The Impact of Anonymity in Online Communities

Eli Omernick
Computer Science Dept., Pomona College
185 East Sixth Street
Claremont, CA 91711
elijah.omernick@pomona.edu

Sara Owsley Sood
Computer Science Dept., Pomona College
185 East Sixth Street
Claremont, CA 91711
sara@cs.pomona.edu

ABSTRACT
The scale of participation on social news sites has challenged community managers, leaving them unable to detect and remove all inappropriate content by hand. Automated insult and profanity detection systems have helped, but have failed to address the problem of why this content is contributed in the first place. That is, what implications do interface design choices have on the content being generated? One such design choice is whether or not a site allows anonymous comments. What impact does allowing anonymity have on the quality or quantity of participation on a site? This case study analyses the impact of anonymity on a technology social news site, TechCrunch.com. TechCrunch is ideal for this study in that it underwent a shift from allowing anonymous comments (using the Disqus commenting platform) to disallowing them (using the Facebook commenting platform) in March of 2011. We compare the quality of *anonymous* and *real identity* comments through measures of reading level, relevance to the target article, negativity and presence of swear words and anger words. We couple this qualitative analysis with a quantitative analysis of the change in participation to give a complete picture of the impact of anonymity in this online community, with the end goal of informing design on similar social news sites.

Author Keywords
Online communities; comment threads; user-generated content; anonymity; community management; negativity; profanity.

ACM Classification Keywords
H.5.3. Information interfaces and presentation: Group and Organization Interfaces.

General Terms
Design; Experimentation; Human Factors.

INTRODUCTION
Social news sites have become a great way to encourage participation on the web. However, commenting on a social news site differs greatly from a face-to-face conversation with a friend. Imagine walking into a party with 2 Billion people (the 2010 estimated number of internet users (Miniwatts Marketing Group 2011)). Some might find this intimidating. Some might feel discomfort or awe from the sheer magnitude of the group. Some may have some sensation of empowerment from the lack of recognition from within the group (feeling anonymous); there is no apparent social hierarchy to submit to.

Several psychological models explain the effects of anonymity, namely "deindividuation," which originated in the famous works of Stanley Milgram and Philip Zimbardo in the 1960's and 70's, and "communication bandwidth," which is associated with these sort of behavioral changes as early as the 1890's with the technological advancement of telegraphs (Watt et al. 2002). Deindividuation is characterized by a reduction in private self-awareness and accountability, resulting in lower self-regulation and concern for the reactions of others and is brought about by an individual not being identifiable or distinguishable in a group (Johnson, Cooper, and Chin 2009). It's clear that the Internet is a place where such social conditions are present, and there even seems to be some disassociation of self-awareness just from working on computers, even without interactions with others. Survey takers who responded via a computer terminal reported less social anxiety, more self-esteem, and fewer adherences to social norms than their counterparts who took the same survey with pen and paper (Joinson 1999).

Given these effects of anonymity, it is clear why many Internet users prefer to remain anonymous, or even won't participate on sites where they are required to share their real identity. Many social news sites offer users the ability to participate under the cloak of anonymity. However, a competing concern is that social news sites are often plagued with negative content of malicious intent and cyberbullying (Boyd and Marwick 2011; Dinakar, Reichart, and Lieberman 2011; Li 2006). Studies have shown that large portions of social news site comments contain profanity and/or insults and personal attacks. The scale of these sites (the sheer volume of comments) makes them hard, if not impossible for community managers to moderate. The recent development of automated moderation tools - systems that detect insults and profanity automatically (Boyd and Marwick 2011; Dinakar, Reichart, and Lieberman 2011; Li 2006; Sood, Churchill, and Antin 2011; Sood, Antin, and Churchill 2012; Yin et al. 2009) - attempt to aide community managers in this daunting task. Additionally, social moderation efforts allow users to 'flag

978-0-7695-5137-1/13 $26.00 © 2013 IEEE
DOI 10.1109/SocialCom.2013.80

as inappropriate' and even attempt to resist collusion (Lou, Chen, and Lei 2009).

While monitoring and removing inappropriate content is necessary, one should also question the environmental characteristics the lead to the generation of this content. That is, what implications do interface design choices have on the content being generated? One such design choice is whether or not a site allows anonymous comments. Are anonymous users more likely to be malicious? What impact does allowing anonymity have on the quality or quantity of participation on a site? These are questions that cannot only be addressed empirically, but one must also consider the political implications of these design decision – the basis of the "nymwars" debate from participants in response to Google's 2011 decision to require real names on Google+ (Boyd 2012).

This case study addresses these questions through analysis of a technology social news site, TechCrunch.com. TechCrunch is ideal for this study in that it underwent a shift from allowing anonymous comments (using the Disqus commenting platform) to disallowing them (using the Facebook commenting platform) in March of 2011. By creating a large data set of comments (on TechCrunch news articles) from before and after the change, we compare the quality of *anonymous* and *real identity* comments through measures of reading level, relevance to the target article, negativity and presence of swear words and anger words. We couple this qualitative analysis with a quantitative analysis of the change in participation, along with a discussion of limitations, to give a complete picture of the impact of anonymity in this online community, with the end goal of informing design on similar social news sites.

