

Anonymous Quorans are still Quorans, just anonymous

Malte Paskuda
ICD, HETIC, Tech-CICO
Troyes University of Technology
UMR 6281, CNRS
12 Rue Marie Curie - CS 42060
10004 Troyes Cedex - France
malte.paskuda@utt.fr

Myriam Lewkowicz
ICD, HETIC, Tech-CICO
Troyes University of Technology
UMR 6281, CNRS
12 Rue Marie Curie - CS 42060
10004 Troyes Cedex - France
myriam.lewkowicz@utt.fr

ABSTRACT

This article presents a study that investigates how anonymity influences user participation in an online question-and-answer platform (Quora¹). The study is one step in identifying hypotheses that can be used to address a research and design issue concerning the role of anonymity in online participation, particularly among older informal caregivers. We present here a model that describes the factors that influence participation, which we based on the literature. These factors were used when analyzing the answers to questions in the health category on Quora. The results of this study complement an earlier study that we conducted on YouTube comments. On Quora, there was only one significant difference between anonymous and non-anonymous answers: with anonymous answers, social appreciation correlated with the answer's length.

Categories and Subject Descriptors

H.5.3 [Group and Organization Interfaces]: Web-based interaction

Keywords

Anonymity; social support; Quora

1. INTRODUCTION

The context of our research is an ambient assisted living project called TOPIC² (The Online Platform for Informal Caregivers), which aims to build an online platform for social support among elderly informal caregivers (people providing help to their ailing relatives).

The illness of someone close is obviously a burden for the relatives who are caring at home. However, how burdensome the situation can become might be a surprise. [41], a study of the burden met by informal caregivers of depressed family members, gives such an impression. It describes the frustration with the changed life situation, the discrepancy between the expectations for life before the

¹<https://www.quora.com/>

²<http://www.topic-aal.eu/>

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

C&T '15, June 27 - 30, 2015, Limerick, Ireland.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3460-0/15/06\$15.00

DOI: <http://dx.doi.org/10.1145/2768545.2768551>

illness and the reality after, and the common feeling of disconnection with the primary doctor. The depressed people change and can become aggressive and have mood swings. This might be especially hard for family members because they are surprised by that new behavior.

Informal caregivers search for ways to ease their burden, including ways to obtain information on the Internet. However, the information found is often contradictory. Another issue is the desire to conceal, which collides with the desire to share. Informal caregivers would like to share their situations with others, discuss them, and obtain advice and help [41]. However, they do not want to speak ill of their loved ones, and they do not want to fall under the stigma attached to those illnesses. A number of local initiatives exist that offer support groups and respite care services. However, many informal caregivers do not use these services because they are not aware of their existence, because they cannot organize to attend the support groups, or because they do not dare to talk openly about their emotions and mixed feelings (which is especially true for older husbands who are caregivers for their wives).

One solution for these problems is to design adapted online platforms [34]. When designing the TOPIC platform, we pursued the goal of offering informal caregivers a place to find validated information, hold support group meetings, seek help and share experiences in online discussion forums. The idea is that being online could help caregivers ease their burdens by providing them with social support in three dimensions: informational, emotional and tangible [33].

For such a platform, it is crucial that people take part. The more they participate, the higher the chance is that they will receive support. Experience sharing is especially important because it is a key factor in giving and getting social support [34]. However, it is difficult to know beforehand the critical factors that influence participation, experience sharing and general user satisfaction with online collaborative systems. Moreover, older caregivers are the target group of the TOPIC platform, but they use the web or even just computers less often than others [21]. An online platform for older informal caregivers must not only overcome this but also ensure trust. First, the caregivers need to view the information on the platform as credible. Second, they need to feel comfortable discussing something as private as life with an ailing relative.

To sum up, our goal is that such a platform will become more than a place to find information. Members should discuss with each other and help each other by sharing their experiences and giving

advice. In short, this place should become a home for a community of informal caregivers (one as described in [38]).

The question that we are focusing on in this article is the role that anonymity could play in this process of community building. In fact, this question intrigued us because there is no easy answer about the identity model that the platform should allow. On the one hand, user anonymity can diminish credibility [25], and anonymity could make it more difficult to organize people who live in the same area to later help each other offline. On the other hand, it could make it easier for informal caregivers to talk about the more intimate aspects of their role and not follow their desire to conceal. In consequence, we focus on both sides of anonymity: being anonymous and interacting with anonymous people.

In this context, our research question is then as follows: what is the influence of anonymity on user participation in general and experience sharing in particular? Specifically, what would be the best option when designing a tool for a community such as the one wanted in the TOPIC project (elderly people who are experiencing stressful situations)? Should users be able to choose their own names, stay completely anonymous, or use their real names?

To address this research and design question, the first goal is to define our expectations for the user identity model of the TOPIC platform. *Anonymous users will participate more in this social support platform* would be an example. To test this type of hypothesis, we defined the following process:

1. Conduct a twofold review of the literature in the field of computer-mediated communication, HCI and CSCW, to first define the factors that influence participation and then to search for factors that anonymity influences.
2. By looking at the overlapping factors, select the ones to be tested. It will be necessary to determine how to identify the factors in a text (here: online messages).
3. Collect messages on existing platforms with varying degrees of anonymity and analyze them for the factors defined in the earlier step.

