notices in order to emphasize unexpected privacy practices.

3.1 Short-form Notice Design

We created four prototype short-form privacy notice designs, and conducted a survey to assess the effect of design on awareness of Fitbit’s privacy practices. The designs are shown in Figure 2: table format, bulleted icon format, bulleted format, and icon format. Table formats have been used successfully in standardizing bank privacy policies [18, 31] and in privacy nutrition labels [29]. Fitbit’s illustrated privacy notice uses icons with text and Fitbit’s full legal privacy policy includes bulleted text [19]. While our four formats had different layouts and graphical elements, they all contained the same text. We designed our first study to test which of these formats led to the greatest awareness of Fitbit privacy practices.

3.1.1 Study Design

In summer 2015 we conducted a 200-participant survey on Amazon Mechanical Turk, using a between-subjects design with 50 participants per format. We chose 200 participants after conducting a power analysis using Cohen’s medium effect size to ensure that we achieved 80+% power, even with study drop outs. Participants were paid \$0.60 for completing the survey. Only US Turkers with 95% or higher HIT acceptance were recruited. To reduce bias, the survey was marketed as a survey on fitness wearables: no recruitment information indicated the survey was related to privacy.

After being asked a set of demographic questions, participants were shown one of the four short-form privacy notice designs and instructed to read it carefully as they may be asked questions about it. The goal was to create a best-case scenario in which all participants would pay attention to the notice, so that we could assess differences in awareness based on notice design, rather than due to

Question	Correct(%)	Incorrect (%)	Unsure (%)	In Short Notice	In Medium Notice
Collect					
Steps	94	3	3		
Distance	94	4	1		
Info Posted to Profile	93	6	1		
When Exercising	93	6	1		
Heartrate	93	6	1		
Stairs Climbed	88	11	1		
Name	81	16	3		*
Sleep	76	20	4		*
Exercise Comp. to Friend	73	22	5		*
Weight	72	24	4		*
Height	70	25	5		*
Location Specific (Q. 20 in Appendix)	31	56	13	*	*
Share With					
Fitbit Friends (Q. 16 in Appendix)	76	20	4		*
Companies Providing Services	72	22	6		*
Directed Organizations (e.g. Facebook)	67	26	7	*	*
Government	29	66	5	*	*
Misc.					
Where to Find Privacy Policy	88	12	0	*	*
Use Fitbit Without an Account	31	53	16	*	*
Selling Data Conditions	23	57	20	*	*
Data Retention Policy	22	47	31	*	*

Table 1: Results from our second MTurk study (70 participants). Shows the % correct/incorrect/unsure for Fitbit privacy practices without any form of privacy notice. Using Fitbit without an account denotes the functionality a Fitbit maintains without a connection to a Fitbit account (and thus without any form of data collection). We use asterisks to indicate which practices we displayed in our short and medium notices; all practices were displayed in our long notice.

different levels of attention. Participants could move on to the next survey page as soon as they wanted but were not able to return to the notice after that. They were then asked questions to test their awareness of the Fitbit privacy practices. After answering these questions, participants were again shown the assigned privacy notice format, and asked to rate its helpfulness on a 5-point Likert scale (not very helpful to very helpful), and to evaluate how comfortable they were with Fitbit’s collection of location data, storage practices, and sharing practices on a 7-point Likert scale (very uncomfortable to very comfortable). We asked these questions to get a sense of a participant’s feelings towards the privacy notices, as well as their feelings towards some of Fitbit’s privacy practices.

3.1.2 Study Results and Conclusions

We found no statistically significant differences between formats in awareness of Fitbit’s privacy practices. Additionally, using Kruskal-Wallis tests, we found no difference between privacy notice format in terms of how helpful participants found notices ($H(3,197)=.3326$ $p=.95$), or how they felt about collection of location data ($H(3,197)=.7017$ $p=.87$), storage practices ($H(3,197)=.0816$ $p=.99$), or sharing practices ($H(3,197)=.4961$ $p=.51$). In past studies that have found differences in the performance of privacy policy format variants [29, 34], the tested formats varied in wording, length, and layout, while our formats varied only in layout.

We selected the bulleted icon format (second from the left in Figure 2) for our final study because it was in line with Fitbit’s general design motif of mixing icons and text [19].

3.2 Baseline Knowledge of Privacy Practices

One of our key hypotheses was that removing commonly expected pieces of information from a privacy notice would increase awareness of the information contained in the privacy notice, since

there would be less information for people to read and understand. We conducted a study to identify which privacy practices described in the Fitbit privacy policy were commonly expected.

3.2.1 Study Design

We designed a survey asking participants questions about Fitbit’s privacy practices without showing them any privacy notice. In addition, we let participants know at the beginning of the survey that they would not be penalized for wrong answers, so as to discourage them from searching for this information in Fitbit’s privacy policy.

We recruited 70 Turkers from the US with 95% or higher HIT acceptance during Fall 2015. The survey was marketed as a survey on fitness wearables, with no recruitment information indicating the survey was related to privacy. Participants were paid \$0.60 for completing the survey. After answering basic demographic questions, participants were directed to visit the Fitbit Surge page on Fitbit’s website and could not move on from this page for 2.5 minutes. We included this provision because a potential buyer of a Fitbit device would likely spend some time looking at its webpage before purchasing the device. However, we did not enforce that participants look at the Fitbit Surge page, only that they wait 2.5 minutes before advancing in the survey.

Participants were then asked questions about 30 collection, sharing, selling, and data retention practices, specifically pertaining to the Fitbit Surge watch. These questions included 20 practices actually included in Fitbit’s privacy policy (shown in first column of Table 1), as well as questions regarding ten additional fictitious practices. Examples of fictitious practices include collecting perspiration, altitude, and mood; and sharing with researchers, Facebook friends, and the public. We included fictitious practices in order to ensure that participants did not believe that all practices mentioned were performed by Fitbit. Participants were then asked

a series of multiple choice questions related to Fitbit policy details. Because we were interested in baseline knowledge of actual privacy practices, we report only these results (see columns 2-4 of Table 1).

3.2.2 Study Results and Conclusions

As shown in Table 1, there was a wide range of participant awareness. 94% of participants knew that the Fitbit Surge collected steps, whereas only 22% were aware of Fitbit’s data retention policy. Many of these questions were based on a likert scale, as can be seen in questions 12 and 14 in Appendix. For our results, we aggregated any choice (from might to definitely) to a binary collect/did not collect. Overall, participants were more knowledgeable about data collection practices, somewhat less knowledgeable about sharing practices, and least knowledgeable about specific policies such as data retention or using the Fitbit Surge without a Fitbit account.

We used our results to inform our decisions about what practices to omit in our shorter notices. We wanted to remove practices only when a strong majority could answer questions relating to those practices correctly. We determined that removing items that 70% or more and 85% or more of participants were able to answer correctly allowed for the removal of two clear clusters of information. The data practices that were retained in the medium- and short-length notices are shown in the right two columns of Table 1.

4. NOTICE FRAMING AND LENGTH

Our two preliminary studies informed the design of the short-form privacy notices that we used to test our hypotheses relating to effects of the framing and length of the notice. We considered three forms of framing (positive, negative, neutral), and three notice lengths (short, medium, and long). This led to a 3x3 experimental design, with a tenth condition as control.

In the positive framing conditions we included positive reasons for Fitbit to engage in some of its practices, namely sharing and data retention. In the negative framing conditions we included potential drawbacks/risks related to the same practices. Figure 3 provides the positive and negative framing text. The neutral condition did not include any framing. What practices were included in which notice length can be seen in Table 1. Figures 4, 5, and 6 show the long, medium and short length notices. The figures show examples from different framing conditions. All use the bulleted with icons design from our design format study.

