

Retrospective Privacy: Managing Longitudinal Privacy in Online Social Networks

Oshrat Ayalon
Department of Industrial
Engineering
Tel Aviv University
oshratra@post.tau.ac.il

Eran Toch
Department of Industrial
Engineering
Tel Aviv University
erant@post.tau.ac.il

ABSTRACT

Online social networks provide access to the user's information for long periods of time after the information's initial publication. In this paper, we investigate the relation between information aging and its sharing preferences on Facebook. Our findings are based on a survey of 193 Facebook users, in which we asked users to specify their sharing preferences and intentions towards posts that were published in different periods of time (from the time of the survey and up to 24 months prior to the time of the survey.) Our results show that willingness to share significantly drops with the time passed since publishing the post. The occurrence of life changes, such as graduating from college or moving to a new town, is correlated with a further decrease in the willingness to share. We discuss our findings by relating it to information aging theories and privacy theories. Finally, we use our results to reflect on privacy mechanisms for long-term usage of online social networks, such as expiry date for content and historical information reviewing processes.

Categories and Subject Descriptors

H.5.3 Group and Organization Interfaces evaluation: collaborative computing.

General Terms

Human Factors; Design; Measurement.

Author Keywords

Online social networks; privacy; information aging.

1. INTRODUCTION

Some scholars argue that nothing is forgotten in online social networks [1, 2, 7, 11]. Personal information shared and stored by users in online social networks (OSNs) is not deleted by default. New and powerful mechanisms, such as Facebook Timeline, Facebook Graph Search and Google social website indexing, are making historical information in Facebook and in other OSNs readily available to the user's social network. The pervasiveness of OSNs is raising crucial questions with regard to the impact of accessibility of historical OSN information on people's social lives. The Facebook Timeline, introduced at the end of December 2011, allows viewers to easily browse through the user's actions on Facebook, jumping back to earlier events with the click of a button [8]. Criticisms by the public and the media press were not delayed, highlighting privacy concerns. As a Facebook user quoted in USA Today: "*I don't think something I did four years*

ago is really representative of whom I am today" [8]. The accessibility of historical information raises new and thought provoking questions: How would people cope with embarrassing information posted when they are in college when they move into the workplace? What kind of information would children see about their parents? How would people manage their privacy when the future audience of their information is still unknown? All of these questions will have an outstanding and long term impact on people's social lives.

Research suggests that while users are increasingly using privacy tools in online social networks [3, 12, 17], they still face serious problems in understanding, expressing and managing their privacy preferences [13, 14, 15, 16]. Almost all research in the field of privacy in social networks is focused on privacy preferences in the context of information shared in the user's present time. In other words, in previous research in OSN privacy, the studies assume that the outcome of the user's privacy actions is apparent roughly at the same time that the OSN interaction is carried out. However, as information ages, the privacy preferences that relate to the information can change, as the information becomes irrelevant, embarrassing or un-noteworthy. It is unclear what is going to be the effect of time on sharing preferences and whether there is any temporal dimension to sharing preferences at all. In this study, we wish to narrow this crucial knowledge gap and to understand the temporal dimension of privacy preferences.

We present a study that investigates the privacy preferences and manifested behavior of Facebook users with regard to historical information. We investigate *retrospective privacy*: how sharing preferences correlate with the time passed since publishing the information. We conducted a survey of 193 Facebook users, who have been using Facebook for at least two years and have their profiles in a timeline form. The participants were surveyed regarding four posts published at different times: in the current time, a month prior to the survey, a year prior to the survey and two years prior to the survey. Our results show the extent and the nature of the relation between time and sharing preferences. We find that there is a negative correlation between willingness to share a post and the time passed since its first publication. Our analysis reveals other ties between time and user's preferences: a positive correlation between willingness to share the post and the post's relevance, as it is perceived by the user. Finally, we find that major changes in the user's life reduce the likelihood that a user would share a post that was published prior to the change.

Our main contributions are twofold: (1) we quantitatively describe the relation between the publication time and the privacy preferences of the social network information; (2) we analyze several aspects that can impact longitudinal privacy, including life-changes and information relevance. We discuss the theoretical framework necessary for understanding the relation between time and privacy, and demonstrate how the knowledge generated in our

study can guide the development of proactive and retroactive privacy management mechanisms.

2. RELATED WORK

The relation between time and privacy was discussed in two main domains: digital memory theories and privacy theories. In his book “Delete”, Mayer-Schönberger discussed the impact of digital memory on the individual and on the society [11]. To Mayer-Schönberger, digital memory negates time and threatens humans’ ability to make decisions, in both the individual and the societal level. Humans’ forgetting acts as a mechanism that puts past events in a temporal perspective:

“Forgetting plays a central role in human decision-making. It lets us act in time, cognizant of, but not shackled by, past events. Through perfect memory we may lose a fundamental human capacity to live and act firmly in the present” (p. 12)

When reviewing past information via digital memory, this perspective is lost and might cause indecision or incorrect judgment. Forgetting has also a societal function, accepting people’s ability to evolve and change by forgetting past events such as bankruptcies and previous relationships [22]. As forgetting in the digital age is non-trivial, and requires special mechanisms and processes, the current default in many information systems is to retain all the information collected about users. As a response, several scholars, from diverse domains such as Law [22] and Pervasive Computing [23], argued for the adoption of mechanisms that would allow information systems to forget information, for the sake of supporting individual and societal dynamics.

Palen and Dourish described the relation between time and privacy, focusing on the individual identity. Basing their framework on Altman’s boundary regulation theory [5], they described privacy management via three boundaries: disclosure, identity and temporal, and refer to the possible impact of information persistence on one’s online identity [4]. Since the user has a little control over how the information would be accessed in the future, temporal aspects are becoming part of any action people carry out, as personal information is almost never deleted from databases by default [11]. People’s identities are based on “artifacts of simply having been somewhere or done something at a particular time —such as visiting a cookie-enabled web page, or as being listed as a member of an email distribution list.” Westin emphasized the aspects of the individual’s control in data privacy: “The claims of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others” [24]. Due to the vast data collection and retention, one’s information can be found in several databases, allowing cross-referencing. As a result, it is easy to create an individual’s personality picture, reducing her control over her own information [25]. Zaho et al. performed one of the few studies which investigate users’ personal information management in OSNs *over time*. They used Goffman’s theatrical “front stage/back stage” metaphor and Hogan’s “exhibition in a museum” metaphor as their guidelines. These theories, together with their results of a qualitative study (n=13) brought them to define three functional regions in OSNs: performance, exhibition and personal. The last two regions are highly related to temporal aspect, reflecting the user’s usage of OSNs as a platform for long term presentation and archiving personal data, respectively [33].

Other studies specifically referred to privacy risks in OSNs and their relation to the durability of information. Gross and Acquisti referred to the data durability in the social network [1]. Together with the users’ information disclosure, it is possible to build a digital dossier, which might hurt users in the long run. Rosenblum mentioned the lack of realistic sense of how the shared information is public and permanent [2]. In the domain of social networks, Wang et al. found that a user’s common regret when publishing information on Facebook is unforeseen or ignored consequences [11]. For example, participants reported on unintended audiences which led to consequences that were not foreseen while publishing the information. Our paper complements these research works by specifically investigating the temporal aspect of privacy and sharing in OSNs.

