

Seven Words You Can't Say on Answerbag: Contested Terms and Conflict in a Social Q&A Community

Rich Gazan
University of Hawaii
2550 McCarthy Mall, HL 2A
Honolulu, HI 96822
gazan@hawaii.edu

ABSTRACT

On a social Q&A site with thousands of transactions per day, what constitutes inappropriate content is not always obvious. In virtual communities such as this, users largely define and continuously renegotiate what constitutes appropriate participation, and moderators must allow healthy debate while curbing conflict. This paper presents an empirical analysis of Answerbag, a social Q&A site where moderators combined content analysis and transaction log analysis with information retrieval principles to identify non-obvious words associated with reported and unreported instances of conflict on the site. Content and transaction log analysis revealed the processes by which Answerbag users negotiated the meaning of contested terms, and suggested instances when conflict was a positive force for the community.

CCS Concepts

• **Human-centered computing** → **Collaborative and social computing** → **Empirical studies in collaborative and social computing**

Keywords

Online communities; Community Q&A; Community management; Boundary objects.

1. INTRODUCTION

This paper's title is adapted from George Carlin's monologue "Seven Words You Can Never Say on Television," [1] in which he questioned why seven widely used profanities were essentially unspeakable in certain public situations. After a performance of the monologue in Milwaukee, Carlin was arrested for disturbing the peace, and an unedited radio broadcast of the routine by station WBAI-FM in New York City led to a listener complaint and a series of legal actions, culminating in the 1978 US Supreme Court ruling that government can censor "indecent" material on public broadcasts. The goal of this paper is to unpack how members of an online community negotiated and renegotiated the meaning of words and their appropriateness in changing contexts, and how moderators of a high-traffic social Q&A site attempted to identify these areas of conflict in parallel with user reports.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org. *HT '16*, July 10 - 13, 2016, Halifax, NS, Canada
Copyright is held by the owner/author(s). Publication rights licensed to ACM

Social Q&A sites such as Yahoo Answers [2] and Quora [3] offer some of the same affordances as social networking sites, such as the ability to friend or follow certain topics or users, to "like" or "favorite" content, and to curate profile pages of their past contributions. Through these affordances, individuals and groups can support and promote what they feel is appropriate behavior and contributions, and publicly question, downrate or disparage content and actions they feel are inappropriate. Areas of the site where these negotiations occur are critical to identify, both for moderators to intervene if discussion degenerates into personal attack, but more importantly to identify points where diverse ideas converge, and the community evolves.

Relying solely on user reports of inappropriate content in a high-traffic environment presents several problems. Wrongly reported content consumes scarce moderator resources. Unreported conflict can fester on the site, and following the "broken window" theory [4], can yield the impression of an undermoderated site where personal attacks go unpunished, or worse, are selectively punished based on moderator whim, increasing the likelihood of more misbehavior. Guiding the few site moderators to areas of the site where conflict or other inappropriate behavior is brewing requires a quick, near-real-time organization of user-generated content into appropriate and possibly inappropriate content, the latter with a fairly high level of confidence to warrant moderator attention.

This participant observation describes how moderation staff at the Answerbag social Q&A site attempted to design an algorithm to detect unreported conflict instances following a site redesign that led to large-scale user conflict [5]. The author supervised paid and volunteer moderators, and had access to their reports and site transaction logs. Content analysis reveals some of the processes by which users debated the meaning of contested words, reshaped community norms, and helped moderators rethink their definition of conflict.

2. BACKGROUND

Online communities are as diverse as the theories and methods used to study them [6]. Lampe et al. [7] studied user motivation for online community participation through the lens of both uses and gratifications theory [8] and organizational commitment theory [9], and found that social and cognitive factors such as a sense of belonging, connection and importance to the community are more important than site usability in predicting long-term participation and continued contribution. Similarly, Jin et al. [10] employed theories of social capital, social exchange and social cognition, and found that user self-presentation and peer recognition were associated with continued participation and contribution.