In addition to notice content, for all notice lengths we included at the end of the first two sections of the notice the phrase “Find further [collection/sharing] practices at Figbit.com/privacy.” At the bottom of the policy we included the text “Full Fitbit Privacy Policy www.fitbit.com/privacy.” We did this to avoid the perception that the absence of well-known practices from the shorter notices indicates that these practices do not occur.

In January 2016 we recruited 400 Turkers from the US with 95% or higher HIT acceptance, approximately 40 per condition in a between-subjects study design. We chose 400 participants as a result of a power analysis using Cohen’s medium effect size to ensure that we achieved 95+% power, even with study drop outs. Due to randomized condition assignment and some participants failing to complete the survey after being assigned to a condition, actual conditions ranged in size from 33 to 42 participants (Mean=38.7 SD=3.71). The survey was marketed as a survey on fitness wearables, with no recruitment information indicating the purpose was related to privacy. Additionally, we noted within the survey that participants would not be penalized for incorrect answers, as we were more interested in their opinions and knowledge level than achieving the best answers. This was done to reduce the likelihood of Turkers looking up answers in the survey.

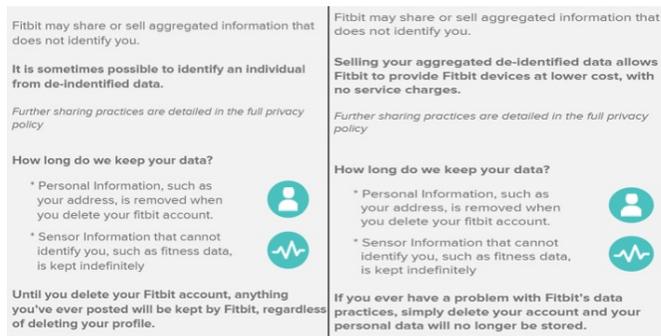


Figure 3: Negative (left) and positive (right) framing statements (in bold) for data sharing and retention practices.

This survey followed the same study design as our baseline knowledge survey until after participants were directed to view the Fitbit Surge’s webpage (survey can be found in Appendix). At that point, participants were shown a notice (or shown nothing in the control), based on their condition. In order to make participants’ interaction with the notices realistic, participants were allowed to skip to the next page of the survey without spending any time looking at the notice. We recorded the time participants spent on the notice page. We then presented questions relating to Fitbit privacy practices as we had in the baseline knowledge survey. Participants were not able to return to the notice while answering these questions.

In addition to the questions relating to Fitbit’s privacy practices, we also asked participants to rate their general concern with Fitbit’s privacy practices, as well as to answer the 10-item variant of the IUIPC privacy concerns scale [32]. We did this to measure what we expected to be the mechanism (concern) by which our framing conditions affected participant awareness of Fitbit privacy practices. To account for this longer survey length, we compensated the participants \$1.50.

5. RESULTS

In the following sections, we aggregate our results by type of practice and overall awareness. For the purposes of aggregation, we count a participant’s correct answer about each of 20 data practices as 1, an incorrect answer as -1, and unsure answers as 0. The question categories are shown in Table 1 and the questions are shown in the Appendix. Our metric led to a non-normal distribution of awareness for certain conditions. As a result, we performed non-parametric statistical tests (Kruskal-Wallis).

Our short-form notices increased awareness of privacy practices over the control condition (no notice). However, framing did not have a statistically significant effect on privacy practice awareness or concerns. Additionally, the shortest notices performed worse in terms of privacy practice awareness than the medium and long notices, particularly on practices removed from the short notices. Age and Gender were related to awareness, but there was no interaction effect between these factors and condition. Participants who visited the Fitbit website during the survey had significantly higher awareness scores than those who did not, and those in the control condition benefited most from visiting the website. Additionally, we found no significant difference in time spent reading notices between conditions. However, we found that longer reading times, concern about Fitbit privacy practices, and high IUIPC scores were associated with greater awareness of privacy practices. We discuss the results in detail below.

With an account, Fitbit will collect:

- * Your **location**, when location features, such as maps, are active
- * Your **name, height, and weight**
- * Your **steps, distance and stairs climbed**
- * When and how long you **exercise**
- * When and how long you **sleep**
- * Your **heartrate** throughout the day
- * Exercise **compared with Friends**
- * **Information posted** to your profile

You can track your heartrate, distance and step count with your Fitbit, without needing an account.

Find further collection practices at [Fitbit.com/privacy](https://www.fitbit.com/privacy)

With whom do we share data?

- * **Government Entities**
- * **Companies** providing services to Fitbit
- * **Organizations** you specifically direct Fitbit to share data with (e.g. Facebook)
- * **Fitbit friends** you've listed (opt-out of sharing with friends in your profile settings)

Fitbit may share or sell aggregated information that does not identify you.

It is sometimes possible to identify an individual from de-identified data.

Find further sharing practices at [Fitbit.com/privacy](https://www.fitbit.com/privacy)

How long do we keep your data?

- * Personal Information, such as your address, is removed **when you delete your fitbit account**.
- * Sensor Information that cannot identify you, such as fitness data, is **kept indefinitely**

Until you delete your Fitbit account, anything you've ever posted will be kept by Fitbit, regardless of deleting your profile.

Full Fitbit Privacy Policy:
www.fitbit.com/privacy

Figure 4: Long length notice (negative framing): Includes all Fitbit privacy practices relevant for using a Fitbit Surge, as well as negative framing statements for certain practices.

5.1 Participants

We initially recruited 400 participants through Amazon Mechanical Turk. Nine participants were removed when our survey tool (SurveyGizmo) indicated they were connecting from outside the US, despite being identified as US MTurkers.

Our sample was fairly diverse. The median age was 29, with a range of 18-69. 193 (49.4%) of our participants were male, 196 (50.1%) female, with two participants not reporting their gender. As shown in Table 2, most of our participants reported currently or previously using a fitness wearable device.

5.2 Effectiveness of Notices

Our short-form privacy notices led to increased participant awareness of privacy practices. Performing a Mann-Whitney U test, we found participants who saw one of our short-form privacy no-

With an account, Fitbit will collect:

- * Your **location**, when location features, such as maps, are active
- * Your **name, height, and weight**
- * When and how long you **sleep**
- * Exercise compared with Friends

You can track your heartrate, distance and step count with your Fitbit, without needing an account.

Find further collection practices at [Fitbit.com/privacy](https://www.fitbit.com/privacy)

With whom do we share data?

- * **Government Entities**
- * **Companies** providing services to Fitbit
- * **Organizations** you specifically direct Fitbit to share data with (e.g. Facebook)
- * **Fitbit friends** you've listed (opt-out of sharing with friends in your profile settings)

Fitbit may share or sell aggregated information that does not identify you.

Selling your aggregated de-identified data allows Fitbit to provide Fitbit devices at lower cost, with no service charges.

Find further sharing practices at [Fitbit.com/privacy](https://www.fitbit.com/privacy)

How long do we keep your data?

- * Personal Information, such as your address, is removed **when you delete your fitbit account**.
- * Sensor Information that cannot identify you, such as fitness data, is **kept indefinitely**

If you ever have a problem with Fitbit's data practices, simply delete your account and your personal data will no longer be stored.

Full Fitbit Privacy Policy:
www.fitbit.com/privacy

Figure 5: Medium length notice (positive framing): Has had relevant Fitbit privacy practices which 85% or more individuals assume are true removed.

tices had significantly higher overall privacy practice awareness (Mean=12.06, SD= 5.89) than control participants (M=9.54, SD= 5.86), with ($U(1,390)=-3.03, p=.002, r=.153$).

We examined whether our hypotheses relating to framing and length of the notice led to significant changes in awareness. Performing a Kruskal-Wallis test, we found there was no statistically significant interaction between the framing and length conditions ($H(8, 343)=14.26, p=0.08$) on overall privacy practice awareness. Therefore, when conducting further analysis on each of these variables individually, we aggregate conditions by their framing or length.