Several recent studies investigated how access-control mechanisms can aid users better manage their privacy in OSNs. Vaniea et al. tested “proximity interfaces”, access-control interfaces embedded in the interaction of an online photo-sharing website [28]. They found that placing the access-control information below the photo enhance the participants’ attention to privacy policy. Their results highlight the importance of locating the privacy management task as close as possible to the users’ primary task [28]. Kairam et al. studied Google+, which allows users to share content with selected ‘Circles’ of people. Their results reflect a need to share information separately with people from different *life facets* and according to the social *tie strength* [29]. Sleeper et al. explored Facebook users’ self-censorship behavior. Despite the ability to selectively share content, users choose not to share some of the content. They suggested mechanisms to improve the existing tools for selective sharing. One of their results revealed that users are willing to share with or block users according to the post’s relevance: sharing with people which may find the information as relevant and blocking those who may find it as irrelevant. Therefore, they suggested implementing a system in which users will self indicate themselves as interested in a certain topic, similar to Twitter’s Hummingbird [30]. Mazzia et al. also explored groups’ privacy management and suggest *PVIs*, an interface that visually displays the user’s privacy policy with regard to her social groups [31]. The temporal dimensions of privacy can have significant effects on the effectiveness of access-control mechanisms. In this work, we review several options for enhancing existing information sharing interfaces to take into account longitudinal privacy management.

Acknowledging the challenges in managing historical data, several studies suggested technical solutions for handling information permanence on the web. Barua et al. presented a theoretical foundation for the design of mechanisms for forgetting in storages of personal information, enabling users to control their own information, including deleting, compacting, blocking and archiving information [6]. Geambasu et al. developed *Vanish*, a system that makes all copies of a certain data unreadable after a user-specified period of time. The self-destructing mechanism is aimed to help users controlling the lifetime of their Web objects such as emails, private messages on Facebook or private photos on Flickr [18]. Patil & Kobsa investigated privacy in Instant Messaging and developed *PRIVacy-Sensitive Messaging system (PRISM)*. The system allows a negotiation between the conversation contacts over setting of an expiration date for the conversation [32]. There are several mobile applications which also suggest a self-destruction mechanism for the objects which

are sent through them, such as Snapchat¹, Facebook's Poke, Wickr² and Vidburn³. As Peter Deng, Facebook's director of product management was quoted: "*The demand comes from real life,*" he said. "*People want something that is more lightweight than a message and less permanent.*" [26] Conley suggested ways of implementing deletion mechanisms using technical tools, legal regulation, social norms and market forces. The "Web 2.0 Suicide Machine" is another example for some progress in the direction of automated personal information deletion. The machine allows users to erase their records and profiles from multiple social networking sites using a single click [7]. In our work, we hope to inform and guide the design of technological solutions for supporting long-term usage of online social networks and other Web services.

3. RETROSPECTIVE PRIVACY

In this study we aim to understand the temporal aspects of information aging on sharing preferences in Facebook, currently, the most popular and widely-used OSN. The work explores *retrospective privacy*: the relation between information sharing preferences and the time passed since first publishing it. To understand retrospective privacy, we ask two research questions. First, what is the relation between the publication time of posts published on Facebook and the privacy associated with the information? Second, what is the impact of the user's life-changes and information relevance on sharing preferences?

We expect that people, who regularly publish information on an OSN, would have different sharing preferences for old and new content. This broad observation, leads us to introduce the following hypotheses concerning the impact of information age on sharing preferences:

H1. Facebook posts' relevancy to the user's social network decreases with time.

H2. Willingness to share information with the user's social network decreases with time.

H3. Willingness to alter the Facebook post (delete, change or to restrict access) increases with time.

H4. Major life changes increase the decline in willingness to share.

Section 4 describes the methodology we used to investigate the hypotheses, and the results are described in Section 5. Understanding the impact of time on sharing preferences will help us to outline possible design solutions to longitudinal privacy in OSNs. One of the goals of this study is to help designers analyze and evaluate tools that allow users to look back to what they did in the past, and to be better prepared to the future. Section 6 discusses the implications of the study to the design of privacy management tools.

4. METHODOLOGY

To answer the research questions and to test the hypotheses we designed and ran a within-subjects user study ($n = 193$), which was executed using Amazon's Mechanical Turk crowdsourcing

service. In the following sub-sections, we describe the variables, the research tool that was used, and the user study.

4.1 Survey

Participants were surveyed using a questionnaire that included four parts: screening, demographics, sharing preferences, and general Facebook usage. The survey took place in March 2012. In the sharing preferences part, we wanted to ask the participants about four posts: the latest post which they published, a post they published one month ago, one year ago and two years ago or earlier than that. In practice, the participants chose the posts on their own. Therefore, in order to achieve maximum consistency we guided the participants regarding the way to choose the posts. We asked them to choose them as follows: 1) the most recent post; 2) the first post that was published in the former month (in relation to the day they answered the survey), or earlier, in the case that they did not publish a post in the specific month; 3) the last post that was published at the end of the year 2010; 4) the last post that was published at the end of year 2009. If the account did not exist at that time, we asked to refer to the first published post. In the last two cases choosing the last post was the easiest access from the point of view of the participants, requiring them to view the first post which appears after choosing the requested year. As a result of instructing the participants to choose the posts as previously described, the participants were asked about posts which were approximately published 1.25 and 2.25 years before taking the survey. We asked the participants to enter the posts' publication dates. The time since publishing the post is the main independent variable of the study. Figure 1 provides a screenshot of our survey and details how participants are guided to find their posts in the requested time. See the appendix section for the entire questionnaire.

The dependent variables in the sharing preferences part were based on five-point Likert scale questions. Throughout the survey we did not use privacy wordings in order to avoid biases, as suggested by Braunstein et al. [34]. Instead, we asked the participants about issues related to information sharing. The questions included three questions about intended behavior: whether the participant intended to delete, change or block the content in the future. Following these questions, the participants were asked about their approaches towards the post: their satisfaction with the post's content, its relevancy and how willing they are to show the post on their timeline and to show it to their friends. Since Facebook's Timeline feature was new when running the survey (approximately 3 months), we asked both about simply sharing the post with Facebook's friends and also about keep publishing the post on timeline. The sharing preferences included questions that elicit further explanations from users and a set of controlling questions in which the participant categorized the post according to content type (text, photo, video, a link to an article, a link to a web site, a link to an application, and others), and according to the subjective intent of the post (expression of personal feelings, information about other people, information about work or company, information about relationships, political views, opinion about a certain product/service, a request for some kind of help, a question, and others.) The post's content type, its subjective intent and the user's demographic information were later used as control variables.

4.2 Recruitment

We recruited adult Facebook users using Amazon's Mechanical Turk (MTurk), a crowdsourcing service that is becoming more popular for use in HCI research [11]. We presented our survey as

¹ <http://www.snapchat.com>

² <https://www.mywickr.com>

³ <http://www.vidburn.com>

When viewing your personal profile, you will find a timeline at the right side of the screen, please click on "2010". The following questions are referring to the first status update you posted at the end of 2010.