This process generates hypotheses about how anonymity influences the factors selected in step 2. The hypotheses will then be tested in an experiment that will test anonymity on the TOPIC platform.

In this paper, we describe how we performed these three steps, in particular the study we conducted, which is our second study. The first was a study on YouTube: we looked at YouTube's move to Google+. In 2013, Google moved YouTube's comment system from the integrated one that allowed pseudonyms (and thereby a degree of anonymity) to Google+. This move gave us the chance to study the effect of changes in the degree of anonymity. Here, we present the study we conducted on the Q&A platform Quora. Quora has a function for answering anonymously and is thus a mixed environment. We analyzed these answers for differences between anonymous and non-anonymous Quorans.

In the following sections, we describe the related work on the influence of anonymity. We then present the model of participation factors that we built to generate hypotheses about our research question and that we used in our study. We then go on to present and

discuss our findings, explain the limitations of this study, and, finally, conclude.

2. RELATED WORK

In this section, we present the related literature about the effect of anonymity. We feature preferentially the effects on online interaction. However, we begin with the two fundamental theories about anonymity in general.

According to the deindividuation theory, a member in a group loses his self-awareness and thus loses his social conscience [22]. He can behave in ways he would normally never allow himself to behave because he is no longer bound by the norms of general society. In the framework of this theory, anonymity could fortify or enable this loss of control. In practice, the tone of the communication becomes less polite. Applying this theory to online interaction was coined *the online disinhibition effect* [31].

The reduced social cues approach is another perspective on this theory. In that model, electronic means do not transport the social elements used in face-to-face conversations, which leads to the deindividuation and deregulation of behavior [30].

The social identity model of deindividuation effects (SIDE) regards group behavior more positively, insofar as the actions taken by the members are in line with what is good for the group as a whole because group members identify with the group norms [9]. Anonymity is then twofold: are the others known, and do the others know me? If the others are anonymous to the member, this increases the likelihood of identifying with the group because individuality—which could hinder that identification—stays hidden. In practice, a user might feel more connected with the group (and this should improve the user experience).

Following these theories, there have been a number of studies on anonymity online. Kilner et al. conducted an important analysis about an online forum for soldiers that gradually changed its account model from anonymity with pseudonyms to asking for the full civil identity [14]. Kilner et al. analyzed the comments in the different stages and found that removing anonymity options led to fewer antisocial comments and fewer comments in total. This work heavily influenced our analysis in selecting possible hypotheses.

Another influence comes from a study on the move of the tech site TechCrunch from Disqus to Facebook as a comment system [20], thereby disabling the option to comment anonymously or under a pseudonym. By comparing comments from before and after the change, Omernick and Sood found evidence for a negative influence of anonymity on comment quality and politeness (what [14] would have classified as antisocial, thus underlining this result). However, those changes did not result in a significant decrease in participation; there were fewer comments, but they were longer.

In the area of behavioral science, a number of experiments (such as [16]) have attempted to find effects of anonymity. [40] measured the effect of having moderation features and changes in the response rate on the intent to participate. This experiment is noteworthy especially because it constructs a connection between interactivity and the reduced social cues model.

In computer science, theoretical models of anonymity and anonymity on a technical level play a more important role [10], and research exists on the link among politeness, civility and anonymity, analyz-

ing that link from a political angle [27].

There is also a broad amount of literature describing the factors influence participation. Although anonymity is seldom the focus, it is occasionally mentioned. An example of that is at the same time a main thread in the literature: the common identity and bond theory used by Kraut et al., as described in [26]. The theory assumes two connections between the members of a community—identity and bond—that are influenced by different factors, *social categorization*, *interdependence* and *intergroup comparisons* for identity and *social interaction*, *personal information* and *personal attraction through similarity* for bond. Some of these can be linked to anonymity, and SIDE theory does that explicitly with *personal attraction through similarity*.

Although the research cited here paints a mostly negative picture of the effect of anonymity on online interaction, one should not forget its positive effects. An example is the broad range of different and justified motivations to stay anonymous [13]. These effects are also deductible by following SIDE theory and noting the positive effects that are possible with easier group identification. The negative results, such as more rude comments in the case of available anonymity, do not unanimously show that this leads to less participation. Additionally, our own study of YouTube's change away from anonymous comments found more rude comments after the change, thus contradicting this specific expectation.

Those contradictions show that we are not able to simply use existing literature to paint a clear picture of the effect of anonymity.

We needed to build a model to show the factors that we found in the literature and extract those factors to use in our studies. The building of the model and how we plan to use it are described in the following section.

3. BUILDING A MODEL TO DESIGN STUDIES

The current section describes the background of the study that we conducted and explains the choice of measured factors.

To see the effect of anonymity on participation, we needed a model of factors that influence participation. In fact, we assumed that anonymity cannot be the only factor that influences participation. We claim, supported by the literature, that anonymity influences the factors that influence participation. This is the reason the first step in generating our model was to review the existing work to find general factors that influence participation. After having identified these factors, we made a second literature review focusing on the effects of anonymity.

3.1 Factors influencing Participation

We found many factors that might influence participation.

3.1.1 Anonymity itself

In [14], a platform moved stepwise from anonymous participation possibilities to making it necessary to disclose one's full civil identity. Kilner et al. observed that although many of the metrics that measured participation did not change, what did change was the number of comments posted.

