

A comparison of users' perceptions of and willingness to use Google, Facebook, and Google+ single-sign-on functionality

Lujo Bauer

Cristian Bravo-Lillo

Elli Fragkaki

William Melicher

Carnegie Mellon University
Pittsburgh, PA
{lbauer, cbravo, elli, billy}@cmu.edu

ABSTRACT

Identity providers such as Google and Facebook are increasingly used to sign in to third-party services like Flickr and USA Today. For users, this can increase convenience (e.g., fewer passwords to remember) and security (e.g., service providers need not keep passwords). At the same time, relying on identity providers that have rich information about users (e.g., all information in a Facebook profile) creates the risk that users will lose oversight or control over the access that service providers are given to this information. To address such concerns, identity providers show users consent interfaces at sign on and provide audit tools for post hoc review.

In this paper we report on a 424-participant on-line study through which we seek to understand the effectiveness of these interfaces: We induced participants to log in with one of three identity providers, and measured their awareness of the information that was being sent by identity providers to service providers, their awareness of identity providers' audit tools, and their sentiment about various aspects of single sign-on. Participants logged in under one of two treatments: a basic treatment, which requested a minimum of personal data; and an invasive treatment, which requested data that most people would find invasive to their privacy. We found that participants' understanding of the information identity providers shared with service providers was based on preconception rather than the content of informational dialogs displayed by the identity providers; and that they were almost uniformly unaware of audit tools offered by identity providers. At the same time, many participants exhibited strong preferences and concerns about data sharing, several of which did not match current data-sharing practices.

Categories and Subject Descriptors

[Security and Privacy]: Security Services—*Authentication*
; [Human Centered Computing]: Interaction Design—*Empirical Studies in Interaction Design*

Keywords

Facebook; Google; single-sign-on; identity providers; privacy

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

DIM'13, November 8, 2013, Berlin, Germany.

ACM 978-1-4503-2493-9/13/11.

<http://dx.doi.org/10.1145/2517881.2517886>

1. INTRODUCTION

Internet users are accumulating more and more identities. A seminal study by Florêncio and Herley found that a typical internet user has 25 different identities, each of which have different credentials [4]. In part because managing these identities and credentials is difficult for users and encourages behaviors like password reuse, the US Government has declared the creation of a digital identity ecosystem a national security priority.¹

In such an ecosystem, *single sign-on* (SSO) systems allow users to authenticate to an *identity provider* (IdP); the IdP in turn vouches for the user to multiple *service providers* (SPs), absolving them of the need to authenticate users themselves. This frees users from remembering many sets of credentials, and service providers from the need to maintain their own authentication mechanisms.

An authentication process with a single sign-on system begins with the enrollment of a user with an identity provider. When a user logs into a service provider using an identity provider, the latter sends or authorizes the former to access a set of attributes about the user. Entities like Facebook and Google are starting to be widely used as identity providers by service providers such as Flickr and USA Today. Both Facebook and Google have social networking capabilities that make them uniquely qualified to provide rich information about users to service providers. Examples of attributes that may be conveyed from an identity provider to the service provider are age, gender, friend list, email address, current location, photos, and relationship status. The convenience to the user of using identity providers hence comes with a potential cost to privacy: the identity provider may send a service provider data that the service provider otherwise would not have known, and identity providers may learn yet more about users by keeping track of which service providers they access.

To address the former concern, identity providers like Google and Facebook explain to the user at (first) login what types of information about the user will be sent to a service provider. The types of information to be transmitted are typically displayed in a *consent dialog*, which gives the user a chance to abort the process of logging in to a service provider if she does not want to share the described information with the service provider. A question that naturally arises in this context is to what extent identity providers in fact succeed at conveying to users what personal information they will share with service providers and give users the ability to make an informed choice.

This paper describes a 424-participant on-line study designed to test the effectiveness of these consent dialogs as implemented by Facebook, Google, and Google+. We induced participants to log

¹<http://www.nist.gov/nstic/>

in to a survey site with one of these identity providers and measured their awareness of the information that was being sent by the identity providers to the survey site, their awareness of identity providers' audit tools, and their sentiment about various aspects of single sign-on. Participants logged in under one of two treatments: a basic treatment, which requested a minimum of personal data; and an invasive treatment, which requested data that we expected most people would find invasive to their privacy.

We found that participants' understanding of the information identity providers shared with service providers was based on preconception rather than the content of consent dialogs shown by the identity providers. Participants were equally likely to consent to logging in regardless of how much information the identity provider was going to send to the service provider. Participants' beliefs about what information identity providers were sending were in general similarly unaffected by the consent dialogs, and they had little understanding of how long and on what occasions the service provider was going to have access to this data. We also found that participants were almost uniformly unaware of audit tools offered by identity providers.

At the same time, many participants exhibited strong preferences and concerns about data sharing, some of which were clearly at odds with current common single sign-on practices. For example, participants overwhelmingly expressed a desire to be at least reminded of the information sent by identity providers to service providers at every login to the service provider; common practice is for identity providers to display this information only once, on initial login. Reported concerns about data sharing were substantiated by participants reporting that they frequently used multiple SSO identities to help control what personal information identity providers shared with specific service providers.

In summary, our study reveals that several aspects of how user information is handled in single sign-on systems are currently largely opaque to users; users neither understand in detail what information about them is sent by identity providers to service providers, nor do they believe they have control over this process. On the other hand, both self-reported data and users' actions indicate a need for better insight and control—including in some specific, relatively easy to implement ways—and suggest that addressing this need would encourage greater adoption of, and satisfaction with, single sign-on.

The remainder of this paper proceeds as follows. We first discuss related work, in Section 2. Next, in Section 3, we describe the design of the study. We present our results in Section 4. We discuss the limitations of our study in Section 6, and conclude in Section 7.

2. RELATED WORK

A number of researchers have investigated tools for privacy-preserving identity management, including in the context of SSO systems. Mowbray and Pearson, for example, proposed an architecture for a cloud computing privacy manager which obfuscates client data before it reaches the server for processing [10]. Weyl et al. proposed a tool for privacy-preserving identity management for multiple services that allows users to have multiple identities and to choose which identity is exposed to which service provider [14]. Steuer et al. proposed a tool for identity verification that allows a semi-trusted identity manager to authenticate the user [12]. This tool allows proving an identity claim to a service provider without fully trusting an identity provider or service provider with all of the contents of an identity claim.

In contrast to these approaches, Facebook's and Google's identity-management systems organize information about users into sometimes coarse-grained categories, and allow service providers to specify which categories of data they wish to access. Each identity



Figure 1: Log-in-with buttons, displayed by web applications (service providers) that provide a single-sign-on service.

provider has a core set of information that is always sent to any service provider that makes use of its services. For example, Facebook's "public profile" includes a user's name, profile picture, age range, gender, and other demographic information, as well as the user's friend list. Google sends a similar set of information when requested to authenticate a user. The types of information that we experiment with in our study are shown in detail in Table 2.

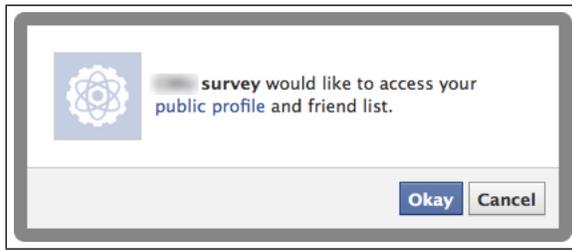
The sets of information that an identity provider will allow a service provider to access are displayed to the user when the user tries to log in to the service provider for the first time. To obtain informed consent from the user, an effective user interface must communicate what private information the service provider can access. Privacy-preserving user interfaces have been the subject of past research that informs our work. Lau et al., for example, have advocated for interfaces that allow the user to create active privacy policies that are automatically applied [7]. However, privacy policies for single sign-on systems require the user to make a decision about each service provider individually.

