# Analyzing and Optimizing Access Control Choice Architectures in Online Social Networks

RON HIRSCHPRUNG, ERAN TOCH, HADAS SCHWARTZ-CHASSIDIM, TAMIR MENDEL, and ODED MAIMON, Tel Aviv University

The way users manage access to their information and computers has a tremendous effect on the overall security and privacy of individuals and organizations. Usually, access management is conducted using a *choice architecture*, a behavioral economics concept that describes the way decisions are framed to users. Studies have consistently shown that the design of choice architectures, mainly the selection of default options, has a strong effect on the final decisions users make by nudging them toward certain behaviors. In this article, we propose a method for optimizing access control choice architectures in online social networks. We empirically evaluate the methodology on Facebook, the world's largest online social network, by measuring how well the default options cover the existing user choices and preferences and toward which outcome the choice architecture nudges users. The evaluation includes two parts: (a) collecting access control decisions made by 266 users of Facebook for a period of 3 months; and (b) surveying 533 participants who were asked to express their preferences regarding default options. We demonstrate how optimal defaults can be algorithmically identified from users' decisions and preferences, and we measure how existing defaults address users' preferences compared with the optimal ones. We analyze how access control defaults can better serve existing users, and we discuss how our method can be used to establish a common measuring tool when examining the effects of default options.

CCS Concepts: • **Security and privacy** → **Access control**; **Social network security and privacy**; **Privacy protections**; *Social aspects of security and privacy*; *Usability in security and privacy*; • **Human-centered computing** → **Social content sharing**; *Empirical studies in collaborative and social computing*; User studies; • **Social and professional topics** → **Privacy policies**; *Governmental regulations*; • **Applied computing** → Sociology

Additional Key Words and Phrases: Access control, privacy, choice architecture, social networks

**ACM Reference Format:**
Ron Hirschprung, Eran Toch, Hadas Schwartz-Chassidim, Tamir Mendel, and Oded Maimon. 2017. Analyzing and optimizing access control choice architectures in online social networks. ACM Trans. Intell. Syst. Technol. 8, 4, Article 57 (May 2017), 22 pages.
DOI: http://dx.doi.org/10.1145/3046676

## 1. INTRODUCTION

Human factors influence how end-users interact with cyber-security systems and are the cause of many successful cyber-attacks [Dutt et al. 2013]. While cyber-security technologies provide a powerful technical solution, users' failure to comply with security guidelines is the cause of the majority of breaches in enterprise computing [Pfleeger et al. 2014; Deloitte 2013]. An important way in which users profoundly affect cyber

Authors' addresses: R. Hirschprung, E. Toch, H. Schwartz-Chassidim, T. Mendel, and O. Maimon, Department of Industrial Engineering, The Iby and Aladar Fleischman Faculty of Engineering, Tel Aviv University, P.O. Box 39040, Tel Aviv, 6997801, Israel; emails: {ronyh, erant, hadasc}@post.tau.ac.il, tamirmendel@mail.tau.ac.il, maimon@eng.tau.ac.il.

DOI: http://dx.doi.org/10.1145/3046676

security is by managing access to information and resources in a large variety of systems [Benantar 2006]. For example, making bad access control decisions in firewalls can allow access to internal organization networks [Wool 2004] or make one vulnerable to privacy threats when browsing the Web [Friedman et al. 2002]. In another example, bad access control decisions on Facebook can lead to over-sharing of personal information that can jeopardize the user's privacy [Madejski et al. 2012]. Access management is conducted using a *choice architecture*, which is a concept from behavioral economics that describes the way decisions are framed to users, that is, how they are presented, designed, and explained to users.

Research in behavioral economics has repeatedly shown, both experimentally [Thaler and Sunstein 2008; Adjerid et al. 2013; Knijnenburg et al. 2013a; Korff and Böhme 2014] and observationally [Palen 1999; Stutzman et al. 2013], that the way default options are selected in the choice architecture has a profound effect on users' ongoing behavior, nudging users toward specific choices. In response, default options are increasingly gaining the attention of regulators and legislators, who attempt to regulate choice architectures in online services. For example, a California bill (which did not pass) required that Online Social Networks (OSNs) establish a default privacy setting that prohibits the display of most personally identifiable information [California Bill 2011].

Policy makers and technology designers look at a problem such as defaults in OSNs in a very different way. Lawmakers use societal terms to suggest specific values that should influence or dictate design, promoting privacy and security [EU Directive 2002, 2011; California Bill 2011]. On the other hand, many system designers invoke basic usability principles and other values, such as information sharing, to justify specific designs [Buchanan 2011]. Balancing opposing values in a value-sensitive design process [Friedman et al. 2002, 2013] can take many forms, but it first requires some common language and assumptions to be handled correctly. In light of the growing awareness to the effect of the choice architecture on users' decisions and the overall privacy and security of the system, it has become crucial to better understand how systems' choice architectures serve the preferences of their users and the effect the architecture has on usability. This understanding is particularly important in rich choice architectures, where users make decisions constantly about how to share their information. The richness of the access control choice architecture varies between social networks and systems, from a binary choice of public or private (such as in Twitter) to nuanced control over each post (such as in Facebook or LinkedIn). This study takes an empirical and algorithmic approach, describing analysis evaluation methods and testing them on real-world data in rich choice architectures.

To evaluate choice architectures, this article suggests measuring the usability coverage of the default options. The definition of usability coverage is based on the observation that the system's choice architecture provides some social welfare to the population of users. A social welfare function ranks common alternatives (such as different sets of default options) in the system [Sen 1970], and in our case, it is interpreted as the proportion of users who would find at least one usable option to be satisfactory. To apply this concept, usability coverage is measured in two ways: in retrospect, by analyzing how real users have employed choices over time, and by asking users to indicate their preferences. Alternatively, we discuss how our method can be used by eliciting the preferences of individuals who did not use the system in the past. We collected previous choices that were made by users and current preferences of people and test how well an existing choice architecture complies with those choices and preferences (the usability coverage). We examined the gap between the usability coverage of the existing choice architecture and that of the optimal choice architecture. We also tested how usability coverage changes over time. Finally, we measured the bias the system

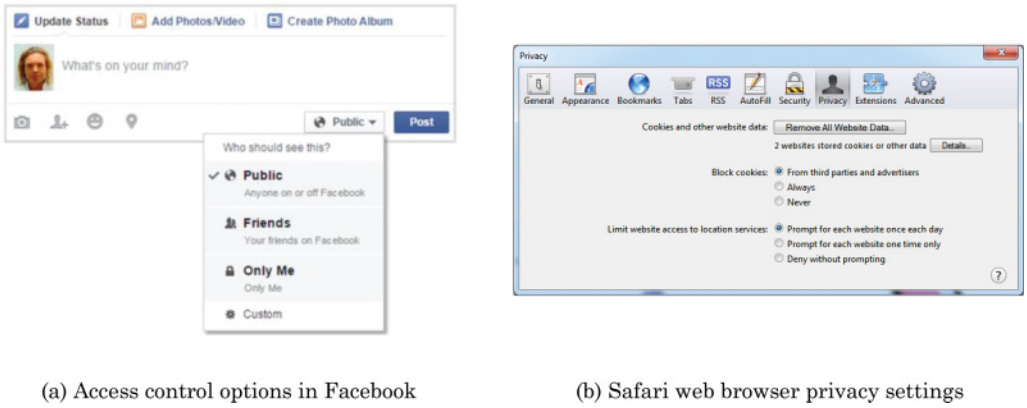


(a) Access control options in Facebook (b) Safari web browser privacy settings

Fig. 1. Two examples of user interfaces for access control: The left screenshot (a) shows Facebook's privacy control user interface, which is used to decide the audience for a published post. The user can select one of the displayed options, which include *Public*, *Friends* and *Only-Me*. The user can also abandon the default options by selecting *Custom* and decide the audience using fine-tune settings. The right screenshot (b) shows the Safari web browser privacy settings. These settings are configured by default when the application is installed. By using this interface, the user can manually set authorizations to use cookies and location services.

applies to its choice architecture by estimating whether the system nudges users to disclose more or less information.

To empirically test our method, we examined actual privacy settings for 266 Facebook users who published a total of 21,950 posts, and we surveyed 533 participants who were asked to elicit their preferred set of default options. Facebook was chosen as our research case study because it provides a rich interface to manage fine-grained access control (see Figure 1(a)) and is an information-sharing framework in which users are active in managing their privacy [Stutzman et al. 2013]. The findings of this research describe how access control choice architectures can be improved to maximize the users' social welfare. Our results from both user studies show that usability coverage can be used as an intelligent tool in analyzing choice architectures, determining how they contribute to the experience of existing users, how they perform over time, and in which way they nudge their users. We end the article by discussing how our methods can be used as a design and regulation tool and how choice architectures can be used to enhance the security and usability of information systems.

