

Privacy Concerns vs. User Behavior in Community Question Answering

Imrul Kayes
University of South Florida
Tampa FL, USA
imrul@mail.usf.edu

Nicolas Kourtellis
Telefonica Research
Barcelona, Spain
nicolas.kourtellis@telefonica.com

Francesco Bonchi
Yahoo Labs
Barcelona, Spain
bonchi@yahoo-inc.com

Adriana Iamnitchi
University of South Florida
Tampa FL, USA
anda@cse.usf.edu

Abstract—Community-based question answering (CQA) platforms are crowd-sourced services for sharing user expertise on various topics, from mechanical repairs to parenting. While they naturally build-in an online social network infrastructure, they carry a very different purpose from Facebook-like social networks, where users “hang-out” with their friends and tend to share more personal information. It is unclear, thus, how the privacy concerns and their correlation with user behavior in an online social network translate into a CQA platform. This study analyzes one year of recorded traces from a mature CQA platform to understand the association between users’ privacy concerns as manifested by their account settings and their activity in the platform. The results show that privacy preference is correlated with behavior in the community in terms of engagement, retention, accomplishments and deviance from the norm. We find privacy-concerned users have higher qualitative and quantitative contributions, show higher retention, report more abuses, have higher perception on answer quality and have larger social circles. However, at the same time, these users also exhibit more deviant behavior than the users with public profiles.

Keywords—Community question answering, privacy concerns, crowdsourcing.

I. INTRODUCTION

Community-based Question-Answering (CQA) platforms, such as Yahoo Answers (YA), Quora and Stack Overflow, are online platforms where community members ask and answer questions. For example, YA, launched in December 2005, has more than one billion posted answers [24]. Liu et al. [23] found that about 2% of web searches performed by users of YA lead to a question posted to the community.

Such communities have a social network component, where users can follow other users’ activity via updates. Privacy settings are typically available for users to personalize. Two conflicting goals in privacy setting configurations emerge: on one hand, the platform is most useful when user-generated content is publicly available. On the other hand, various studies on general-purpose online social networks (such as Facebook) showed that the users who exercise their privacy rights (specifically, by restricting the visibility of their content) are more engaged and thus contribute more to the community.

To the best of our knowledge, this is the first study of the association between users’ privacy concerns and contribution behavior in CQA platforms via analysis of user activity logs. In our previous work on cultures in YA [17], we found that users’ privacy concerns vary across cultures: users from individualistic countries are more concerned about their privacy compared

to collectivistic countries. However, this work doesn’t explore the relationship between users’ privacy concerns and their contribution behavior. In this work, we analyzed more than a year of activity traces from 1.5 million users from YA to answer the following questions related to users’ contribution behavior:

- (1) Are there quantitative and qualitative differences in user contributions between user groups with private vs. public settings?
- (2) Is user engagement (measured by frequency of contributing content and number of social contacts) correlated with user privacy settings?
- (3) Do users with privacy settings enabled tend to violate community norms more than users with public content?

Our study makes two main contributions. First, while the previous related studies [6], [28], [32] on Facebook were based on self-reported data (shown to be subject to bias [10], [25]), this study uses modifications of the privacy settings as a proxy of privacy concern, and users’ recorded activity logs to infer their behavior. This is the first data-driven study that shows correlation between privacy controls and online user behavior. Second, this study is the first that characterizes a CQA platform from the privacy perspective. Our study finds that privacy-concerned users contribute more to the community. They are more engaged, having higher retention and larger social circles, and have higher perception on answer quality. However, they also exhibit more violations of platform rules in asking and answering questions than the users with public profiles.

The paper overviews related work in Section II, describes the YA platform and our dataset in Section III and presents our data analysis in Section IV. We conclude with a discussion of results in Section V.

II. RELATED WORK

Community-based Question Answering has attracted much research interest from diverse communities as web science, HCI and information retrieval. We divide research on CQA in four categories: content perspective, user perspective, system perspective and social network perspective. Content perspective research focuses on various aspects of questions and answers such as answerability of questions [9], [29], question classification (e.g., factual or conversational) [12], [13], quality of questions [20], [34] and answers [2], [30]. Kucuktunc et al. [19] investigate the influence of gender, age, education level, and topic on sentiments of questions and answers.

User perspective research sheds light on why users contribute content: that is, why users ask questions (askers are failed searchers, in that, they use CQA sites when web search fails [23]) and why they answer questions (e.g., they refrain from answering sensitive questions to avoid being reported for abuse and potentially lose access to the community [7]). Moreover, Liu et al. [22] explore the factors that influence users' answering behavior in YA (e.g., when users tend to answer and how they choose questions). Pelleg et al. [26] investigate truthfulness of users and offer a quantitative proof that users post sensitive and accurate information to fulfill specific information needs.

System perspective research develops techniques and tools to improve platform usability. It includes routing questions to expert users [27], [35], extracting factual answers from QA archives [4] and reusing the repository of past answers to answer new open questions [31]. Weber et al. [39] derive "tips" (a self-contained bit of non-obvious answer) from YA to address "how-to" queries. Social network perspective research attempts to understand the interplay between users' social connections and Q&A activities such as analyzing the social network of Quora [37], using social network properties and contribution behavior for content abusers detection [18].