A study by Sun et al. found that users would value the convenience provided by SSO systems but have privacy and other concerns about adopting SSO systems. The authors found a large number of usability problems with OpenID², a distributed SSO system. They built and tested an identity-enabled browser tool based on their findings. Participants that used the tool made fewer mistakes in a number of tasks that involved logging into websites, compared to participants that used the unmodified version of OpenID [13].

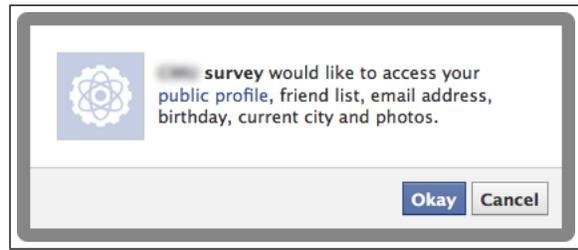
This agrees with the findings of Egelman's study on the privacy and convenience tradeoff of single sign-on systems [3]. This study revealed some aspects of how users react to the privacy statements when signing into a service provider. Participants in the study were observed in a laboratory environment where they were given a choice to complete tasks either by using a SSO system or by creating a new account. The study's findings are generally consistent with ours: users were found to be somewhat aware of the types of data that may be sent to service providers; however, in general users did not notice the details of which data is sent to service providers as conveyed to them by informational display. Although our study has the same focus, it differs from this work in a number of ways: We conduct our study in an online environment, which arguably more faithfully reproduces the context in which users would normally interact with single sign-on systems. We also test three sets of consent interfaces from different identity providers, and examine other aspects of users' interactions with and concerns about single sign-on systems.

An eye-tracking study by Arianezhad et al. sought to find the effect of technical expertise on how users make decisions about SSO by relying on the browser's security indicators [1]. The study found that users with more technical expertise are more likely to look at security indicators than those without technical expertise; however, this alone is not correlated with good use of security indicators. This type of demographic correlation in the use of SSO systems may also hold for reacting to privacy indicators for SSO systems. Although the focus of our work is different, we also examine the relationship between self-reported privacy concern level and behavior in a single sign-on context.

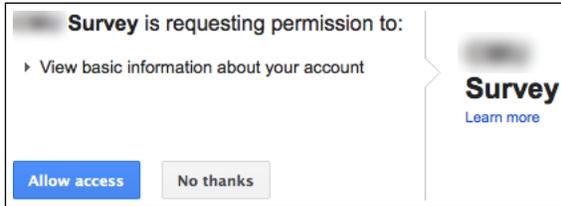
²<http://openid.net>



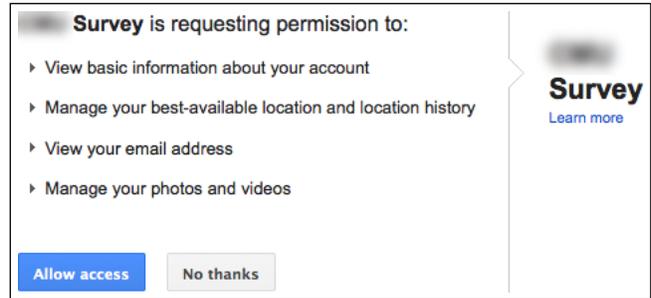
(a) Facebook “basic” condition.



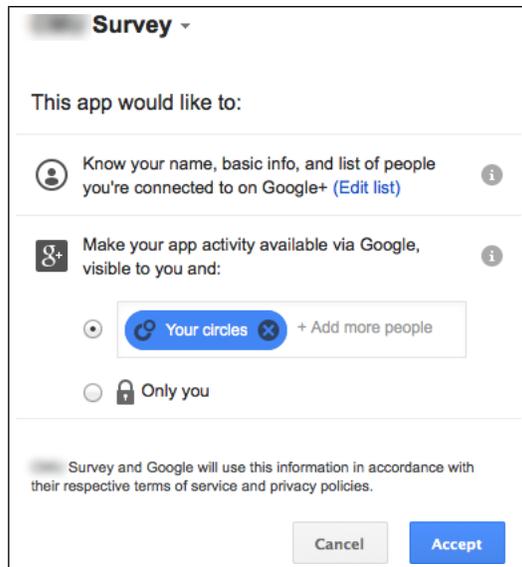
(b) Facebook “invasive” condition.



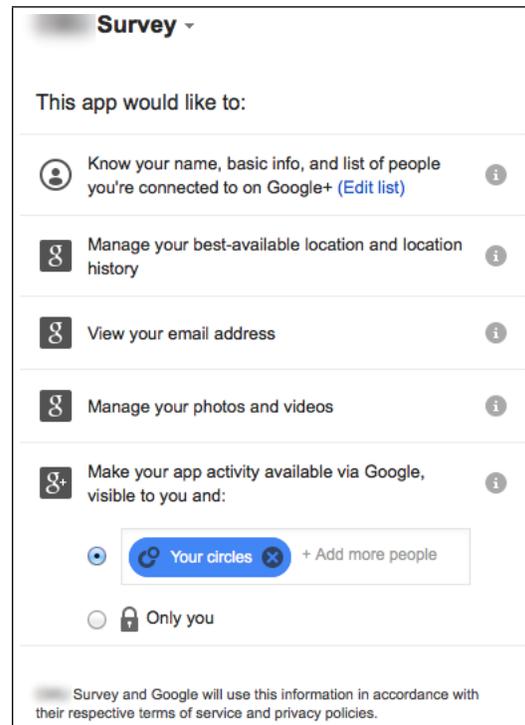
(c) Google “basic” condition.



(d) Google “invasive” condition.



(e) Google+ “basic” condition.



(f) Google+ “invasive” condition.

Figure 2: Consent dialogs presented to participants after they clicked the log-in-with button.

3. METHODOLOGY

We conducted a between-subjects online study to investigate the effectiveness of identity providers’ informational and consent dialogs shown during single sign-on (“consent dialogs” throughout this paper). Participants were asked to log in to our survey using a popular single sign-on provider (Facebook, Google, or Google+), which involved consenting to the release of information by the identity provider to the site on which the survey was hosted. The survey asked participants a series of questions about their understanding of what information about them was released by the identity provider; their attitudes towards single sign-on; their understanding of the privacy features offered by identity providers;

which involved consenting to the release of information by the identity provider to the site on which the survey was hosted. The survey asked participants a series of questions about their understanding of what information about them was released by the identity provider; their attitudes towards single sign-on; their understanding of the privacy features offered by identity providers;

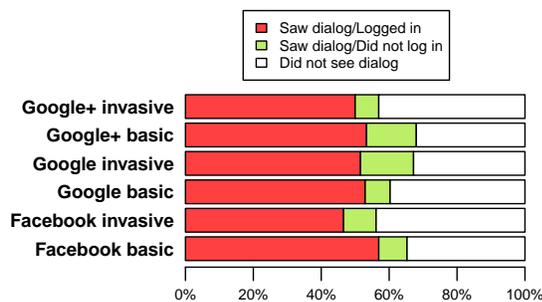


Figure 3: Participants’ responses to consent dialogs, per condition. In the beginning of our survey, a participant chooses either to login with the corresponding service provider or not. If she does, she is presented with a consent dialog, in which she either allows passing the information to the service provider or not. The graph shows these three different outcomes.

and demographic information. Participants were assigned to conditions that differed in the choice of identity provider and the amount of information that was released by the identity provider to the survey site. Our study design was reviewed and approved by our institutional review board (IRB).