3.1.2 Common identity

In [26], it was argued that attachment to the group influences participation. Kraut et al. mentioned two theories to explain that attachment: common identity and common bond.

Common identity theory makes predictions about the causes and consequences of people's attachment to the group as a whole [26, p. 377].

The authors also highlighted a number of factors that might help achieve a group identity and therefore foster participation:

Social Categorisation. Simply declaring that people are in a common group based on arbitrary criteria.

Interdependence. Being dependent on others members to achieve a common goal or by a shared fate. An example is described in [18]. In that experiment, users contributed more work when they were told that their work was unique and thus necessary to achieving the group's goal.

Intergroup Comparisons. Comparing group members with other groups.

3.1.3 Common Bond

Its definition:

Common bond theory makes predictions about the causes and consequences of people's attachment to individual group members [26, p. 377].

The theory highlights the following factors [26]:

Social Interaction. [8] described that the expression of negative emotions led to high interactions between users and to a high amount of participation in threads on the BBC forum. [12] also showed the success of personalized invitations, stressing the social aspect of a forum. However, this success was not observed in [28]; on the contrary, social aspects in invitations led to fewer registrations with less filled profiles. In [37], interaction was rather observed as a metric of a successful community. [32] described that depending on the community, direct interaction can be necessary to have an effect, in contrast to only creating social awareness.

The different studies show how difficult it is to distinguish between cause and effect in this area.

Personal Information. Opportunities for self-disclosure. For example, in [37], the option to have an avatar photo seemed to increase the number of messages and forum threads created.

Personal Attraction through Similarity. Because people like people they share similarities with.

The factors linked to common bond and common identity profit from a design that enables *social presence*, which is defined as making other users visible while using a system. An example would be integrating teams with visible and shared progress [11].

3.1.4 Other Factors

Many other factors exist that are not easy to integrate into common identity and common bond theory.

First to mention are the factors of the *community activity framework* [37]. Some factors are content based, such as the use of graphical emotions in posts and having rules to guide group discussions. Others are more functional, such as email notifications and posting counts next to posts or related news sections.

A number of authors describe *politeness* as a factor. [6] analyzed a small sample of messages from discussion groups and measured their perceived politeness with an Internet survey. Polite messages received three times more replies in technical groups, but impolite messages received more replies in political groups.

Introductions and *requests* are rhetorical strategies that are analyzed for their impact on responses in [5]. There, they increased the likelihood of replies by 7% and 6%. However, other rhetorical features such as the use of *self-references* also elicited responses:

Posts that included testimonials or requests were more likely to receive a reply. Including self-references (“I”), third-person pronouns, describing cognitive states and processes, and expressing either positive or negative emotions all increased the likelihood that a message received a response. The topical coherence of a message with respect to other recent discussions in the community also affected the likelihood of getting a reply [2, p. 959].

Timely and positive *feedback* increases the effort put into the task at hand or the general motivation [43]. Strong negative feedback decreases motivation to participate [42].

The *behaviour of the founder* of a group can influence its chance of success. For example, groups founded by very controlling people die early [15]. Similarly, in the context of a learning community, the number of prompts in the course material to answer questions by the organizers led to greater learner participation [1].

How to activate already present members in a community who do not actively participate (lurkers) receives special attention in the literature. [23] presented some factors that could activate these members, and these factors were divided into the categories usability and sociability. In its essence, it follows the thought that *easy access to the means to contribute* and *social appreciation of the contribution* will activate lurkers. In contrast, inactive users have a variety of reasons to remain inactive, including privacy concerns [19], and lurking is sometimes simply viewed as a metric that shows that the community does not fit the nonparticipants [24]. Nonetheless, lurkers are sometimes considered a strong negative factor for the survival of online communities, as in [29]. [29] also proposed that *perceived risks* and *social ties* are sufficient to explain lurking behavior.

3.2 Influence of Anonymity

It is a common thought that anonymity can change situations and that it influences various factors.

A survey of 44 people on the Internet with various backgrounds focused on the self-perceived merits of being anonymous. One such

merit is the *emotional benefit*. Additionally, the participants perceived anonymity as something that enables more honest ratings or recommendations [13].

Thus, anonymity appears to also influence *Credibility*. Although there are theories in both directions—both more and less credible—[25] observed that perceived anonymity decreased credibility.

Conformity in a group situation appears to be at least minimally affected by perceived anonymity [35].

[17] attributed *Uncivility* and *Impoliteness* to anonymity.

[29] stated that anonymity will result in *stronger social ties*, thus minimizing lurking behavior.

In contrast, [7] suggested that anonymity leads to more *antisocial behavior* in the context of grieving in online games. This phenomenon had already been mentioned in [14], where the removal of anonymity options led to fewer *antisocial comments*.

However, in [14] it was also measured whether the change had an *effect on participation*, and the authors found less direct participation but the same number of logins and page views.

An experiment with two groups tested the differences between groups with anonymous and identified members. They found the following:

Depersonalization was associated with greater attitude differentiation than individuation was [22, p. 11].

Participants also *identified more strongly* with their own groups. In a similar vein, as a fitting summary [3] states,

The attributes of anonymity, including minimal accountability, disinhibition, and deindividuation, can encourage robust political speech, provide safety from reprisal, permit the freedom to speak freely, and create a strong sense of group identity [3, p. 30].