A number of studies [6], [28], [32] on social networks like Facebook have shown the correlation between users' self-reported privacy concerns and their self-reported behavior. For example, Staddon et al. [32] showed that users who express concerns on Facebook privacy controls and find it difficult to comprehend sharing practices also report less engagement such as visiting, commenting, and liking content. At the same time, users who report more control and comprehension of privacy settings and their consequences are more engaged with the platform. Similarly, the frequency of visits, type of use, and general Internet skills are shown to be related to the personalization of the default privacy settings [6]. Acquisti and Gross' [1] survey on Facebook finds that a user's privacy concerns are only a weak predictor of his joining the network: that is, despite expressing privacy concerns, users join the network and reveal great amounts of personal information. Young et al. [40] used surveys and interviews on Facebook users to show that Internet privacy concerns and information revelation are negatively correlated. Tufekci's study [36] on a small sample (704) of college students shows that students on Facebook and Myspace manage privacy concerns by adjusting profile visibility but not by restricting the profile information.

Wang et al.'s [38] demographic study on privacy concerns among American, Chinese, and Indian social network users shows that American respondents are the most privacy concerned, followed by Chinese. However, there has been no research on privacy concerns and user behavior in CQA platforms. Our previous work on cultures in YA used Geert Hofstede's cultural dimensions [14], such as individualism index, and showed that users from higher individualism index countries exhibit higher level of concern about their privacy compared to the users from collectivistic countries. In this study, we focus on understanding how the users' behavior, characterized by broad engagement, accomplishments and deviance metrics, relates to their privacy concerns.

III. DATASET DESCRIPTION

Launched by Yahoo! in 2005, YA is available in 12 languages and has 56M monthly visitors in U.S. alone¹. The functionalities of the YA platform and the dataset used in this analysis are presented next.

A. The YA Platform

YA is a CQA platform in which community users ask and answer questions on various topics in predefined taxonomies, e.g., *Business & Finance*, *Cooking*, and *Politics & Government*. A question consists of a title and a body (typically, additional details). Members can find questions by searching or browsing through the hierarchy of categories.

The goal of asking a question is to find a best answer for the question. Users can write one answer per question and a question remains open for four days to answer. The asker can extend the answering duration for an extra four days. If the asker of the question selects a best answer within this time period, YA archives it as a *reference* question and only comments can be added to a reference question. The asker can rate a best answer between one to five, which is known as *answer rating*. However, if the asker doesn't select a best answer, community members get an opportunity to vote for a best answer. YA deletes all unanswered questions when the answering duration expires.

Users in YA can flag content (questions, answers or comments) that violates the Community Guidelines and Terms of Service using the "Report Abuse" functionality. YA requires its users to follow the Community Guidelines that forbid users to post spam, insults, or rants, and the Yahoo Terms of Service [3] that limits harm to minors, harassment, privacy invasion, impersonation and misrepresentation, and fraud and phishing. Users click on a flag sign embedded with the content and choose a reason between violation of the community guidelines and violation of the terms of service. They can select between two reasons: violation of the community guidelines (e.g., chat or rant, adult content, spam, insulting other members, etc.), or violation of the terms of service (e.g., harm to minors, violence or threats, harassment or privacy invasion, impersonation or misrepresentation, fraud or phishing, etc.). Reported content is then verified by human inspectors before it is deleted from the platform.

There is a point system in YA to encourage and reward participation. In short, a user is given two points for answering a question; ten points for a best answer. However, the user is penalized five points for asking a question, but if she chooses a best answer for her question, three points are given back. Users are ranked daily on a *leaderboard* based on their points. The points are also used to split users into seven levels (e.g., 1-249 points: level 1, 250-999 points: level 2, ..., 25000+ points: level 7). YA uses the levels to limit user actions, such as posting questions, answers, comments, follows, and votes: e.g., first level users can ask 5 questions and provide 20 answers in a day.

YA users follow each other and create a Twitter-like follower-followee relationship. Users are free to follow anyone. The followee's questions, answers, ratings, votes, best answers

¹<http://www.listofsearchengines.org/qa-search-engines>

and awards are automatically disseminated to the followers' newsfeed. In addition, users can follow questions, in which case all responses are sent to the followers of that question. Users can control the exposure of their information using two options in the privacy settings: they can choose to hide their content (questions and answers), and they can also choose to hide their network (followers and followees) from other users.

B. Dataset

We studied a sample of 1.5 million users, who were active between 2012 and 2013. These users are connected via 2.6 million follower-followee relationships in a social network that has 165,441 weakly connected components. The largest weakly connected component has 1.1M nodes (74.32% of the nodes) and 2.4M edges (91.37% of the edges). Our study includes all the users.

The available privacy configurations allow 4 user groups:

- 1) Public: all information is publicly visible (87.20% of users). This is the default setting.
- 2) QA-private: only Q/A information is private (2.23% of users), i.e., their questions and answers are visible only to their followers.
- 3) Network-private: only network information is private (0.81% of users), i.e., only their followers see the network.
- 4) Private: Q/A and network information is private (9.74% of users), thus only visible to the user's followers.

In the rest of the study we collectively refer to the QA-private and network-private users as semi-private. The default privacy in YA is public. It might be possible that many of the users in the public group are dormant: users who signed up, asked and answered some questions, and disappeared quickly. These users might skew the results of our study, thus, we only consider active users, who have asked and answered more than 10 questions. The active users are about 68% of the population and out of them 84.43% are public, 2.50% are QA-private, 0.89% are network-private, and 12.16% are private. We note that our observations remain the same even if we consider more active users by filtering-in users who have asked and answered more than 20 questions.