## 2. BACKGROUND

Our work was inspired by scholarly research on access control decision-making and behavioral biases. In the following sections, we relate each field to our study.

### 2.1. Decision-Making in Access Control

The disclosure of sensitive data can be initiated by the user (like in OSNs) or by the system (like in cyber-security mechanisms). Cyber-security systems regularly monitor network traffic, device use and personal communications in order to cope with a variety of vulnerabilities (e.g., intrusion detection, malware detection, data leakage prevention, and phishing identification). However, these cyber-security mechanisms raise a new challenge: balancing cyber-security risks against privacy and civil liberties concerns [Thang and Nguyen 2016; Landau 2014]. Thus, privacy concerns might reduce the acceptance and use of cyber-security systems by organizations and individuals, leading to decreased security for everyone's activities [Pfleeger et al. 2014; Warkentin and

Willison 2009]. An access control mechanism can provide the user a way to balance privacy disclosure vs. social benefits and security.

How are access control decisions conducted? In many online services, the answer to this question depends on the context of the information, resource sharing, and the interaction model of the users with the services. In a growing number of online services, users are offered a high level of control by using sophisticated access controls. Generally speaking, access control is the selective restriction of access to a place, information or other resource, based on a formal model that describes the type, characteristics, destinations and extent of information disclosure [Sandhu et al. 1996]. The basic mechanism of access control is similar for many types of information security system scenarios, including access to physical objects and places [Bauer et al. 2005], firewall management [Wool 2004], browser cookie management [Friedman et al. 2002], and OSN sharing [Madejski et al. 2012].

However, the usability of access control mechanisms in modern distributed systems has been widely criticized, mainly because of the limited technological literacy of the average user and the significant effort required on the user-side to determine how to implement the desired access rules [Tolone et al. 2005; Moyer and Abamad 2001]. This phenomenon is an outcome of the reality that most users are not experts at this task and that access management actions are almost always secondary to the collaborative task at hand [Cao and Iverson 2006].

Access control decision-making is conducted in the presence of uncertainty, where the benefits and probability of harm are weighted against each other [Dinev and Hart 2006]. Uncertainty inherently complicates any decision-making process, making it prone to human biases [Kahneman and Tversky 1979]. Behavioral economics research shows that choice architecture, the way choice alternatives are framed and contextualized, significantly and strongly affects the actual user's choices. For example, users tend not to change their initial calendar-sharing options [Palen 1999], decisions on receiving electronic mail from Websites [Johnson et al. 2002], retirement plans [Madrian and Shea 2000], or even decisions regarding organ donation [Johnson and Goldstein 2003].

Default choices play on people's tendencies to stick to the status quo, and therefore, if sharing is the default, people will tend to share more [Schweitzer 1994; Smith et al. 2013]. Additionally, defaults can be considered to be the service's official or unofficial recommendation, pushing users toward a particular choice [Sher and McKenzie 2006; McKenzie et al. 2006]. Research on privacy has shown that users' decisions in disclosing information can be affected by the wording of the options [Adjerid et al. 2013; Staddon et al. 2013] and by the number and granularity of the choices [Knijnenburg et al. 2013a, 2013b].

Several studies have shown the potential of reconstructing default options to improve privacy decision-making and to reduce the user's burden. For example, algorithms for finding defaults for location sharing applications were evaluated by Ravichandran et al. [2009], and clustering-based algorithms were suggested by Toch et al. [2010] and Hirschprung et al. [2015]. Knijnenburg and Kobsa [2014] experimentally evaluated different default options, showing how users' sharing tendency can be increased without increasing their privacy concerns. However, to the best of our knowledge, there were no works that aimed at evaluating and optimizing the default options of existing systems that have an existing user base. Analyzing existing user behavior raises new challenges in understanding the interaction between users' decisions and the context in which they are made, challenges we aim to address in this article.

To demonstrate the place of choice architecture in access control, let us look at two examples. Figure 1(a) depicts Facebook's privacy controls, which allow users to control the audience for a published post. This interface allows users to choose who will have access to a specific piece of information by choosing among a small set of possibilities, which are ordered according to the level of sharing (*Public*, *Friends*, and

*Only-Me*). If the user is not satisfied with the available possibilities, there is an option to choose "*Custom,*" which provides a fine-grained configuration. Figure 1(b) depicts the Safari web browser privacy settings. These settings are configured by default when the application is installed. By using this interface, the user can manually set general authorizations to use cookies, can watch the list of web sites that stored cookies in the browser and selectively delete them, and can limit web sites from accessing location services. The short list of initial possibilities that is presented to users comprises the set of default options that the designers assume the users will choose from. We distinguish between *default options*, which are the set of initial possibilities introduced to the user, and the *default choice*, which is the predefined choice that is selected by the service if the user does not explicitly select another choice.

### 2.2. Choice Architectures for Access Control

The power of default options and choices has encouraged policy-makers to regulate those defaults to discourage or encourage certain practices by electronic services. For example, the 2002 EU Opt-In Directive states that marketing email messages can be sent only to recipients who have given their prior consent, thus adopting an opt-in approach [EU Directive 2002]. Similarly, Canada's 2010 Anti-Spam Legislation establishes that all commercial electronic messages can be sent only to recipients who have given their prior consent [Canada's Justice Laws 2010]. On the other hand, the United States 2003 CAN-SPAM Act allows direct marketing email messages to be sent to anyone, without permission, until the recipient explicitly requests that they cease, thus allowing an opt-out approach [US Public Law 2003]. The 2011 EU Consumer Rights Directive subjected electronic commerce websites to a regulation that bans pre-ticked check boxes on websites for charging additional payments [EU Directive 2011]. For example, it is now illegal for a European website to add items to consumers' shopping carts by default.

Attempts to regulate OSN services were even more specific in describing how privacy defaults should be engineered. The EU Article 29 Working Party recommended that OSN providers offer default privacy settings that restrict viewing the user's profile to self-selected contacts [EU Directive 1995]. The 2011 California Bill S.B. 242 goes even further. It requires that OSNs set defaults to limit access in such a way that the users must choose the information that is to be made public and the OSN would have to ask users to establish their privacy settings when they register to join the site instead of after they join [California Bill 2011]. The bill's author, California senator Ellen Corbett, explains the logic behind the bill [Buchanan 2011]:

> "You shouldn't have to sign in and give up your personal information before you get to the part where you say, 'Please don't share my personal information.' "

This quotation highlights the trend in which regulation attempts to oversee the user's experience in OSNs and guides the interaction between the user and the service when using the default options and choices, according to privacy values.

The response of OSNs to the bill contains language that highlights usability rather than privacy. According to a letter sent by the Internet Alliance, a trade association that includes Facebook and other OSN websites, the bill [Buchanan 2011]:

> "would force users to make decisions about privacy and visibility of all information well before they even used the service for the first time, and in such a manner that they are less likely to pay attention and process the information."

The response highlights several key issues in usability engineering: the user's cognitive state, context of use, and user knowledge at the time of the decision. In this debate, privacy and usability are in direct conflict. However, to productively discuss both

concepts together, we must define the common language that combines these concepts as well as metrics that would provide a clear goal for effective privacy controls.

## 3. MODELING ACCESS CONTROL MECHANISMS

### 3.1. Access Control Evaluation

An access control mechanism usually relies on a choice architecture that offers the user few canonical options (default options). If $C$ is the set of all possible options, and $C'$ is the set of the default options, then $C'$ must be a subset of $C$ ($C' \in C$). Usually, the number of options offered in the default set is significantly smaller than that of the full set ($|C'| \ll |C|$). To formalize the coverage index, we indicate the satisfaction of user $u$ by a specific option $c$ with the function $f_s(u, c)$, which returns a numeric value of 1 if satisfied, otherwise 0, as given by:

$$f_s(u, c) = \begin{cases} 1, & \textit{user } i \textit{ is satisfied by option } c \\ 0, & \textit{otherwise} \end{cases}, \tag{1}$$

while $f_s(u, c)$ provides the satisfaction by a specific option, the satisfaction of user $u$ by the default set $C'$ (which includes few options) is given by:

$$S(u, C') = \begin{cases} 1, & \exists c \in C' : f_s(u, c) \neq 0 \\ 0, & \textit{otherwise} \end{cases} \tag{2}$$