RELATED WORKS

The effects of deindividuation have been studied outside of the contexts of computer-mediated communication. In an interesting modification of the Ultimatum game, players were found to be significantly more generous when the names of the players were revealed, indicating a pro-social behavioral shift from a less anonymous condition (Charness and Gneezy 2008). A large number of studies have also examined self-disclosure (spontaneously revealing personal details or experiences) as a result of deindividuation from anonymity in online contexts (Qian and Scott 2007).

Hartnett and Seligsohn observed disinhibition on survey responses positively correlated with increasingly anonymous conditions (Hartnett and Seligsohn 1967). This deindividuation effect was also observed in participants taking computer-moderated surveys over pencil and paper, demonstrating that the computer-mediated communication is also a factor in causing deindividuation, not just anonymity (Chester and Gwynne 1998).

According to Wallace, subjects participating in computer-mediated communication are less likely to conform to confederates responding incorrectly to obvious questions than when put in the same situation in a FTF context (Wallace 2001, 60–61). She goes on to observe that when in more social scenarios, while individuals are more likely to deviate from normative behavior in computer-mediated communication, they are less likely to influence the majority they are deviating from (Wallace 2001, 83).

The disposition of commenters posting under varying anonymity conditions has also been studied. Kilner & Hoadley observed that the prohibition of anonymous posts on a military professional education site almost completely eradicated negative comments, and that the subsequent prohibition of pseudonym use had no noticeably effect (Kilner and Hoadley 2005). In South Korea the *Real Name Verification Law* was introduced in 2005, requiring users to have their identities verified before posting on popular sites. It was concluded that the implementation of this law lead to decreases in uninhibited behavior in a site that did not have similar measures in place previously while having no effect on a site that did; it was also shown that the participation on the previously anonymous site dropped somewhat immediately after the switch, but then returned to normal levels, again with the other site being unaffected (Cho, Kim, and Acquisti 2012).

COMMENTING PLATFORMS FOR SOCIAL NEWS SITES

With the web2.0 rise in blogs and forums, tools to support and enhance such sites have been developed. In particular, many tools have arisen to support commenting on social news sites. For example, Disqus and Facebook have released commenting platforms that integrate smoothly into these news sites, and manage contributed comments. Though they support the same task, the Disqus and Facebook commenting platforms have one fundamental difference. Disqus allows users to identify themselves with pseudonyms, or to contribute anonymously, while Facebook requires users to give a 'real identity' by linking their comment to a Facebook account (or to a verified Hotmail, AOL, or Yahoo! email address).

When a user makes a comment on a site using the Disqus commenting platform, they are associated with one of the following three categories: *Anonymous*, *Pseudonym*, and *Identified*. *Anonymous* users do not have any personal or account information associated with the comments they make. *Pseudonym* users identify themselves with a 'Disqus account' - an account, with a visible username, created for making comments on the Disqus platform. *Identified* users are those who post comments under their real name by linking their comments to their Facebook account (via the Disqus platform).

Similarly, when a user makes a comment on a site using the Facebook commenting platform, they are associated with one of the following two categories: *Account Identified*, and *Email Verified*. *Account Identified* users have a Facebook account associated with their post, and are identified by name. *Email Verified* users provide a valid Yahoo, AOL or Hotmail email address to associate with their comments. Their posts are identified with whatever alias is associated with the account used.

As the result of an internal study Disqus released a promotional piece arguing that Pseudonym users were the driving force behind their virtual communities (danielha 2012). In it they state that the average *pseudonym* user contributes 6.5 times the amount of content as the average *anonymous* user and 4.7 times as much as the average *Identified* user. They also claim that 61% of *Pseudonym* posts are positive, vs. 51% for *Identified* posts and 34% for *Anonymous* posts.

DATASET

Our dataset consists of all TechCrunch articles posted from Jan 1, 2010 through May 30, 2012. On March 1, 2011 TechCrunch switched from using the Disqus commenting platform to the Facebook commenting platform. An approximately equal window of time was sampled before and after that date, generating two datasets: the Facebook dataset (collected from TechCrunch after the switch to the Facebook commenting platform), and the Disqus dataset (collected from TechCrunch before the switch, while comments were managed using the Disqus platform).

In the *Facebook* dataset, a total of 14,895 TechCrunch articles were collected, yielding a total of 295,495 comments (286,742 made from Facebook *Account Identified* users (97%), and 8,753 from *Email Verified* users (3%)), from 91,303 unique Facebook *Account Identified* users. *Email Verified* users were not uniquely identified in the Facebook API responses used to generate our dataset. Thus, no approximation can be made as to how many unique *Email Verified* users are represented.

In the *Disqus* dataset, a total of 7,344 articles were collected, yielding 194,835 comments, 104,161 from *Anonymous* users (53%), 77,058 by *Pseudonym* users (40%), and 13,616 from *Identified* users (7%). The dataset includes the following numbers of unique users: 23,077 unique *Pseudonym*s, 45,823 unique *Anonymous* identifiers, and 6116 *Identified* Facebook users. Unique users were determined by the name listed in association with the comment in the case of *Pseudonym* and *Anonymous* users, and the url of the associated Facebook account for *Identified* users.