We next describe the design of the study, including recruitment and study conditions, as well as the analyses we performed.

3.1 Recruitment

We created three Human Intelligence Tasks (HITs) in Amazon’s Mechanical Turk (mturk) crowdsourcing service. Each HIT was advertised as a study of “*identity-provider* users’ habits while surfing online,” where *identity-provider* was Facebook, Google, or Google+. Potential participants were told they may be asked to evaluate a web site, and were required to be over 18 years old and have an account with the identity provider. Participants were paid \$1.00, which is a typical payment on mTurk for a survey of similar length.

3.2 Study procedure

After participants accepted our HIT and agreed to a study consent form, we directed them to a web page where we displayed the message “Please log in with *identity-provider* to confirm you are an *identity-provider* user,” along with the corresponding log-in-with button (Figure 1), and the message “... or click [here](#) to continue to the survey without logging in.” The latter message was displayed in a smaller font than the former. If a participant clicked on the log-in-with button, she was redirected to the identity provider, which asked her to log in and informed her about the information about her that would be shared with the study web site. These consent dialogs for the three providers are shown in Figure 2. The dialogs were implemented by the identity providers, and we did not modify them. Each consent dialog offered the participant two options: allow the identity provider to share the described information with the survey web site, or cancel sign on. Agreeing to share would redirect the participant to our survey; canceling would redirect her to the previous page, where she would choose between logging in with the identity provider or taking the survey without logging in.

Survey.

The survey contained two sections: data sharing and demographics. In the first section, we asked our participants what types of information about them were known by their identity providers; what part of that information was in their public profiles; and what infor-

mation was sent to researchers through the single sign-on process. Also in this section were questions about how comfortable participants were with sharing different types of information with different websites, and participants’ preferences and history with logging into websites with the identity provider from their condition. Questions about the participants history and preferences included whether the participant had used this identity provider in the past, whether the participant had ever started logging in with this identity provider and decided not to, and what factors influenced this decision. Notably, this section also asked participants about how much control participants felt they had over their data. In general, the data-sharing section measured participants knowledge about the SSO process, and their preferences for using SSO systems. The purpose of these questions was to reveal whether participants understood the SSO process and what data will be sent about them to the service provider, and to elicit participants’ privacy preferences both with direct questions and indirectly by asking them about their previous behaviors.

In the second section of the survey we asked typical demographic questions such as participant’s gender, age, race/ethnicity, occupation, and level of education. In addition, we tried to gauge their technical knowledge and depth of their online experience in general. In this section we also measure participants’ inclinations towards privacy, which we call *privacy concern level* (PCL), via a Likert scale composed of five questions that explore which privacy-preserving online behaviors participants had engaged in, and to what extent. We used this metric to understand how privacy-inclined our participants were, and whether this tendency had any influence on their behavior. We provide more details about this metric, including the questions we used, in Appendix B.3.

The survey included two questions to give us an idea of a participant’s technical savviness: whether the participant knew a programming language, and whether the participant had studied or worked in a computer-science-related field. While questions like whether the participant has ever read a website’s privacy policy, and whether the participant has ever used a social networking site, describe their general online experience.

3.3 Conditions

The conditions we tested differed in two dimensions: choice of identity provider, and the set of attributes that the identity provider would release to the survey site. We tested three identity providers (Facebook, Google, and Google+) and two levels of information sharing (basic and invasive). Our study had a full-factorial design, resulting in six conditions. Each participant saw exactly one dialog, and responded to the survey only once.

We tested multiple identity providers for several reasons. One reason was to examine and (partially) control for differences in participant behavior caused by the specific way in which each identity provider informs users what information about them will be released. A second reason was to examine whether participants behaved differently depending on whether the identity provider they were using was an online social network.

We tested two levels of information release: the minimum supported by the identity provider (we call this the *basic* treatment); and a richer set of information that included categories of information that we expected users to be less willing to release (the *invasive* treatment). We made each treatment as uniform as possible across the three identity providers. The identity providers supported slightly different types of information release, which resulted in some differences between identity providers. In the invasive Facebook condition the identity provider would ask to reveal the participant’s email address, birthday, current city, and photos;

while in the invasive Google and Google+ conditions the identity provider would ask to reveal email address, best available location and location history, and photos and videos.

Table 2 shows in detail the sets of attributes that each condition included, and Figure 2 shows the identity providers’ consent dialogs for all six conditions.

3.4 Analysis

We conducted a number of statistical tests to explore our participants’ answers. Throughout Section 4 we describe briefly each test where applied, and we report the corresponding statistics in those cases where our tests were significant or where the lack of statistical significance was notable. All tests were run at the usual $\alpha = .05$ significance level.

4. RESULTS

In this section we present the results of our study. We first discuss participant demographics (Section 4.1) and then detailed results about participants’ understanding of consent dialogs and attitudes towards various aspects of single sign-on (Section 4.2).

4.1 Participants

We recruited 482 participants from Amazon’s Mechanical Turk, in three separate Human Intelligence Tasks (HITs), each advertising a different identity provider. After removing participants who reported they used a “fake” single sign-on account or failed to correctly answer a survey question designed to reveal whether participants were paying attention³, we were left with 424 participants; the rest of the paper focuses exclusively on these. Participants were 30.7 years old on average ($\sigma = 9.7$); 49.5% were female; 64.4% were Caucasian; 84% were from the United States; and the top two reported occupations were “student” (13.7%) and “science, engineering and IT professional” (12.3%). 23.3% of participants reported having knowledge of a programming language. In comparison with typical demographics measured for mturk workers, our sample is more educated and shows a higher proportion of participants from the United States. [2, 5, 6, 8, 9].

4.2 Understanding Consent Dialogs

We found participants to be somewhat aware of the range of attributes passed by the identity provider to the service provider, but the factors that affected their awareness were largely not exposed by this experiment. Participants appeared to have a preconceived idea of which attributes would be sent based on the identities of the identity provider and service provider. Their precise understanding of what is sent, and their willingness to log in, was not significantly affected by consent dialogs. Instead, it was affected by their privacy concern level (see Appendix B.3). This suggests that users have already made a decision about whether to log in with the identity provider before viewing the dialogs. However, participants who see a consent dialog alerting them that a larger set of attributes will be sent to the service provider (our invasive conditions) do realize that more attributes are being sent, even if not which ones.

Login rates per service provider.

Figure 3 shows the proportion of participants who decided to log in with each service provider, per condition. We were interested in understanding whether the information displayed in the consent dialogs influenced participants’ logging behavior. Since only those participants who clicked on the log-in-with button actually saw the

³Question 39 in Appendix B.

| Factor | Estimate | Std. Error | z-value | p-value |
|----------------|----------|------------|---------|---------------|
| is.invasive | -0.2521 | 0.2064 | -1.2210 | 0.2220 |
| is.socialnet | 0.0053 | 0.2565 | 0.0210 | 0.9836 |
| is.facebook | -0.1881 | 0.2539 | -0.7410 | 0.4589 |
| has.itdegree | -0.0333 | 0.2904 | -0.1150 | 0.9086 |
| priv.con.level | -0.4326 | 0.1572 | -2.7520 | 0.0059 |
| fam.with.SP | 0.4444 | 0.2142 | 2.0750 | 0.0380 |
| know.prog.lang | 0.5756 | 0.3014 | 1.9100 | 0.0561 |

(a) Results of a logistic regression applied to participants’ decision of whether to click on the log-in-with button. Values in bold indicate (marginal) statistical significance.

| Factor | Estimate | Std. Error | z-value | p-value |
|----------------|----------|------------|---------|---------|
| is.invasive | -0.1362 | 0.3357 | -0.4060 | 0.6849 |
| is.socialnet | 0.0706 | 0.4025 | 0.1750 | 0.8608 |
| is.facebook | 0.1210 | 0.4197 | 0.2880 | 0.7730 |
| has.itdegree | -0.3755 | 0.4588 | -0.8180 | 0.4131 |
| priv.con.level | -0.1525 | 0.2492 | -0.6120 | 0.5405 |
| fam.with.SP | 0.2603 | 0.3481 | 0.7480 | 0.4547 |
| know.prog.lang | 0.8068 | 0.4993 | 1.6160 | 0.1061 |

(b) Results of a logistic regression applied to participants’ decision of whether to agree to the consent dialog. Participants included in this regression were only those who saw the consent dialog, that is, only those who clicked on the log-in-with button.