Assuming that we have $N$ users, the coverage index for the default set $C'$ is defined as the proportion of users who comply with $S(u, c') = 1$, which is given by

$$\textit{Coverage}(C') = \frac{1}{N} \sum_{u=1}^{N} S(u, C') \tag{3}$$

### 3.2. Social Network Posts Publishing Access Control

To empirically test our definitions, we look at OSNs, in which users make many access decisions, such as setting the audience for each published post. Specifically, we formalize Facebook's access control model. Let $A_u$ be the number of actions user $u$ made, which is, in this specific example, the number of posts he published, and let $p_{ua}$ be the configuration user $u$ chooses for her action $a$ ($a = 1, 2, 3 \ldots Au$). In this case, we define the satisfaction of user $u$ for a specific post that she published (for a single action $p_{ua}$) by the set of default options $C'$ as follows: $f_p(u, \mathrm{a}, \mathrm{C}') = \exists \mathrm{c} \in \mathrm{C}' : p_{ua} = c$. When measuring the coverage, we can refer to all of the actions of all of the users homogeneously or examine how the default options address each user's action decisions. Therefore, we define two evaluation indexes (two types of coverage) for the usability coverage of a set of default $C'$: *choice coverage* and *user coverage*. The *choice coverage* measures the coverage by calculating the proportion of the actions' decisions that are matched with at least one of the options in the default set (without distinguishing among the users) and is given by:

$$\textit{choice coverage}(C') = \frac{\sum_{u=1}^{N} \sum_{a=1}^{A_u} f_p(u, a, C')}{\sum_{u=1}^{N} A_u} \tag{4}$$

The *user coverage* measures the coverage by calculating for each user a grade, which is the proportion of her actions' decisions that are matched with at least one of the options in the optimal set and then averaging all of the user's grades, which is given by:

$$\textit{user coverage}(C') = \frac{1}{N} \sum_{u=1}^{N} \frac{\sum_{a=1}^{A_u} f_p(u, a, C')}{A_u} \tag{5}$$

The *choice coverage* can be derived from the *user coverage* by giving the result of each user's coverage a weight, which is equal to the percentage of posts (out of all posts) the user published. To illustrate how these indexes work, let us take two Facebook users who have published posts. In this case, each posting is an action: User A has 3 posts, and 2 of them can be satisfied by the set of default options (e.g., one of the options in the set would include the user's choice). User B has 18 posts, and 16 of them are satisfied by the set of default options. We have a total of $2 + 16 = 18$ covered posts out of 21 posts. Thus, *choice coverage* $= \frac{2+16}{21} = 86.7\%$. The grade for user A is $\frac{2}{3}$, and the grade for user B is $\frac{16}{18}$, and therefore, *user coverage* $= \frac{\frac{2}{3}+\frac{16}{18}}{2} = 77.8\%$. These two indexes represent two different approaches to characterizing and evaluating the default options in this case. The *choice coverage* addresses the overall publishing activities (all of the posts of all of the users in the above Facebook example), ignoring its distribution across users, which gives "heavier" users who published more posts a higher weight. The *user coverage* addresses the proportion of users who are covered by the default set, normalizing each user by the number of her publishing actions (the posts of each user in the above example), and thus gives each user an equal weight. In the above example, it can be observed that because user B is dominant (because of her quantity of posts compared to user A), she significantly increased the *choice coverage*, while maintaining an effect equal to that of the other user on the *user coverage*.

Setting a criterion for the *user coverage* and *choice coverage* indexes in order to estimate normative approaches to what are "good" and "bad" options is not a straightforward task, because each choice has a complicated tradeoff between different values. For example, if we increase the number of options in the default set, we will most probably increase the *user coverage* and the *choice coverage*, but the choice architecture will provide lower usability to the user. Theoretically, it is possible to set boundaries to the coverage indexes. However, we think that a more proper approach is to bound security issues such as privacy violation (as described in Section 5.3) and then optimize the coverage given these boundaries as constrains.

### 3.3. Analysis of Access Choice Architectures

To evaluate how access control choice architectures fare, we compared the coverage of users' actual choices by the existing choice architecture and an optimal choice architecture that we generated (heuristically). The difference between the performance of the two architectures reflects how well the existing choice architecture answers its existing users' preferences. To generate the optimal set, we build on previous studies that allocated options by applying a general algorithm for optimizing the choice architecture [Olson et al. 2005; Watson et al. 2015; Hirschprung et al. 2015]. The algorithm extracts $k$ canonical options ($k$ is a parameter) from the full decision space based on users' actual choices or preferences.

The choice reduction algorithm's objective function can be one of the options: optimization by actions (maximal *choice coverage*) or optimization by user (maximal *user coverage*). The algorithm can be customized to a pre-reduction of the configuration space in such a way that not all of the theoretical options are included in the process. This feature is useful when, in a preliminary analysis, it is found that some options were not selected by any of the users and, thus, can be eliminated to reduce complexity and increase the algorithm's performance. Additionally, the algorithm has the flexibility of defining a threshold on two options to reflect a situation when two different options can be similar such that if a specific user is satisfied by one of them, she will be satisfied by the other. However, in our case, we set the threshold to 0, which means we ask for complete equality. Thus, the reduced configuration space can be limited to only those configurations that were chosen by at least one user. Because the algorithm

has factorial complexity, this preliminary process is significant as a means of reducing the computational effort.

The algorithm can work in two modes: *standard mode*, seeking the best $k$ configurations to maximize the objective function, and *preset mode*, in which instead of generating the optimal configurations, it is possible to insert $k'$ preset configurations. The preset mode is useful for measuring the performance of a current set of default options in an existing system. For example, when $k' = 3$, it reflects the current number of Facebook defaults. The algorithm can run for all of the posts or only for posts that were customized by the user (not configured by one of the three defaults). The full pseudo-code of the algorithm applied for the Facebook post publishing example (standard mode for all posts) is given in Algorithm 1.

**ALGORITHM 1:** The pseudo-code of the adopted algorithm for optimizing Facebook post publishing privacy settings (describe the standard mode)

```
% Eliminate non-relevant configurations
==========================
conf_options ⟵ all possible configurations    % the combinatorical combinations
                                              of all parameters
for x=1 to | conf_options |
   conf_eliminate_flag ⟵ TRUE
   for u =1 to N
      for a=1 to A_u
         if conf_options[x] == p_ua then
            conf_eliminate_flag ⟵ False
      if conf_eliminate_flag
         eliminate conf_options[x] from conf_options

% Find optimal configuration
====================
conf_space ⟵ all k combinations out of conf_options      % (|conf_options| choose k)

for x=1 to | conf_space |
   for u =1 to N
      User_Coverage_count ⟵ 0
      for a=1 to A_u
         conf_fit_flag ⟵ FALSE
         for s=1 to k
            if conf_space[x]_k == p_ua then conf_fit_flag ⟵ TRUE
         if conf_fit_flag
            Coice_Coverage[x] ++
            User_Coverage_count++
   User_Coverage[x] ⟵ User_Coverage[x] + ( User_Coverage[x] / A_u )
Case Objective Function
      Optimize by Actions:  SELECTED CONFIGURATION ⟵ configuration with
                            MAX (Coice_Coverage)
      Optimize by Users:    SELECTED CONFIGURATION ⟵ configuration with
                            MAX (User_Coverage)
```

The algorithm includes two main phases. In the first phase, the algorithm generates all possible configurations of the choice architecture space. Then, the algorithm eliminates from the configuration space each configuration that was not used by at least a single user. In the second phase, the algorithm generates all possible combinations of $k$ configurations out of the reduced configuration set. The algorithm calculates the

choice coverage and the user coverage for each set and selects the optimal configuration according to the objective function, which is the maximal *choice coverage* or maximal *user coverage*.

## 4. METHOD

To evaluate the methodology, we first applied it to real data of post sharing configuration decisions made by 266 users. These data were collected directly from Facebook. Then, we launched a survey with 533 participants who were asked to indicate their preferred three defaults. In the survey, the participants were not bound to Facebook's choice architecture, and they had the freedom to select their preferred set of defaults. We calculated the optimized configuration set and compared it with current Facebook defaults. Furthermore, we evaluated the amount of bias from the optimized set by quantifying the openness level of sharing.

### 4.1. Data Collection

We collected the data of 266 users (21,950 posts) by using Facebook's application programming interface (API), which accesses and analyzes privacy settings on Facebook. The application first asks for the participant's consent and then asks the user to grant specific access permissions to Facebook data. Afterward, the application accesses the participant's old posts using Facebook's Graph API. The application also collects general information about the participant, such as the number of friends and some demographic information (e.g., age, gender, and education).

### 4.2. Survey

We collected the answers of 533 participants in an online questionnaire that was not framed under Facebook's environment by using a survey engine. The participants were asked to select their 3 preferred default options out of a set of 11 options. The questionnaire of this survey is available in Appendix A. The options were shuffled and displayed in random order. The survey set of options is similar to the Facebook set of options and includes the option: "Share with people who live in your city or in your area," which can be selected in Facebook by using the smart-list[1] feature.