24. When did you post it? *

27. On a scale from 1 to 5, how much do you agree or disagree with the following statements? (1 - I highly disagree, 5 - I highly agree) *

	1	2	3	4	5
I'm satisfied with the status *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I will delete the status in the future *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I will change the status' content in the future *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I will change the people who can view the status *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
The status is relevant today *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would like the status to be seen in my timeline *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I would like that my Facebook's friends will see the status today *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 1: A screenshot of the survey. The participants were asked several questions about each post. We added explanations throughout the survey, making sure the participants are able to locate the posts from the necessary time frames. For example, above screen shot directs the participants to a post they published a year before taking the survey.

a task on MTurk and presented few qualifications required for our survey. Participants were required to have an active Facebook account for the last two years, to have their Facebook account in the timeline mode, and to reside in the U.S. to ensure English proficiency. In the MTurk task description we did not mention privacy to avoid biasing our population towards more privacy sensitive people. We asked 200 participants to take our survey, calculating this number according to a confidence Level of 95% and margin of error of 1%.

Participants who answered our requirements were asked to follow a link to our survey on the SurveyGizmo commercial web survey service. They had electronically authorized the IRB-approved consent form, and the limitations on participation. At the end of the survey each participant received a unique number. We then compared between this number and the user's ID received from MTurk, allowing us to compensate only the participants who answered the survey seriously. We paid each participant \$0.75 for completing the survey. The survey took approximately 20 minutes

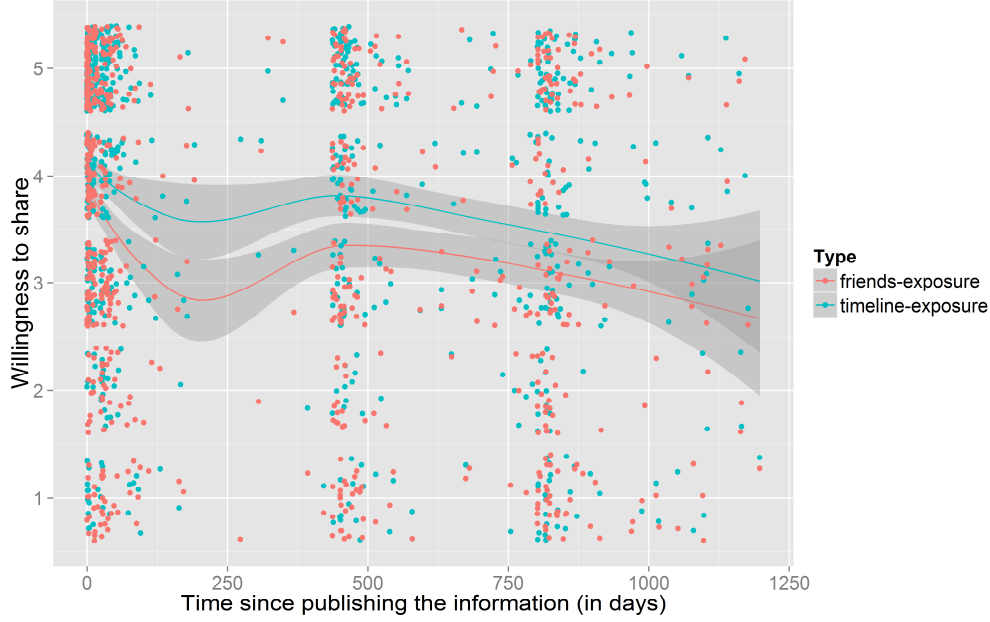


Figure 2: Timeline exposure (willingness to share on timeline) and friends exposure (willingness to share with friends) versus time since publishing the information (in terms of days). The grey areas represent the 95% confidence intervals around the mean

to complete and our compensation rate was about \$2 per hour, which is on par with the normal hourly pay on MTurk [11].

In order to determine whether our participants were taking their tasks seriously, we used a combination of measures. We gave a reading apprehension test, which included reading a short paragraph and answering two questions about it (screening task), as suggested by Downs et al. [10]. We also paid attention to inconsistent answers. Some of the questions were ranking questions and with contrary directions. An example for inconsistency is a participant who reflected the intention to delete a post and to keep it on timeline at the same time. If the participant chose the wrong answer in both of the screening task questions and provided inconsistent answers, we then excluded the participant's results. After manually checking the answers and filtering out suspected participants, we removed 7 responses out of 200. 59.07% of the participants were male, 39.9% were female and 2 participants chose not to expose their gender (1.03%). Fifty-one participants were between the ages of 18-24 (26.42%), 98 participants were between the ages of 25-34 (50.78%), 38 participants were between the ages of 35-54 (19.69%) and 6 participants were at the age of 55 or older (3.11%).

4.3 Data Analysis

In order to analyze the data, we first used the Shapiro-Wilk test for normal distribution. We used this test for all the dependent variables and we found that they are not normally distributed ($p < 0.0001$). Therefore, we used non-parametric tests for the next data analysis. We used three tests, according to the type of analysis: 1) the Spearman correlation test for testing correlations, marked as "p" in the results section. 2) the Kruskal-Wallis rank sum test for comparing between groups, marked as "H". 3) the Tukey method, a post-hoc test for finding which groups are significantly different, after performing the Kruskal-Wallis test.

5. RESULTS

Most of the variables were measured using the Likert scale, ranging from 1 to 5 (from *I highly disagree* to *I highly agree*). Therefore, when using the terminology "score" we are referring to the score received in these ranked answers. We controlled the results for the content type, the post subjective intent and demographics. Per each variable we compared the different groups (e.g., males versus females for gender variable) checking for significant differences in sharing preferences, using Kruskal-Wallis test. Only in the case of content type we found significant differences ($H = 14.01$, 6 df, $p = 0.03$). A post-hoc Tukey test showed that two types of content differ significantly: photo and text, with higher willingness to share in the case of photo. However, the content types distribute similarly between the epochs ($p = 0.31$, using chi-squared test).

5.1 Impact of Time

In our first set of results, we looked at the relation between the participants' approach regarding their published posts and the time passed since first publishing them. We looked at two variables that reflect two possible user approaches towards sharing information: *timeline-exposure* that reflects the willingness to show the post on the timeline (the phrasing was: "I would like the status to be seen in my timeline"), and *friends-exposure* that reflects the willingness to show the post to the participant's friends today (the phrasing was: "I would like that my Facebook friends will see the status today").