Table 1: Results of logistic regressions. Variables used in the tables are defined as follows: *is.invasive*: true when a participant is assigned to an invasive treatment (see Section 3.3); *is.socialnet*: true when a participant is assigned to a condition in which the identity provider is either Facebook or Google+; *is.facebook*: true when a participant is assigned to a condition in which the identity provider is Facebook; *has.itdegree*: true when the participant answered affirmatively “Do you have a degree in an IT field...?”; *priv.con.level* (privacy concern level): Likert scale composed of 5 questions described in Appendix B.3, that measure participants’ level of privacy awareness; *fam.with.SP* (familiarity with the service provider): binary variable; true when a participant answered the question “Have you ever used ‘log-in-with service-provider’ to log into a website (other than this survey)?” with “A few times” or “Multiple times” (other possible answers were “Never,” “Once,” and “I’m not sure”).

dialog, we had two binary outcomes: participants either clicked on the log-in-with button or not, and those who did either agreed to the dialog (i.e., clicked “Okay,” “Allow access,” or “Accept,” depending on identity provider) or not.

Tables 1a and 1b show the results of logistic regressions applied to both of the previously described participants’ decisions.

There was no statistically significant difference by condition in which participants logged in to our study via an identity provider (see Figure 3 and the first three rows in Tables 1a and 1b). Closest to significant was that participants in invasive conditions that used a social network identity provider had a marginally higher tendency to log in compared to participants in other conditions ($p = 0.08$). This disproves our hypothesis that participants would log in less if that required revealing more information.

As seen in Table 1a, two variables had a significant influence on participants’ willingness to click on the “log-in-with” button: the privacy concern level ($e = -0.433$, $p = 0.0059$), and Familiarity with the service provider ($e = 0.444$, $p = 0.038$). Self-reported knowledge of a programming language had a marginally significant ($p = 0.0561$) but strong, positive influence ($e = 0.5756$) on the decision.⁴

⁴That is, participants who self-reported as having knowledge of a programming language were more likely to click on the “log-in-with” button than those who did not; however, the probability that this observed effect is due to pure chance is higher than for the

| Data type | Facebook | | Google | | Google+ | | Comments |
|--------------------------------|----------|----------|--------|----------|---------|----------|--|
| | Basic | Invasive | Basic | Invasive | Basic | Invasive | |
| Full name | Y | Y | Y | Y | Y | Y | 1,2,3 |
| User ID | Y | Y | | | Y | Y | 3; User ID formats differ between providers |
| Profile Picture | Y | Y | Y | Y | Y | Y | 1,2,3 |
| Profile URL | | | Y | Y | Y | Y | 2,3 |
| Age range | Y | Y | | | Y | Y | 1,3 |
| Gender | Y | Y | Y | Y | Y | Y | 1,2,3 |
| Language | Y | Y | Y | Y | Y | Y | 1,2,3 |
| Country | Y | Y | Y | Y | Y | Y | 1,2,3 |
| Friend list | Y | Y | | | Y | Y | Facebook "friend list" and Google+ "circles" |
| Email address | | Y | | Y | | Y | |
| Birthday | | Y | Y | Y | Y | Y | 2,3 |
| Current city | | Y | | | | | |
| Location | | | | Y | | Y | "Best-available current location" and "Best-available location history" in Google Latitude |
| Photos | | Y | | Y | | Y | |
| Photo tags | | | | Y | | Y | 3 |
| Videos | | | | Y | | Y | |
| Video tags | | | | Y | | Y | 3 |
| Timezone | | | Y | Y | Y | Y | 2,3 |
| App activity | | | | | Y | Y | Default is to share it to circles; can be restricted to only the person |
| Other public, unspecified info | Y | Y | | | Y | Y | 1 |

Table 2: Attributes used in this study, per identity provider and study treatment (see Section 3.3 for a description of conditions and treatments). Notes in the rightmost column correspond to: [1] Included in Facebook "Public profile." [2] Included in Google/Google+ "Basic information." [3] Must open a secondary dialog during consent process to see that this information is sent.

Remarkably, all factors stop being significant for the second decision (see Table 1b). Since participants could only find out what attributes would be passed to a service provider after examining the consent dialog that popped up during the second decision, this strongly suggests that the knowledge of what information was passed to the service provider did not influence participants' decision about whether to use the log-in-with functionality.

Awareness of transmitted information.

In order to determine whether participants were aware of the information being passed to the survey site when they consented to the consent dialog, we first controlled for what information was actually being sent; the list of attributes sent is shown in Table 2. We compared this list to participants' answers about whether each particular type of information was passed to the service provider.

Figure 4 shows the proportion of participants who believed that a specific attribute was shared with the service provider, per condition, as well as whether the attribute was in fact shared. If consent dialogs were effective at communicating which attributes were passed to the survey website for each provider, we would expect strong correlation between an attribute being shared and participants indicating that it was shared. In general, however, we found that the majority of participants believed that various attributes were not shared even when they were.

Between 22% and 42% of participants were able to ascertain correctly that their full name was being passed to the service provider. While in conditions where Facebook or Google were identity providers this can be explained by noticing that full name was not mentioned explicitly in the corresponding consent dialogs (see Figures 2a through 2d) and by assuming that most participants did

other two factors.

not follow any links within dialogs to inspect carefully which data types were passed to the service provider, in conditions where Google+ was the identity provider this was not the case, as can be seen in Figures 2e and 2f.

While between 69% and 80% of participants indicated correctly that their email addresses were not being passed in the basic conditions, only between 25% and 43% were correct in the invasive conditions. The same tendency can be observed for photos/pictures. Here, the large discrepancy between the basic and invasive conditions is due to participants' beliefs about information transmitted being relatively similar across conditions; sometimes these beliefs were aligned with reality (e.g., email addresses were not sent in basic conditions), and sometimes it was not (e.g., email addresses were sent in invasive conditions).

Since there is little correspondence between actual data being passed (Table 2) and the data participants believed was passed (Figure 4), this suggests that participants simply acknowledged the dialog without examining it, and later made an educated guess about what information was being passed.

An alternate explanation is that participants understood the consent dialogs and chose whether to proceed based on this understanding, but forgot the content of the consent dialogs before answering the survey question about what attributes they thought were shared. Our next finding strongly suggests that this explanation is not likely, however.

Trust and willingness to share.

We asked those participants who thought that a specific attribute was being sent to the service provider how comfortable they would feel if they had to share that attribute with "a trusted website" and an untrusted one. Figure 5 includes our participants' answers for the same attributes shown in Figure 4.