5.2.1 Framing

Our positive and negative framing statements (shown in Figure 3) had no noticeable effect on participants' awareness of Fitbit's privacy practices. Performing a Kruskal-Wallis test, we found no significant differences in overall privacy practice awareness based on the framing of the notice ($H(2,349)=2.643, p=.267$).

5.2.2 Length of Notice

We found that our shortest notice resulted in lower privacy practice awareness than longer notices, and that this was particularly

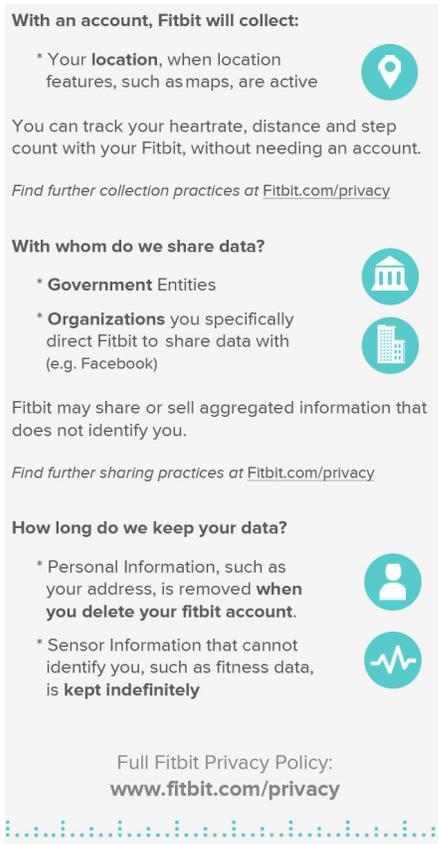


Figure 6: Short length notice (no framing): Has had relevant Fitbit privacy practices which 70% or more individuals assume are true removed. No framing statements included.

Category	Percent
I currently use a wearable Fitness device	30.2
In the past, I regularly used a wearable fitness device, but I no longer do so	10.5
I have tried out a wearable fitness device, but have never regularly used one	17.1
I have never a wearable fitness device, but am familiar with the concept	40.2
I was unfamiliar with wearable fitness devices, before taking this survey	2.0

Table 2: Participant experience with fitness wearables.

true in the case of practices removed from the shorter notices (see 7). Demonstrating this, we ran a Kruskal-Wallis test and found significant differences in awareness of privacy practices ($H(2,349) = 10.42, p = .005$) based on length.

Performing post-hoc Mann-Whitney U tests with Tukey correction, we found that long notices (Mean = 12.52, SD = 5.98) and medium length notices (Mean = 12.65, SD = 5.14) outperformed short notices (Mean = 11.05, SD = 5.82) in terms of overall awareness of privacy practices and collection practices with ($U(1,232) = -2.909, p = .012, r = .191$) and ($U(1,238) = -2.604, p = .027, r = .168$), respectively. We found no significant difference between long and medium length notices in terms of overall awareness of privacy practices.

While important in aggregate, we also examined whether the

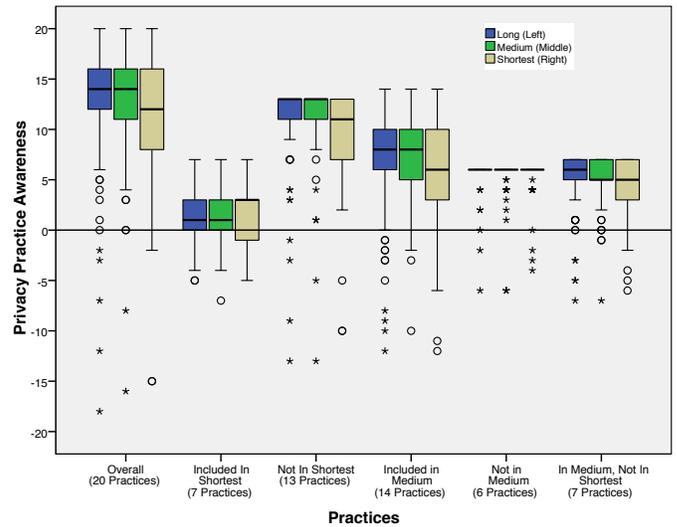


Figure 7: Privacy practice awareness by Length. Strong similarity in performance between long and medium length notices. Significantly worse performance for the shortest notice on privacy practices overall, and specifically on practices removed from the shortest notice. Medium length notices performed similarly to long length notices for practices both left in and removed from the medium length notice.

change in awareness between notice lengths was focused on practices that remained in the shortest version of the notices, or practices that were removed from the shortest version of the notice. We originally postulated that participants in shorter length conditions would perform less well on practices removed from their notices, and potentially better on practices that remained in their notices.

Performing a Kruskal-Wallis test, we found significant differences in awareness by length when considering practices that had been removed from the shortest notices ($H(2,349) = 22.439, p < .0005$). Performing post-hoc Mann-Whitney U tests with Tukey correction, we found long and medium length notices (Long: Mean = 11.05, SD = 4.14; Medium: Mean = 11.27, SD = 3.66) outperformed short length notices (Mean = 9.50, SD = 4.36) in terms of practices removed from the shortest notice, with ($U(1,232) = -3.891, p < .0015, r = .255$) and ($U(1,238) = -4.127, p < .0015, r = .267$) respectively. We found no significant differences between long and medium length notices.

Additionally performing a Kruskal-Wallis test, we found no significant difference in awareness of practices remaining in the shortest notice by length.

While we found no difference in the performance of long and medium length notices overall, we also analyzed whether there was a difference in performance when considering practices left in and removed from the medium notices independently. Performing a pair of Kruskal-Wallis test, we found a significant difference in awareness of practices remaining in the medium notice (with $H(2,349) = 10.126, p = .005$, and not significant difference in awareness of practices removed from the medium notice. Performing post-hoc Mann-Whitney U tests with Tukey correction, we found no significant difference between long and medium notices in awareness of practices remaining in the medium length notice ($p = .882$). Instead, we found that both the medium and long notices outperformed the shortest notice when considering practices remaining in the medium length notice, with (Long vs. Short: $U(1,232) = -2.726, p = .018, r = .181$) and (Medium vs. Short: $U(1,238) = -2.756, p = .018, r = .178$).

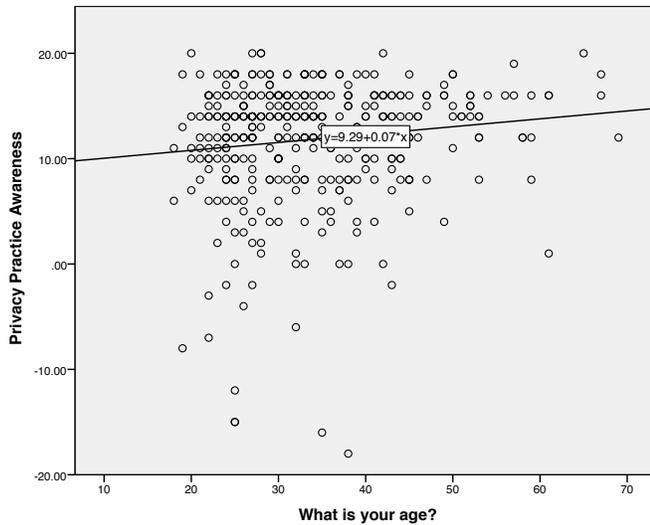


Figure 8: A statistically significant relationship between age and privacy awareness, with each year of age being associated with a .07 increase in awareness score.