3.3 The Model

From the factors that influence participation and the factors that are influenced by anonymity, we kept the intersecting factors, i.e., the ones influenced by anonymity that influence participation. Of those, we kept the factors for which we were able to find visible markers in text (see fig. 1).

We observed that a large part of the literature assumes that anonymity influences politeness (see [17]). Politeness appears to influence participation, and [8] showed that impolite comments provoked other comments.

The relationship between anonymity and intergroup comparisons and social interaction is indirect via social presence. [11] described that the factors linked to common bond and common identity could both profit from social presence, and [36] showed that anonymity influences social presence.

Anonymity can change the perception of contributions and can lead to less social appreciation [25]. However, social appreciation and specific types of feedback foster participation [23, 42].

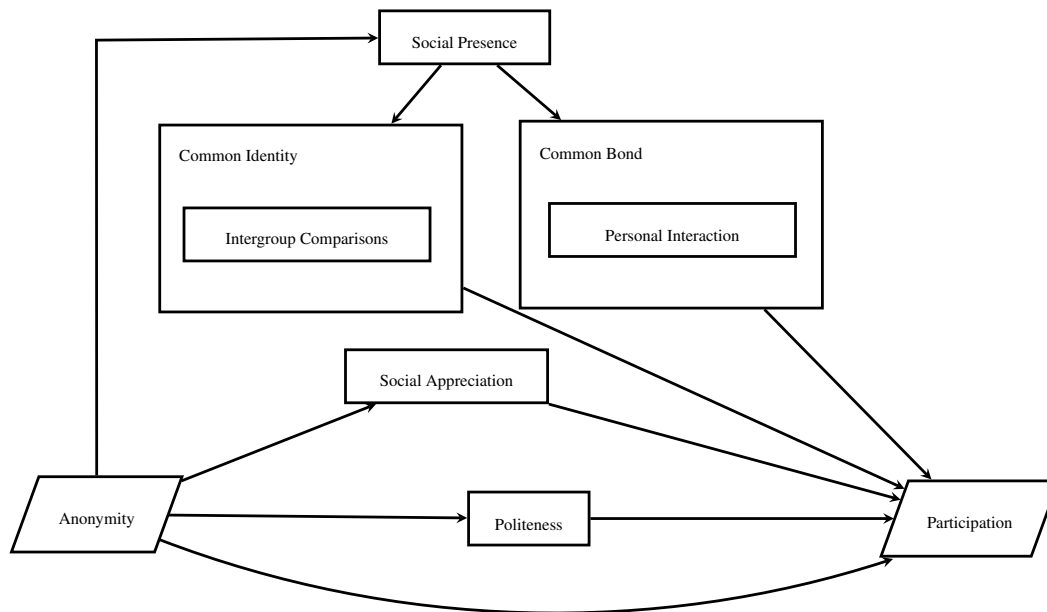


Figure 1: Simplified model showing anonymity interacting with participation

To synthesize the links that we identified in the literature between these factors, we built a model (see fig. 1). It serves as an anchor in our approach to quantifying the effect of anonymity on interaction in an online community. If the model's factors are valid in influencing participation and anonymity truly influences these factors, we can use that model—or, rather, measure markers of the model's factors—to search for differences in anonymous versus non-anonymous contributions. If these differences exist, we can assume that there will be a difference if we allow anonymity in the participation on the TOPIC platform, or in general. If there are no differences, then anonymity should not significantly influence participation in our community.

3.4 Using the Model

One place that appeared to be fit for testing the model was YouTube. This is a place where users interact that hosts videos related to the topic of informal caregiving. It also provides a chance to study a change in anonymity because YouTube moved from an internal comment system with pseudonyms to Google+, where users are required to give their full civil names. To capture the effect of moving away a degree of anonymity, we collected comments from before and after the change from a number of videos related to informal caregiving and Alzheimer's disease (because most of the future users in the city where we are conducting our fieldwork are informal caregivers of people suffering from this disease). We then searched for signs of changes in the model's factors.

As a result, we found a degree of evidence for a link between anonymity and politeness, although it was a positive one: anonymous comments were more often polite. We also observed a strong link between anonymity and less personal interaction but no significant influence of anonymity on the use of intergroup comparisons.

The YouTube study did not answer the research question adequately. Its conclusion of an increase in rude comments after moving away from anonymity collides with the literature, which led us to ex-

pect a positive influence on politeness. The study was of a move from one environment to another, which made it difficult to pinpoint anonymity as the deciding factor in the change. To minimize such other influences, it seemed to be a good idea to repeat a similar study in an environment that by default mixes anonymous and non-anonymous user participation, such as Quora. On Quora, anonymous responses can be given by regular users, who must still be logged in, by toggling the anonymous mode before sending.

Thus, the study on Quora that we describe in the next section had multiple goals:

1. It was another opportunity to apply the model and determine whether the factors' markers are detectable and whether it is possible to identify any differences.
2. It was a control for the results of the YouTube study. The generated hypotheses would be stronger if they applied here, and they would be weakened if they did not hold up. Because the final experiment required hypotheses to be possibly disproved, this study helped in choosing them.
3. Because the situation on Quora is different because it already allows anonymity rather than changing its anonymity conditions, it provided us with the chance to study the effect of anonymity without having to fear that other changes would influence the measured factors

In a direct comparison with the YouTube study, we were able to see whether the change in the level of politeness would be the same as it was on YouTube, where fewer pseudonymous comments were rude, contradicting the literature. The number of intergroup comparisons between the groups (anonymous and not) could have been the same, as on YouTube, or it could have differed this time. The number of personal interactions with anonymous authors—replies to answers or comments—could have been lower as well.