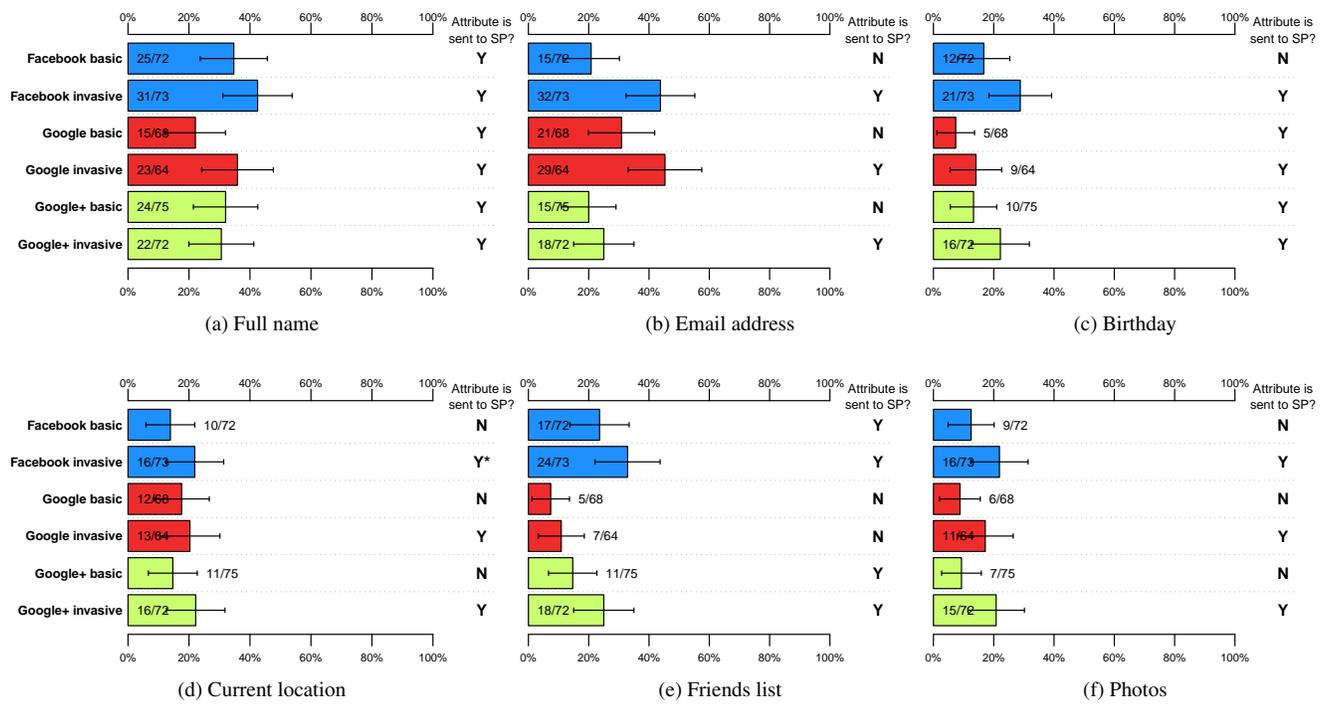


Figure 4: Proportion of participants who answered “Yes” when asked whether an attribute would be passed to the service provider, per condition. The letter “Y” or “N” to the right of each bar indicates whether the corresponding attribute is actually passed to the service provider. Whiskers to the right of each bar correspond to 95% confidence intervals. In Figure 4d, the attribute shared by Facebook is “current city,” which is Facebook’s closest equivalent to Google’s “current location.”

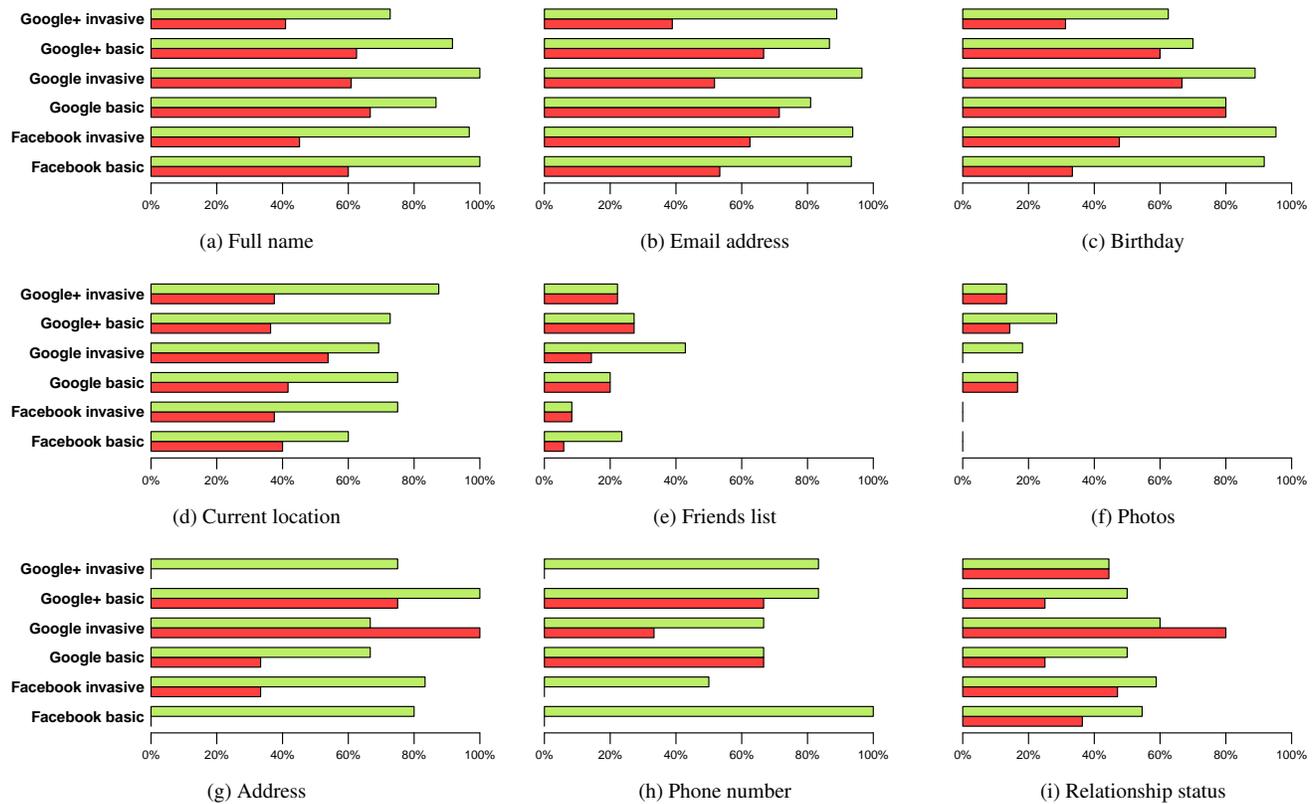


Figure 5: Proportion of participants that reported to feel either “comfortable” or “very comfortable” about sharing a specific attribute with a trusted website (green; the upper bar of each pair) and an untrusted website (red).

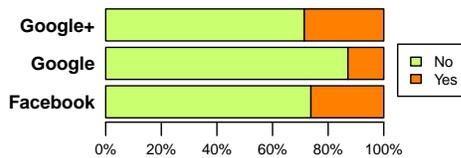


Figure 6: Answers to the question “Would you choose something else if you had the option to log in with another provider (like Y) instead of X?” X corresponded to the provider assigned to each participant (either Facebook or Google), whereas Y corresponded to the other provider (i.e., Facebook if X was Google and vice versa.) An omnibus χ^2 test is significant ($\chi^2 = 11.09, p < 0.004$) indicating that participants in the Google condition were significantly more likely to answer “no” than participants in the other conditions.