These results prompted us to examine the performance of the various notice lengths on awareness of those 7 practices that were removed from the shortest notice, but were retained in the medium and long notices. Note that these practices were expected by between 70 and 85% of participants in our baseline knowledge study, while the other 6 removed practices were expected by over 85% of participants in that study. Performing a Kruskal-Wallis test, we found that there was a significant difference in awareness of these practices by notice length, with $H(2,349)=14.268$ $p=.001$. Performing post-hoc Mann-Whitney U tests with Tukey correction, we found significant differences in awareness of these practices between both the long and shortest notice lengths, and the medium and shortest notice lengths with (Long vs. Short: $U(1,233)=-3.037$ $p=.006$ $r=.199$) and (Medium vs. Short= $U(1,238)=-3.435$ $p=.003$ $r=.222$). Indeed the participants in the shortest notice conditions performed similarly to those in the control condition on these 7 practices, while participants in the medium and long conditions became more aware of these practices. This suggests that 70 to 85% awareness of practices may not be high enough for successful removal from a privacy notice.

5.3 Impact of Demographic Factors

Interestingly, age and gender both had significant effects on participants' overall privacy practice awareness, although we did not find any interaction between these factors and participant condition. This means that our conclusions regarding our notice conditions are generally applicable across these demographic factors.

We performed a linear regression, and found that for each year of age, participants had a .075 higher awareness score (see Figure 8), with $t=2.558$, $p=.011$. Performing a Mann-Whitney U test, we found that women (Mean= 11.13, SD=6.31) had higher overall privacy practice awareness than men (Mean=10.50, SD=7.77), with ($H(1,387)=-2.104$, $p=.035$, $r=.109$). We removed two participants who chose to not share their gender from this analysis.

5.4 Impact of Participant Behavior

We examined the relationship between participant behavior in our survey and privacy practice awareness in two ways. First, as mentioned in the methodology, we indicated to participants that they should visit the Fitbit Surge page on the Fitbit website as if

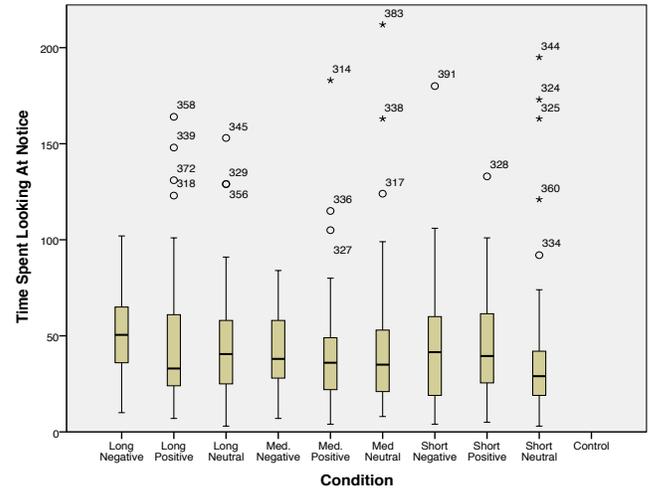


Figure 9: Time spent on notice by condition. No significant difference in time spent on notice by condition, with an average between 45–55 seconds for each condition.

they were shopping for a Fitbit. However, we did not force our participants to do so. While we did not record whether participants visited the website, we asked participants how much the Fitbit Surge costs (\$250). This acted as a knowledge check to determine who had at the very least visited the page, as the cost was prominently displayed at the top right corner of the page. We did this to get a measure of participants' commitment to researching the device, and the extent to which a privacy notice would help those more or less likely to examine a fitness wearable's details on their own.

Additionally, we tracked how long participants spent on the page of the survey that showed them our privacy notice before moving on. We hypothesized that participants could spend less time on our shorter notices while maintaining at least similar performance.

5.4.1 Knowledge Check

Our analysis showed a strong majority, 339 (86.7%) participants, knew the cost of the Fitbit Surge, as compared to 52 who didn't know(13.3%). We performed a Mann-Whitney U test showing that participants who knew the cost of the Fitbit Surge had significantly higher overall privacy practice awareness (Mean=11.70, SD=5.47) than participants who did not (Mean=8.25, SD=6.85) with ($U(1,390)=-3.719$, $p<.0005$, $r=.188$).

Examining the data more closely, we found that there was a major jump in overall privacy practice awareness for participants in our control condition, from (Mean=1.50, SD=8.22) to (Mean=10.46, SD=4.90) for those who passed the knowledge check, whereas the increase in awareness for those who passed this knowledge check in the treatment conditions (with notices) was not as dramatic going from (Mean=9.58, SD=6.41) to (Mean=12.46, SD=5.47). This may be due to the fact that the Surge page contained information about its functionality, which included mention of the data it collects. Participants in the control condition were not presented with information about data collection except on this page, whereas participants in the other conditions received this information both on the Surge page and in the privacy notice.

5.4.2 Time Spent on Notice

We found a number of interesting results regarding time spent looking at our notices. We found no significant differences between time spent reading the notices in each condition. However,

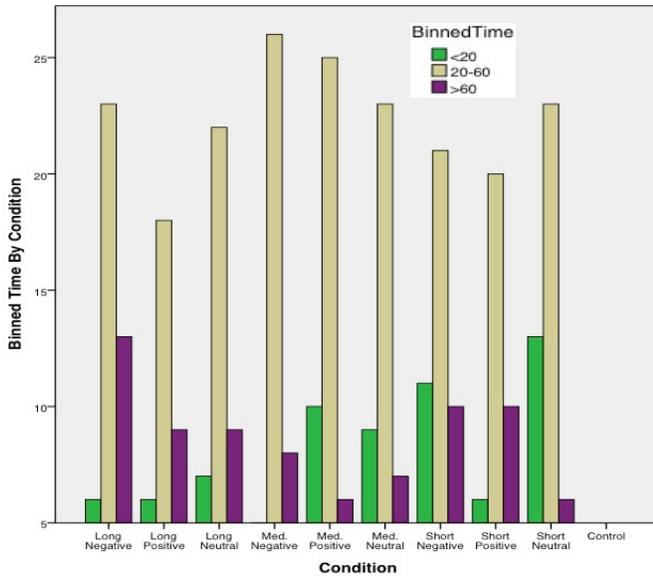


Figure 10: Binned time spent on notice by condition. We found no significant difference in binned time by condition.

we did find some relationships between time spent and overall privacy awareness.

In addition to analyzing time as a continuous variable, we also binned time into three segments: less than 20 seconds, between 20 and 60 seconds, and more than 60 seconds. We chose these bins as we did not think someone could read through the entire privacy notice in less than 20 seconds, but that almost anyone could read through the notice in 60 seconds, and would be examining it closely (or were distracted by another task) if they looked at it for longer.

We found that regardless of whether time was measured as continuous or binned, there was no difference in time spent on notice between length conditions. The distribution of participants by condition in each bin is shown in Figure 10. Using Pearson's Chi-Square test, we found no relationship between binned time and length of notice. The overall length of time spent on notice by condition is shown in Figure 9. Performing a Kruskal-Wallis test, we found no statistically significant differences in the length of time spent on the notice, and condition.

Performing a linear regression, we did not find a statistically significant relationship between time spent on the notice and overall privacy practice awareness. However, performing a Kruskal-Wallis test, we did find that binned time had an effect on overall privacy practice awareness ($H(2,349)=26.89, p<.0005$). Using Tukey correction for multiple testing, we compared each binned time with Mann-Whitney U tests. We found that bin 0 (<20 seconds) was significantly outperformed (Mean=8.59, SD=8.09) by both bin 1 (20-60 seconds, Mean=12.87, SD=4.61) and bin 2 (>60 seconds, Mean=13.26, SD=4.03), with ($U(1,274)=-4.839, p<.001, r=.292$) and ($U(1,150)=-4.431, p<.001, r=.361$) respectively, in terms of privacy practice awareness. We found no significant difference between bins 1 and 2. This suggests that there is a difference between glancing at a notice and reading the notice, but how much time is spent reading or studying the notice may not matter as much.