However, there was also the chance to generate new hypotheses given that we were able to include social appreciation through the number of upvotes (which was not possible on YouTube because there were no parsable upvotes). An upvote on Quora usually means that the person who is upvoting an answer believes that the question was properly answered and that the answer contributed in a meaningful way to Quora’s repository of knowledge. However, many user base their upvotes on their own interpretations of what an upvote should represent. Quora’s administration provided guidelines for upvoting, but it appears that many users have their own personal upvoting philosophies³.

We present the findings in the second next section, but social appreciation proved to be an important metric, which hints at a measurable influence of the reduced number of social cues provided by anonymity.

4. DATA COLLECTION AND ANALYSIS

We manually gathered 3626 answers from 337 questions on Quora, of which 293 answers were anonymous. We selected them by taking the then popular questions in the health category (fig. 2). Only questions without answers or those that merged with multiple other question threads were skipped. The obtained HTML was then parsed, and the generated data were saved in a database.

Depression: What is the diagnosis process like for someone to be diagnosed with depression?

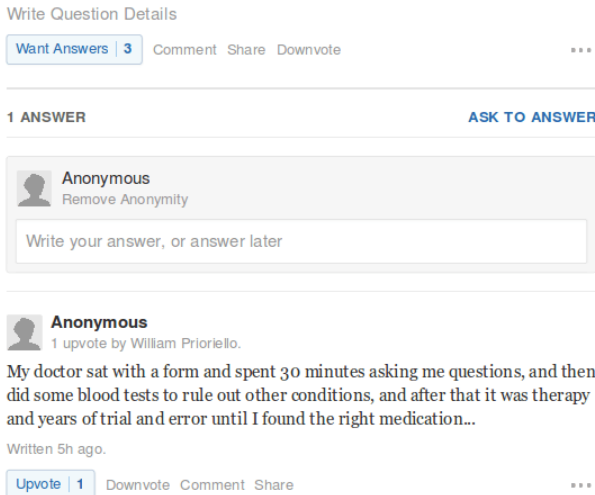


Figure 2: Example of an anonymous question and answer on Quora, <http://goo.gl/md4WJ3>

The data were then analyzed with a number of scripts, in particular, calling a Bayes classifier⁴ and a statistic toolkit⁵.

We searched for a number of factor markers from our participation model:

1. *Politeness*. How polite the message was. We used a Bayes classifier to attempt to categorize the answers into the cat-

³<http://goo.gl/qf6hBj>

⁴<https://github.com/jekyll/classifier-reborn>

⁵<https://github.com/clbustos/statsample>

egories *polite*, *neutral* and *rude*. This was based on reports that algorithmic approaches can work acceptably well for detecting politeness [39] and our own good experience with the method in the YouTube study.

2. *Intergroup Comparisons*. We searched for the words "we/us/our/them", which show that a group of people is being mentioned [4, p. 86]. In the model, anonymity influences the use of intergroup comparisons through *social awareness*.
3. *Personal Interaction*. To approximate personal interaction, we used the number of comments on an answer. In the model, anonymity influences this through *social awareness*.
4. *Social Appreciation*. The number of upvotes reflected this.

The scripts used and the generated database are available at <https://db.tt/EeG0s9sx>.

5. FINDINGS

We analyzed the answers, which means that we searched for significant differences in the selected markers between anonymously and non-anonymously posted answers. The result mainly showed that the two groups did not differ greatly, with one noteworthy exception described in section Social Appreciation.

5.1 Politeness

It was not possible for us to algorithmically analyze the answers for politeness as we had done previously. The algorithm failed to distinguish among the three categories, categorizing nearly all answers as either all rude or all polite while almost ignoring the much more fitting neutral category. This was a surprise given that the same software and workflow were used in the earlier YouTube study, in which we found 80% accuracy.

A manual examination showed that except for literally one answer, all of them followed a specific tone that appeared to be common on Quora. That is not to say that all answers were equal; there was a great range of quality and length in the sample data. Many responses were factual, and others were filled with pathos, but they all lacked easily distinguishable indicators of politeness. These were present in the comments on YouTube, where it was easy to categorize an insult as rude and many best wishes as polite. We assume that this impression is correct and that there was no difference between the groups. We discuss this further in section Limitations.

5.2 Intergroup Comparisons

There was only a small difference in the number of intergroup comparisons made by anonymous and non-anonymous users (fig. 3, table 1).

The difference was not significant with a t-test, which suggests the conclusion that the preference to post anonymously on an otherwise non-anonymous platform does not influence identification with the group, at least on Quora.

Table 1: Amount of comparisons

Group	mean	sd	median	n
Known	0.5426	1.8984	0	3332
Anonymous	0.4573	1.0114	0	293




Figure 3: Difference comparisons made

5.3 Personal Interaction

A t-test showed no significant difference in the number of comments received for the two groups (fig. 4, table 2). In general, comments to answers are not overly common on the platform. Quora uses UI elements to not highlight them: they are not visible by default, and they must be made visible by clicking on a small grey-colored link. As such, an average of roughly one comment for every second answer was already unexpectedly high.