A limitation with this mechanism of identifying users is that multiple users could post under the same alias. It is unlikely, for instance, that all posts made by “John” originate from the same user. It is also possible that any one user may have posted under multiple aliases, but at an intuitive level this seems somewhat less likely. We therefore consider this user count to be a lower bound on the actual number of users represented, but even with this bound we observe meaningful trends in the data.

IMPACT OF ANONYMITY ON COMMENT QUALITY

How do real identity, pseudonym and anonymous comments compare? The debate over anonymity and its place/role in online communities centers largely on a notion of observable behavioral differences between groups of anonymous and identified participants. To that end we have considered a variety of dimensions for comparing the quality of comments made including relevance, readability and vocabulary.

Relevance

Relevance is meant to be a measure of how *on topic* comments are. Relevance is measured in relation to the article a comment was made in response to. A list of relevant terms was generated for each article, comprised of title words, keywords, and tags. Keywords were defined as words in the main text of the article that linked to other online resources. Some articles also provided tags, a short list of terms, which applied to the content of the article. The relevance of a comment was then calculated as the proportion of words in the comment (after stop words were removed), which were present in the relevant terms list.

User Group	Mean	StDev	Median
Disqus-*Anonymous*	0.106	0.076	0.098
Disqus-*Pseudonym*	0.110	0.085	0.098
Disqus-*Identified*	0.112	0.112	0.091
Facebook-*Email Verified*	0.085	0.117	0.043
Facebook-*Account Identified*	0.101	0.078	0.090

Table 1: Relevance Distributions

In Table 1, we see that in both the Facebook and Disqus datasets, comments from the *Identified* (Disqus *Identified* and Facebook *Account Identified*) users were more relevant than comments from their more anonymous counterparts (Disqus *Anonymous* or Facebook *Email-Verified*). This trend (based on relevance means) is more significant in the Facebook dataset, and it is also interesting to note that the median relevance for the *Email-Verified* users is considerably lower than the other groups. Also, the Facebook dataset as a whole is somewhat less relevant than the Disqus dataset. The mean relevance of each account type (Table 1) was paired individually with each of the remaining account types and tested for significance using a t-test. The pair Disqus *Pseudonyms* and Disqus *Identified* was found to be significantly different with $p < 0.5$. The remaining pairs of means are significantly different from each other with $p < 0.01$. In conclusion, our analysis shows that, within both the Disqus and Facebook commenting platforms, more relevant comments are associated with more revealed identity.

User Group	Mean	StDev
Disqus ***Anonymous***	63.5	192.7
Disqus ***Pseudonym***	59.0	125.0
Disqus ***Identified***	63.1	65.5
Facebook ***Email Verified***	67.6	62.1
Facebook ***Account Identified***	69.3	73.0

Table 2: Flesch Reading Ease

Reading level

There are a variety of useful metrics for assessing the reading grade level or general readability of documents, usually based on some combination of word length (syllables or characters), sentence length, and frequency of complex words (words with three or more syllables). We evaluated the comments collected with several of these metrics, including *Flesch Reading Ease*, *Flesch-Kincaid Grade Level*, *RIX*, *Coleman-Liau Index*, *Gunning Fog Index*, *ARI*, *SMOG Index*, and *LIX*. All tests yielded the same pattern of results, so for the sake of brevity, we consider only Flesch Reading Ease in this discussion (Flesch 1948).

Flesch Reading Ease is a metric for estimating the difficulty of a text to be read based on the number of words per sentence and syllables per word. A higher Flesch score indicates easier readability (a lower reading grade level) (Flesch 1948).

In Table 2, we see that Flesch Reading Ease scores between groups within each dataset exhibit few trends, but the difference between the two corpora is notable. Within the Disqus dataset, the Disqus and *Anonymous* and *Identified* comments were both significantly 'easier' than the Disqus *Pseudonym* comments (by t-test with $p < 0.001$). Overall, comments in the Facebook dataset were significantly 'easier' than in the Disqus dataset (by t-test with $p < 0.001$).

Table 3 gives a summary of several common (and immediately intuitive) readability features. The only one that shows a significant observable difference between groups is *Words per Sentence*. The biggest difference here is between datasets again, with the Facebook dataset using an average of 2.8 fewer words per sentence. The other values are all similar across all divisions of the data, so it seems as if the nature of the words being employed is relatively constant, but that the length of sentences is a more volatile feature.

Word Usage

Linguistic Inquiry and Word Count (LIWC) is a system developed and refined over the last twenty years that categorizes words into functionally and psychologically significant categories. Based on a predefined dictionary, it returns the percentages of a document's words that are categorized into each group. It has been used to provide insight and results in a number of linguistic and psychological studies (Francis and Pennebaker 1993).