Figure 2 depicts the relation between time and willingness to share the post. The results show that both variables decrease with time. We found significant negative correlation between time and timeline-exposure ($\rho = -0.19$, $p < 0.0001$, using the Spearman correlation test) and with friends-exposure ($\rho = -0.22$, $p < 0.0001$). The average timeline-exposure score for new posts, which were published in one month since answering the survey, is 4.02 (SD =

Table 1. Mean scores and standard deviations of timeline-exposure, friends-exposure and relevance

	Current	1 m	12 m	24 m
Timeline-exposure	4.02 (1.19)	4.03 (1.14)	3.76 (1.35)	3.78 (1.46)
Friends-exposure	3.87 (1.27)	3.51 (1.37)	3.36 (1.48)	2.99 (1.54)
Relevance	3.62 (1.41)	3.06 (1.47)	2.71 (1.58)	2.55 (1.48)

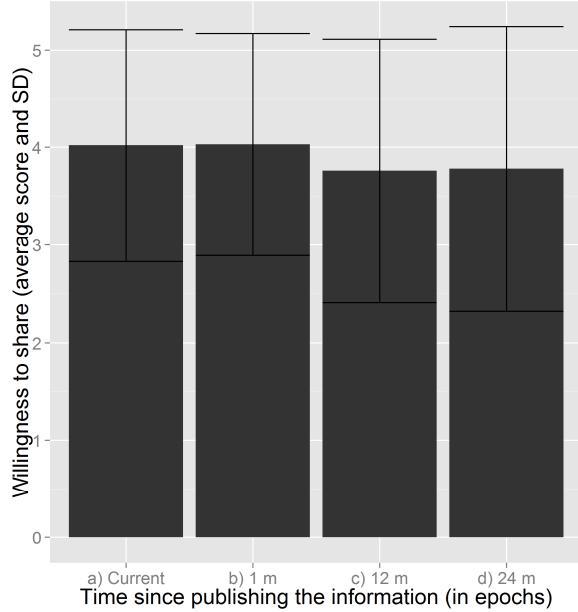


Figure 3: Average and SD of willingness to share per each epoch (time passed since publishing the information).

1.19), while for older posts, published at least two years earlier, is 3.38 (SD = 1.46). Timeline-exposure and friends-exposure reflect very similar privacy semantics, and therefore it is not surprising that the two variables are highly correlated ($\rho = 0.68$, $p < 0.0001$). However, throughout the tested period of time, friends-exposure score is lower than timeline exposure. A possible explanation can be the more concrete and tangible phrasing of the friends-exposure question. We also analyzed the results according to four discrete time frames, defining them as epochs: "currently", 1 month, 12 months and 24 months respectively. The mean and variance results of timeline-exposure and friends-exposure are presented in table 1. In the following sections we will refer to timeline-exposure (the willingness to show the post on timeline) as "*willingness to share*".

A second trend is observed when comparing the standard deviations of the means of willingness to share between the epochs. The Levene's test showed a significant difference in the variances between the epochs ($F = 10.69$, $p < 0.00001$). The results are illustrated by figure 3, showing a SD that increases with time. Although the diversity is larger for the current time rather than 1 month, there is a general increase over time. This may be derived by contradicting factors which impact on the users' willingness to share. For example, content may be irrelevant for one user, but nostalgic for another user.

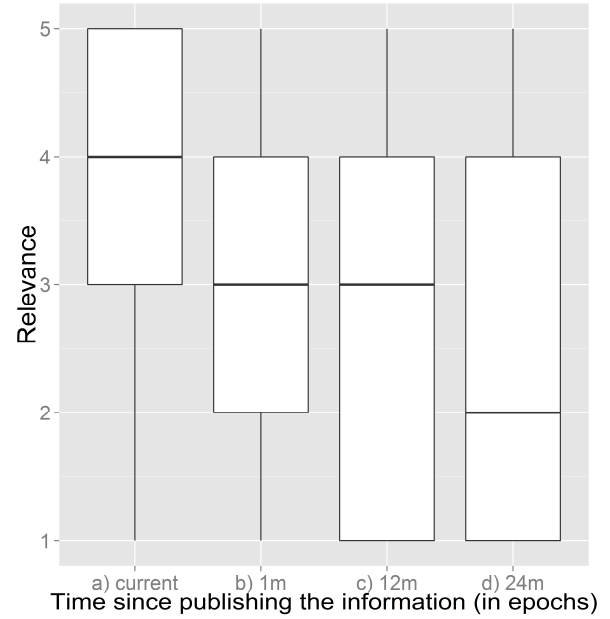


Figure 4: Relevance versus time since publishing the information (in terms of epochs).

5.2 Relevancy

In order to find possible explanations for the correlation between time and sharing preferences, we asked participants to rank the post's relevance (the phrasing was: "The status is relevant today"). We found a significant positive correlation between the post's relevance and the user willingness to share ($\rho = 0.47$, $p < 0.0001$). In addition, we found a significant negative correlation between the post's relevance and the time passed since publishing it ($\rho = -0.28$, $p < 0.0001$). The average relevance score for posts which were published one month before taking the survey is 3.06 (SD = 2.17) while the average score for posts which were published two years or earlier is 2.55 (SD = 2.20), reflecting about 20% decrease. Figure 4 displays the differences in the median relevance scores between the epochs. The means of relevance scores and SDs are presented in table 1.

Another aspect of the post's relevance can be seen in the reasons participants gave for deleting a post. The participants were asked whether they would consider deleting a post and if so, why they would do that (the phrasing was: "If you would consider deleting the status, why would you do so?"). In 230 cases out of 772 (29.79%) the participants would consider deleting it. The distribution of the reasons can be seen in Figure 5. Irrelevance of the post was found as the most popular reason for considering deleting the post (the reason was irrelevance for 127 cases out of 230 or 55%) The second reason was a change in the point of view of the participant since the post was published (55 out of 230 or 24%). The other given reasons (offending, inappropriate and other) answer for 27, 11 and 10 of the other cases, respectively.

We also asked the participants to rank their satisfaction with the post (the phrasing was: "I am satisfied with the status"). We found a significant positive correlation between the user's satisfaction and the user willingness to share ($\rho = 0.25$, $p < 0.0001$). A positive correlation was found between satisfaction and the post relevance ($\rho = 0.21$, $p < 0.0001$). The more the post is relevant, the more the user is satisfied with the post. We did not find a significant correlation between satisfaction and time.

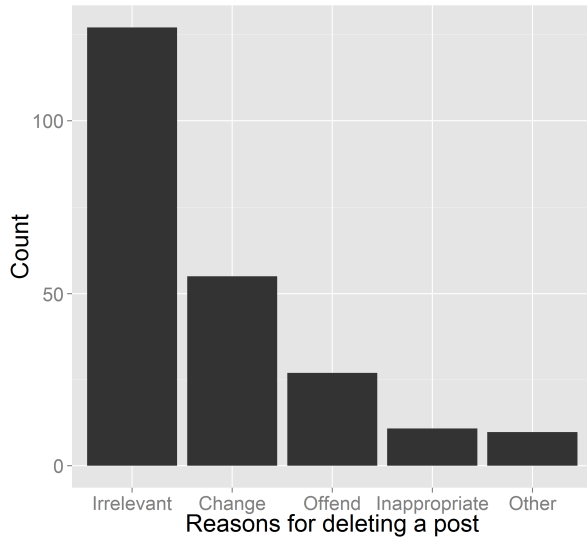


Figure 5: Number of instances for each reason given to delete a post. Irrelevancy is the most popular reason.