We should note here that the amount of personal interaction that occurred through direct messages was invisible to us (see also section Limitations).




Figure 4: Average amount of received comments

The insignificance of the difference between the groups was also unexpected. We observed a high increase in personal interaction in the YouTube study after the change to less anonymous comments. That there was no difference here suggests another explanation: that the change on YouTube was not caused by the change in anonymity but by the change in the comment UI and the link with the social network Google+.

Table 2: Received comments

Group	mean	sd	median	n
Known	0.6747	5.0662	0	3332
Anonymous	0.4778	1.8048	0	293

5.4 Social Appreciation

The number of received upvotes did not differ significantly between anonymous and non-anonymous questions. However, anonymous answers also received less feedback (fig. 5, table 3).




Figure 5: Average upvotes received

There was sufficient literature with different results and theories to expect a stronger difference. In particular, [25] led us to expect that anonymous answers would be less appreciated. In that study, an experiment measured the assigned persuasiveness of responses that were either linked to a participant or were anonymous. These anonymous responses were viewed as being inferior to the non-anonymous ones, as less trustworthy and less persuasive. We expected that the same would happen here, i.e., that the anonymous responses would receive fewer upvotes.

Table 3: Received upvotes

Group	mean	sd	median	n
Known	9.1267	44.8019	1	3332
Anonymous	7.4539	59.6217	1	293

However, another factor that was measured was the length of the answer, which by itself was nearly significant in a t-test, with $p = 0.08$ (table 4). Using a Pearson correlation, we found a positive correlation $r = 0.383$ with 291 degrees of freedom ($p < 0.01$) between the length of the answer and the number of upvotes, but only for anonymous answers (fig. 6).

Table 4: Comment length

Group	mean	sd	median	n
Known	609.9988	938.1715	336.5	3332
Anonymous	741.8771	1273.7350	346.0	293

This means that for anonymous answers only, the number of upvotes increased with the length of the answer. This is surprising given that that correlation did not exist for the other answers.

It seems plausible to expect that in general, longer answers will receive more upvotes on Quora. They take longer to write, they can contain more relevant information, and they show that an effort was made. That there was no correlation between an answer's length and its number of upvotes shows that this is not the case. Perhaps longer answers are more cumbersome to read, or perhaps short answers better convey the needed information to answer a typical Quora question.

However, as soon as answers were anonymous, the correlation became significant: Longer answers by anonymous users received significantly more upvotes. Why is that?

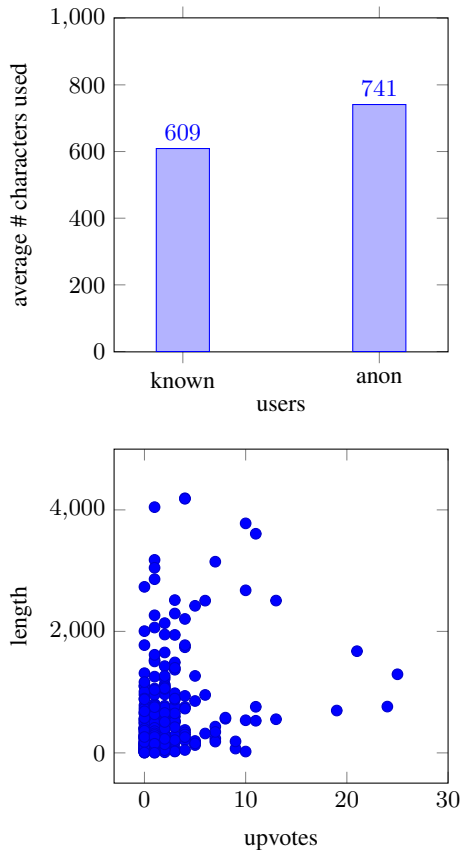


Figure 6: Correlations of answers length and upvotes for anonymous users

A possible explanation is that without the added social cues provided by username and avatar image, readers resorted to comment length as a signal of comment quality.

We base that hypothetical explanation on the assumptions of the reduced social cues approach, as exemplified in [40]. A social signal used in non-digital communication can in that model be replaced by a digital signal. Here, it could be the social dimension provided by username, attributability and avatar image that normally influences upvote behavior more than the answer’s length. When those signals are missing, a number of things could happen. It could be that the answer’s quality becomes more relevant and that quality is linked to length, or it could be that comment length works as a social signal for an answer’s credibility, a function that was previously covered by the now hidden peripheral elements.

Other explanations are possible. One could transform the removal of social cues into a positive and argue that, as in [9], removing the avatar image leads people to identify more strongly with those than with anonymous members. Comment length would again be a more objective factor of comment quality, a factor that was previously covered by negative identification through identity elements such as username and password.

We will not uncover the cause here; the finding is only that the correlation exists.

6. LIMITATIONS

Quora did not give us raw sample data, and we did not have access to an API. We collected our data manually and then parsed it with a handwritten parser. Thus, the answers that we collected had already been filtered by Quora’s moderation, with potential consequences for the validity of the question selection. If, for example, anonymous answers were in general less polite, that effect could have been invisible to us if the rude comments had already been deleted.

Moreover, all of the information that is not available to the public, such as direct messages between members, was also invisible to us. Having this as metadata would have allowed us to measure personal interaction more accurately, and without those data, one should judge the personal interaction metric as an estimate.