5.5 Impact of Privacy Concern

We measured participants' privacy concern in two ways. First, we asked participants to rate their concern with Fitbit's privacy practices at the aggregate levels of collection practices, sharing

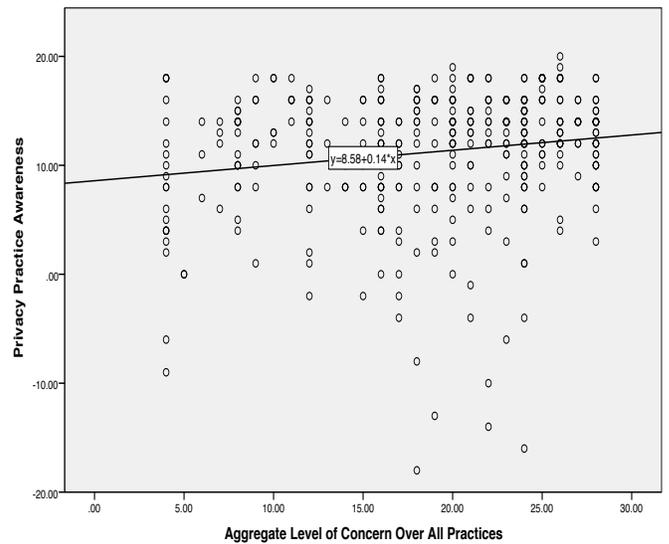


Figure 11: Relationship between awareness of privacy practices and overall concern with Fitbit practices: for every point of concern, there is an increase of .14 in awareness of privacy practices (or .14 more questions answered correctly).

practices, selling practices, and storage practices on a 7-point Likert scale from not very concerned to very concerned (see Q. 30 in Appendix). Second, participants completed the 10-item IUPC questionnaire [32], which results in scales for awareness, collection, and control on a 7-point Likert scale from strongly disagree to strongly agree (see Q's 44-53 in Appendix A).

5.5.1 Concern With Fitbit Privacy Practices

We found that participants were most concerned with Fitbit's sharing practices, participants with greater concern had greater privacy practice awareness, and there was no significant relationship between framing and concern.

We performed a Friedman test, finding participant concern was greatest for sharing practices (Mean=5.06, SD=1.89), compared to collection (Mean=4.52, SD=1.84), selling (Mean=4.72, SD=1.93), and storage (Mean=4.70, SD=1.93) with ($\chi^2(3,388)=74.32, p<.0005$). Performing pair-wise Wilcoxon tests with post-hoc correction, we found that concern with sharing practices was significantly higher than collection practices ($Z(1,390)=-7.968, p<.003$), concern with storage practices was significantly higher than collection practices ($Z(1,390)=-2.782, p=.030$), concern with sharing practices was significantly higher than concern with selling practices ($Z=-5.051, p<.0030$), concern with sharing practices was significantly higher than concern with storage practices ($Z=-5.696, p<.0030$).

We had originally hypothesized that framing would lead to greater concern, causing participants to pause to reflect on the practices in the policies to a greater extent. The second part of this hypothesis appears to be correct, as can be seen in Figure 11. Performing a linear regression, we found that for every increase in overall concern over privacy practices, there was a .146 increase in overall privacy practice awareness, with ($t=3.488, p=.001$). However, performing a Kruskal-Wallis test, we found no relationship between condition and concern, as can be seen in Figure 12. We additionally tested whether aggregating notices by their framing and excluding the control condition made any difference. However, a Kruskal-

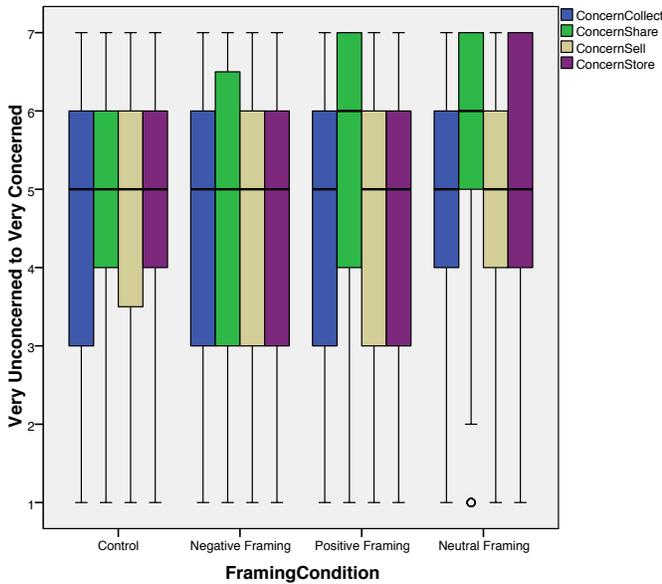


Figure 12: Concern with Fitbit privacy practices by framing condition. While concern over sharing personally identifiable information was slightly greater, we found no statistically significant differences between conditions.

Wallis test did not reveal significant differences. It seems that our framing conditions fail to impact overall privacy practice awareness because they fail to impact participant concern.

5.5.2 IUIPC Concern

Prior work has shown a significant relationship between IUIPC scores and putative online privacy behavior [32]. Therefore, we examined the relationship between the IUIPC scales and participants’ awareness. Performing a linear regression, we found the IUIPC awareness scale was positively associated with awareness of privacy practices, with every point of agreement with the IUIPC awareness questions leading to (on average) an increase of 2.221 in overall awareness of privacy practices ($p < .0005$), see Figure 13.

However, performing a Kruskal-Wallis test, we found no relationship between condition and any of the IUIPC scales, suggesting agreement with IUIPC variables was not noticeably affected by notices, framing, or length of notices, see Figure 14.

On the whole this confirms the effectiveness of the IUIPC scales to predict overall privacy concerns of participants. It also demonstrates that our framing did not affect participant concern about on-line privacy in general, as measured by IUIPC questions.

6. DISCUSSION

We explored the idea that shorter short-form privacy notices focusing on less expected privacy practices might lead to greater awareness of privacy practices. We specifically investigated this approach in the context of fitness wearables’ privacy practices. We measured success by participant awareness of Fitbit’s privacy practices.

We first discuss potential limitations of our study design. We then discuss the effectiveness of privacy notices, the specific effects (or lack thereof) of our enhancements, as well as explanations for these effects from the data, and implications for notice design.

6.1 Limitations

It is unclear how generalizable our results are, as our surveys focused on a single context, the privacy policy of a single company,

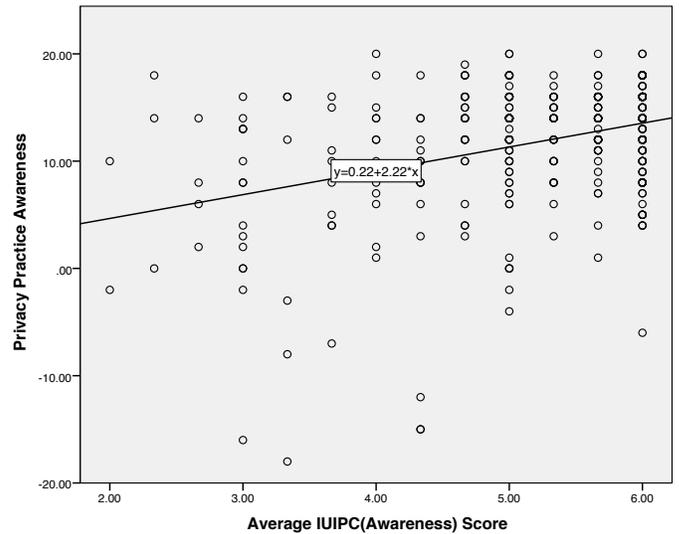


Figure 13: Relationship between the IUIPC awareness scale and awareness of Fitbit’s privacy practices. For every point of agreement with an IUIPC awareness question, awareness of Fitbit privacy practices increased by 2.22.

one specific wording of that policy, and one specific device. We chose to focus on Fitbit as it is the market leader in fitness wearables [25], and the Fitbit Surge as it was Fitbit’s newest product at the time our research commenced. Our examination of other fitness wearable manufacturers (e.g. Jawbone, Misfit), found their policies to be functionally similar to Fitbit’s [26,36]. More importantly, the focus of our research – improving privacy notice design through framing or length – is not specific to Fitbit or even fitness wearables, except for the privacy practices we displayed in the tested notice formats. Our notice-development process could be applied to any company’s privacy policy.