Participants’ self-expressed willingness to share specific attributes differed depending on their level of trust in the service provider. These attributes include address (Figure 5g), birthday (Figure 5c), phone number (Figure 5h), and current location (Figure 5d).

Notably, there is a mismatch between participants’ comfort level with sharing certain attributes and the sets of attributes that are usually shared. E.g., participants are very uncomfortable sharing their friend lists with a service provider, yet these are always shared.

Participants using Google as their identity provider were significantly less likely to answer “No” to the question “Would you choose something else if you had the option to log in with another provider (like Facebook) instead of Google?” (Figure 6). This suggests that participants using a specific social network as their identity provider were more likely to want to use a different identity provider for specific sites, or, alternately, to be influenced by the choice of identity provider when deciding whether to create an account with the service provider or sign in via the identity provider.

Participants had little insight into the level of access service providers received from identity providers to user attributes. 38% of participants (161 of 424) erroneously thought that the service provider could access the attribute exactly once; 45% (192 of 424) were unsure.

These findings all strongly suggest that we need better mechanisms to convey to users what attributes an identity provider is about to send to a service provider.

4.3 Need for Control

About half (210 of 424) of participants reported feeling that they had no control over the information passed by an identity provider to the service provider. Only 4.5% (19 of 424) reported that they had “a lot of” control. At the same time, 94% of participants (399 of 424) thought it was “very” or “extremely” important to them to have such control, with another 5% calling it “important.” In fact, a minority of participants (13%) reported creating different identity providers accounts purely for controlling the amount of information sent by an identity provider to a service provider.

Two thirds of participants (283 of 424) expressed a desire to have some level of control over or insight into which information is sent by an identity provider to a service provider during every transaction between the two. Of these, 78% (221 of 283) expressed a preference for being reminded rather than being asked for a decision each time they log into the service provider. Participants’ preferences about how to log in to a service provider were split: 29% (125 of 424) preferred to do so via an identity provider and 44% (185 of 424) by creating a new account with the service provider; 27% (114 of 424) would make this choice depending on context. For most (77%, 327 of 424), this preference does not depend on

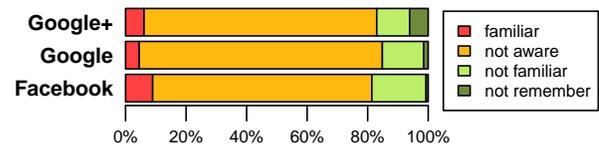


Figure 7: Participants’ answers to the question pertaining awareness with the audit tool. An omnibus χ^2 test was significant ($\chi^2 = 14.07, p = 0.029$); however, since we were not expecting this value to be significant we did not conduct further analysis.

the identity of the identity provider.

Participants had strong but inconsistent preferences about whether they preferred using a single identity provider for most purposes (37%, 157 of 424 participants), or preferred using multiple identity providers (48%, 203 of 424 participants).

A large minority (43%, 183 of 424) of participants reported that they had at least once started to log in to a service provider via an identity provider, but then changed their mind.

4.4 Knowledge of Audit Tools

Participants had minimal familiarity with out-of-band tools to audit and control the flow of information between IdPs and SPs; at the same time, most participants reported that the availability of such tools would increase their willingness to use an identity provider to log into a service provider.

The vast majority of participants (74%) was not aware of audit tools provided by the identity provider; most others (16%) were aware of but not familiar with these tools. At the same time, almost half (48%) of the participants reported that the availability of an effective audit tool would cause them to log in more often with an identity provider.

Similarly, most participants (84%) were not aware of or were not familiar with tools that allowed them to change post-hoc the sharing decisions they had previously made. Most participants (71%) thought the availability of such a feature would make them more likely to log in with the identity provider.

Although generally poor, participants’ awareness of these tools was significantly different between providers (see caption of Figure 7). However, since we were not expecting this result, we did not conduct further analysis.

4.5 Login Time

Finally, we measured the time that participants took between the moment when they clicked the log-in-with button, and when they entered the survey (either by consenting to the dialog or not).

There was no significant difference between basic and invasive treatments in how long participants took to log in (a Wilcoxon rank sum test yielded a statistic $W = 24563, p = 0.097$). However, participants in the Google+ conditions took about twice as long to log in as participants in conditions with other identity providers, which is a statistically significant difference (Kruskal-Wallis rank sum test, $\chi^2 = 36.25, p < 0.001$). Although interesting, this is not surprising, as Google+ dialogs were considerably longer than the other two consent dialogs (see Figure 2). Perhaps more surprising is that though participants took longer to read Google+ dialogs, this was not reflected (as previously described) in a better understanding of the content they convey.

5. DISCUSSION

Our results show that participants’ understanding of which information is passed to service providers largely was not affected by

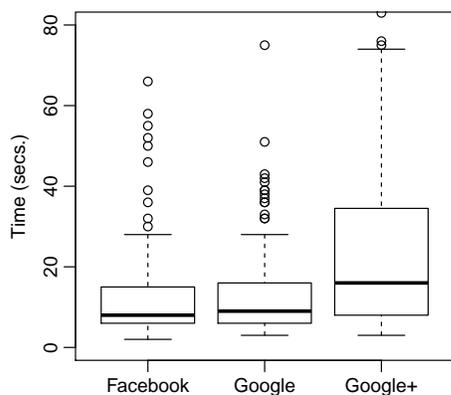


Figure 8: Box plot summarizing participants’ response times to consent dialogs, per identity provider. Boxes show the interquartile range (i.e., 75th-25th percentiles), thick black lines within the boxes show the median response times, and whiskers show the minimum and maximum values that fall within the range enclosed by the median ± 1.5 times the interquartile range. Dots represent outliers, i.e., values that fall outside of the aforementioned range. The Y-axis was constrained to the range 0–80 seconds, omitting a number of outliers. Median times were 8 (Facebook), 9 (Google), and 16 (Google+) seconds, while average times were 16.2, 14.5, and 32.3 seconds, respectively. A Kruskal-Wallis rank sum test shows that the timing distribution for Google+ is significantly different from the other two identity providers ($\chi^2 = 36.25, p < 10^{-8}$).

the consent dialogs shown by identity providers. However, participants did have a general idea of the types of data that can be sent. Few participants saw the dialogs and chose not to proceed with the login, which strongly suggests that this general idea is formed largely before the participant is actually shown the specifics of the information that will be sent. Because our study did not find a statistically significant difference by condition in whether a user chose to log in via an identity provider, neither the amount of information requested nor the specific identity provider (and the corresponding styles of presenting information in consent dialogs) significantly affects whether users will choose to create a new account with the service provider or use a SSO system. However, participants did convey that the level of trust they have in the service provider would affect whether they share data with that service provider.

Communicating with users.

In general, users’ lack of understanding of what data is sent to service providers and the ineffectiveness of consent dialogs at communicating this information suggest that new methods for communicating this to users are required. Users’ preconception of the data that is sent to service providers likely complicates effective communication through consent dialogs. Participants did notice the difference between more categories of data (i.e., the invasive conditions) and fewer (the basic conditions); this suggests that users can notice certain aspects of consent dialogs. A possible way to make dialogs more effective would be to highlight categories of data that are particularly privacy-sensitive by using special attractors or by reordering the list of categories by putting the most significant categories first. These methods may ameliorate the effect of users’ preconceptions and allow users to be alerted to specific privacy issues. More research would be required to reveal exactly what methods of communicating with users are most effective.

The lack of awareness of resources to control and audit data

that passes between identity providers and service providers contributes to participants’ feelings of lacking control over this information flow. Because participants stated that knowledge of these tools would make them more likely to log in with SSO systems, communicating the existence of these tools may result in higher feelings of control and increased use of these systems. This could benefit both users, who may get more control over this process, and identity providers, which may get more people to use their systems.