User Group	Chars Per Word	Syllables Per Word	Words Per Sentence	Complex Words (%)
Disqus All	4.60	1.44	18.25	11.54
Disqus ***Anonymous***	4.58	1.42	17.99	11.39
Disqus ***Pseudonym***	4.61	1.42	18.36	11.50
Disqus ***Identified***	4.65	1.43	19.65	11.57
Facebook All	4.52	1.40	15.45	11.14
Facebook ***Email Verified***	4.65	1.38	17.02	10.96
Facebook ***Account Identified***	4.51	1.40	15.40	11.12

Table 3: Readability Features

We utilized the LIWC2007 system to analyze the word usage by commenters in our two datasets. To eliminate bias (to the extent possible) between the topics of comments in our two datasets, topical keywords (keywords and tags from the focal article) were removed from comments before analysis by LIWC. The LIWC analysis of a document covers 80 output variables including:

"4 general descriptor categories (total word count, words per sentence, percentage of words captured by the dictionary, and percent of words longer than six letters), 22 standard linguistic dimensions (e.g., percentage of words in the text that are pronouns, articles, auxiliary verbs, etc.), 32 word categories tapping psychological constructs (e.g., affect, cognition, biological processes), 7 personal concern categories (e.g., work, home, leisure activities), 3 paralinguistic dimensions (assents, fillers, nonfluencies), and 12 punctuation categories (periods, commas, etc)" (LIWC)

While all categories of comparison are of interest to us in this case study, we found the *psychological construct* variables to be most interesting. In particular, because of the known impact on behavioral norms in online communities, we were most interested in the variables that analyzed the proportion of swear words, affect words, and anger words used.

Table 4 gives the LIWC reported proportion of words in each category per dataset. The data in Table 4 sends a consistent message about comment quality – groups with more identity associated with their comments user less swear words, less anger words, more affect words, more positive emotion words, and less negative emotion words. One-way anova tests reveal that nearly all of these differences between groups are significant. All comparisons between the Disqus and the Facebook datasets as a whole are significant ($p<0.05$ for 'swear words' and $p < 0.001$ for all others). Within the Disqus dataset, differences in proportion of 'swear words' and 'negative emotion words' is only significant between the *anonymous* and *pseudonym* users with $p<0.001$, there are no significant differences in the use of 'anger words', and use of 'affect words' and 'positive emotion words' differs significantly between all groups with $p<0.01$. Within the Facebook dataset, differences in use of 'affect words' is not significant, but all other comparisons are significant with $p < 0.01$.

User Group	Swear words	Anger words	Affect words	Pos emo words	Neg emo words
Disqus All	0.20	0.57	6.17	4.58	1.56
Disqus *Anonymous*	0.22	0.59	6.09	4.43	1.62
Disqus *Pseudonym*	0.17	0.55	6.31	4.76	1.51
Disqus *Identified*	0.21	0.58	6.99	5.43	1.51
Facebook All	0.17	0.48	6.84	5.43	1.38
Facebook *Email Verified*	0.22	0.60	6.71	5.20	1.49
Facebook *Account Identified*	0.15	0.44	6.88	5.51	1.34

Table 4: LIWC analysis on comments from user groups within the Facebook and Disqus datasets.

Another indication of comment quality is how other users actually responded to posted comment. Both Facebook and Disqus provide the ability to "like" comments. Disqus also provides users the opportunity to "dislike" comments. The dislike functionality appears to be seldom used, however, as only 1 comment in our dataset had a "dislike score" other than 0 (it was 1).

Within the Disqus data, the mean number of likes is smallest for *anonymous* users (0.928) and greatest for *pseudonym* users (1.363) (all differences between groups within the Disqus dataset are significant from a t-test with $p<0.01$). The mean number of likes between the Facebook dataset (1.074) and Disqus dataset (1.113) are relatively similar, but the standard deviations are much higher in the Facebook data (6.30 in Facebook data and 3.19 in Disqus data), indicating that there is much more variation in the distributions of the number of likes on comments.

IMPACT OF ANONYMITY ON COMMENT QUANTITY

One major criticism of prohibiting anonymous users in online communities is that it limits participation. Users who are unwilling or unable to identify themselves are excluded from the conversation, and it follows that the community is smaller because of it. It is important to attempt to quantify this effect, therefore, to gauge the weight of this claim. We consider the problem from two perspectives: on a per article basis and on a per user basis.

Per Article Analysis

From the Disqus dataset a total of 1,029 out of 7,344 articles went uncommented (14%). In the Facebook group 3,933 out of 18,757 articles were uncommented (21%). At first glance this seems quite significant, but we must make the observation that while the number of articles appearing on TechCrunch more than doubles from the Disqus to the Facebook dataset (again, datasets spanning two different periods of time), the number of observed users only increases from 75,016 to 91,303 (22% growth). The 91,303 (observed users) figure is conservative, as it includes none of the Email-Verified users who contributed during this time, but that group represented only 3% of the contributions of that dataset, so it is unlikely that that group contributes a grossly larger proportion of the numbers of contributors. We also observe that this number will be offset to some degree by the overlapping nature of users' identities (discussed above and in further detail in the 'Per User Analysis' section that follows).

If we are to assume the same ratio of users to commented articles in both groups the analysis changes somewhat. In the Disqus group the ratio of users to commented articles is 75,016 / 7,344 = 10.2 users per commented article, while in the Facebook group it is 91,303 / 14,824 = 6.16 users per commented article. This shows us that while the percentage of articles that went uncommented increased in the Facebook data, that users in the more recent body were making contributions on a larger spread of the articles released relative to their size as a group.