5.3 Users' Behavior

In our second set of results, we looked at the effect of time on the manifested behavior of the participants, their indication to alter the post: to delete the status, change its content or block it to their friends (the phrasings were, respectively: “I will delete the status in the future”; “I will change the status' content in the future”; “I will change the people who can view the status”). When analyzing the results we did not find significant results for any of the mentioned actions; therefore, we normalized the participants' answers. The idea of normalizing the results is to remove biases which might stem from using a Likert scale. Such biases may occur due to variance between the participants, ranking their answers differently. For example, one participant's high score might be another participant's medium score. Therefore, we look at comparative results, based on the participants' answers. Per each participant we calculated the mean score and SD for all the answers provided in all epochs, providing us with a standard score. The interpretation of the resulted score was defining a direction, positive score for increased willingness to share and negative for a decrease, compared to the average of the particular users. Indeed, after normalizing the results, we found significant correlation between the likelihood to delete the post and the time passed since publishing the information. The correlation is weaker than the correlation between willingness to share and time (delete: $\rho = 0.08$, $p = 0.03$; willingness to share: $\rho = -0.19$, $p < 0.0001$). The weakness of the correlation may explain the existence of significant results only in the case of normalized results. When not normalizing the results, biases which may occur due to Likert scale overshadowed the weak correlation between time and deletion, thus resulting in insignificant results. The likelihood to change the people who can access the post or to block it increases over time but in a very mild manner.

5.4 Life Changes and Social Circles

Life changes, major changes occurring in users' lives, may impact longitudinal privacy for many possible reasons: changes in social circles that happen due to changes in life

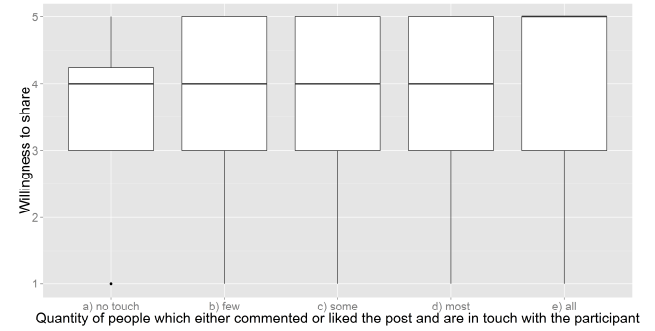


Figure 6: Willingness to share versus quantity of people the participant is in touch with and commented or liked the post.

Table 2. Mean scores and standard deviations of willingness to share, depending on the occurrence of life changes.

Changed occurred	Yes	No
Mean and SD	3.61 (SD 1.35)	3.91 (SD 1.25)
n	301	471

circumstances, changes in individual preferences and so on. We asked the participants to reflect on the people who commented or liked the post and choose the best description of the relationship between them (the phrasing was: “review the people who commented or liked or reshared the status. What best describes your relationship with them today?” Possible answers, for example: “I'm in touch with **all** of them”, “I'm in touch with **few** of them”). We found significant differences in the willingness to share between the groups ($H = 26.966$, 4 df, $p < 0.0001$). A post hoc Tukey test shows that the following pairs of groups differ significantly: 1) few and all ($p = 0.01$); 2) most and all ($p = 0.02$); 3) some and all ($p < 0.0001$). In other combinations of pairing of groups we did not see significant difference between them. Figure 6 shows that in the case when the participants are in touch with all of those who commented or liked the post, the score for willingness to share is significantly higher than in the other categories.

Finally, we asked the participants to indicate whether any life-changes had occurred since publishing a given post (the phrasing was: “indicate the changes that happened in your life since you had posted the status”). We offered several options, such as moving to a new town, changing relationship status, getting a new job, having a baby and so forth. The participants could choose all kinds of options at once, resulting in many combinations available. Therefore, we first divided the results to two categories: “changes occurred” and “no changes occurred”. A Kruskal-Wallis rank sum test showed significant differences in the willingness to share scores between the two groups ($H = 7.98$, 1 df, $p = 0.0047$), after controlling for time. The results are represented in Table 2, showing higher willingness to share when no changes occurred. We find significant differences in the scores between the changes types ($H = 16.19$, 6 df, $p < 0.01$), nevertheless, a post-hoc Tukey test did not show significant difference between specific groups.

6. DISCUSSION

Prior research has investigated privacy concerns, approaches and behaviors and demonstrated the challenges users face when trying to manage their privacy in OSNs [1, 11, 12, 14]. Our research adds the time dimension to the discussion, by analyzing retrospective privacy and contextualizing the temporal dimension of privacy approaches. When first facing retrospective privacy, we have defined four hypotheses, as previously mentioned: 1) Facebook posts' relevancy decreases with time; 2) willingness to share decreases with time; 3) willingness to alter the Facebook post increases with time; 4) major life changes increase the decline in willingness to share. Our findings confirm hypotheses 1, 2 and 4. Hypothesis 3 is not confirmed, showing insignificant ($p = 0.085$) temporal effect on deletion and significant temporal effect with regard to comparative approaches towards deletion, but not on content changing and blocking. Our findings demonstrate the decaying characteristics of Facebook posts with regard to relevancy and sharing preferences. About 89% of our participants are willing to share a recent post on Facebook, 81% are willing to share a post published a year ago and 73% are willing to share a post published two years ago. Why do participants change their sharing preferences? The analysis points to two possible explanations, which we call *staleness* and *anachronism*. In 51% of the surveyed posts, we witness information staleness: an information post that was published at least a year before the survey is defined as irrelevant by the participants, which is significantly correlated with the decline in willingness to share the information. In 20% of cases of irrelevant posts, users wished to delete the post. In their work, Zhao et al. also found that their participants are managing their profile to present relevant information [33].

Anachronism is expressed in the impact life changes and social circles have on sharing approaches. Sixty-three percent of the participants had reported at least one life-change, 79% of these had happened more than a month before answering the survey. As for social circles, 71% of the participants whose Facebook friends commented or liked their posts had a downward change in the relation with those friends since publishing the information (i.e., they are not in touch with *all* of them). The results may point to a situation in which social links have grown weaker over time, in the online or offline domain, thus resulting in a decrease willingness to share updates with them. The occurrence of life changes (such as moving to a new town, having a baby, starting a new work, changing relationship status and graduating from college) reduces the participants' willingness to share information that was published before the change occurred. We explain this result by the relation between life changes and publication context. When a life-change happens, there are high chances that the user's social circles will change as well. The new social relations might not understand the context of the old information, making the old information anachronistic. For example, co-workers in a law firm might not understand drunk photos taken during student times.

There is a gap between users' manifested approaches towards sharing the information and their manifested behavior. The likelihood of users to delete old posts is weaker than manifested sharing preferences. The likelihood to change the people who can access the post or to block it slightly increases over time. It is clear that there is a gap between users' sharing approaches and their plans for acting upon these approaches. This result might have several explanations: the well-known gap between

privacy preferences and actual privacy-related behavior on OSNs [1], personal uses people have for old information, or people's limited beliefs about their ability to control the information's distribution. Another explanation can be the users' tendency to stick to default settings, as described by Mackay [27]. In this case the default is keeping the post as it is on timeline. The provided options in the questionnaire (deleting, changing content and blocking people) did not cover all the options. Zaho et al. referred to Facebook's hide feature, which removes content from timeline, but keeps it accessible to other contexts, such as photo albums [33]. Offering a reversible option (i.e., hiding the post), might have been more widely accepted by the participants. The unwillingness to block the content to part of the audience was also previously observed, pointing at the complexity of such action, thus causing the users not to use it [33].