IV. PRIVACY AND USER BEHAVIOR

Our goal is to study the association between privacy concerns and behavior in YA. Previous works [6], [32] on Facebook have inferred users' privacy concerns using their self-reported feedback on privacy. Rather than self-reporting, which is subject to bias [10], we use modifications on privacy settings as a proxy for privacy concern. We measure several characteristics of user behavior that are related to CQA such as engagement, retention, accomplishments, abuse reporting, and deviance. We ask the following research questions:

- 1) **Is privacy preference associated with user engagement?**

We consider two metrics of user engagement: retention, which measures the average interval time between consecutive user contributions (addressed in Section IV-A), and social engagement, given by the number of followers and followees (Section IV-B). This question aims to

investigate the pattern identified in survey-based Facebook studies, but using CQA-specific and more nuanced engagement metrics on longitudinal activity traces.

- 2) **Do privacy-concerned users contribute differently to the community than public users?**

Users contribute by posting questions and answering others' questions. The quality of user-generated content is measured in the number of best answers and the askers' satisfaction with the answer received. The overall activity is measured in points. We characterize user contributions quantitatively and qualitatively in Section IV-C.

- 3) **Do privacy-concerned users have different perception on answer quality than public users?**

Users can themselves select best answers for their posted questions or they can rely on community voting to mark the best answers. In Section IV-D, we look at how the community sees the best answers selected by the users who received them. Specifically, we compare the quality of best answers selected by privacy-concerned users with those selected by public users in terms of the number of thumbs-up and thumbs-down given by the community.

- 4) **Are privacy-concerned users also more abuse-conscious?**

Intuitively, engagement is also correlated with the desire to keep the community free of unethical users (who, for example, may post spam in violation of the community rules). The related analysis is presented in Section IV-E.

- 5) **Are privacy-concerned users more likely to violate community rules?**

Intuitively, reduced visibility can give a false sense of confidence that might lead to violations of community rules. One study [5] in online gaming social networks shows that newly found and banned cheaters are more likely to change their profile to a more restrictive privacy settings than non-cheaters. In YA, we ask, is this observed more with privacy-concerned users than public users? This question is studied in Section IV-F.

A. Privacy and Retention

We define retention as the inverse of the average time difference between two actions not marked as abusive (i.e., fair). We consider two types of retention, based on questions and answers. For both types, if a user has a high average time difference between two fair actions, her retention is low.

Figures 1(a) and (b) show the medians and CCDF of the question inter-event time for the different groups, respectively. On average, private users have lower question inter-event time (thus higher retention) than public users. The answer inter-event time in Figures 2(a) and (b) show similar patterns. It seems semi-private (QA-private and network-private) users have higher average inter-event time, compared to private users, but similar to public users.

We performed a Kruskal-Wallis test to assess the difference among privacy groups in terms of retention. The test shows that at least one of the groups is different from at least one of the others for question ($\chi^2 = 458.83, df = 3, p < 2.2e - 16$) and answer retention ($\chi^2 = 119.32, df = 3, p < 2.2e - 16$). All-pairwise comparison tests after the Kruskal-Wallis test show that besides the QA-private and network-private for question retention and network-private and private for answer retention, all others are different for questions and answers retention

($p < 0.05$). These results show that privacy-concerned users are more retained than others.

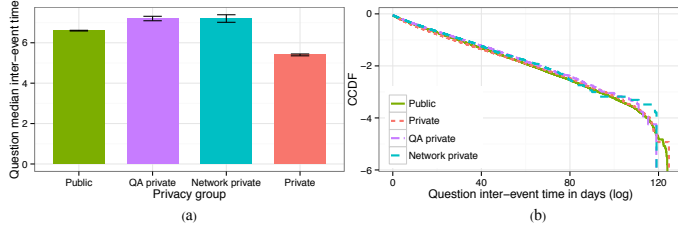


Fig. 1: (a) Median of question inter-event time in days with standard error bars; (b) CCDF of question inter-event time in days.

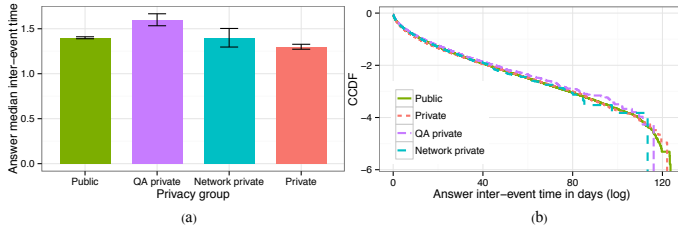


Fig. 2: (a) Median of answer inter-event time in days with standard error bars; (b) CCDF of answer inter-event time in days.

B. Privacy and Social Circles

YA users can follow each other, thus, we compute the indegree (total number of followers) and outdegree (total number of followees) of different privacy group users. Figures 3(a) and 4(a) show the median of indegree and outdegree, respectively, for the four privacy groups. The CCDF of indegree and outdegree of them are shown in Figures 3(b) and 4(b), respectively. While 20.56% of private users have more than 5 followers, only 4.42% of public users do. However, 15.33% of network-private and 14.48% of QA-private users have more than 5 followers. Alternatively, while 14.79% of private users follow more than 5 users, only 5.85% of public users do. For network-private and QA-private users, these numbers are 12.92% and 9.79%, respectively.

The results indicate that more restrictive private settings users have richer social circles. Indeed, Kruskal-Wallis tests show that at least one of the privacy groups is different from at least one of the other groups, for both the indegree ($\chi^2 = 29383.67, df = 3, p < 2.2e - 16$) and outdegree ($\chi^2 = 2913.63, df = 3, p < 2.2e - 16$). All-pairs comparison tests between the privacy groups show that all pairwise privacy groups are different ($p < 0.05$) for indegree, and only network-private and private users are same for outdegree ($p < 0.05$).