Table 5 shows a similar story. The Disqus data exhibits significantly more comments per article. The trend is even more apparent after removing articles that had no comments from the analysis, despite the Facebook data having the larger proportion of uncommented articles. Both trends are significant by t-test with $p < 0.001$. This analysis has the same issue of unequal relative sizes of the commenting body and comment space discussed above, however, so this higher density of content in the Disqus data is not surprising.

Dataset	Mean	StDev	Median	Mode
Disqus-All Articles	26.5	29.2	16	0
Facebook-All Articles	15.8	19.4	9	0
Disqus-Commented Articles	30.9	29.3	20	7
Facebook-Commented Articles	19.8	19.8	13	1

Table 5: Number of Comments Per Article

Dataset	Mean	StDev	Median	Mode
Disqus – All	24.9	21.7	17	7
Disqus – *Anonymous*	13.5	13.6	9	3
Disqus – *Pseudonym*	9.6	11.0	5	0
Disqus – *Identified*	1.8	2.4	1	0
Facebook - All	15.2	13.2	11	1

Table 6: Unique Commenters Per (Commented) Article

In Table 6 we again see that in the Disqus data the users are more densely concentrated in articles. Note that this is a per article measurement, and that any single user will be counted in every article they post in. If we divide the number of comments per commented article (Table 5) by the number of unique commenters (Table 6) we generate the number of comments per user on the basis of article. The Disqus group gives a value of 1.24 comments per user per article and the Facebook group gives a value of 1.30 comments per user per article. These values are not significantly different.

It is also interesting to note that in the Disqus group, the sum of average *Pseudonym* Users and average *Identified* Users per article is less than the average number of *Account Identified* Users per article in the Facebook data, and that in the average article in the Disqus group there are 40% more *Anonymous* commenters than *Pseudonym* users; in the promotional piece mentioned above, Disqus has claimed that *Pseudonym* users represent the core of their user contributions (danielha 2012).

We also need to consider these values normalized to the relative sizes of each group in the Disqus dataset. To do this we divide each of the means in Table 6 by the proportions of the entire user base represented by each category of user. This gives us values of 22.1 for *Anonymous*, 31.2 for *Pseudonyms*, and 21.8 for *Identified* users. These values indicate that *Pseudonym* users are more diversely represented in threads with respect to their size, and that there isn't a substantial difference between *Anonymous* and *Identified* users in this group (recall that this is only the Disqus data).

Due to the similarity of the comments per user per article values, and the discrepancy between the growth of the user base and comment space we cannot conclude that the participation of users, from a strictly quantitative perspective, differs significantly between the time before and time after the switch from Disqus to Facebook.

Another dimension of the issue is the length of the comment threads (number of comments posted in response to a given article). Longer comment threads may indicate a higher level of discourse on an issue, as it is more likely for large numbers of comments to be made in response to each other as the relative amount of content swings from favoring the original article to the responses, and readers tend to consider comments as equally valid sources of content as the original article (Reader 2012). Figures 1 and 2 summarize the trend of average proportion of Disqus *Anonymous* (or Facebook *Email-Verified*) users observed in a thread versus the length of that thread (number of comments made).

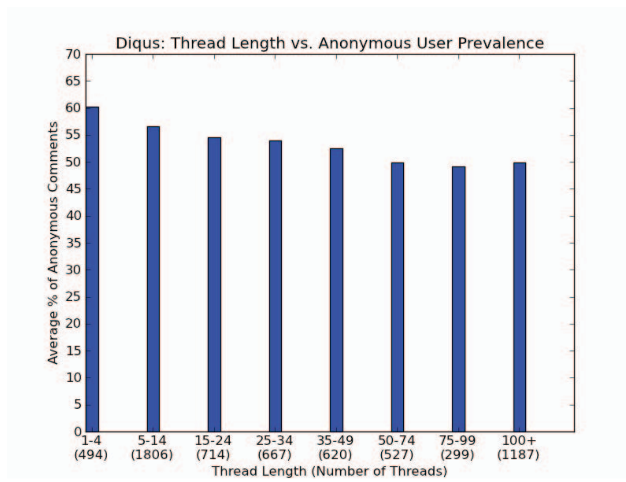

Figure 1: Disqus dataset: Thread Length vs. Anonymous User Prevalence

Note that these figures are made as a result of bucketing the threads into categories based on number of comments. To determine if there is a statistically significant correlation between thread length and proportion of anonymous commenters, we do a Pearson's Correlation on both of these distributions (without bucketing). For the Disqus data (Figure 1) – comparing the trends of thread length and proportion of Anonymous comments - $r = -.58$ with 98 degrees of freedom ($p < .01$). The Facebook data (Figure 2) yields $r = -.85$ with 128 degrees of freedom ($p < .01$).