6.1 Theoretical Implications

This paper draws upon several theoretical threads to suggest a framework that model retrospective privacy. Mayer-Schönberger points at several elements that may cause historical information to be more privacy sensitive than present information: loss of control and limitations of accessibility and comprehensiveness. Facebook Timeline does provide accessibility and control of historical information (or at least the illusion of control), but we were able to prove and characterize the effect of durability on sharing historical information. To Mayer-Schönberger, older information may be perceived as less comprehensible to the user's current social network than new information. The high correlation between relevancy and sharing strengthen the validity of this observation.

In the boundary regulation theory of privacy, Altman describes privacy as a dynamic process, in which a person is constantly adjusting information available to others [5]. Palen and Dourish emphasize the fundamental change that persistent electronic information poses to privacy. They predict that people would rely on new ways to write and publish information in a way that is geared towards the future [4]. Our results provide an empirical basis for this view, characterizing the challenges users face when regulating privacy over long periods of time. Our results show that the variability in privacy behaviors grows over time. Some information remains relevant, and users are willing to keep sharing it. On the other hand, some is becoming irrelevant with time, and users wish to hide, and sometimes delete the information. Time complicates privacy concerns, resulting in varied sharing preferences, according to the specific sensitivity of the content and the context in which it was published. Furthermore, we see that irrelevancy is the primary reason for wishing to hide or delete an older piece of information. This phenomenon provides a clue to how users would actually future-proof their published information: by writing and framing it for a broader context than the immediate one. Whether people are writing for the future, what type of future audience people are writing for, and how people can actually write for the future, are fascinating questions that await future research.

Changing life circumstances, which makes old information anachronistic and inappropriate in the current period, plays a significant role in users' approaches towards historical information. Specific life changes, such as moving to a new town, having a baby, finding a new job or graduating, decrease the chances that users will want to share information from the time before the change. Privacy theories emphasize the

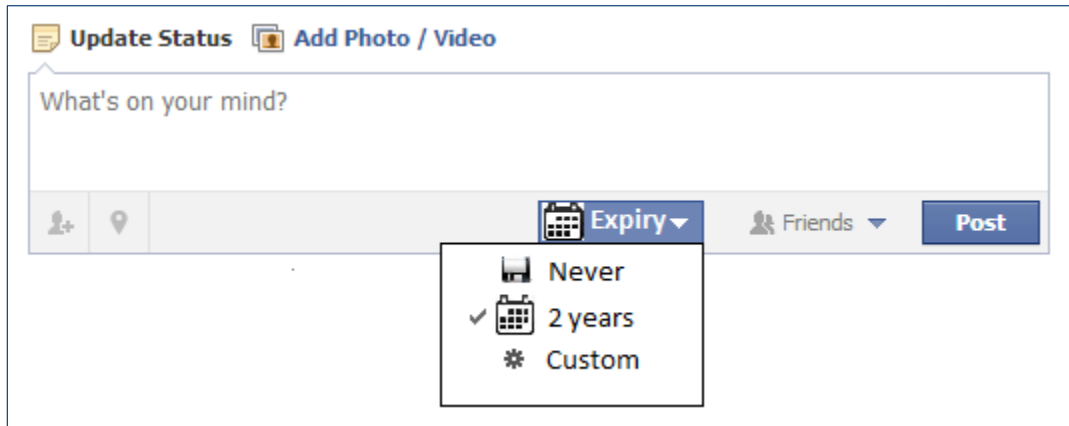


Figure 7: mockup implementation of expiry date for Facebook posts

importance of shared context in which information can be communicated and understood [21]. Life changes can bring on a change in the user's social circle and possibly in other spheres of life, and can represent a deep context switch. Users cannot be sure that information published in a different context, e.g., when the user was in college, will be understood in the same way two years later. As long-term usage of OSNs make it inevitable that some life changes would occur during this time, making long-term privacy management a complex and unchanging problem.

Another theoretical aspect of our results is the dynamics of information sharing and self representation in online social networks. Our results point to Facebook as a dynamic self-representation space, with time as a mediating factor in the characteristics of the presentation. Zaho et al. qualitative analysis of self-representation in online social network makes use of two impression management theories: Goffman's theatrical "front stage/back stage" metaphor and Hogan's "exhibition in a museum" approach [33]. Goffman's metaphor is used to describe users' profile management with relation to their recent published information, which was defined by Zaho et al. as the *performance region*. They refer to Hogan's approach with relation to the steps taken by users in order to manage longer term presentation, defined as *exhibition region*. Instead of keeping all bits of information, users are selectively choosing which items to keep, creating an "exhibition" of the self. Lastly, there exists a *personal region* for archiving meaningful facets of life. The users would like to keep this data, but not necessarily to share it with everyone else. Our results reflect the dynamic nature of temporal representation in Facebook, which can be explained as the tension between the different regions. The growing variability in sharing preferences in time can be explained by the transition from the exhibition region to the personal region, a process that contains some inherent uncertainty. Staleness and anachronism are expressed in the exhibition region, where users are less willing to share information due to irrelevancy and the out-of-date context.

6.2 Interfaces for Retrospective Privacy

One of our motivations for this paper was to inform the design of user interfaces and tools that help users manage their privacy in long-term usage of online social networks. One proposed mechanism is to reintroduce forgetting into the electronic domain by establishing an expiration date for information [9]. Expiration dates associate information stored in digital memory

with expiration dates that users set and that trigger the destruction of the information. Figure 7 displays our suggestion of expiry date implementation in Facebook. We suggest adding a feature which allows the user to determine the post's expiry date when publishing it, similar to the ability to determine its publicity, as exists today. We provide three options: 1) never, for those who wish to keep it for good (or at least they wish to do so at the publication point); 2) suggested default option, two years after publication date; 3) custom, for those who wish to set their own expiry date. The technical implementation of such a mechanism is non-trivial, but new approaches, based on cryptographic methods may provide the technical foundations for such a technology [18, 19].