C. Privacy and Accomplishments

We consider two accomplishments that measure the quantity and quality of user contribution, through the point system described in Section III-A. Quantity of contribution is measured by the points users earn for their activities. To measure

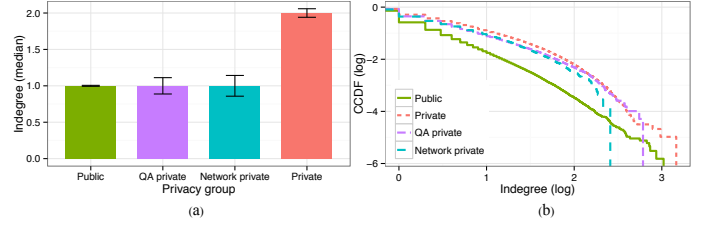


Fig. 3: (a) Median of indegree with standard error bars; (b) CCDF of indegree.

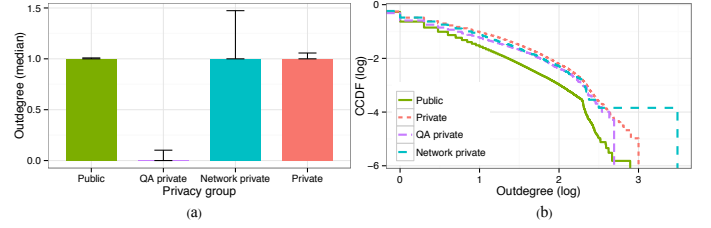


Fig. 4: (a) Median of outdegree with standard error bars; (b) CCDF of outdegree.

quality of contribution we use two metrics: Best Answer Percentage (BAP) and Award Rating Percentage (ARP). BAP is the percentage of a user's answers that are selected as best. ARP measures how satisfactory a user's best answers are. A YA asker can rate a best answer from 1 to 5 to declare how satisfied she is with the answer. ARP_j is the average rating a user j receives for her best answers:

$$ARP_j = \frac{\sum_{i=1}^{\# \text{best answers of } j} \text{Award rating for best answer } i}{\# \text{Total answers of } j * 5} * 100$$

Figure 5(a) shows median points with standard error for different privacy group users. It appears that median points of private and semi-private users are higher than public users. In fact, the CCDF of points in 5(b) shows that while 53.28% of private, 52.35% of QA-private and 45.51% of network-private users have more than 1000 points, only 14.14% of public users have more than 1000 points.

A Kruskal-Wallis test shows at least one of the privacy groups is different from at least one of the other groups for award points ($\chi^2 = 75884.12, df = 3, p < 2.2e - 16$). Moreover, all-pairs comparison tests between the four privacy groups show that besides private and QA-private, all others are different ($p < 0.05$). These results indicate that privacy-concerned users contribute more in YA from a quantitative point of view.

However, unlike quantitative contributions where public users are far behind the private ones, we found smaller, albeit significant, difference in the qualitative contributions among the four privacy groups. Figures 6(a) and 7(a) show the medians of best answer percentage (BAP) and award rating percentage (ARP) of different privacy group users, respectively. Although in both cases, private and semi-private group users have higher percentage than public users, the difference

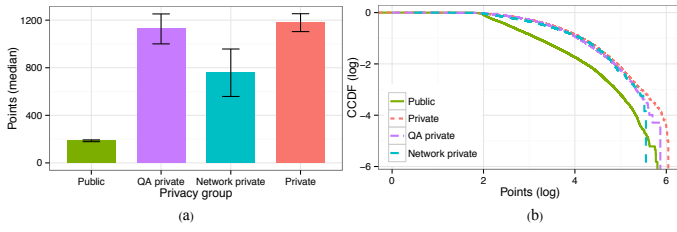


Fig. 5: (a) Median of points with standard error bars; (b) CCDF of points.

is less compared to points (even by a visible inspection on CCDF of BAP (Figure 6(b)) and ARP (Figure 7(b)) shows no difference across all privacy groups). Analyzing the CCDFs we get 27.96% of public users have best answers percentage more than 20, and 34.91% of private, 35.42% of network-private and 37.17% of QA-private users have best answers percentage more than 20. On the other hand, 27.10% of public, 33.56% of private, 34.03% of network-private and 36.01% of QA-private users have award rating percentage more than 20.

For both BAP and ARP, we notice that all privacy groups' numbers (median or CCDF) are close, especially private and network-private. So, one important question is how different privacy groups are in terms of users' qualitative contribution. We conducted a Kruskal-Wallis test on both BAP and ARP. The test results show that at least one of the privacy groups is different from at least one of the other groups for BAP ($\chi^2 = 5832.93, df = 3, p < 2.2e - 16$) and also for ARP ($\chi^2 = 5604.056, df = 3, p < 2.2e - 16$). Moreover, all-pairs comparison tests between the four privacy groups show that only private and network-private groups are the same ($p < 0.05$), and all other pairwise privacy groups are different. Thus, we confirm that privacy-concerned users have higher quantitative and qualitative contributions than others.

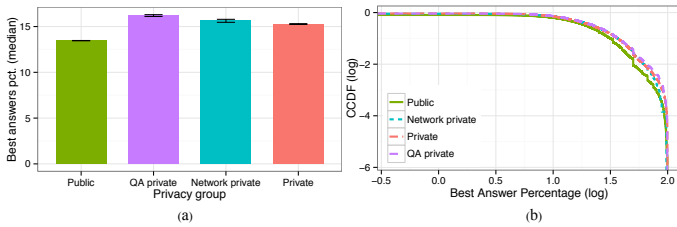


Fig. 6: (a) Median of best answers percentage with standard error bars; (b) CCDF of best answers percentage.