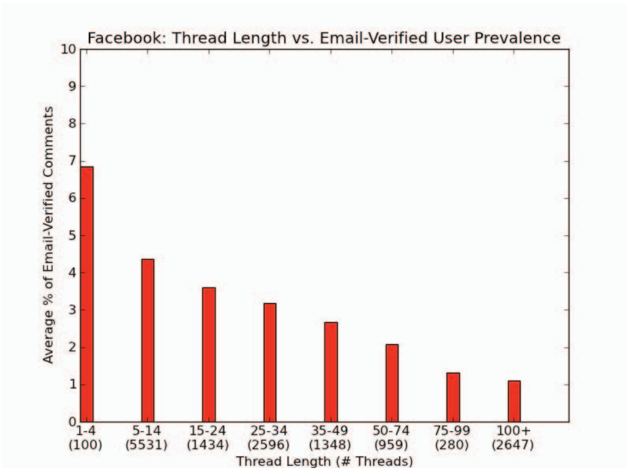

Figure 2: Facebook dataset: Thread Length vs. Email-Verified User Prevalence

It is quite clear that longer threads become increasingly less dominated by *Anonymous* users (in the case of Disqus), or *Email-Verified* users (in the case of Facebook). This suggests that these users are less likely to contribute to quantitatively fruitful discourse. In other words, more identity revealed is correlated with longer threads.

Per User Analysis

An intuitive mechanism for comparing user behavior is considering the trends in individual behavior over time. In this section we analyze the patterns of user activity over the history of our case study.

Group	Mean	StDev	Median	Mode
Disqus– *Anonymous*	2.3	18.5	1	1
Disqus– *Pseudonym*	4.3	15.4	1	1
Disqus– *Identified*	2.2	5.5	1	1
Facebook	3.1	14.5	1	1

Table 7: Comments Per User

Table 7 holds the average total number of comments posted by users over the course of the experimental period. One of the flaws of our classification becomes painfully obvious here. Since users are only differentiated by the name associated with the post it is possible for multiple actual users to be binned together if they share a screen-name. For instance, the top anonymous contributors include: 'Guest' (which is the default suggestion for anonymous posts), 'John', 'Mike', 'Tom', and 'Anon'. Presumptuous as it may be, we assert that the entirety of these users' contributions do not originate from the same real individuals. Conversely, the top several Pseudonym users are: 'MrGamma', '*** I own the © of myself ***', 'MG Siegler', 'Universal_Mind', 'daniel14214', 'Michael Hart', and 'BrianD'. Though it is still possible, it seems considerably less likely that these aliases map to as many individuals as those listed above. Facebook users were differentiated on the basis of their Facebook page's url, so there is no potential overlap there.

There is also the reverse issue of one anonymous individual being able to post under many different names. It seems somewhat less likely that this behavior would be more prevalent than that noted above, so we have chosen to consider our count of unique anonymous users to be a lower bound. That being said, the mean in Table 7 under *Disqus – Anonymous* is likely higher than the actual value.

One conclusion that can be made based on this analysis is that the average *pseudonym* user does seem to contribute more than the average *anonymous* or *identified* user in the Disqus group.

Another factor we considered was duration of activity, which is defined as the number of days between a given user's first and last post in the dataset as a whole. This data is summarized below in Table 8.

Group	Mean	StDev	Median	Mode
Disqus– *Anonymous*	21.4	75.3	0	0
Disqus– *Pseudonym*	28.2	65.7	0	0
Disqus- *Identified*	22.2	61.9	0	0
Facebook	39.1	90.2	0	0

Table 8: Duration of Activity

We observe that the issue of multiple individuals using the same alias will have the same effect on this piece of analysis; considering the *Anonymous* data point an upper bound again, however, only validates the result more in this case. *Pseudonym* users have a longer lifespan within the community than *Anonymous*, but the Facebook users after the switch show substantially more longevity. In the Facebook group, it follows that an individual is more likely to remain an active part of the community for longer. All pairs of means in Table 8 are significantly different (t-test with $p<0.001$) excluding the Disqus *Identified* and Disqus *Anonymous* pair.

It is curious to note that the Facebook users in the Disqus group (the *Identified* subset) do not exhibit the same trends in our quantitative analysis. We observe that those users tend to be more statistically similar to *anonymous* users. One possible explanation to this is that most users fall onto the extremes of either being consistently active members of the community or only posting once. This is evident from the repeated median/mode scores of 0 and 1 in the two above analysis. When making a comment on the Disqus platform the user was prompted to give identification. They had to choose from a list of options, so choosing the

Facebook option is equally as convenient as the Anonymous option.

User Group	Mean	StDev	Median	Mode
Disqus All	40.5	59.1	23	10
Disqus *Anonymous*	40.0	53.7	23	8
Disqus *Pseudonym*	40.7	63.5	24	10
Disqus *Identified*	42.3	57.2	24	8
Facebook All	35.0	48.0	20	7
Facebook *Email Verified*	44.2	57.6	28	10
Face *Account Identified*	34.7	47.6	20	7

Table 9: Comment Length (Words)

There are a few trends we would like to point out in Table 9, on average comment length per user group. We again see the discrepancy between the two corpora where the Facebook group provides shorter responses than the Disqus group. This is made curious, however, by the observation that Facebook *Identified* users within the Disqus group use slightly longer sentences on average than those from the Facebook platform, while there is no substantial difference between the *anonymous* and *pseudonym* users of Disqus. All pairwise comparisons of means in Table 9 are significant (by t-test where $p < 0.05$).