2.1 Social Q&A

Describing a social Q&A site, where people, ask, answer and rate content while interacting around it [11], as an online community is not accidental. While some Q&A sites are designed around anonymous, fact-based interactions, social Q&A sites allow or require users to create an online identity and user profile, where their past contributions can be browsed by others as a means to identify the trustworthiness of the individual contributor.

The importance of social factors found in online community research more generally has also been observed in social Q&A sites. Raban and Harper [12] found evidence that extrinsic factors such as ratings and social recognition motivated user participation, and intrinsic motivations such as pride in content ownership—distinct from the quality of that content—and altruism were important as well. Designing in affordances to support and encourage these kinds of social interactions yielded social Q&A sites where conversations around content became arguably more important to users than the content itself, even in topic areas where accuracy is paramount, such as health information [13, 14].

2.2 Theoretical framework

The literature of online communities mapped prior research on social processes into the online environment to understand why people would be motivated to contribute and participate. In response to decades of literature in library and information science conceptualizing information seeking in terms of one-to-one, patron-to-librarian reference transactions, Shachaf [15] proposes an input-process-output theoretical framework for information seeking via social Q&A and related services. Her framework for understanding, analyzing, and evaluating social reference includes (a) answer quality, (b) user satisfaction, (c) service viability, and (d) a digital repository. While the majority of social Q&A research has focused on answer quality, and to a lesser extent user satisfaction [11], the viability of the social Q&A service and the community surrounding it have been studied far less often. According to Shachaf, viability is measured primarily through continued participation, but an impediment has been the scarcity of researcher access to backend user data, which tends to be proprietary [16]. This study utilizes backend data from Answerbag, the first social Q&A site in the US, and is conducted as a participant observation, with access to both user and moderator transaction data, to investigate user behavior surrounding potential instances of conflict.

Just as in physical communities, members of online communities are invested in their identities, support their friends, and take ownership of the community by legitimate participation as they define it, all of which can result in conflict. Savolainen [17] adapted the theoretical model of argumentation developed by Toulmin [18] to the Yahoo! Answers social Q&A site, and found that 24% of Q&A contained oppositional or mixed argument patterns, and only 2% of the grounds used in argumentation were emotive appeals such as ad hominem attacks, the level of conflict where moderator review would be appropriate.

2.3 Q&A retrieval and classification

The goal of any information retrieval system is to match query representations with document representations. How documents and queries should be best represented for optimal retrieval has been the subject of decades of research, but at its core, the problem is little changed from the days of the earliest information

retrieval experiments [19, 20]. The relevance of a document to a query has been operationalized along two major dimensions: the query term must be common within a document (term frequency), but rare across the collection (inverse document frequency). While the relevance of a given document to a given query is highly contextual [21, 22, 23], and information retrieval is best understood as a probabilistic endeavor, the core challenge is to identify which words across all documents in a collection, or pages on a Website, best represent the content.

Initial information retrieval models reflected the technological realities of their day, assuming a relatively static database, updated in occasional batches, and searched only by professionals as intermediaries for user requests. The rise of networked computing made it possible for users to search collections directly, and was quickly followed by Web 2.0 models of user generated content, then social computing models where users could rate, comment and interact around that content. These new affordances necessitated the creation of new models of social information seeking and retrieval.

Harper, Moy and Konstan [24] developed algorithms to distinguish informational and conversational questions on three social Q&A sites: Yahoo Answers, Answerbag and Ask Metafilter, positing that conversational questions were likely to be more ephemeral, and less relevant for future retrieval. Liu and Jansen [25] implemented and evaluated several classification algorithms for social Q&A along similar lines, based on question subjectivity vs. objectivity, and found fundamental differences in response time and respondent characteristics between the two question types. Similarly, Mendes Rodrigues and Milic-Frayling [26] integrated content analysis and network analysis in their study of the MSN QnA community, and proposed a typology of Q&A user intent, first suggesting a broad typology of social vs. non-social question types, and develop measures of engagement to quantify user participation. They developed a metric of user behavior summarized by a “social score” (p. 1136) designed to be used in parallel with social network analysis metrics, and suggest that future research focus on more fine-grained sense of user intent within the broader categories of social and non-social questions.