Matching user preferences.

Some specific preferences of our participants appear to be out of sync with current implementations of the single sign-on process. In particular, participants’ ideas of what is very private does not seem to match with identity providers’ handling of personal information. An example of this is that Facebook *always* sends service providers the user’s friend list, even when sending the least amount of information, although participants stated that they were very uncomfortable sharing this information. Better aligning what is commonly shared with users’ notions of privacy sensitivity would likely increase users’ comfort with the single sign-on process.

Participants also expressed a strong desire to be kept better informed about the information that was being sent to service providers on subsequent logins, during which identity providers typically send information to service providers without giving any notice or control to users. It seems likely that adding such functionality would further increase users’ comfort without negatively impacting how much information they are willing to allow to be shared.

6. LIMITATIONS

As with other studies, our work has some limitations that are important to keep in mind when interpreting the results. We next discuss the most significant ones.

The details of the user interfaces for the three single sign-on systems we tested change frequently, and we tested with only three identity providers. Hence, one concern is whether our results hold only for exactly the interfaces that were in use while we ran our study, or perhaps only for the specific identity providers. However, while the interfaces will undoubtedly change, the three sets we tested were similar in the functionality they offered, indicating some degree of convergence. Each had a “log in with X” button where X is Facebook, Google, or Google+. After clicking this button, each presented a dialog which described what categories of data the service provider could access and asked for consent from the user (see Figure 2). Furthermore, we observed surprisingly few differences in participants’ understanding of the information presented in these consent dialogs, suggesting that this finding is robust to minor changes in the appearance of dialogs and the choice of identity provider.

Our population was recruited from Amazon’s Mechanical Turk, and thus, like most study populations, does not match the population of all users likely to encounter single-sign-on systems. However, mturk has been used in other studies and been proven to be, with appropriate care, a good source of high-quality data [2, 5, 6, 8, 9]. Furthermore, mturk has better population diversity than populations typically found in on-campus laboratory environments, and allowed us to have a high volume of participants.

Our study was performed online as opposed to in a laboratory environment. The study by Egelman that also investigates users’ understanding of single sign-on was performed in a laboratory environment, which may allow more insight into user’s thoughts and actions [3]. However, a laboratory setting may also influence participants’ behavior, who often feel more trusting and safe in such a context and hence may behave differently than they would nor-

mally. In this situation, an online environment is also arguably closer to the actual experience that users have when interacting with SSO systems. The choice of an online environment also allowed us to have a greater volume of participants and examine more aspects related to participants' understanding of and attitudes towards the single sign-on process.

Finally, our participants, by selecting one of the three HITs we were advertising on mturk, could effectively select the identity provider and potentially introduce some self-selection bias, which may reduce the internal validity of the results. However, we found no significant demographic differences between participants assigned to different identity providers. Additionally, the main effect of a potential self-selection bias by identity provider would have likely been to lower confidence in results that show differences in behavior or opinion by identity provider; however, we found no such differences.

7. CONCLUSIONS

Our participants were unable to recognize what data types were passed from identity providers to service providers during the login process, meaning current consent dialogs which are meant to convey this information are ineffective at doing so. One possible explanation for why consent dialogs are ineffective at informing users is that identity providers are not motivated to effectively communicate these privacy concerns to users. We hypothesize that identity providers like Facebook and Google do not make money by providing better privacy settings, but by having people use these platforms. Our results show, however, significant misalignment between current single sign-on processes and users' expectations and needs. Our results further suggest that this misalignment discourages some users from participating in single sign-on, or causes them to develop workarounds like maintaining "fake" identities. We hope that the results of this study can inform the design of better consent dialogs that are more effective at communicating privacy concerns to users.

8. ACKNOWLEDGMENTS

This work was supported in part by a grant awarded to the University Corporation for Advanced Internet Development (Internet2) under the sponsorship of the U.S. Department of Commerce, National Institute of Standards and Technology; and by NSF grants CNS0831428, CNS1116934, CCF0917047, and DGE0903659.

9. REFERENCES

- [1] M. Arianezhad, L. J. Camp, T. Kelley, and D. Stebila. Comparative eye tracking of experts and novices in web single sign-on. In *Proc. of 3rd ACM Conf. on Data and application security and privacy*, CODASPY '13, 2013.
- [2] J. S. Downs, M. B. Holbrook, S. Sheng, and L. F. Cranor. Are your participants gaming the system? Screening Mechanical Turk workers. In *Proc. of SIGCHI Conf. on Human Factors in Computing Systems*, CHI '10, 2010.
- [3] S. Egelman. My profile is my password, verify me! The privacy/convenience tradeoff of Facebook Connect. In *Proc. of SIGCHI Conf. on Human Factors in Computing Systems*, CHI '13, 2013.
- [4] D. Florencio and C. Herley. A large-scale study of web password habits. In *Proc. of 16th International Conference on WWW*, WWW '07, 2007.
- [5] J. K. Goodman, C. E. Cryder, and A. Cheema. Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *J. Behav. Dec. Making*, 2012.

- [6] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with Mechanical Turk. In *Proc. of SIGCHI Conf. on Human Factors in Computing Systems*, CHI '08, 2008.
- [7] T. Lau, O. Etzioni, and D. S. Weld. Privacy interfaces for information management. *Commun. ACM*, 42(10):88–94, Oct. 1999.
- [8] W. Mason and S. Suri. Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44(1):1–23, Mar. 2012.
- [9] T. K. Michael Buhrmester and S. D. Gosling. Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, pages 3–5, 2011.
- [10] M. Mowbray and S. Pearson. A client-based privacy manager for cloud computing. In *Proc. of the 4th International ICST Conf. on COMMunication System softWare and middlewaRE*, COMSWARE '09, 2009.
- [11] R. Rosenthal and R. L. Rosnow. *Essentials of Behavioral Research*. McGraw Hill, 2008.
- [12] K. Steuer, Jr., R. Fernando, and E. Bertino. Privacy preserving identity attribute verification in Windows CardSpace. In *Proc. of the 6th ACM Workshop on Digital Identity Management*, DIM '10, 2010.
- [13] S.-T. Sun, E. Pospisil, I. Muslukhov, N. Dindar, K. Hawkey, and K. Beznosov. What makes users refuse web single sign-on? An empirical investigation of OpenID. In *Proc. of the 7th Symposium on Usable Privacy and Security*, SOUPS '11, 2011.
- [14] B. Weyl, P. Brandao, A. Gomez Skarmeta, R. Marin Lopez, P. Mishra, H. Ziemek, and C. Hauser. Protecting privacy of identities in federated operator environments. In *Proc. of the 14th IST Mobile & Wireless Communications Summit*, June 2005.

APPENDIX

A. MTURK SOLICITATION TEXT

Researchers at Carnegie Mellon University are conducting a study to determine some of the habits of [Facebook|Google|Google+] users. You may be asked to evaluate a website, then answer a few questions about the website. The whole survey should take you between 10 and 15 minutes in total. We will pay you \$1 for your participation.

Requisites to participate:

1. You must be 18 years old or older.
2. You must have a [Facebook|Google|Google+] account.
3. You must not take this survey (including other versions of it) twice.

To be paid, follow these steps:

1. Go to: [URL shown here]
2. After completing the survey you will receive a confirmation code in the last page. Enter the code in the box below and we will approve your payment. Please do not enter the code more than once. If you are not sure about having entered the code correctly, please send us a message instead of trying to send the HIT twice. Please do not make up codes. If you make up a code to obtain the payment, we will reject your HIT.