### 4.3. Participants

The participants were Amazon Mechanical Turk (Mturk) workers who were Facebook users. The participants were recruited via Mturk, which is a common tool for running behavioral studies [Mason and Suri 2012; Komarov et al. 2013] and has been used extensively in the field of privacy [Ayalon and Toch 2013; Kelley 2010]. Also, "MTurk especially is suitable to conduct survey research if Internet users are the intended population" [Schaarschmidt et al. 2015]. American MTurk workers, who were the population of our study, have a similar amount of personal information online as the general American population [Kang et al. 2014]. In the data collection study, for an approximately 10-minute task, the participants were paid $1.00 with a bonus of $0.25 for providing an additional detailed explanation with regard to sharing their decision-making process (which was not relevant to this particular study). In the survey study, the participants were paid $0.25 for approximately 3 minutes of work. Those rates are within the standard hourly compensations in MTurk studies [Ross et al. 2010]. Participants were required to be over 18 years old, have an Amazon MTurk HIT rate

[1]Facebook Smart lists use the information the user and his friends added to the Education, Work, and Current City of their profile to automatically create a sub-list of friends. For example, if a user lists Los Angeles as his current city, he will have a list with all of his friends who also listed Los Angeles as their current city.

of 90% or higher, and be from the U.S. (to ensure that they understood the survey at a native-tongue level).

The sample of the data collection study is heavily biased toward women: 87 of the participants were male, and 177 were female (2 preferred not to report). With regard to age, 49% of the participants were above age 35, 36% were between 25 and 34, and 15% were between 18 and 24. Most of the participants (93%) had at least a bachelor's degree, while 7% graduated high school. These results are in line with the studied demographic properties of Mechanical Turk crowd workers [Tomlinson et al. 2010]. The average number of Facebook Friends was 400 for males and 382 for females, with no significant differences.

The study was authorized by the institutional ethics review committee. We took several steps toward ensuring the participants' privacy: data were collected and surveyed in an anonymized and secure fashion (we did not record Facebook identity data) only after receiving the participant's consent; the data were accessed only at the time at which the participant granted consent; the post's actual text and content were not analyzed; and the data were accessed and stored in a secure and encrypted way.

## 5. RESULTS

The results of the studies are depicted in Figure 2. In both the behavior and preferences, sharing the post with Friends was the most common choice, having been used in 48% of the posts and 79% of the preferences. Public was the second most commonly used option in both studies, with 28% of the posts and 58% of the defaults. Only-Me was used in only 2% of the posts and only 31% of the defaults.

In the data collection study, custom options were used in 23% of the posts. This means that for a certain post the user did not choose one of the default options but instead used the advanced interface to customize the sharing settings. In addition, many of the users used more than one sharing option (for different posts they published). Seventy percent of the users used Friends as one of their options, while Public was used by 48% and a Custom option by 33% of the participants. It is interesting to note that 60% of the users who had used Public had also used Friends, and 48% of the users who had used Custom had used Friends. However, only 21% of users used both Friends and Public. Gender was correlated with a difference in publishing behavior: females published more posts than males, with an average of 122.0 for females and 72.6 for males ($t(198) = 3.669$, $p < .001$).

In the survey study, "share with specific people" was chosen as one of the options in the default (with a variable number of specified people) by 43% of users, Friends-of-Friends by 34%, "include Friends but exclude specific people" (with a variable number of specified people) by 30%, and "share with people in your area" by 12%. It can be noted that the results of both studies generally line up. However, when comparing them, it is important to indicate that in the data collection study, users made actual choices regarding a sharing option (single one each time), which reflects their behavior, while in the survey study, they elicited their preferred set of choices (three options at once), which reflects their preferences.

### 5.1. Choice Coverage

In the data collection study, where users' behavior was sampled, we applied the optimizing algorithm (Algorithm 1, as described in Section 3.3) to generate optimized privacy options. Then, we synthesized only those options that were actually chosen by

[2]Since each action in the data collection study included exactly one option, the frequencies in graph (a) sum to 100% (only the top 8 are displayed). However, in the survey study, each default includes 3 options, and thus, the sum exceeds 100%.

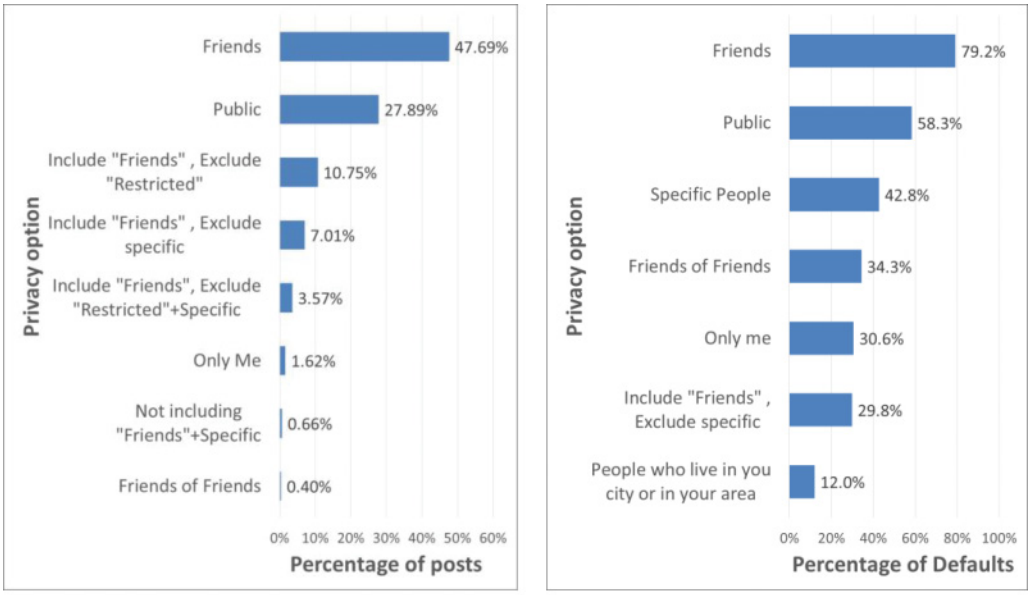


(a) Options in the data collection study Users' behavior

(b) Options in the survey study Preferences

Fig. 2. How participants use and select privacy options. The left graph (a) depicts the distribution of the top 8 most popular access options that were chosen by the users in the data collection study (a total of 21,950 posts, with a single action per post). The right graph (b) depicts the distribution of access options that were included by the participants in their preferences (a total of 533 participants, with 3 options per participant). The Y-axis describes the various privacy options that were chosen, while the X-axis describes for graph (a) the percentage of posts that were configured with that option and for graph (b) the percentage of defaults that included this option[2].

the users, and we discarded the theoretical options that belong to the configuration space. As a result, out of 21,950 posts, only 45 privacy options were actually in use. The algorithm was configured to run for one to four choices ($k \in \{1, 2, 3, 4\}$) and for both objective functions: optimized for posts and optimized for users. For each objective, both the *choice coverage* and the *user coverage* were calculated by introducing the optimized set of defaults to the whole set of options. Afterward, the algorithm was run in pre-set mode, with Facebook's three default options inserted (Public, Friends, Only-Me). The optimized default options that our algorithm produced for $k = 3$ configurations, when optimized for *choice coverage,* are the following:

1. Friends
2. Public
3. Friends except restricted (restrict friends from the pre-defined list)

We applied the same optimizing algorithm to the survey data where users' preferences were sampled. Due to the nature of this study, which did not sample specific Facebook actions (and each participant had only one result), the only relevant mode was optimized for users, and the algorithm was configured to run for three choices ($k = 3$). The *user coverage* in this case was calculated by the rate of users who had at least one of their preferred choices included in the optimized set. The default options our algorithm produced are the following:

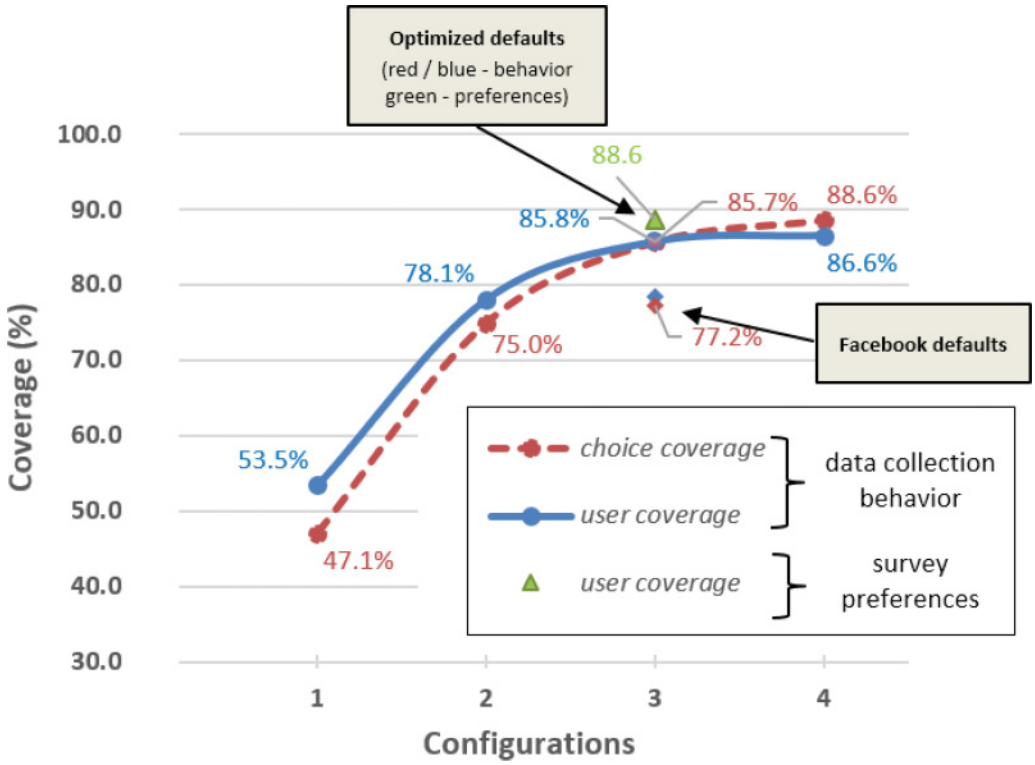


Fig. 3. Coverage rates. The x-axis stands for the number of options ($|C'|$), while the y-axis stands for the coverage rate. The line graphs depict the *choice coverage* (red) and *user coverage* (blue) that were achieved by the algorithm in the data collection study for all of the posts when the subjective function was maximizing the *choice coverage*. The triangle depicts the *user coverage* of the survey study, while the rhombuses depict the coverage that is achieved by Facebook defaults.[3]

1. Friends
2. Public
3. Share with specific people

Figure 3 depicts the coverage according to the number of options and type of coverage. It can be seen that for three default options, when optimized for posts, the *choice coverage* and *user coverage* were 85.8% and 85.7%, respectively, for the data collection study, and the *user coverage* was 88.6% for the survey study. This result means that approximately 86% of the users would be satisfied with at least one of our optimally set options. The asymptotic nature of the graph suggests that there is no significant benefit from increasing $k$ above 4 options. This result sets a clear bound on the optimality of the default options in the given distribution of privacy decisions. The conclusion is that we would probably not find any number of options that would satisfy the whole user population without sacrificing too much of the user's burden. The gap between the

[3]As can be seen in Figure 3, the differences between the *user coverage* of Facebook defaults and the two configurations optimized set are less significant than with the choice coverage. The *choice coverage* is exactly the same in the Facebook defaults and the two configurations optimized set. This is a result of the internal overlap of users' choices, so that the addition of the *Only-Me* option does not contribute much to the *user coverage*.

*choice coverage* and the *user coverage* graphs is an outcome of the dispersion of satisfied and unsatisfied posts among the users.

To understand how the algorithm can address sophisticated custom options, we applied it exclusively to custom options (in the data collection study). When optimizing the *choice coverage* for only the set of customized posts ($k = 3$), the algorithm produced the following options:

1. Friends excluding 1 specific friend
2. Friends, but Restricted (restrict friends from the pre-defined list)
3. Friends, but Restricted + 1 specific friend

These options yield coverages of 67.6% and 48.9% for *choice coverage* and *user coverage,* respectively. When optimized for *user coverage*, we obtained 57.8% and 62.9%, respectively. The reason for the lower coverage rates in comparison with coverage for all of the posts is the difference in homogeneity between the user's choices. We see sufficient overlap between the user's choices when choosing general options, but Custom options are used much more sporadically by the users. There are very few cases in which the same user uses more than two custom options.

The current defaults of Facebook can be compared to the optimal defaults that are generated by our algorithm. As Figure 3 shows, the Facebook default of *choice coverage* is 77.2% and of *user coverage* is 78.4%, while the optimal generated options cover 85.7% and 85.8%, respectively. The line graph depicts the coverage that is achieved by the algorithm (for 1 to 4 options), while the rhombuses depict the coverage that is achieved by Facebook defaults (there were three options). The vertical distance between them is the gap between the optimal defaults and Facebook's current defaults, and thus, it reflects the contribution of the algorithm to the coverage. The proposed methodology can immediately increase the *choice coverage* by 8.5% and the *user coverage* by 7.4% solely by altering the proposed defaults. The reasons for the better coverage of our defaults is simple: Options such as Only-Me are hardly used by users. Additionally, specific custom options (e.g., Friends except for particular Friends) are a common choice among users but are difficult to define and, therefore, may be inaccessible to the users.

### 5.2. Temporal Analysis

When looking at changes through time, we see more positive improvements in the coverage and user burden. Before 2012, Facebook's privacy defaults included also Friends-of-Friends as a default option [FTC-USA 2011], and because that configuration had a very low coverage, it downgraded the overall coverage of the options. Facebook has a "built in" default choice for post publishing as well as for other information items. This default is the privacy level that will be allocated to a post when a new user has registered to Facebook, and has not changed the default suggestions while going through the registration process, and has not changed the post's privacy setting when publishing the post. It was found that users publish significant amounts of information on Facebook, and nearly half of them adopt the default privacy choice [Liu et al. 2011]. Facebook is changing those defaults continually; for example, by 2005, the default was Friends,[4] by 2009 Friends-of-Friends, by 2010 Public, and today (May 2015) Friends again [Loewenstein et al. 2015].

If we apply the results from our experiments to the changing choice architecture settings, we can calculate the coverage of Facebook defaults across time, as depicted in

[4]The definition at that time was "network," which is ranked one level above friends in the order of sharing levels, but it is not used today of Facebook.

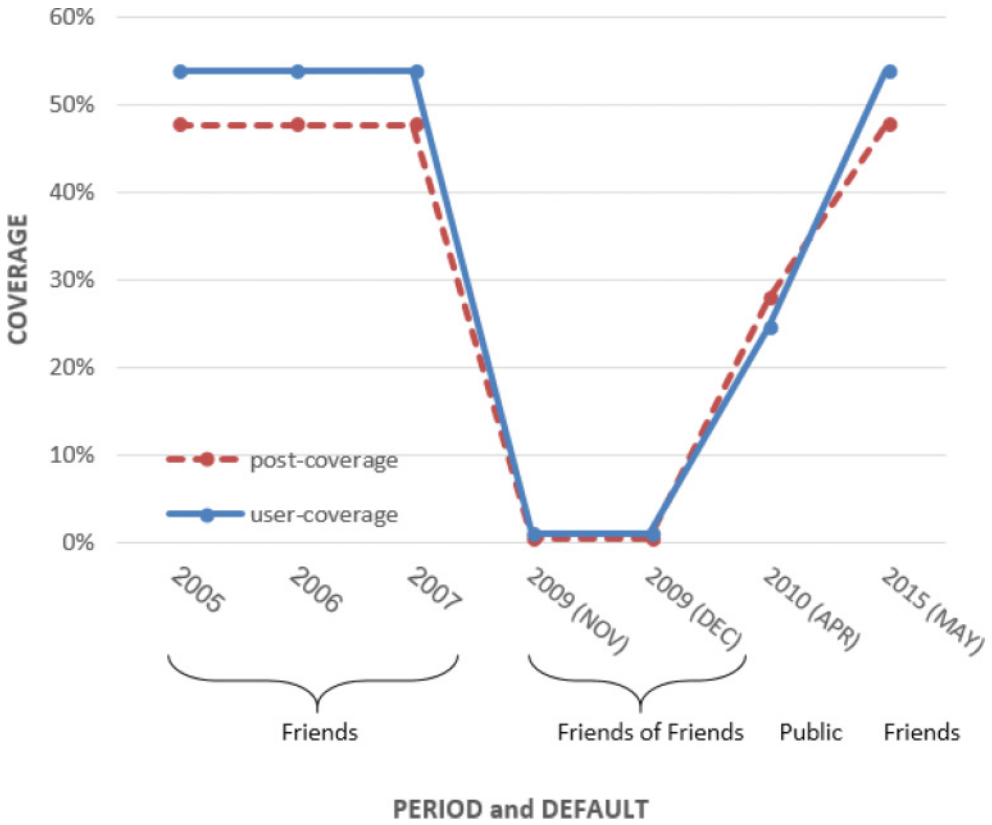


Fig. 4. The coverage of the Facebook default choice for a post's privacy across time.

Figure 4. It can be observed that from 2005 to 2007, Facebook provided defaults with coverages of approximately 50%. Then, the default was changed to Friends-of-Friends, which caused the coverage to drop to nearly 0% because this setting is not popular among Facebook users. Since December 2009, Facebook has improved the coverage and abandoned the Friends-of-Friends setting. Those changes across time refer to Facebook defaults, i.e., when a user is publishing a post without even selecting any of the three canonical options. In any case, a single choice has a relatively small coverage of approximately 50%, but a set of three optimized options provides a coverage of 85%, according to our experiment results. For that reason, we argue that Facebook should optimize the major choices of the privacy configuration; selecting the single best default is not sufficient.