The factor that does seem to correlate with the length of responses is the relative plurality of the size of the poster's group. In the Facebook corpus, *Email Verified* posts make up only 3% of the total, with *Account Identified* Users contributing 97%. In the Disqus group, *Anonymous* comments are 53%, *Pseudonyms* 40%, and *Identified* 7%. In the Disqus groups there is a far less profound disparity between sizes. We see the smaller groups tending to be more verbose here and in Table 3, Words per Sentence. This deviation from the norm is a phenomenon predicted by the Social Identity Model of Deindividuation Effects (Christopherson 2007).

LIMITATIONS

There are a few factors that impact the quality of analysis in this case study. At a high level, these factors relate to temporal changes (on TechCrunch and on the Web as a whole), identification of anonymous users and equity in the creation of our datasets.

There are many temporal changes that could influence our analysis. To reiterate, these temporal changes are important because our datasets do not span the same time period. First, there are participation changes for the focal site, TechCrunch.com – these differences include both the volume of participation and site traffic, as well as changes in *who* is participating (on TechCrunch as well as the Web as a whole). Next, clearly the articles in the two datasets span different topics that trended in the news during the two time periods – topics that might shift not only due to events in the world, but also because of editorial or writing staff changes at TechCrunch. Finally, other interface changes to the site or commenting systems during these time periods can influence participation.

As thoroughly described in the "per user analysis" subsection of the "Impact of Anonymity on Comment Quantity" section above, there are complications in the task of uniquely identifying anonymous users given that we are completely reliant on the "provided name". That is, two users could use the same name or one user could write comments under multiple names. These two scenarios are discussed thoroughly in the previous section, and do have potential to impact the validity of our quantitative analysis though we do not think this limitation effects the conclusions of our analysis.

Finally, one detail of our data collection resulted in a minor imbalance. When gathering comments for the Disqus dataset, the Disqus platform only allowed access to the first 100 comments in each thread. While in most cases this is the entire thread, this does leave our dataset slightly incomplete.

CONCLUSIONS

This case study was aimed at the following question:

What impact does allowing anonymity have on the quality or quantity of participation on a site?

Through analysis of a technology social news site, TechCrunch.com, we addressed this question by comparing two datasets of TechCrunch articles and comments. In the first dataset, TechCrunch allowed anonymous comments (using the Disqus commenting platform). In the second, following a change to the Facebook commenting platform in March of 2011, anonymous comments are no longer allowed.

Through our qualitative analysis, we have many findings that support the claim that real identity comments are of higher quality. Through relevance analysis, we found that users who reveal more of their identity write comments that are more relevant to the focal news story (Table 1). Similarly, through analysis using the Linguistic Inquiry and Word Count tool, we see that more identity revealed yields less swearing, less anger, more affect words, more positive emotion words and less negative emotion words in comments (Table 4). Finally, within the Disqus dataset specifically, we see that *anonymous* comments were "liked" less than *pseudonym* comments.

In our quantitative analysis, we cannot conclude that the participation of users, from a strictly quantitative perspective, differs significantly between the time before and time after the switch from Disqus to Facebook. In the Disqus dataset, there are on average more comments per article, but the relative size of users to articles is smaller in the Disqus dataset. In the Facebook dataset, users comment on a larger number of different articles, while Disqus users focus more intently on specific articles. In favor of revealing identity, we found that longer thread lengths are associated with less anonymous participation (see Figures 1 and 2), a quantitative measure that favors disallowing anonymity. Finally, also in favor of revealing identity, more identity revealed yields longer duration of participation in the community (Table 8).

With an increase in some qualitative measures of participation and no significant decrease in participation (in fact, we see increases in quantity by some measures), the results of this study lean in favor of disallowing anonymity on similar sites. However, in making design decisions with regard to anonymity on a website, one must couple this analysis with knowledge of the community at hand and the political implications that such a decision might have (see discussion of "nymwars" in (Boyd 2012)).

FUTURE WORK

This case study focused on the impact of allowing anonymity in online communities. Future work will continue to call to question other design choices that impact participation quantity and quality as the user-generated web continues to grow. While work in natural language and image processing are critical in allowing sites to manage (keep to the standards and norms of a site) large amounts of user generated content, equally important is an understanding of what environments, and what qualities of these environments foster and promote (perhaps unintentionally) negative content of malicious intent, content that is destructive to these online communities.

REFERENCES

Boyd, Danah. 2012. "The Politics of 'Real Names'." *Commun. ACM* 55 (8) (August): 29–31. doi:10.1145/2240236.2240247.

Boyd, Danah, and Alice Marwick. 2011. "Why Cyberbullying Rhetoric Misses the Mark." *The New York Times*, September 22, sec. Opinion. http://www.nytimes.com/2011/09/23/opinion/why-cyberbullying-rhetoric-misses-the-mark.html.