It is possible that some of our results can be attributed to the wording of the short-form privacy notice we tested. For example, we included wording intended to inform participants that the short-form notices did not contain all of Fitbit’s data collection or sharing practices. However, we did not directly investigate whether participants understood that. In addition, we tested only one set of words for our framing conditions. It is possible that other approaches to framing might have produced different results.

Another potential limitation is the use of MTurk for conducting surveys. Some prior work has shown that MTurkers can differ from the general population, and that individuals may interact with a survey differently than they would in reality [23]. Other research has shown that MTurkers constitute a reasonably good sample of the general population [11]. We addressed this potential problem in two ways: first, our survey was consistently designed to elicit natural reactions to privacy notices. Our recruitment materials did not mention privacy or security, participants were informed at the beginning of the survey that they would not be penalized for wrong answers, and at no point did we force participants to look at privacy notices, but instead we let them click through to the next page of the survey if they so chose. These design decisions were meant to, as closely as possible, mirror a participant’s actual interaction with privacy policies and privacy notices. Secondly, we examined relative effectiveness of our various design decisions, with a control group included, which should mitigate biasing effects.

A related potential limitation is the direct confrontation of participants with a privacy notice. We chose this approach to reduce

variations in participants’ attention. This provided us with a best case scenario for a comparative assessment of how notice length, framing, and other characteristics impact participants’ awareness of privacy practices. We expect that under real conditions, participants would likely perform worse, due to distractions and lack of attention to the notice. Since we did not observe framing effects in our study, it is unlikely that they would surface in a field study with the type of privacy notice we focused on.

6.2 Privacy Notices Can Be Effective

An important result from our work is demonstrating that short-form privacy notices uniformly led to significantly higher awareness than the control. This result, while a reconfirmation of the basic effectiveness of privacy notices [43] is important for two further reasons. First, fitness wearables generally collect data that is inherent to their function (e.g., steps, distance, heart-rate). It was therefore possible that since many of Fitbit’s privacy practices would be linked to the function of the device, participants might have had a higher awareness of such practices without seeing a privacy notice. This was not the case.

Second, Fitbit does not currently have comparable short-form privacy notices. Our results show a practical method by which Fitbit and other fitness wearable manufacturers could increase user awareness of their privacy practices by integrating privacy notices similar to ours into their mobile companion apps or websites.

6.2.1 Framing Did Not Affect Concern

The results from our analysis of participants’ reported concern over Fitbit’s privacy practices provide a potential explanation for the lack of significant difference we found between framing conditions. We found no significant difference in concern with Fitbit’s privacy practices or general privacy concern (IUIPC) and the framing conditions. In other words, framing some practices in a positive or negative light did not seem to make a difference in how concerned participants were about them. However, the lack of change in level of concern suggests that this was due to a lack of effectiveness of our chosen framing technique, and not a failure in the underlying concept of framing itself. Including framing statements that lead to greater or lesser participant concern might very well lead to greater or lesser awareness of policies. This could be done through heightening the focus on risk and implications, or by including personalized information, such as the data that could be re-identified for the particular user receiving the notice.

6.2.2 Shortest Notices Led to Less Awareness

Our results show that removing well-known privacy practices to make short-form notices even shorter actually led to similar or worse participant awareness of privacy practices. Our intuition was that further condensing a short-form privacy notice would lead to even better performance, provided that the practices removed were well known. However, this intuition proved false, as our results show no increase in awareness of the practices remaining in the notice when some practices are removed.

Our medium length notices did not result in significantly different performance compared to our longest notices. This suggests that removing some of the most known practices had little effect on participant awareness, but that there may be some benefits of using such medium notices when space is constrained.

Our shortest notices performed significantly worse than our longest notices, suggesting that there may be a lower bound to the length of an effective privacy notice. In addition, the awareness threshold we selected for removing practices from the shortest notice may have been too low.

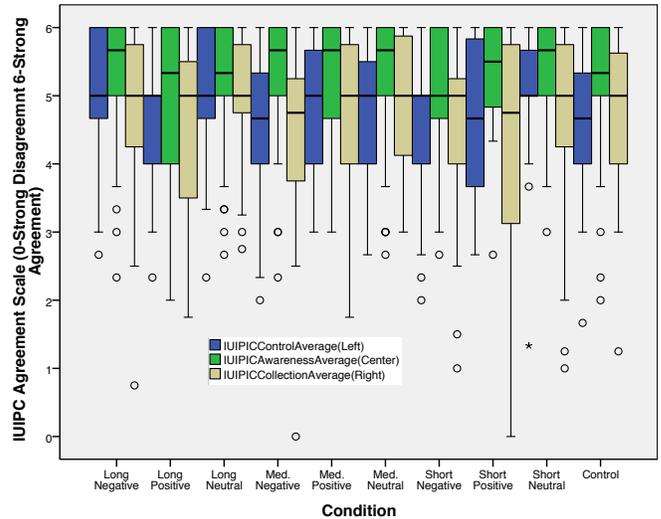


Figure 14: Average agreement with IUIPC scales. No significant difference between conditions. Awareness questions had the highest agreement, followed by control and collection.

Analyzing the time spent on notices does not make the picture clearer. As part of our study design, we did not force participants to look at our notices for a set period of time, instead they could click through to the next page immediately if they so chose. We made this design decision to increase ecological validity, since in the real world users can generally quickly click through a privacy policy. We did, however, record the amount of time participants stayed on the page with the notice. We examined this time as both a continuous variable, as well as binning it into three time lengths based on our estimation of the time necessary to read through the notice. We found that there was no significant relationship between continuous time and participant awareness. However, binned time showed that participants in the larger time bins had significantly higher awareness of privacy practices. Our results also showed that there was no significant difference in time spent on notices by condition (either length or framing). Given the length disparity between our long and short notices (see Figures 4 and 6), we expected participants to be able to spend far longer on the remaining privacy practices in the short notice, and therefore have better awareness of these practices. However, we did not find such a difference.

It is possible that participants spent their time looking at those practices that were unknown or alien to them, with only very brief confirmation of those practices which they assumed or were well known. Participants therefore would have spent a roughly equivalent amount of time on the lesser known privacy practices regardless of length, and participants in the short length notice conditions did not have the benefit of quick confirmations of practices they were already aware of, leading to worse awareness of these practices. It is also possible that even the long privacy notice we created was short enough to achieve all of the gains from condensing a privacy notice, and that shortening a notice for a more complex and lengthy privacy policy could achieve better results.

6.3 Importance of Participant Factors

Our sample was diverse with respect to age, gender, and experience with fitness wearables. Interestingly, we found that each of these participant factors had at least some statistically significant effect on awareness of privacy practices; with older participants and

women having significantly higher awareness of Fitbit’s privacy practices. While not the focus of our study, these results are important as they showcase that awareness of privacy practices varies based on demographic factors. This demonstrates that user studies on the effectiveness of a privacy notice should be conducted with a diverse sample in order to account for demographic differences or should target specific audiences with a specific notice design.

7. CONCLUSIONS AND FUTURE WORK

We presented in this paper a series of three MTurk user studies. Our first survey was focused on determining an effective design format for a Fitbit short-form privacy notice. Our second survey focused on determining participant awareness of each of 20 Fitbit privacy practices. Our final study examined the potential for removing generally expected privacy practices from notices, as well as including framing statements in notices, to increase participant awareness of privacy practices.