### 5.3. Evaluating the Choice Architecture Bias

To analyze the direction in which choice architectures lead their users, we define an openness *index,* which is a normative model that relies on a notion of openness vs. closedness in a user's access decision. The index is based on the following assumption: When a user sets the access permissions to a Facebook post, she can experience some privacy loss due to unexpected use of the data, or due to disclosure to an unintended population. For example, if Alice published a post using the Friends option, and Bob, who is a Facebook friend of Alice, publishes this post to his Facebook friends, then

Table I. An Example of Possible Measurements of Openness of Facebook Post Sharing Decisions

| Parameter | Rank (noted OP) | Explanation | Example Value |
|---|---|---|---|
| Public | *Cp* | If Public is selected, then the rank is *Cp* | $Cp = 20$ |
| Only-Me | *0* | If Only-Me is selected, then the rank is zero because there is no disclosure. | *0* |
| Friends | *Cf* | If the Friends option is selected, then the rank is *CF* | $Cf = 10$ |
| Friends-of Friends | *Cfof* | If Friends-of-Friends is selected, then the rank is *CFoF* | $Cfof = 15$ |
| Restricted | *Cr* | This parameter is defined in Custom mode and indicates whether to exclude a pre-defined list of friends. If used, it reduces openness, and thus, its value is negative. It should be carefully handled when added to other values to avoid $OP<0$. | $Cr = -5$ |
| Include (n) | $fi(n)$ | This parameter defines the number of friends who will be disclosed to the post. Here, $n$ is the number of friends who are explicitly included in the audience of the post. We know that $fi(n)$ is a monotonically increasing function (because when more friends are disclosed to the post, there might be a higher rank of openness), and we can assume asymptotic behavior (because there might be a significant difference between 3 and 4 friends and a less significant difference between 50 and 51). Such a function defined by the regulator might be: $$f_i(n) = a\left(1 - \frac{1}{bn}\right)$$ | $fi(n) = n$ |
| Exclude (n) | $Fe(n)$ | This parameter defines the number of friends who will not be disclosed to the post. Here, $n$ is the number of friends who are explicitly excluded from the audience of the post. We know that $Fe(n)$ is a monotonically decreasing function (because when more friends are excluded from the post, there might be a lower openness), and we can assume asymptotic behavior (as in Include). It should be carefully handled when added to other values, to avoid $OP<0$. | $Fe(n) = -n$ |

*OP* represents the openness level; the higher *OP* is, the higher the openness is.

Alice is disclosed at a Friend-of-Friends level. To quantify the privacy loss, we created an index to measure the openness level of a configuration, as described in Table I. The designer of the choice architecture can set a subjective rank that represents the hypothetical openness for each parameter of the post disclosure configuration. The ranks have no units and can be used mainly as a comparison between two or more configurations. For example, if Friends has a rank of 10 and Public has a rank of 20, then Public openness is double that of Friends. Because Public is the maximal disclosure, the rank as exemplified in Table I cannot exceed the rank of Public (i.e., 20). Thus, for example, if a user included 22 explicit names of friends, then the rank would be 20 rather than 22.

Given the openness index, we can measure the openness of each specific Facebook post-sharing configuration. If we use the example values of Table I (right column), we obtain an openness value of 20 for the current Facebook default, which is Public. However, when applying the openness index to the post's sharing decisions of the users in our data collection study, the average openness of all of the posts is 11.6 and only 1.3 for customized posts. An immediate conclusion of this result is that it appears that

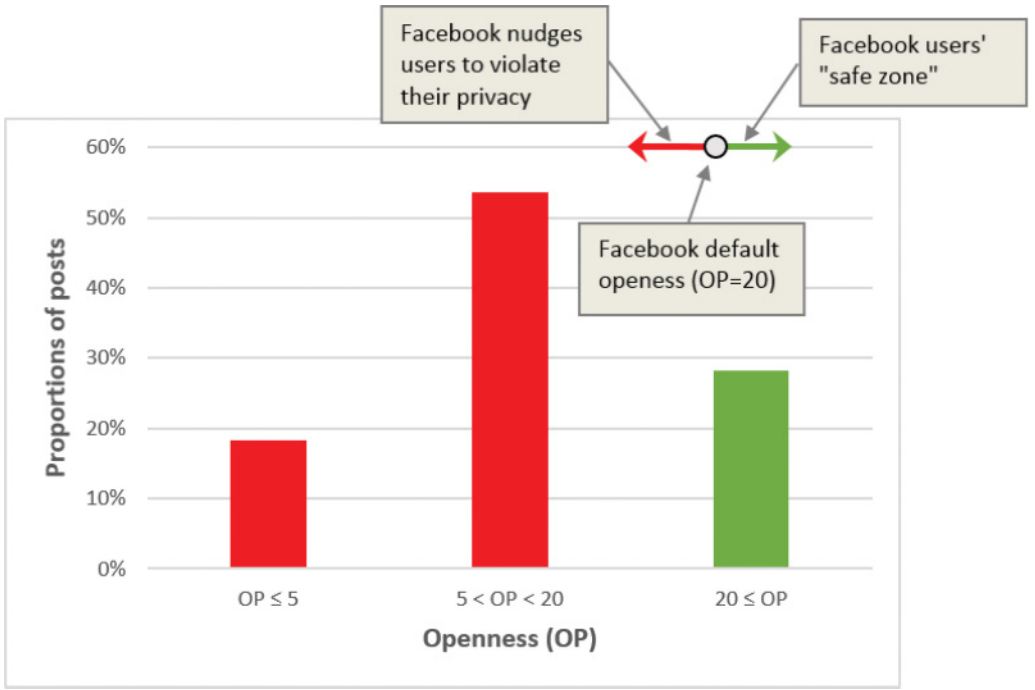


Fig. 5. The distribution of the openness (*OP*) across all of the published posts. The X-axis is the rank of openness, while the Y-axis is the proportion of posts that have that rank value. The red bars indicate the population that is nudged by the Facebook default choice (Public) to higher openness than their preferences and, thus, will not be satisfied by this default. The green bar indicates the population that will be satisfied.

Facebook introduces to its users a set of options that are more open compared to the user's actual choices. Figure 5 depicts the distribution of the openness across all of the posts that were published according to the user's choices. It can be observed that most users will find the default choice too open, and only 28% of users will find it safe.

To measure the fitness of the canonical configurations to a specific choice according to the rank of openness, we define an *openness distance* between them. The *openness distance* is the minimum of the absolute differences between each of the options in the canonical set and the preference, i.e., the distance between the preference $p$ and the canonical set $C'$ (with $i$ options) is *openness distance* $= \min_i |OP(p) - OP(c_i)|$. The average *openness distance* between the posts in the user study dataset and the Facebook defaults is 0.89 and between the optimized defaults is 0.29, which indicates that the optimal set better fits users' preferences and, thus, has a lower potential of violating users' privacy. Figure 6 depicts the accumulated *openness distance*s across the various *openness distance* levels. It can be noted that the optimal configuration defaults (orange line) have better fitness than the current Facebook defaults (blue line).

## 6. DISCUSSION

The design of access control choice architectures can reflect users' preferences to varying degrees, and thus is a critical issue in protecting users' security and privacy. Our methodology provides a way to analyze, monitor, and assess how the architecture operates. To provide a meaningful and intelligent analysis, two capabilities are required: (a) the ability to measure the fitness of an existing choice architecture to the user's preferences and (b) the ability to optimize the choice architecture to ensure minimal

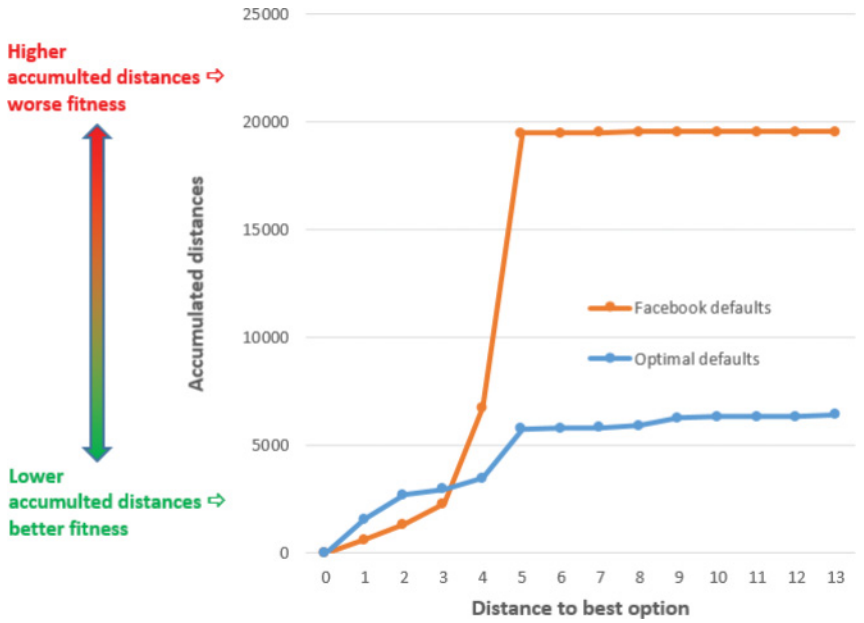


Fig. 6. The accumulated distance between users' choices and: Facebook defaults (orange line) and the optimal defaults (blue line). The X-axis describes the distance to the best option of the set, and the Y-axis is the accumulated distances. The higher the accumulated distance, the worse the fit to the user's preferences.

deviation from the user's preferences. In this study, we show that by applying an algorithmic approach, the choice architecture can be evaluated and optimized while relying on the existing user preferences.