Enter your code here: [Textbox here]

For questions and problems, please contact us through Mechanical Turk's contact functionality.

B. SURVEY QUESTIONS

B.1 Data sharing section

1. For the following types of information about you, check all of the following to indicate whether Facebook/Google/Google+ knows this information, this information is part of my Facebook/Google/Google+ public profile, Facebook/Google/Google+ sent this information:
 - (a) My Facebook/Google/Google+ username
 - (b) My Facebook/Google/Google+ password
 - (c) My full name
 - (d) My email
 - (e) My address
 - (f) My phone number
 - (g) My birthday
 - (h) My current location
 - (i) My relationship status
 - (j) My friend list
 - (k) My photos
2. Imagine you are visiting a news or a shopping website. How comfortable do you feel sharing the following information with the website?
 - (a) Your full name
 - (b) Your email
 - (c) Your address
 - (d) Your phone number
 - (e) Your birthday
 - (f) Your current location
 - (g) Your relationship status
 - (h) Your friend list
 - (i) Your photos
3. Imagine you are visiting the website of a trusted organization (for example, your bank or your health insurance company). How comfortable do you feel sharing the following information with the website?
 - (a) Your full name
 - (b) Your email
 - (c) Your address
 - (d) Your phone number
 - (e) Your birthday
 - (f) Your current location
 - (g) Your relationship status
 - (h) Your friend list
 - (i) Your photos
4. Which of the following statements do you think is true for the information Facebook/Google/Google+ sent information to the website?
 - (a) Facebook/Google/Google+ sent my information only once
 - (b) Facebook/Google/Google+ may send my information multiple times, until I exit the study website
 - (c) Facebook/Google/Google+ will send my information each time I log in to the study website
 - (d) Facebook/Google/Google+ may send my information multiple times, as long as the study website is up
 - (e) I am not sure
5. Suppose that you used 'log in with Facebook/Google/Google+' to log in to a website (e.g., library website, online game, etc). As part of this process the website receives some information from Facebook/Google/Google+. Which of the following best describes your preference?
 - (a) I want to approve what information is sent each time I log in
 - (b) I want to approve what information is sent only the first time
 - (c) I don't want to approve what information is sent, it's always fine
 - (d) I don't care
6. Would you prefer to be reminded of what information is sent from Facebook/Google/Google+ to the website instead of approving it every time?
7. Have you ever used 'log in with Facebook/Google/Google+' to log in to a website (other than this survey)?
8. Suppose that you want to visit a website that offers two choices for logging in: to create a new account on the website or to 'log in with Facebook/Google/Google+'. What would you prefer to do?
 - (a) Create a new account
 - (b) Log in with Facebook/Google/Google+
 - (c) It depends on the website
9. Why?
10. Would you choose something else if you had the option to log in with another provider (like Facebook/Google/Google+⁵) instead of Facebook/Google/Google+?
11. Have you ever started 'logging in with Facebook/Google/Google+' to a website but decided not to?
12. How important were the following factors in influencing your decision in the previous question?
 - (a) How much I trust or distrust Facebook/Google/Google+
 - (b) How much I trust or distrust the external website
 - (c) How well I was explained what/when my information would be sent to the website
 - (d) Whether I could choose what information would be sent to the website
 - (e) How much information the website was asking forAnswers were on a 5 point scale from 'Not important' to 'Extremely Important'
13. Please explain any other factors that may have influenced your decision.
14. Do you think that 'logging with Facebook/Google/Google+' to a website gives you any level of control over the information that is passed from Facebook/Google/Google+ to the website?

Answers: I have a lot of control; I have some control; I do not have control
15. How important is it for you to have control over the information that is sent by Facebook/Google/Google+ to the websites to which you logged in with Facebook/Google/Google+?

Answers were on a 5 point scale from 'Not important' to 'Extremely Important'

⁵This is the "opposite" provider than the user's condition; i.e., "Google" if the participant was assigned Facebook and vice versa.

16. Facebook/Google/Google+ has a tool that allows you to see what information has been sent to the websites to which you have logged in with Facebook/Google/Google+. Prior to taking the survey, how familiar were you with this tool?
Answers were one of four choices: I was not aware that this tool exists; I was aware that this tool exists but I was not familiar with it; I was familiar with this tool; I don't remember.
17. How often do you use this tool?
Answers were one of four choices: I use it often; I use it occasionally; I have used it once or twice; I have never used it.
18. Please explain briefly situations where you used this tool.
19. If Facebook/Google/Google+ didn't provide this tool, would you log in with Facebook/Google/Google+ less often?
20. If you knew that you can always easily check which websites access your private information, would you log in more often with Facebook/Google/Google+?
21. The Facebook/Google/Google+ tool also allows you to change your preferences so that websites to which you logged in with Facebook/Google/Google+ can no longer access your information. Prior to taking the survey, how familiar were you with this feature?
22. How often do you use this feature?
Answers were one of four choices: I use it often; I use it occasionally; I have used it once or twice; I have never used it.
23. Please explain briefly situations where you used this feature.
24. If Facebook/Google/Google+ didn't provide this feature, would you 'log in with Facebook/Google/Google+' to other websites less often?
25. If you knew that this feature exists, would you 'log in with Facebook/Google/Google+' to other websites more often?
26. Many people use different email addresses for different purposes (for example, one address for work and another for personal life).
If you could do the same with identity providers (like Facebook or Google) to log into websites (like Amazon, Expedia, banking websites, etc), what would you do?
Answers were one of three: I would use the same identity provider to log into all websites, I would use different identity providers for different purposes, I'm not sure.
27. Why?
28. Do you have any additional Facebook/Google/Google+ accounts that you use for the purpose of logging in to different websites
29. Did you use your real (i.e. main) Facebook/Google/Google+ account to log in to this survey?

B.2 Demographic and knowledge section

A selection of the questions included in the exit survey that were reported in this study follows below.

31. What is your gender?
32. What is your age?
33. What is your race/ethnicity?
34. What is your current occupation?
35. What is the highest level of education you have completed?
36. Do you know any programming languages?
37. Do you have a college degree or work experience in computer science, software development, web development or similar computer-related fields?
38. How much do you agree or disagree with the following statements:
 - (a) It bothers me when websites ask me for personal information
 - (b) I am concerned that websites are collecting too much personal information about me
 - (c) It bothers me to give personal information to so many websites
 - (d) Consumers have lost all control over how personal information is collected and used by companies
 - (e) I feel that as a result of my visits to websites, others know more about me than I am comfortable with
 Answers were on a 5 point scale from 'Strongly disagree' to 'Strongly Agree'
39. The power switch on a computer is used to:
 - (a) Print documents to a laser printer
 - (b) Run an anti-virus program
 - (c) Install new software from a DVD
 - (d) Send email messages
 - (e) Turn the computer on and off
 - (f) Call customer support for help
 - (g) I don't know

B.3 Privacy concern level

The privacy-concern level is a Likert scale that we designed to capture the level of privacy awareness of a person. It is composed of the five statements included in question 38 of our exit survey (see Appendix B.2).

A Likert scale is a set of questions that measures participants' attitudes toward a specific notion or idea [11]. Participants have to report a level of preference or agreement with each sentence, usually by picking one option among an odd number of answers that ranges from one extreme (e.g., "completely disagree") to the other (e.g., "completely agree"). Each answer is assigned a numeric value, and the sum of the values of the answers yields a single numeric value per participant that represents their privacy concern level.

The internal reliability of our Likert scale, measured with Cronbach's alpha metric, was $\alpha=83\%$, $N=424$. Values of α over 80% are considered indicative of a high degree of internal consistency.