Charness, Gary, and Uri Gneezy. 2008. "What's in a Name? Anonymity and Social Distance in Dictator and Ultimatum Games." *Journal of Economic Behavior & Organization* 68 (1) (October): 29–35. doi:10.1016/j.jebo.2008.03.001.

Chester, Andrea, and Gillian Gwynne. 1998. "Online Teaching: Encouraging Collaboration Through Anonymity." *Journal of Computer-Mediated Communication* 4 (2): 0–0. doi:10.1111/j.1083-6101.1998.tb00096.x.

Cho, Daegon, Soodong Kim, and Alessandro Acquisti. 2012. "Empirical Analysis of Online Anonymity and User Behaviors: The Impact of Real Name Policy." In *Hawaii International Conference on System Sciences*, 3041–3050. Los Alamitos, CA, USA: IEEE Computer Society. doi:http://doi.ieeecomputersociety.org/10.1109/HICSS.2012.241.

Christopherson, Kimberly M. 2007. "The Positive and Negative Implications of Anonymity in Internet Social Interactions: 'On the Internet, Nobody Knows You're a Dog'." *Computers in Human Behavior* 23 (6) (November): 3038–3056. doi:10.1016/j.chb.2006.09.001.

danielha. 2012. "Pseudonyms." *Disqus*. http://blog.disqus.com/post/15638234811/pseudonyms.

Dinakar, Karthik, Roi Reichart, and Henry Lieberman. 2011. "Modeling the Detection of Textual Cyberbullying." In *Proceedings of International AAAI Conference on Weblogs and Social Media, Workshop "Social Mobile Web."*

Flesch, Rudolph. 1948. "A New Readability Yardstick." *Journal of Applied Psychology* 32 (3): 221–233. doi:10.1037/h0057532.

Francis, M. E., and J. W. Pennebaker. 1993. "LIWC: Linguistic Inquiry and Word Count." *Dallas, T.: Southern Methodist University*.

Hartnett, Rodney T., and Harriet C. Seligsohn. 1967. "THE EFFECTS OF VARYING DEGREES OF ANONYMITY ON RESPONSES TO DIFFERENT TYPES OF PSYCHOLOGICAL QUESTIONNAIRES." *Journal of Educational Measurement* 4 (2): 95–103. doi:10.1111/j.1745-3984.1967.tb00576.x.

Johnson, Norman A, Randolph B Cooper, and Wynne W Chin. 2009. "Anger and Flaming in Computer-mediated Negotiation Among Strangers." *Decis. Support Syst.* 46 (3) (February): 660–672. doi:10.1016/j.dss.2008.10.008.

Joinson, A. 1999. "Social Desirability, Anonymity, and Internet-based Questionnaires." *Behavior Research Methods, Instruments, & Computers: a Journal of the Psychonomic Society, Inc* 31 (3) (August): 433–438.

Kilner, Peter G., and Christopher M. Hoadley. 2005. "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!*, 272–280. CSCL '05. International Society of the Learning Sciences. http://dl.acm.org/citation.cfm?id=1149293.1149328.

Li, Qing. 2006. “Cyberbullying in Schools.” *School Psychology International* 27 (2) (May 1): 157 – 170. doi:10.1177/0143034306064547.

LIWC. . http://www.liwc.net/howliwcworks.php.

Lou, J. K, K. T Chen, and C. L Lei. 2009. “A Collusion-resistant Automation Scheme for Social Moderation Systems.” In *IEEE Consumer Communications and Networking Conference, 2009.*, 571 – 575.

Miniwatts Marketing Group. 2011. “Internet World Stats.” *Internet World Stats*. http://www.internetworldstats.com/.

Qian, Hua, and Craig R. Scott. 2007. “Anonymity and Self-Disclosure on Weblogs.” *Journal of Computer-Mediated Communication* 12 (4): 1428–1451. doi:10.1111/j.1083-6101.2007.00380.x.

Reader, Bill. 2012. “Free Press Vs. Free Speech? The Rhetoric of ‘Civility’ in Regard to Anonymous Online Comments.” *Journalism & Mass Communication Quarterly* (May 22). doi:10.1177/1077699012447923. http://jmq.sagepub.com/content/early/2012/05/22/1077699012447923.

Sood, Sara Owsley, Judd Antin, and Elizabeth Churchill. 2012. “Profanity Use in Online Communities.” In *In Proceedings of ACM SIGCHI.*

Sood, Sara Owsley, Elizabeth Churchill, and Judd Antin. 2011. “Automatic Identification of Personal Insults on Social News Sites.” *Journal of the American Society for Information Science and Technology.*

Wallace, Patricia M. 2001. *The Psychology of the Internet*. Cambridge University Press.

Watt, S. E., M. Lea, R. Spears, and P. Rogers. 2002. “How Social Is Internet Communication? Anonymity Effects in Computer-mediated Groups.” *Virtual Society*: 61–77.

Yin, Dawei, Zhenzhen Xue, Liangjie Hong, Brian Davison, April Kontostathis, and Lynne Edwards. 2009. “Detection of Harassment on Web 2.0.” In *Proceedings of the Content Analysis in the WEB 2.0 (CAW2.0) Workshop at WWW2009*. Madrid, Spain.