We applied the proposed methodology to real preferences of users who configured their post-sharing options in Facebook, and we show that the Facebook defaults can be improved from a coverage of approximately 75% of the users to 85%. Our assessment echoes the discussion around Facebook's default privacy options, including the criticism against the existence of the Friends-of-Friends option [FTC-USA 2011]. Long-term analysis of OSN sharing behavior shows that between 2005 and 2012, public sharing of almost every information item on Facebook dropped from 95%–85% to approximately 20%-10% [Stutzman et al. 2013], and from early 2010 to 2012, people became dramatically more private on Facebook [Dey et al. 2012]. This result is relevant to the interpretation of our findings, because in this period, Facebook's default choice was public for most of the data types. Therefore, we assume that experienced users have settings that are very different from the initial defaults that were offered by Facebook. Our study does not provide a silver bullet for designing the default options; however, our conceptualization and evaluation can provide clear design guidelines.

Using similar algorithmic approaches combined with information-flow analysis, user interfaces of applications can be examined with no human intervention. Therefore, massive numbers of applications on the Web or available through app stores can be regulated with relatively little effort. For example, a regulator can set ranks to the default options of web browser security, evaluate those defaults and set rules to restrain those options. The methodology can also be applied in a mobile applications store. By using these approaches, regulators can ensure that all of the app store applications have a specific privacy default option with regard to location API calls, for example. Because the methodology relies on the user's actual preferences, it can also address different sub-groups of the users, who can differ in their preferences, such as groups

based on geographic regions [Krasnova and Veltri 2010]. We can see the beginning of these abilities in automatically enforced filtering procedures in the Apple App Store. This vision of wide-scale policy enforcement can be seen to be a powerful tool for policy makers, continuing the current trend of regulating human-computer interactions by enabling new types of legislative and administrative tools. On the other hand, the vision of automatic enforcement is also disturbing, since it practically removes control from the user.

Our method can be applied to other security domains in which users make decisions within a choice architecture. For example, the deployment of organizational firewalls or data leakage prevention systems can lead to conflicting interests of employees and organizations, where organizations close access to websites and services that employees still want to access. This situation can lead employees to refrain from using the organizational channels and to switch to their own personal devices [Pfleeger et al. 2014]. In browser security and privacy interfaces, users can manage whether cookies from second-party and third-party websites are accepted, which can lead to conflicting interests between website functionality and privacy [Friedman et al. 2002]. Our method can be applied to the analysis of access control choice architectures, providing a way to adjust these architectures accordingly and to improve their design in order to encourage users to make use of these systems. For example, we can measure the preferences of users and compare them to the actual behavior to see whether users' actions are more correlated with the system's choice architecture or with their preferences.

### 6.1. Limitations and Future Work

The analysis relies on data collected by users' behavior in an existing choice architecture. We cannot estimate the effect of the architecture and whether and how it has biased any user choices. Probably, people's choices will be different if the choice architecture is different. It is important to point out that if a system's choice architecture does not fit users' choices made within the boundaries of that architecture, then the architecture is definitely lacking. This limitation was mitigated by conducting the second user study, when users elicited their preferences outside of the Facebook choice architecture environment. Moreover, we do not know whether our optimal set of options is the most fitting for the entire Facebook user base. However, the privacy choices made by our participants are similar to the choices obtained by representative U.S. surveys [Madden 2012] and long-term observational studies [Stutzman et al. 2012]. A further study that validates the advantage of the proposed (optimal) choice architecture over Facebook's architecture could better establish the results. This study can be performed, for example, by creating an environment that simulates the Facebook configuration mechanism, which will be introduced with Facebook options to one group of users and to another group with the optimal set of options. Another aspect of the configuration setting is the detailed list of friends who are included or excluded from a post disclosure. In our work, we only referred to the number of friends (and not to friends by name), since in the vast majority of the cases, they were the same friends across the posts a specific user published. Further research could investigate the cases in which a user list of included or excluded friends is not constant.

Naturally, this research opens up new questions and challenges. For example, it is possible to increase the average fitness of the access control mechanism to users' preferences by improving the usability to a large degree for a small subset of users or by improving the usability to a small degree for all of the users [Sen 1970]. However, if we want to optimize the social fairness across users, i.e., to strive for higher equity, this approach can reduce the average fitness for all users because fairness and average utility are in a trade-off relation [Blackorby and Donaldson 1977]. Finally, it is also important to emphasize that our work assumes that choice architectures are static:

they do not adapt to particular users and do not change over time. Most configuration systems follow this assumption, but some new interfaces, such as Facebook's current privacy settings, are adaptive and dynamic. Our method could serve as the basis for future studies that evaluate the fairness of adaptive interfaces by analyzing the options available to each user individually and comparing them to previous behavior or preferences.

## 7. CONCLUSIONS

In this article, we describe a method for evaluating and optimizing access control to better protect a user's desired privacy according to their preferences and, specifically, to quantify how architecture designers aspire to nudge users toward specific choices. We empirically tested our model on Facebook post sharing behaviors of users and on the results of a survey on default preferences. Our results paint a complicated picture. Facebook's defaults at the time of the study comply for the most part with users' behavior and stated preferences. At the same time, these default options provides a coverage that is lower by approximately 8% than the optimal set of default options. Overall, the current choice architecture is biased toward a higher openness level, a design that can reflect the software company's interests rather than the user's preferences. Our proposed methodology can quantify this type of bias and can suggest an optimal choice architecture that answers users' preferences and best represents their behaviors.

## APPENDIX A — THE QUESTIONNAIRE OF THE SURVEY STUDY

The survey study includes these 11 options, which were shuffled and displayed in random order to the participants:

1. Public (anyone on or off Facebook)
2. Friends (your friends on Facebook)
3. Friends of Friends
4. Only me
5. Share with people who live in you city or in your area
6. Share only with 1 specific person (which you choose)
7. Share only with 2–4 specific people (which you choose)
8. Share only with a group of 5 or more specific people (which you choose)
9. Share with all my friends except 1 specific person (which you choose)
10. Share with all my friends except 2–4 specific people (which you choose)
11. Share with all my friends except with 5 or more specific people (which you choose)

## REFERENCES

Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347 (6221): 509–514.

Idris Adjerid, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2013. Sleights of privacy: Framing, disclosures, and the limits of transparency. In *Proceedings of the 9th Symposium on Usable Privacy and Security.* ACM, New York, 9.

Oshrat Ayalon and Eran Toch. 2013. Retrospective privacy: Managing longitudinal privacy in online social networks. In *Proceedings of the 9th Symposium on Usable Privacy and Security.* ACM, New York, 4.

Lujo Bauer, Scott Garriss, and Michael K. Reiter. 2005. Distributed proving in access-control system. In *Proceedings of the 2005 IEEE Symposium on Security and Privacy.* 81–95.

Messaoud Benantar. 2006. *Access Control Systems: Security, Identity Management and Trust Models.* Springer Science & Business Media.

Charles Blackorby and David Donaldson. 1977. Utility vs equity: Some plausible quasi-orderings. *Journal of Public Economics* 7, 3, 365–381.

Wyatt Buchanan. 2011. Social-networking sites face new privacy battle. Retrieved from http://www.sfgate.com/bayarea/article/Social-networking-sites-face-new-privacy-battle-2371641.php.

California Bill. 2011. California Bill S.B. 242-Privacy Control Requirements for Social Networks. Retrieved from http://www.leginfo.ca.gov/pub/11-12/bill/sen/sb_0201-0250/sb_242_bill_20110525_amended_sen_v96.html.

Canada's Justice Laws. 2010. Canada's Anti-Spam Legislation.

Xiang Cao and Lee Iverson. 2006. Intentional access management: Making access control usable for end-users. In *Proceedings of the 2nd Symposium on Usable Privacy and Security.* ACM.

Deloitte. 2013. *2013 TMT (Technology, Media, and Telecommunications) Global Security Study.* Deloitte Touche Tohmatsu Limited (DTTL).

Ratan Dey, Zubin Jelveh, and Keith Ross. 2012. Facebook users have become much more private: A large-scale study. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops).* 346–352.