Our results reconfirmed the utility of short-form privacy notices, as all notice conditions outperformed the control. However, we also found that while condensing long legalistic privacy policies into succinct privacy notices does increase awareness, taking this a step further by further condensing privacy notices to succinctly include only practices that users are not generally aware of, had the opposite effect. Participants with shorter notices had similar performance on practices that were left in the notice, but performed significantly worse on practices that were removed. Additionally, incorporating positive and negative framing statements into our privacy notices did not bear fruit, with no statistically significant difference in performance. Our analysis of participant concern over Fitbit privacy practices suggests that this lack of effect was due to insufficient differences in the level of concern between framing conditions to elicit significant changes in awareness.

Given these results, we suspect that a lower bound for the potential to compress privacy notices exists, and that further research should focus on personalization of privacy notices [5, 22, 30, 48], or in the timing of the notices (e.g. just-in-time notification, or notification on a regular basis rather than on purchase/install) [7, 17, 30]. That said, further studies investigating the effectiveness of generic short-form privacy notices may be able to address some of the limitations of our study and shed additional light on ways notices may be shortened effectively.

8. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under grants CNS-1012763, CNS-1330596, SBE-1513957 and CSR-1526237, as well as by DARPA and the Air Force Research Laboratory, under agreement number FA8750-15-2-0277. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA, the Air Force Research Laboratory, the National Science Foundation, or the U.S. Government.

The authors would like to thank Blase Ur and Ariel Polakoff for their input and feedback.

9. REFERENCES

- [1] A. Acquisti, I. Adjerid, and L. Brandimarte. Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security & Privacy*, (4):72–74, 2013.
- [2] A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
- [3] I. Adjerid, A. Acquisti, L. Brandimarte, and G. Loewenstein. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proc. SOUPS ’13*. ACM, 2013.
- [4] Y. Agarwal and M. Hall. Protectmyprivacy: Detecting and mitigating privacy leaks on iOS devices using crowdsourcing. In *Proc. MobiSys ’13*, pages 97–110. ACM, 2013.
- [5] H. Almuhammedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal. Your location has been shared 5,398 times!: A field study on mobile app privacy nudging. In *Proc. CHI ’15*, pages 787–796. ACM, 2015.
- [6] I. Ayres and A. Schwartz. No-Reading Problem in Consumer Contract Law, The. *Stanford Law Review*, 66:545, 2014.
- [7] R. Balebako, J. Jung, W. Lu, L. F. Cranor, and C. Nguyen. Little brothers watching you: Raising awareness of data leaks on smartphones. In *Proc. SOUPS ’13*. ACM, 2013.
- [8] R. Balebako, F. Schaub, I. Adjerid, A. Acquisti, and L. Cranor. The impact of timing on the salience of smartphone app privacy notices. In *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*, SPSM ’15, pages 63–74, New York, NY, USA, 2015. ACM.
- [9] O. Ben-Shahar and A. S. Chilton. ’Best Practices’ in the Design of Privacy Disclosures: An Experimental Test. SSRN ID 2670115, Oct. 2015.
- [10] R. Calo. Against Notice Skepticism In Privacy (And Elsewhere). SSRN ID 1790144, Mar. 2011.
- [11] K. Casler, L. Bickel, and E. Hackett. Separate but equal? A comparison of participants and data gathered via Amazon’s MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*, 29(6):2156–2160, 2013.
- [12] F. H. Cate. The limits of notice and choice. *Security & Privacy, IEEE*, 8(2):59–62, 2010.
- [13] Center for Information Policy Leadership. Ten Steps to Develop a Multilayered Privacy Notice. White paper, Mar. 2007.
- [14] L. F. Cranor, P. Guduru, and M. Arjula. User interfaces for privacy agents. *ACM Trans. Comput.-Hum. Interact.*, 13(2):135–178, June 2006.
- [15] J. B. Earp, Q. He, W. Stufflebeam, D. Bolchini, C. Jensen, and others. Financial privacy policies and the need for standardization. *IEEE Security & privacy*, (2):36–45, 2004.
- [16] Federal Trade Commission. Protecting Consumer Privacy in an Era of Rapid Change, Mar. 2012.
- [17] Federal Trade Commission. Internet of Things: Privacy & security in a connected world, 2015.
- [18] Federal Trade Commission and Kleimann Communication Group. Evolution of a Prototype Financial Privacy Notice: A Report on the Form Development Project, Feb. 2006.
- [19] Fitbit inc. Fitbit Privacy Policy, December 2014. Available at <https://www.fitbit.com/legal/privacy-policy>.
- [20] C. Gates, J. Chen, N. Li, and R. Proctor. Effective risk communication for Android apps. *IEEE Trans. Dependable and Secure Computing*, 11(3):252–265, May 2014.
- [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 *Proc. SOUPS ’05*, pages 43–52. ACM, 2005.

- [22] M. Harbach, M. Hettig, S. Weber, and M. Smith. Using personal examples to improve risk communication for security & privacy decisions. In *Proc. CHI '14*, pages 2647–2656. ACM, 2014.
- [23] D. J. Hauser and N. Schwarz. Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior research methods*, pages 1–8, 2015.
- [24] A. Hiltz, C. Parsons, and J. Knockel. Every Step You Fake: A Comparative Analysis of Fitness Tracker Privacy and Security. 2016. Available at <http://citizenlab.org/2016/02/fitness-tracker-privacy-and-security/>.
- [25] IDC. Worldwide Wearables Market Soars in the Third Quarter as Chinese Vendors Challenge the Market Leaders, Dec 2015. Available at <http://www.idc.com/getdoc.jsp?containerId=prUS40674715>.
- [26] Jawbone. UP privacy policy, December 2014. Available at <https://jawbone.com/up/privacy>.
- [27] C. Jensen and C. Potts. Privacy policies as decision-making tools: An evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 471–478, New York, NY, USA, 2004. ACM.
- [28] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder. A nutrition label for privacy. In *Proc. SOUPS '09*. ACM, 2009.
- [29] P. G. Kelley, L. Cesca, J. Bresee, and L. F. Cranor. Standardizing privacy notices: An online study of the nutrition label approach. In *Proc. CHI '10*, pages 1573–1582. ACM, 2010.
- [30] A. Kobsa and M. Teltzrow. Contextualized communication of privacy practices and personalization benefits: Impacts on users' data sharing and purchase behavior. In *Privacy Enhancing Technologies*, pages 329–343. Springer, 2004.
- [31] A. Levy and M. Hastak. Consumer Comprehension of Financial Privacy Notices. *Interagency Notice Project*, <http://ftc.gov/privacy/privacyinitiatives/Levy-Hastak-Report.pdf>, 2008.
- [32] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.
- [33] A. M. McDonald and L. F. Cranor. Cost of reading privacy policies, the. *ISJLP*, 4:543, 2008.
- [34] A. M. McDonald, R. W. Reeder, P. Kelley, and L. Faith. A Comparative Study of Online Privacy Policies and Formats. *Privacy Enhancing Technologies*, 2009.
- [35] Microsoft. Privacy Guidelines for Developing Software Products and Services. Technical Report version 3.1. 2008.
- [36] Misfit Inc. Misfit Privacy Policy, June 2015. Available at http://misfit.com/legal/privacy_policy.
- [37] OECD. Making Privacy Notices Simple. Digital Economy Papers 120, July 2006.
- [38] S. Patil, R. Hoyle, R. Schlegel, A. Kapadia, and A. J. Lee. Interrupt now or inform later?: Comparing immediate and delayed privacy feedback. In *Proc. CHI '15*. ACM, 2015.
- [39] A. Patrick and S. Kenny. From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *Proc. PET '03*. Springer, 2003.
- [40] President's Council of Advisors on Science and Technology. Big data and privacy: A technological perspective. Report to the President, Executive Office of the President, May 2014.
- [41] A. Rao, F. Schaub, N. Sadeh, A. Acquisti, and R. Kang. Expecting the Unexpected: Understanding Mismatched Privacy Expectations Online. In *Proc. SOUPS '16*. USENIX Assoc., 2016.
- [42] J. R. Reidenberg, T. Breaux, L. F. Cranor, B. French, A. Grannis, J. T. Graves, F. Liu, A. McDonald, T. B. Norton, R. Ramanath, N. C. Russell, N. Sadeh, and F. Schaub. Disagreeable Privacy Policies: Mismatches Between Meaning and Users' Understanding. *Berkeley Tech. LJ*, 30:39, 2015.
- [43] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor. A Design Space for Effective Privacy Notices. In *Proc. SOUPS '15*, pages 1–17. USENIX Assoc., 2015.
- [44] F. Shih, I. Liccardi, and D. Weitzner. Privacy Tipping Points in Smartphones Privacy Preferences. In *Proc. CHI '15*, pages 807–816. ACM, 2015.
- [45] J. Tan, K. Nguyen, M. Theodorides, H. Negrón-Arroyo, C. Thompson, S. Egelman, and D. Wagner. The effect of developer-specified explanations for permission requests on smartphone user behavior. In *Proc. CHI '14*. ACM, 2014.
- [46] T. F. Waddell, J. R. Auriemma, and S. S. Sundar. Make It Simple, or Force Users to Read?: Paraphrased Design Improves Comprehension of End User License Agreements. In *Proc. CHI '16*, pages 5252–5256. ACM, 2016.
- [47] S. Wilson, F. Schaub, R. Ramanath, N. Sadeh, F. Liu, N. Smith, and F. Liu. Crowdsourcing annotations for websites' privacy policies: Can it really work? In *Proc. WWW '16*, 2016.
- [48] M. S. Wogalter, B. M. Racicot, M. J. Kalsher, and S. N. Simpson. Personalization of warning signs: the role of perceived relevance on behavioral compliance. *International Journal of Industrial Ergonomics*, 14(3):233–242, 1994.
- [49] S. Zimmeck and S. M. Bellovin. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. In *USENIX Security Symposium*. USENIX Assoc., 2014.