D. Privacy and Best Answer Quality

In YA, the best answer of a question is selected either by the asker of the question or by the community. If an asker does not select the best answer, the community members do that by voting. We first look at how different privacy groups are in selecting the best answers by themselves. We calculate the percentage of the best answers selected out of the total number of questions asked per user.

Figures 8(a) and (b) show the median and CCDF of the percentage of asker-selected best answers for different

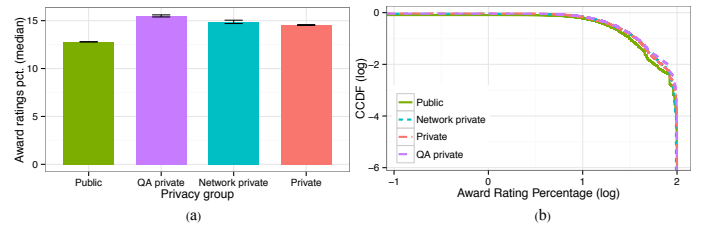


Fig. 7: (a) Median of award rating percentage with standard error bars; (b) CCDF of award rating percentage.

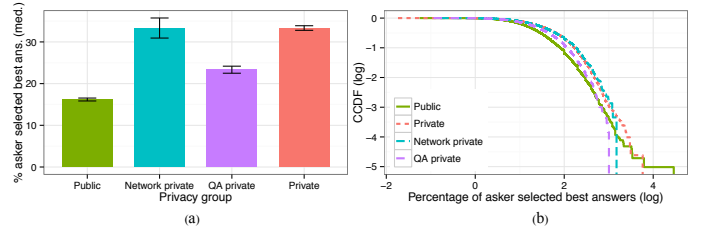


Fig. 8: (a) Median of percentage of asker selected best answers with standard error bars; (b) CCDF of percentage of asker selected best answers.

privacy group users, respectively. Analyzing the distribution, we observe that while 61.35% of private, 57.85% of network-private, 51.61% of QA-private users selected more than 20% of their best answers by themselves, only 38.35% of public users have done the selection by themselves. A Kruskal-Wallis test shows that at least one of the privacy groups is different from at least one of the other groups in terms of asker-selected best answers ($\chi^2 = 9522.60, df = 3, p < 2.2e - 16$). All-pairs comparison tests between the four privacy groups show that besides the network-private and private groups, all other pairwise privacy groups are different ($p < 0.05$).

Next, we focus on the quality of the best answers that users selected by themselves. We measure this quality based on community members' feedback on those answers. Community members can provide feedback on answers by giving either a *thumbs up* or a *thumbs down* (at most one such feedback per answer). For each user j who selected best answers to his own questions, we calculate the average number of thumbs as the ratio between the positive community feedback and the number of asker-selected best answers.

$$\text{AvgThumbs}_j = \frac{\# \text{Thumbs up} - \# \text{Thumbs down}}{\# \text{Best answers selected by } j}$$

Figures 9(a) and (b) show the median and CCDF of the average thumbs on best answers selected by the askers, respectively. The distribution shows that all private group users have more average thumbs on best answers than public users. We observe that while 21.40% of private, 24.62% of network-private, 16.51% of QA-private users have got 5 average thumbs on their best answers, only 11.45% of public users have got 5 average thumbs on the best answers they selected. A Kruskal-Wallis test shows that at least one of the privacy groups is different from at least one of the other groups in terms of

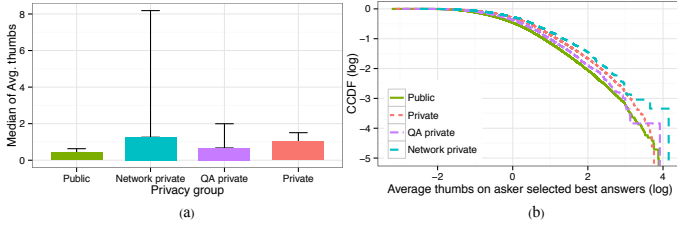


Fig. 9: (a) Median of average thumbs on asker selected best answers with standard error bars; (b) CCDF of average thumbs on asker selected best answers.

average thumbs with $\chi^2 = 5680.47$, $df = 3$, $p < 2.2e - 16$. All-pairs comparison tests between the four privacy groups show all pairwise privacy groups are different ($p < 0.05$).

E. Privacy and Abuse Reporting

As a crowd-sourced community, YA relies on its users for self moderation. Thus, users not only provide questions and answers, but also report inappropriate content using the abuse report functionality. If the report is valid, the content is deleted from the community. In this way, users serve as an intermediate layer in the YA moderation process since these abuse reports are verified by human inspectors. We have already seen that privacy preferences of users have significant association with a number of different dimensions including retention and accomplishments, thus we suspect that privacy is also associated with abuse reporting.

The median and CCDF of the valid abuse reports posted by users are shown in Figures 10(a) and (b), respectively. Although, abuse reports are highly appreciated for maintaining a clean CQA environment, very few people tend to report abuses. We find that 46% of the users reported only one abuse and 90% of abuse reports are contributed by only 7.96% of users. So, it's not surprising that all median values are zero in Figures 10(a). However, the private users have very high variability in abuse reporting compared to the public users.