Our results can be used to characterize the properties of expiration dates, while also highlighting the challenges in implementing it. One of the basic questions regarding expiration dates is how to set the default expiration date. In our analysis, we use linear regression in order to identify the tipping point with regard to the post's relevancy. Approximately 2 years after the initial publication date, the number of users who find the post as irrelevant crosses the 50% threshold, making it an interesting candidate for a default expiration date, as previously mentioned. Figure 8 displays the regression which led us to suggest two years as an optional default expiry date. It can be seen that beyond this point, more than 50% of the participants ranked the post's relevance below 2.5, out of 5. It can also be seen that as for the two sharing variables, timeline-exposure and friends-exposure, we cannot point at such 50% threshold cross point. Current study shows an increase in the proportion of those who are likely to give a low score for willingness to share their post. Further research may reveal the half-life for different types of information and contexts. Existing applications, such as Snapchat and Facebook's Poke, already use the mechanism of expiry date. However, our suggestion solves a different problem. These applications self-destruct messages in the immediate period of time, whereas our expiration date is used for the long term. The content of the messages is different as well: sexting versus everyday content [26]. Therefore, although such services are in use today, they answer a different need. Nevertheless, our results show that a one-size-fits-all solution, in which all messages are temporary, might not be adequate in situations where users wish to selectively share some of the information for short periods of time and some of the information for long periods of time. One important phenomenon to notice is the

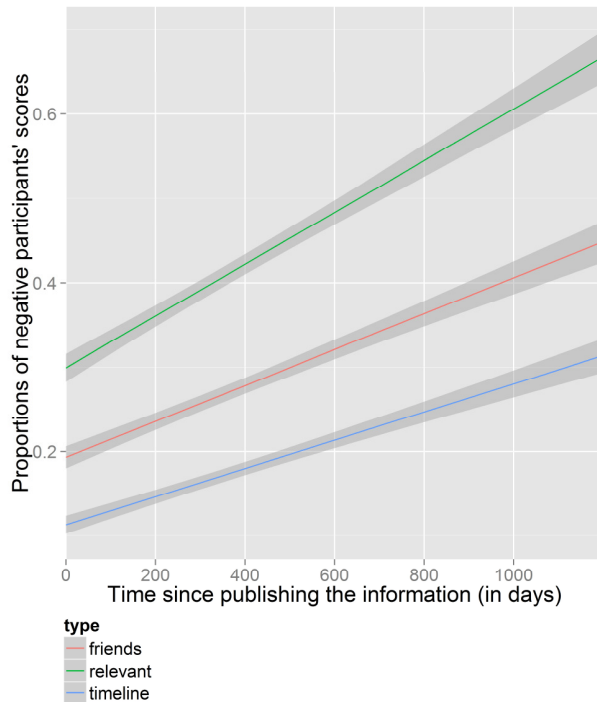


Figure 8: The proportions of participants who are not willing to share the information with their friends, on the timeline, or who believe that the information is not relevant, versus time since publishing the information (in terms of days). Per each variable (relevance, friends-exposure, timeline-exposure), the graph shows the proportion of the participants who ranked it below 2.5.

variance in sharing preferences. A one-size-fits-all expiration date might repulse the 73% of users who do wish to keep exposed a two-year old Facebook post. Moreover, the number of people who actually intended to delete a post was much lower.

We found a gap between the sharing preferences of users and their willing to delete or change a post. Therefore, solutions that completely and permanently delete content, such as the Web 2.0 Suicide Machine [20], might not become a general solution for retrospective privacy. However, there might be other mechanisms for facilitating ongoing privacy management other than deleting or changing the information: a) Archiving: moving irrelevant information to a secondary storage that cannot be accessed directly by current audience; Compaction: summarizing detailed information into a more compact view, and transferring the detailed part to the archive; b) Blocking: limiting the disclosure of the information for some of the people, according to time-based rules. These options keep the information accessible to the users while preventing from others to have an access.

Allowing users to review aged content they published can be a simple, yet effective, method for managing longitudinal privacy on OSNs. The diversity in sharing preferences hints that tools that allow users to review and quickly manage aged content might be useful. The Facebook Timeline is a good start for designing such a tool, as it allows quick access to historical posts. While allowing access to the user's social circle, it can also aid users review and manage their information. The

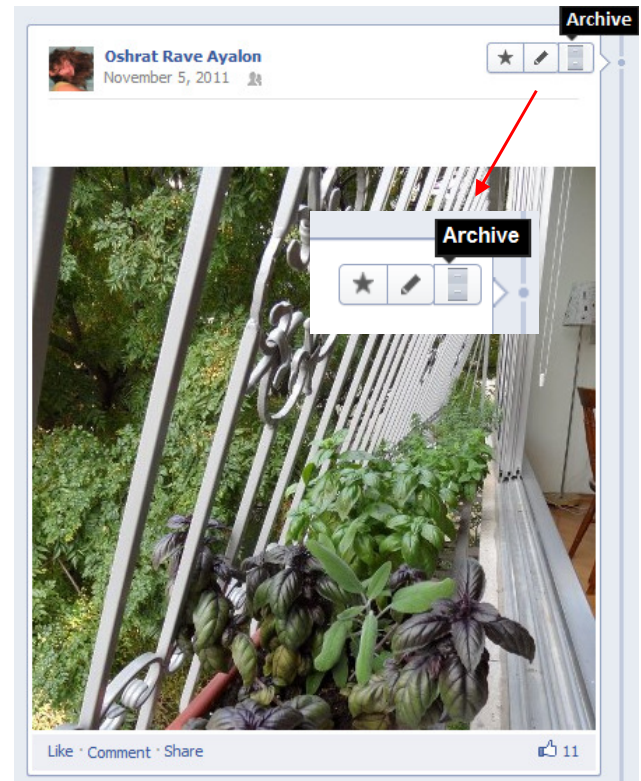


Figure 9: Implementation of highlighting posts which might be problematic in the eyes of the posts' publisher. Hovering on the archive icon reveals the possible options: archive, or using the existing options: edit or remove (pencil) or highlight (star).

correlations reported in this paper can be used as a theoretical basis for an automatic tool that selects potentially irrelevant posts for the user's to review. Figure 9 depicts a sample design, based on the Facebook Timeline, which already visually presents temporal aspect of Facebook information stream. We suggest adding an action, visualized by an icon of a filing cabinet, which will allow the user to archive the post. The icon will appear automatically for posts which are suggested to be more sensitive, such as the correlation between weak tie with those who commented or liked the post and willingness to share. When hovering over the icon the user will have several options: the filing cabinet, standing for archiving the post, and the two options which exists today in Facebook: "highlight" or "edit or remove". Facebook's hide feature provides the option to remove content from timeline, but not from other Facebook's sources, as photo albums. We suggest to fully restricting the access in the case of deciding to archive information.

6.3 Limitations

Our study is subject to several limitations that impact its applicability for design and research. First, our survey respondents were all recruited from Mechanical Turk, and thus our results may not necessarily be representative of the whole Facebook user population. In addition, since we recruited Facebook participants who had Facebook's Timeline feature, our results may not be applicable to other social media such as Twitter. Furthermore, the participants picked Facebook posts according to our instructions, which may introduce hidden biases towards selecting specific posts. In addition, performing a

within-subjects study may cause biases as well, since the users were aware of the changes in time of publications. As described at the results section the most common reason for deleting a post was irrelevance. However, the fact that we asked whether the post is relevant in the former question might have caused a bias. Prior research had demonstrated the gap between users' stated privacy preferences and their actual behavior on ONSs [1]. Since our data is self-reported, it may not always predict users' actual behavior. Lastly, our suggested tools did not cover third parties' ability to keep users' information once they had published. We referred to the users' ability to control the access according to their privacy settings (e.g., friends, public, etc.).

7. CONCLUSIONS

A great deal of past work in social network privacy has focused on examining sharing and privacy behavior in a synchronous time mode, where the information publisher and consumer had interacted with the information in time frames measured by minutes, hours and days. However, to the best of our knowledge, this is the first empirical study that looks into users' concerns and behavior regarding long-term sharing of social network information. To date, the relationship between time and willingness to share information was not captured in a methodological and comparative way.