Organizations you direct Fitbit to share your information with	<input type="checkbox"/>						
--	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------	--------------------------

15) Are there any other groups that you believe that Fitbit shares your personally identifiable information with by default?

16) Do you think Fitbit allows you to control how information is shared with your Fitbit friends? *

- No, anyone you add as a Fitbit friend can see all of your Fitness data
- Yes, you can opt-out of sharing specific forms of data with your Fitbit friends on the Fitbit website
- Yes, you can opt-out of sharing ANY data with your Fitbit friends on the Fitbit website, but it is all or nothing
- No, Fitbit doesn't share your information with Fitbit friends
- None of the above
- I don't know

17) How confident are you in your answer to the question above? (Do you think Fitbit allows you to control sharing information with your Fitbit friends)*

- Very Unconfident
- Unconfident
- Somewhat Unconfident
- Neutral
- Somewhat Confident
- Confident
- Very Confident

18) Under what conditions do you think Fitbit may sell your data?*

- Whenever they want, with no restrictions
- Whenever they want, as long as your real name and address are not attached to the data profile
- They can sell aggregated, de-identified data that does not identify you
- They can sell aggregated, de-identified data that does not identify you, but only if you opt-in (choose to let them do it)
- Never; they cannot sell your data
- None of the Above
- I don't know

19) How confident are you in your answer to the question above? (Under what conditions do you think Fitbit may sell your data)*

- Very Unconfident
- Unconfident
- Somewhat Unconfident
- Neutral
- Somewhat Confident
- Confident
- Very Confident

20) When do you think Fitbit can collect your location?*

- Fitbit can never collect my location
- Fitbit can only collect my location if I choose to let them (opt-in)
- Fitbit will collect my location when location features, such as maps, of my Fitbit device are active
- Fitbit always collects my location
- None of the Above
- I don't know

21) How confident are you in your answer to the question above? (When do you think Fitbit can collect your location?)*

- Very Unconfident
- Unconfident
- Somewhat Unconfident
- Neutral
- Somewhat Confident
- Confident
- Very Confident

22) For how long do you think Fitbit keeps the data it collects?*

- Until that data item has not been accessed for 6 months
- Until you remove an item from your profile or Fitbit device
- Until you fully delete your Fitbit account
- Forever; it never deletes the data even if you delete your account
- None of the above
- I don't know

23) How confident are you in your answer to the question above? (For how long do you think Fitbit keeps the data it collects)*

- Very Unconfident
- Unconfident
- Somewhat Unconfident
- Neutral
- Somewhat Confident
- Confident
- Very Confident

24) In the event of a data breach of some of its consumer data, how soon do you think Fitbit will contact its users to let them know that their data has been stolen?*

- Within 1 week
- Within 1 month
- Within 3 months
- As specified by law
- Never
- I don't know

25) How confident are you in your answer to the question above? (In the event of a data breach of some of its consumer data, how soon do you think Fitbit will contact its users to let them know that their data has been stolen)*

- Very Unconfident
- Unconfident
- Somewhat Unconfident
- Neutral
- Somewhat Confident
- Confident
- Very Confident

26) Do you think you can use a Fitbit device without having a Fitbit account?*

- Yes and the device will function the same way as with an account
- Yes, but only basic functions will work, such as distance, heartrate and step count.
- Yes, but without an account to maintain calibration data, it won't count steps correctly

exercise							
Time of exercise	<input type="radio"/>						
Sleeping habits	<input type="radio"/>						
Information posted to your Fitbit profile	<input type="radio"/>						

33) Please explain your answer(s) to the question above

34) Please indicate the degree to which you are concerned with Fitbit sharing your personally identifiable information with the following groups.*

	Very Unconcerned	Unconcerned	Somewhat Unconcerned	Neutral	Somewhat Concerned	Concerned	Very Concerned
Government entities	<input type="radio"/>						
Organizations providing services to Fitbit	<input type="radio"/>						
Your Fitbit Friends	<input type="radio"/>						
Organizations you specifically direct Fitbit to share data with (e.g. Facebook)	<input type="radio"/>						

35) Please explain your answer(s) to the question above

36) How would you feel about Fitbit collecting and sharing your location while using the device?*

Completely uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very Comfortable

37) How would you feel about Fitbit keeping a copy of all your data, including data you deleted, until you fully delete your entire Fitbit account?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

38) How would you feel about Fitbit sharing all of your fitness data by default, such as exercise and food consumption, with your Facebook friends?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

39) How would you feel about Fitbit sharing all of your fitness data, such as exercise and food consumption, with friends you add on Fitbit?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

40) How would you feel about Fitbit sharing your personally identifiable information with companies providing services to Fitbit, with no limit to what those companies can do with your information, provided they don't share it?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

41) How would you feel about Fitbit selling your personally identifiable information (information that identifies you) to other companies?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

42) How would you feel about Fitbit selling your information as part of a de-identified, aggregated block (does not identify you) to other companies?*

Very uncomfortable Uncomfortable Somewhat uncomfortable Neutral Somewhat comfortable Comfortable Very comfortable

43) How would you rate your desire to buy and use a Fitbit product in the future?*

No Desire Little Desire Some Desire A lot of Desire I already own and use another fitness wearable device

44) Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared. *

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

45) Consumer control of personal information lies at the heart of consumer privacy. *

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

46) I believe that online privacy is invaded when control is lost or unwillingly reduced as a result of a marketing transaction.*

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

47) Companies seeking information online should disclose the way the data are collected, processed, and used.*

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

48) A good consumer online privacy policy should have a clear and conspicuous disclosure. *

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

49) It is very important to me that I am aware and knowledgeable about how my personal information will be used. *

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

50) It usually bothers me when online companies ask me for personal information.*

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

51) When online companies ask me for personal information, I sometimes think twice before providing it. *

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

52) It bothers me to give personal information to so many online companies.*

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree

53) I'm concerned that online companies are collecting too much personal information about me.*

Strongly Disagree Disagree Mildly Disagree Neutral Mildly Agree Agree Strongly Agree