The distributions in Figure 10(b) show that, on average, private users have posted more abuse reports than semi-private and public users. Indeed, all three private groups of users have posted a very large number of valid abuse reports compared to public users. Analyzing the distribution, we observe that 5.93% of private, 3.15% of network-private, 2.73% of QA-private and only 0.20% of public users have posted more than 10 valid abuse reports. A Kruskal-Wallis test shows that at least one of the privacy groups is different from at least one of the other groups in terms of abuse reporting behavior ($\chi^2 = 37647.77$, $df = 3$, $p < 2.2e - 16$). All-pairs comparison tests between the four privacy groups show that besides the QA-private and network-private groups, all other pairwise privacy groups are different ($p < 0.05$).

F. Privacy and Deviance

Deviant behavior is defined by actions or behaviors that are contrary to the dominant norms of the society [8]. Although social norms differ from culture to culture, within a context, they remain the same and they are the rules by which the

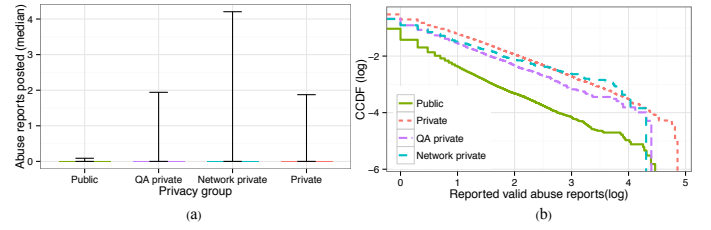


Fig. 10: (a) Median of average valid abuse reports with standard error bars; (b) CCDF of valid abuse reports.

members of the community are conventionally guided. YA has established norms as reflected by its community guidelines and terms of service [3]. We define user behaviors as deviant if they depart from these norms. In our previous work [18] on YA content abusers, we define a *deviance score* metric that indicates how much a user deviates from the norm in terms of received flags considering the amount of the user's activity. In short, we define the deviance score for a user u as the number of correct abuse reports (flags) she receives over the total content (question/answer) she posted, after eliminating the expected average number of correct abuse reports given the amount of content posted:

$$\text{Deviance}_{QA}(u) = Y_{Q/A,u} - \hat{Y}_{Q/A,u} \quad (1)$$

where $Y_{Q/A,u}$ is the number of correct abuse reports received by u for her questions/answers, and $\hat{Y}_{Q/A,u}$ is the expected number of correct abuse reports to be received by u for those questions/answers.

To capture the expected number of the correct abuse reports a user receives for questions/answers, we considered a number of linear and polynomial regression models between the response variable (number of correct abuse reports) and the predictor variable (number of questions/answers). Among them, the following linear model was the best in explaining the variability of the response variable.

$$Y = \alpha + \beta X + \epsilon \quad (2)$$

where Y is the number of correct abuse reports (flags) received for the content, X is the number of content posts and ϵ is the error term. In eq. (1), a positive deviance score reflects deviant users, i.e., those whose deviance cannot be only explained by their activity levels.

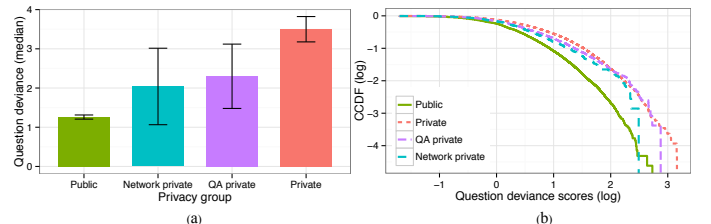


Fig. 11: (a) Median question deviance scores with standard error bars; (b) CCDF of question deviance scores.

Figures 11(a) and (b) show the median and CCDF of the question deviance scores, respectively. In both cases, private and semi-private users' question deviance scores are higher

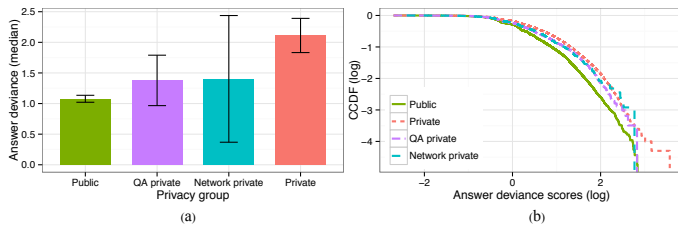


Fig. 12: (a) Median answer deviance scores with standard error bars; (b) CCDF of answer deviance scores.

than the public users. Also private users’ question deviance scores are higher than semi-private users. We reach to the same conclusion for the answer deviance scores from the median and CCDF of the answer deviance scores for all users in Figures 12(a) and (b), respectively. The Kruskal-Wallis test shows that at least one of the privacy groups is different from at least one of the other groups for question ($\chi^2 = 4432.72, df = 3, p < 2.2e - 16$) and answer deviance scores ($\chi^2 = 2662.416, df = 3, p < 2.2e - 16$). All-pairs comparison tests between the four privacy groups show that besides the network-private and QA-private groups, all other pairwise privacy groups are different ($p < 0.05$).

V. SUMMARY AND DISCUSSIONS

By performing a large-scale quantitative study, we have shown how users’ privacy concerns relate to their behavior in Yahoo Answers, a popular community question answering platform. We used users’ modifications on their privacy settings as a proxy of privacy concerns and grouped users into three main categories: private, semi-private (consisting of two groups, QA-private and network-private), and public.

Our study highlighted a number of results. First, we found that 87.20% of user accounts on YA are public, the default privacy setting. This result is similar with Gross and Acquisti’s study [11] on Facebook, where they found that about 90% of user profiles maintained the default, public setting. While expected, this confirmation warns again about the importance of correct default settings (e.g., privacy as contextual Integrity [15], [16]) in online applications.