Answers that are submitted anonymously are anonymous only to other users. Quora itself could know who made which answer because users must always be logged in to answer. The answers are, as such, not fully anonymous in the strictest sense because that would include anonymity to all possible observers.

7. CONCLUSION AND FURTHER WORK

Quora provided us with the option to study a mixed environment of anonymous and non-anonymous user-generated texts. In contrast to the YouTube study, there were fewer differences. On YouTube, the prominent change was in the level of politeness and the increase in social interaction. On Quora, there was only the difference in the correlation between answer length and upvotes for anonymous answers.

Previous work showed the following:

We know that people who have limited motivation to process content are more likely to base evaluations on peripheral cues [40, p. 33]

Something similar could have happened here: Quorans could have normally based their upvotes at least partly on the peripheral social cues provided by username and avatar and resorted to answer length as a relevant factor only when those social cues were not present.

Regarding the effect of anonymity on an online community, we understand the result as an argument for the harmlessness of anonymity. Anonymous answers were, in the eyes of the community in general, not worse, and they did not receive significantly fewer upvotes. They were not shorter but were even slightly longer, which could be a target for community builders. Moreover, in contrast to the expectations generated by [14] and deindividuation theory in general, they were not less polite.

Future work should attempt to test these results in different communities, for example, those with less strict moderation and politeness expectations. Given that SIDE theory predicts an increase in group identity through anonymity, the effect of anonymity could be quite different in communities with different norms. Whether such a group identity for Quora users exists could also be a worthwhile question, and its answer could be useful for interpreting our results.

8. ACKNOWLEDGEMENTS

This work has been supported by European Union, ANR and national solidarity fund for autonomy through AAL program (project AALI 2012-TOPIC).

We thank our colleagues Nadia Gauducheau and Michel Marcoccia for their help in building the model and identifying the markers.

9. REFERENCES

- [1] J. Ahn, C. Weng, and B. S. Butler. The dynamics of open, peer-to-peer learning: What factors influence participation in the p2p university? In *System Sciences (HICSS), 2013 46th Hawaii International Conference on*, pages 3098–3107, Jan 2013.
- [2] J. Arguello, B. S. Butler, E. Joyce, R. Kraut, K. S. Ling, C. Rosé, and X. Wang. Talk to me: Foundations for successful individual-group interactions in online communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06*, pages 959–968, New York, NY, USA, 2006. ACM.
- [3] R. Bodle. The ethics of online anonymity or zuckerberg vs. "moot". *SIGCAS Comput. Soc.*, 43(1):22–35, May 2013.
- [4] N. Bramley. *Pronouns of Politics: The Use of Pronouns in the Construction of 'self' and 'other' in Political Interviews*. Australian National University, 2001.
- [5] M. Burke, E. Joyce, T. Kim, V. Anand, and R. Kraut. *Introductions and Requests: Rhetorical Strategies That Elicit Response in Online Communities*. 2007.
- [6] M. Burke and R. Kraut. Mind your Ps and Qs: the impact of politeness and rudeness in online communities. In *Conference on Computer Supported Cooperative Work*, pages 281–284, 2008.
- [7] V. H.-H. Chen, H. B.-L. Duh, and C. W. Ng. Players who play to make others cry: The influence of anonymity and immersion. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology, ACE '09*, pages 341–344, New York, NY, USA, 2009. ACM.
- [8] A. Chmiel, P. Sobkowicz, J. Sienkiewicz, G. Paltoglou, K. Buckley, M. Thelwall, and J. A. Holyst. Negative emotions boost user activity at {BBC} forum. *Physica A: Statistical Mechanics and its Applications*, 390(16):2936–2944, 2011.
- [9] U. Cress. Why member portraits can undermine participation. In *Proceedings of the 2005 Conference on Computer Support for Collaborative Learning: Learning 2005: The Next 10 Years!*, CSCL '05, pages 86–90. International Society of the Learning Sciences, 2005.
- [10] M. Edman and B. Yener. On anonymity in an electronic society: A survey of anonymous communication systems. *ACM Comput. Surv.*, 42(1):5:1–5:35, Dec. 2009.
- [11] R. Farzan, L. A. Dabbish, R. E. Kraut, and T. Postmes. Increasing commitment to online communities by designing for social presence. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, CSCW '11*, pages 321–330, New York, NY, USA, 2011. ACM.
- [12] F. M. Harper, D. Frankowski, S. Drenner, Y. Ren, S. Kiesler, L. Terveen, R. Kraut, and J. Riedl. Talk amongst yourselves: Inviting users to participate in online conversations.
- [13] R. Kang, S. Brown, and S. Kiesler. Why do people seek anonymity on the internet?: Informing policy and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, pages 2657–2666, New York, NY, USA, 2013. ACM.
- [14] P. G. Kilner and C. M. Hoadley. Anonymity options and professional participation in an online community of practice. In *Proceedings of the 2005 Conference on Computer Support for Collaborative Learning: Learning 2005: The Next 10 Years!*, CSCL '05, pages 272–280. International Society of the Learning Sciences, 2005.
- [15] R. E. Kraut and A. T. Fiore. The role of founders in building online groups. In *CSCW*, pages 722–732, 2014.
- [16] N. Lapidot-Leffer and A. Barak. Effects of anonymity, invisibility, and lack of eye-contact on toxic online disinhibition. *Computers in Human Behavior*, 28(2):434 – 443, 2012.
- [17] S. Levmore, S. Levmore, and M. Nussbaum. *THE OFFENSIVE INTERNET*. Harvard University Press, 2010.
- [18] K. Ling, G. Beenen, P. Ludford, X. Wang, K. Chang, X. Li, D. Cosley, D. Frankowski, L. Terveen, A. M. Rashid, P. Resnick, and R. Kraut. Using social psychology to motivate contributions to online communities. *Journal of Computer-Mediated Communication*, 10(4):00–00, 2005.
- [19] B. Nonnecke and J. Preece. Why lurkers lurk abstract.
- [20] E. Omernick and S. O. Sood. The impact of anonymity in online communities. In *International Conference on Social Computing, SocialCom 2013, SocialCom/PASSAT/BigData/EconCom/BioMedCom 2013, Washington, DC, USA, 8-14 September, 2013*, pages 526–535, 2013.
- [21] Pew Research Center. The Web at 25 in the U.S., February 2014.
- [22] T. Postmes, R. Spears, and M. Lea. Intergroup differentiation in computer-mediated communication: Effects of depersonalization. *Group Dynamics: Theory, Research, and Practice*, 6(1):3, 2002.
- [23] J. Preece and B. Shneiderman. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on Human-Computer Interaction*, 1(1):13–32, 2009.
- [24] J. J. Preece, B. Nonnecke, and D. Andrews. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior*, 20(2):201–223, 2004.
- [25] S. A. Rains. The impact of anonymity on perceptions of source credibility and influence in computer-mediated group communication: A test of two competing hypotheses. *Communication Research*, 34(1):100–125, 2007.
- [26] Y. Ren, R. Kraut, and S. Kiesler. Applying common identity and bond theory to design of online communities. *Organization studies*, 28(3):377–408, 2007.
- [27] A. Santana. Civility, anonymity and the breakdown of a new public sphere. 2012.
- [28] N. Sharma, B. S. Butler, J. Irwin, and H. Spallek. Emphasizing social features in information portals: Effects on new member engagement. *Journal of the American Society for Information Science and Technology*, 62(11):2106–2120, 2011.
- [29] Y.-C. Shiu, C.-M. Chiu, and C.-C. Chang. Exploring and mitigating social loafing in online communities. *Comput. Hum. Behav.*, 26(4):768–777, July 2010.
- [30] L. Sproull and S. Kiesler. Reducing social context cues: Electronic mail in organizational communication.