Tamara Dinev and Paul Hart. 2006. An extended privacy Calculus model for e-Commerce transactions. *Information Systems Research* 17, 1, 61–80.

Varun Dutt, Young-Suk Ahn, and Cleotilde Gonzalez. 2013. Cyber situation awareness modeling detection of cyber attacks with instance-based learning theory. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 55, 3, 605–618.

EU Directive 1995/46/EC. 1995. Directive 95/46/EC of the European Parliament and of the Council: On the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data.

EU Directive 2002/58/EC. 2002. Directive 2002/58/EC of the European Parliament and of the Council: On Concerning the Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector.

EU Directive 2011/83/EU. 2011. Directive 2011/83/EU of the European Parliament and of the Council: On Consumer Rights. *Official Journal of the EU.*

Batya Friedman, Daniel C. Howe, and Edward Felten. 2002. Informed consent in the Mozilla browser: Implementing value-sensitive design. In *Proceedings of the 35th Hawaii International Conference on System Sciences.* IEEE. 10.

Batya Friedman, Peter H. Kahn Jr, Alan Borning, and Alina Huldtgren. 2013. Value sensitive design and information systems. In *Early Engagement and New Technologies: Opening up the Laboratory,* Springer, Netherlands. 55–95.

FTC-USA. 2011. *Facebook Settles FTC Charges that it Deceived Consumers by Failing to Keep Privacy Promises.* Federal Trade Commission. Accessed November 29. http://www.ftc.gov/news-events/press-releases/2011/11/facebook-settles-ftc-charges-it-deceived-consumers-failing-keep.

Ron Hirschprung, Eran Toch, and Oded Maimon. 2015. Simplifying data disclosure configurations in a cloud computing environment. *ACM Transactions on Intelligent Systems and Technology* 6, 3.

Eric J. Johnson, Steven Bellman, and Gerald L. Lohse. 2002. Defaults, framing and privacy: Why opting in-opting out. *Marketing Letters* (Kluwer Academic Publishers) 13, 1, 5–15.

Eric Johnson and Daniel Goldstein. 2003. Do defaults save lives? *Science* 302, 1338–1339.

Daniel Kahneman and Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 2, 263–292.

Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. 2014. Privacy attitudes of mechanical turk workers and the U.S. public. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS'14).*

Patrick Gage Kelley. 2010. Conducting usable privacy and security studies with Amazon's mechanical turk. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS'10).*

Bart P. Knijnenburg, Alfred Kobsa, and Hongxia Jin. 2013a. Dimensionality of information disclosure behavior. *International Journal of Human-Computer Studies 71* 71, 12, 1144–1162.

Bart P. Knijnenburg, Alfred Kobsa, and Jin Hongxia. 2013b. Preference-based location sharing: Are more privacy options really better? In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM.

Bart Piet Knijnenburg and Alfred Kobsa. 2014. *Increasing Sharing Tendency Without Reducing Satisfaction: Finding the Best Privacy-settings User Interface for Social Networks.* AIS Electronic Library (AISeL).

Steven Komarov, Katharina Reinecke, and Krzysztof Z. Gajos. 2013. Crowdsourcing performance evaluations of user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* ACM. 207–216.

Stefan Korff and Rainer Böhme. 2014. Too much choice: End-user privacy decisions in the context of choice proliferation. In *Proceedings of the 10th Symposium on Usable Privacy and Security (SOUPS'14).*

Hanna Krasnova and Natasha F. Veltri. 2010. Privacy calculus on social networking sites: explorative evidence from Germany and USA. In *Proceedings of the 43rd Hawaii International Conference on System Sciences*. IEEE.

Susan Landau. 2014. Highlights from making sense of Snowden, part II What's significant in the NSA revelations. *IEEE Security and Privacy* 12, 1, 62–64.

Yabing Liu, Krishna P. Gummadi, Balachander Krishnamurthy, and Alan Mislove. 2011. Analyzing facebook privacy settings: User expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*. ACM. 61–70.

Mary Madden. 2012. *Privacy Management on Social Media Sites*. Pew Research Center's Internet & American Life Project.

Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. 2012. A study of privacy settings errors in an online social network. In *Pervasive Computing and Communications Workshops (PERCOM Workshops)*. IEEE. 340–345.

Brigitte C. Madrian and Dennis F. Shea. 2000. The power of suggestion: Inertia in 401(k) participation and savings behavior. *National Bureau of Economic Research* w7682.

Winter Mason and Siddharth Suri. 2012. Conducting behavioral research on Amazon's mechanical turk. *Behavior Research Methods* (Springer) 44, 1, 1–23.

Craig R. M. McKenzie Michael, J. Liersch, and Stacey R. Finkelstein. 2006. Recommendations implicit in policy defaults. *Psychological Science* 17, 5, 414–420.

Matthrew J. Moyer and Mustaque Abamad. 2001. Generalized role-based access control. In *Proceedings of the 21st International Conference on Distributed Computing Systems, 2001*. IEEE. 391–398.

Judith S. Olson, Jonathan Grudin, and Eric Horvitz. 2005. A study of preferences for sharing and privacy. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems*. ACM. 1985–1988.

Leysia Palen. 1999. Social, individual and technological issues for groupware calendar systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'99)*. 17–24.

Shari Lawrence Pfleeger, M. Angela Sasse, and Adrian Furnham. 2014. From weakest link to security hero: Transforming staff security behavior. *Homeland Security and Emergency Management 2014* 11, 4, 489–510.

Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman M. Sadeh. 2009. Capturing social networking privacy preferences. In *Privacy Enhancing Technologies*. Springer, Berlin. 1–18.

Joel Ross, Lilly Irani, M. Six Silberman, Andrew Zaldivar, and Bill Tomlinson. 2010. Who are the crowdworkers? Shifting demographics in Mechanical Turk. In *Proceedings of the 28th International Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA'10)*. ACM. 2863–2872.

Ravi S. Sandhu, Edward J. Coynek, Hal L. Feinsteink, and Charles E. Youmank. 1996. Role-based access control models. *Computer* (IEEE) 29, 2, 38–47.

Mario Schaarschmidt, Stefan Ivens, Dirk Homscheid, and Pascal Bilo. 2015. Crowdsourcing for survey research: Where Amazon mechanical turks deviates from conventional survey methods. *Informatic*. University of Koblenz-Landau.

Maurice Schweitzer. 1994. Disentangling status quo and omission effects: An experimental analysis. *Organizational Behavior and Human Decision Processes* 58, 3, 457–476.

Kumar Sen Amartya. 1970. *Collective Choice and Social Welfare*. Vol. 11. Elsevier.

Shlomi Sher and Craig R. M. McKenzie. 2006. Information leakage from logically equivalent frames. *Cognition* 101, 3, 467–94.

N. Craig Smith, Daniel G. Goldstein, and Eric J. Johnson. 2013. Choice without awareness: Ethical and policy implications of defaults. *Journal of Public Policy and Marketing* 32, 2, 159–172.

Jessica Staddon, Alessandro Acquisti, and Kristen LeFevre. 2013. Self-reported social network behavior: Accuracy predictors and implications for the privacy paradox. In *Social Computing (SocialCom)*. IEEE. 295–302.

Fred Stutzman, Ralph Gross, and Alessandro Acquisti. 2013. Silent listeners: The evolution of privacy and disclosure on Facebook. *Journal of Privacy and Confidentiality* 4, 2, 2.

Frederic Stutzman, Jessica Vitak, Nicole B. Ellison, Rebecca Gray, and Cliff Lampe. 2012. Privacy in interaction: Exploring disclosure and social capital in Facebook. In *ICWSM*.

Richard H. Thaler and Cass R. Sunstein. 2008. Nudge: Improving decisions about health, wealth, and happiness. Yale University Press.

Tran Manh Thang and Van Khanh Nguyen. 2016. Synflood spoof source ddos attack defence based on packet id anomaly detection-PIDAD. In *Information Science and Applications (ICISA)*. Springer Singapore, 739–751.

Eran Toch, Norman M. Sadeh, and Jason Hong. 2010. Generating default privacy policies for online social networks. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*. ACM. 4243–4248.

William Tolone, Gail-Joon Ahn, Tanusree Pai, and Seng-Phil Hong. 2005. Access control in collaborative systems. *ACM Computing Surveys (CSUR)* 37, 1, 29–41.

USA Public Law. 2003. The Can-Spam Act 2003.

Merrill Warkentin and Robert Willison. 2009. Behavioral and policy issues in information systems security: the insider threat. *European Journal of Information Systems* 18, 2, 101–105.

Jason Watson, Heather Richter Lipford, and Andrew Besmer. 2015. Mapping user preference to privacy default settings. *ACM Transactions on Computer-Human Interaction (TOCHI)* 22, 6, 32.

Avishai Wool. 2004. A quantitative study of firewall configuration errors. *Computer* 37, 6, 62–67.

Received March 2016; revised December 2016; accepted January 2017