Second, we discovered that users with enabled privacy settings are more engaged with the community: they have higher retention, more social network contacts, they are better citizens in terms of reporting abuses, overall they contribute more and better content, and have higher perception on answer quality. This is in line with Staddon et al.’s study [32] on Facebook, who found that users reporting more control and comprehension over privacy are more engaged with the platform. Therefore, this result is important for two reasons: it applies to a type of online community not previously studied, and it is based on user logs instead of user surveys, prone to self-reporting bias.

Third, we found that, on average, privacy-concerned users show more behavioral deviation in asking and answering questions than users with public accounts. At a first look, this result seems counterintuitive, given that privacy-concerned users keep the environment clean by reporting more abuses.

However, this result is consistent with our previous study [18], which finds that deviance in CQA platforms is not necessarily bad. Deviant users in YA are found to promote user engagement by attracting more users to answer more of their questions.

In addition to characterizing the association between privacy concerns and user behavior, our results may lead to improvements in CQA platforms operation. Whether an expression of privacy awareness or Internet savviness, users who modify their default privacy settings can be expected to be better citizens. If they change their account settings early on in their interaction with the platform, they send a clear signal to platform operators of likely commitment.

CQA platforms could benefit by targeting these users in a number of ways. For example, the indication of changing privacy settings can be used in *question recommendation*, where questions are routed to the most appropriate users who are more likely to answer. To find such answers, typical factors considered are followers, interests, question category, diversity and freshness; privacy settings can also serve as a complementary factor. Also, some of these users could be assigned community moderating duties to monitor community health, as our results show that they report more abuses. However, users who do not change their privacy settings are found to be less engaged. For these users, CQA platforms could provide extra incentives for participation and increased retention.

Our work also shows the importance of user-friendly and more practical design of privacy controls, as we find that increased engagement is associated with the use of privacy controls. For example, the lack of appropriate visual feedback has been identified as one of the reasons of the underutilization of privacy settings [33]. A better interface for setting privacy controls in the CQA platforms can impact users’ understanding of privacy settings and thus their success in exercising privacy controls.

We acknowledge that our study is observational, hence we can only associate privacy concerns with user behavior. In the absence of controlled experimental ground truth data, we cannot draw causal conclusions regarding whether users’ privacy concerns lead to different behavioral patterns in contribution. Understanding what makes users who change their default privacy settings on a CQA platform to also be more engaged in that community is among our future research objectives. The behavioral differences we have found in this paper could be used to create “privacy recommendation” in CQA sites, similar to the work of Li et. al [21]. Our future work includes using machine learning techniques on selected behavioral attributes from this study to predict and recommend privacy on Yahoo Answers.

ACKNOWLEDGMENT

The work was funded by the National Science Foundation under the grant CNS 0952420, and by the Yahoo’s Faculty Research and Engagement Program.

REFERENCES

- [1] A. Acquisti and R. Gross, “Imagined communities: Awareness, information sharing, and privacy on the facebook,” in *6th International Workshop on Privacy Enhancing Technologies*. Springer Berlin Heidelberg, 2006, pp. 36–58.