- Management science*, 32(11):1492–1512, 1986.
- [31] J. Suler. The online disinhibition effect. *Cyberpsychology & behavior*, 7(3):321–326, 2004.
- [32] Y. R. Tausczik, L. A. Dabbish, and R. E. Kraut. Building loyalty to online communities through bond and identity-based attachment to sub-groups. In *CSCW*, pages 146–157, 2014.
- [33] P. A. Thoits. Social support as coping assistance. *Journal of consulting and clinical psychology*, 54(4):416–423, Aug. 1986.
- [34] M. Tixier, G. Gaglio, and M. Lewkowicz. Translating social support practices into online services for family caregivers. In *Proceedings of the ACM 2009 International Conference on Supporting Group Work*, GROUP '09, pages 71–80, New York, NY, USA, 2009. ACM.
- [35] M. Tsikerdekis. The effects of perceived anonymity and anonymity states on conformity and groupthink in online communities: A wikipedia study. *Journal of the American Society for Information Science and Technology*, 64(5):1001–1015, 2013.
- [36] C.-H. Tu. The relationship between social presence and online privacy. *The Internet and Higher Education*, 5(4):293–318, 2002.
- [37] H. van Oostendorp and F. van Varik. Stimulation of activity in online communities. In A. Dittmar and P. Forbrig, editors, *ECCE*, pages 163–170. Universitätsdruckerei Rostock, Germany / ACM, 2011.
- [38] B. Wellman and M. Gulia. Net surfers don't ride alone: Virtual communities as communities. In M. A. Smith and P. Kollock, editors, *Communities in Cyberspace*, pages 167–194. Routledge, 1999.
- [39] F. Wild and C. Stahl. Using latent semantic analysis to assess social competence. In Wild, Kalz, van Bruggen, , and Koper, editors, *Mini-Proceedings of the 1st European Workshop on Latent Semantic Analysis in Technology-Enhanced Learning*, pages 3–4. Open University Netherlands, March 2007.
- [40] K. Wise, B. Hamman, and K. Thorson. Moderation, response rate, and message interactivity: Features of online communities and their effects on intent to participate. *Journal of Computer-Mediated Communication*, 12(1):24–41, 2006.
- [41] N. Yamashita, H. Kuzuoka, K. Hirata, and T. Kudo. Understanding the conflicting demands of family caregivers caring for depressed family members. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2637–2646, New York, NY, USA, 2013. ACM.
- [42] H. Zhu, R. Kraut, and A. Kittur. Effectiveness of shared leadership in online communities. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, CSCW '12, pages 407–416, New York, NY, USA, 2012. ACM.
- [43] H. Zhu, A. Zhang, J. He, R. E. Kraut, and A. Kittur. Effects of peer feedback on contribution: A field experiment in wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 2253–2262, New York, NY, USA, 2013. ACM.