Our findings show that there is a significant negative correlation between the time in which information was published on Facebook and its sharing preferences. The study sheds a partial light on the reasons that lead to this correlation: the relation between relevancy and sharing, and the effect of context. Finally, we highlight how new technological and processes can help users control their privacy over time. As online social networks are playing an increasing role in our lives, understating long-term usage, and even lifelong usage, is crucial. In future research, we plan to further explain the reasons behind information staleness and anachronism, and to design and evaluate tools for helping users manage their sharing preferences for long periods of time.

8. REFERENCES

- [1] Gross, R. and Acquisti, A. (2005). Information revelation and privacy in online social networks. *WPES'05* (pp. 71-80). Alexandria, VA: ACM.
- [2] Rosenblum, D. (2007). What anyone can know: The privacy risks of social networking sites. *IEEE Security & Privacy Magazine*, 5(3): 40.
- [3] Madden, M. and Smith, A. (2010, May 26). *Reputation management and social media*. Retrieved from The Pew Research Center's Internet & American Life Project: <http://www.pewinternet.org/Reports/2010/Reputation-Management.aspx>
- [4] Palen, L. and Dourish, P. (2003). Unpacking "Privacy" for a Networked World. *CHI '03* (pp. 5(1):129-136.). Fort Lauderdale, Florida: ACM.
- [5] Altman, I. (1975). *The environment and social behavior: Privacy, personal space, territory and crowding*. Monterey, CA: Brooks/Cole Pub. Co., Inc.
- [6] Barua, D., Kay, J., Kummerfeld, B. and Paris, C. (2011). Theoretical foundations for user-controlled forgetting in scrutable long term user models. *OZCHI*, (pp. 40-49). Canberra.
- [7] Conley, C. (2010). The right to delete. *AAAI Spring Symposium Series*, 53-58.
- [8] Bond, M. (2011, February 11). *Facebook Timeline a new privacy test*. Retrieved from USA today: <http://usatoday30.usatoday.com/tech/news/internetprivacy/story/2011-11-02/facebook-timeline-privacy/51047658/1>
- [9] Mayer-Schönberger, V. (2009). *Delete: The Virtue of Forgetting in the Digital Age*. Princeton University Press.
- [10] Downs, J. S., Holbrook, M. B., S. Sheng, S., and Cranor, L. F. (2010). Are your participants gaming the system? Screening Mechanical Turk workers. *CHI'10* (pp. 2399-2402). New York, NY, USA: ACM.
- [11] Wang, Y., Komanduri, S., Leon, P., Norcie, G., Acquisti, A., and Cranor, L. (2011). "I regretted the minute I pressed share": A qualitative study of regrets on Facebook. *SOUPS'11*, (pp. 1-13). Newcastle, UK.
- [12] Stutzman, F. and Kramer-Duffield, J. (2010). Friends Only: Examining a privacy-enhancing behavior in Facebook. *CHI'10* (pp. 1553-1562). Atlanta, Georgia: ACM.
- [13] Egelman, S., Oates, A. and Krishnamurthi, S. (2011). Oops, I did it again: mitigating repeated access control errors on Facebook. *CHI'11* (pp. 2295-2304). Vancouver BC: ACM.
- [14] Madejski, M., Johnson, M. and Bellovin, S. M. (February 2011). *The Failure of Online Social Network Privacy Settings*. Technical Report CUCS-010-11, Columbia University.
- [15] Liu, Y., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. (2011). Analyzing Facebook privacy settings: User expectations vs. reality. *IMC'11*. Berlin: ACM.
- [16] Bornoe, N. and Barkhuus, L. (2011). Privacy management in a connected world: 131 students' perception of Facebook privacy settings. *Workshop on collaborative Privacy Practices in Social Media*.
- [17] Lewis, K., Kaufman, J., and Christakis, N. (2008). The taste for privacy: An analysis of college student privacy settings in an online social network. *Journal of Computer-Mediated Communication*, 14(1):79-100.
- [18] Geambasu, R., Kohno, T., Levy, A. and Levy H. M. (2009). Vanish: Increasing data privacy with self-destructing destructing data. *USENIX Security Symp*, (pp. 299-314). Montreal.
- [19] Tang, Y. L. (2010). FADE: Secure overlay cloud storage with file assured deletion. *SecureComm*. Singapore.
- [20] Web 2.0 Suicide Machine, <http://suicidemachine.org/>
- [21] Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(30).
- [22] Blanchette, J.F. and Johnson, D.G. (2002). Data retention and the panoptic society: The social benefits of forgetfulness. *The Information Society*, 18(1):33-45.
- [23] Dodge, M. and Kitchin, R. (2007). "Outlines of a world coming into existence": Pervasive computing and the ethics of forgetting. *Environment and Planning B*, 34 (3):431-445.

- [24] Westin, A. F. (1967). *Privacy and Freedom*. New York: Atheneum.
- [25] Birenhack, M. D. (2010). *Private Space: Privacy, Law & Technology [Hebrew]*. Ramat Gan: Bar Ilan University, Nevo.
- [26] Wortham, J. (2013, February 8). *A growing app lets you see it, then you don't*. Retrieved from The New York Times: <http://www.nytimes.com/2013/02/09/technology/snapchat-a-growing-app-lets-you-see-it-then-you-dont.html?pagewanted=all&r=0>
- [27] Mackay, W. (1991). Triggers and barriers to customizing software. *CHI'91* (pp. 153–160). New Orleans, LA: ACM.
- [28] Vaniea, K., Bauer, L., Cranor, L. F., and Reiter, M. K. (2012). Out of sight, out of mind: Effects of displaying access-control information near the item it controls. *Privacy, Security and Trust (PST), 2012 Tenth Annual International Conference on* (pp. 128-136). IEEE.
- [29] Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. (2012). Talking in circles: Selective sharing in google+. *CHI'12* (pp. 1065-1074). Austin, Texas: ACM.
- [30] Sleeper, M., Balebako, R., Das, S., McConahy, A. L., Wiese, J., and Cranor, L. F. (2013). The post that wasn't: exploring self-censorship on facebook. *The 2013 conference on Computer supported cooperative work* (pp. 793-802). ACM.
- [31] Mazzia, A., LeFevre, K., and Adar, E. (2012). The PViz comprehension tool for social network privacy settings. *The Eighth Symposium on Usable Privacy and Security* (pp. 13). ACM.
- [32] Patil, S., and Kobsa, A. (2010). Enhancing privacy management support in instant messaging. *interacting with computers*, 22(3), 206-217.
- [33] Zhao, X., Salehi, N., Naranjit, S., Alwaalan, S., Volda, S., & Cosley, D. (2013). The Many Faces of Facebook: Experiencing Social Media as Performance, Exhibition, and Personal Archive. *CHI'13* (pp. 1-10). Paris, France: ACM.
- [34] Braunstein, A., Granka, L., & Staddon, J. (2011). Indirect content privacy surveys: measuring privacy without asking about it. In *Proceedings of the Seventh Symposium on Usable Privacy and Security* (pp. 15). ACM.