- [2] E. Agichtein, C. Castillo, D. Donato, A. Gionis, and G. Mishne, "Finding high-quality content in social media," in *Proceedings of the 2008 International Conference on Web Search and Data Mining*. ACM, 2008, pp. 183–194.
- [3] Y. Answers, "Yahoo answers community guidelines," http://answers.yahoo.com/info/community_guidelines, 2013.
- [4] J. Bian, Y. Liu, E. Agichtein, and H. Zha, "Finding the right facts in the crowd: Factoid question answering over social media," in *Proceedings of the 17th International Conference on World Wide Web*. ACM, 2008, pp. 467–476.
- [5] J. Blackburn, N. Kourtellis, J. Skvoretz, M. Ripeanu, and A. Iamnitchi, "Cheating in online games: A social network perspective," *ACM Transactions on Internet Technology*, 2014 (to appear).
- [6] danah boyd and E. Hargittai, "Facebook privacy settings: Who cares?" *First Monday*, vol. 15, no. 8, 2010.
- [7] D. Dearman and K. N. Truong, "Why users of yahoo!: Answers do not answer questions," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2010, pp. 329–332.
- [8] J. D. Douglas and F. C. Waksler, *The sociology of deviance: An introduction*. Little, Brown Boston, MA, 1982.
- [9] G. Dror, Y. Maarek, and I. Szpektor, "Will my question be answered? predicting "question answerability" in community question-answering sites," in *The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*. Springer, 2013, pp. 499–514.
- [10] S. D. Gosling, O. P. John, K. H. Craik, and R. W. Robins, "Do people know how they behave? self-reported act frequencies compared with on-line codings by observers." *Journal of personality and social psychology*, vol. 74, no. 5, p. 1337, 1998.
- [11] R. Gross and A. Acquisti, "Information revelation and privacy in online social networks," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*. ACM, 2005, pp. 71–80.
- [12] F. M. Harper, D. Moy, and J. A. Konstan, "Facts or friends?: Distinguishing informational and conversational questions in social q&a sites," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2009, pp. 759–768.
- [13] F. M. Harper, J. Weinberg, J. Logie, and J. A. Konstan, "Question types in social q&a sites," *First Monday*, vol. 15, no. 7, 2010.
- [14] G. Hofstede, G. J. Hofstede, and M. Minkov, *Cultures and organizations: software of the mind: intercultural cooperation and its importance for survival*. McGraw-Hill, 2010.
- [15] I. Kayes and A. Iamnitchi, "Aegis: A semantic implementation of privacy as contextual integrity in social ecosystems," in *11th International Conference on Privacy, Security and Trust (PST)*. IEEE, 2013, pp. 88–97.
- [16] —, "Out of the wild: On generating default policies in social ecosystems," in *IEEE International Conference on Communications—Workshop on Beyond Social Networks: Collective Awareness*. IEEE, 2013, pp. 204–208.
- [17] I. Kayes, N. Kourtellis, D. Quercia, A. Iamnitchi, and F. Bonchi, "Cultures in community question answering," in *Proceedings of the ACM 26th Conference on Hypertext and Social Media*. ACM, 2015.
- [18] —, "The social world of content abusers in community question answering," in *Proceedings of the 24th International World Wide Web Conference*. ACM, 2015, pp. 570–580.
- [19] O. Kucuktunc, B. B. Cambazoglu, I. Weber, and H. Ferhatosmanoglu, "A large-scale sentiment analysis for yahoo! answers," in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*. ACM, 2012, pp. 633–642.
- [20] B. Li, T. Jin, M. R. Lyu, I. King, and B. Mak, "Analyzing and predicting question quality in community question answering services," in *Proceedings of the 21st International Conference Companion on World Wide Web*. ACM, 2012, pp. 775–782.
- [21] Q. Li, J. Li, H. Wang, and A. Ginjala, "Semantics-enhanced privacy recommendation for social networking sites," in *Proceedings of the 10th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2011, pp. 226–233.
- [22] Q. Liu and E. Agichtein, "Modeling answerer behavior in collaborative question answering systems," in *Advances in Information Retrieval*. Springer, 2011, pp. 67–79.
- [23] Q. Liu, E. Agichtein, G. Dror, Y. Maarek, and I. Szpektor, "When web search fails, searchers become askers: Understanding the transition," in *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2012, pp. 801–810.
- [24] M. McGee, "Yahoo answers hits one billion answers," <http://searchengineland.com/yahoo-answers-hits-one-billion-answers-411671>, 2010.
- [25] D. Nishioka, Y. Murayama, and Y. Fujihara, "Producing a questionnaire for a user survey on anshin with information security for users without technical knowledge," in *45th Hawaii International Conference on System Science (HICSS)*. IEEE, 2012, pp. 454–463.
- [26] D. Pelleg, E. Yom-Tov, and Y. Maarek, "Can you believe an anonymous contributor? on truthfulness in yahoo! answers," in *International Conference on Social Computing (SocialCom)*. IEEE, 2012, pp. 411–420.
- [27] M. Qu, G. Qiu, X. He, C. Zhang, H. Wu, J. Bu, and C. Chen, "Probabilistic question recommendation for question answering communities," in *Proceedings of the 18th International Conference on World Wide Web*. ACM, 2009, pp. 1229–1230.
- [28] L. Rainie, A. Lenhart, and A. Smith, "The tone of life on social networking sites," <http://www.pewinternet.org/2012/02/09/the-tone-of-life-on-social-networking-sites/>, 2012.
- [29] M. Richardson and R. W. White, "Supporting synchronous social q&a throughout the question lifecycle," in *Proceedings of the 20th International Conference on World Wide Web*. ACM, 2011, pp. 755–764.
- [30] C. Shah and J. Pomerantz, "Evaluating and predicting answer quality in community qa," in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2010, pp. 411–418.
- [31] A. Shtok, G. Dror, Y. Maarek, and I. Szpektor, "Learning from the past: Answering new questions with past answers," in *Proceedings of the 21st International Conference on World Wide Web*. ACM, 2012, pp. 759–768.
- [32] J. Staddon, D. Huffaker, L. Brown, and A. Sedley, "Are privacy concerns a turn-off?: Engagement and privacy in social networks," in *Proceedings of the 8th Symposium on Usable Privacy and Security*. ACM, 2012, pp. 10:1–10:13.
- [33] K. Strater and H. R. Lipford, "Strategies and struggles with privacy in an online social networking community," in *Proceedings of the 22nd British HCI Group Annual Conference on People and Computers: Culture, Creativity, Interaction*. ACM, 2008, pp. 111–119.
- [34] K. Sun, Y. Cao, X. Song, Y.-I. Song, X. Wang, and C.-Y. Lin, "Learning to recommend questions based on user ratings," in *Proceedings of the 18th ACM Conference on Information and Knowledge Management*. ACM, 2009, pp. 751–758.
- [35] I. Szpektor, Y. Maarek, and D. Pelleg, "When relevance is not enough: Promoting diversity and freshness in personalized question recommendation," in *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 2013, pp. 1249–1260.
- [36] Z. Tufekci, "Can you see me now? audience and disclosure regulation in online social network sites," *Bulletin of Science, Technology & Society*, vol. 28, no. 1, pp. 20–36, 2008.
- [37] G. Wang, K. Gill, M. Mohanlal, H. Zheng, and B. Y. Zhao, "Wisdom in the social crowd: An analysis of quora," in *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 2013, pp. 1341–1352.
- [38] Y. Wang, G. Norice, and L. F. Cranor, "Who is concerned about what? a study of american, chinese and indian users privacy concerns on social network sites," in *4th International Trust Conference*. Springer, 2011, pp. 146–153.
- [39] I. Weber, A. Ukkonen, and A. Gionis, "Answers, not links: Extracting tips from yahoo! answers to address how-to web queries," in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining*. ACM, 2012, pp. 613–622.
- [40] A. L. Young and A. Quan-Haase, "Information revelation and internet privacy concerns on social network sites: A case study of facebook," in *Proceedings of the Fourth International Conference on Communities and Technologies*. ACM, 2009, pp. 265–274.