However, with the volume and diversity of users and interactions in social Q&A communities, more fine-grained analyses of content types and features have not yet demonstrated reliable predictive power. Burel et al. [27] investigated question selection behavior in the Stack Exchange community, and identified 17 user features (e.g. overall reputation, reputation within a topic, number of answers, etc.), 23 question features (e.g. question age, number of words) and 22 thread features (aggregate of question features with response age normalized) and gauged their relative ability to predict the questions a given user would choose to select based on their history. While they found weak predictive power in question age (newer questions tended to be selected more often), their best results were achieved by combining all the features, indicating that online question answering and the communities around them tend to defy easy distillation into a few relevant features.

Various forms of term analysis have been used to identify undiscovered content in Q&A sites. Yamamoto et al. [28] analyzed a sample word corpus from the Yahoo! Answers and Baidu Zhihao Q&A sites, to present users with adjective facets of noun keywords, allowing them to more easily recognize content of interest, without having to depend on a fortunate choice of query terms.

2.4 Boundary objects

Star and Griesemer conceptualize a boundary object as one that is part of multiple social worlds and facilitates communication between them; it has a different identity in each social world that it inhabits [29].

While boundary objects were initially conceived as tangible items, they are embedded in the context of their social worlds, and in subsequent research the concept has been expanded to include boundary infrastructures [30, 31], and boundary clusters [32], among many others. In a virtual community, where participants share no tangible objects, contested terms may fit the definition of a boundary object as an abstract concept or “ideal type” [29]. Star and Griesemer envisioned boundary objects as the means by which otherwise untenable conflicts between diverse groups might be negotiated; if so, moderators who misinterpret intergroup communication and negotiation around these terms as conflict detrimental to the site, and intervene by removing content, may cause unintentional harm to the community. On the other hand, visibly moderated online communities have been associated with higher information quality [33], and are more self-sustaining [34].

This brief literature review demonstrates that social Q&A sites share some of the same user dynamics that have been observed within online communities more generally, that social factors tend to outweigh content factors when attempting to explain continued user participation, and that analysis of transaction logs and words associated with negotiation and conflict can be productive avenues for research.

3. SETTING

Answerbag was the first social Q&A site in the US. Launched in July 2003, it was initially conceived as a site exclusively for factual Q&A. Its one question—multiple answers architecture acted as a user-generated recommender system, where the answers previous users had rated most helpful were listed first, but all answers were available by scrolling down the page. For its first few years of existence, Answerbag was 100% moderated—every piece of content submitted to the site would go to a moderator for approval, and even errors of spelling and grammar would be corrected. However, as the site grew and social submissions were accepted and quickly overwhelmed factual content, that level of review became impossible, and moderation became crowdsourced.

Registered users could report content they felt was inappropriate via several channels, most commonly via the flag feature. Alongside every Answerbag question, answer or comment was a link labeled “Report,” which routed the content for moderator review. Users indicated whether they felt the content was spam, offensive, miscategorized, nonsense or otherwise inappropriate, though all flagged content appended to the same review queue. Users could also email the 2-3 site moderators, part-time contractors who review flagged content, or send private messages to community leaders, a group of 5-7 volunteer members with limited moderation powers. These few individuals constituted the entire moderation staff during November 2009-February 2010, the period under study.

This particular period was selected because it surrounded a December 2009 major site redesign that led to large-scale conflict and the migration of longtime users to other Q&A sites [5], which directly affected Answerbag’s viability. After roughly three years of increasing traffic and a relatively stable interface, Answerbag

was redesigned by Demand Media, its new owners. The goals were to update the look and feel of the site, create more monetizable ad space, and to add new forms of interaction such as polls and debates. The relaunched site was immediately plagued by malfunctions that limited user interactions, but even when the site functions stabilized, there were far more reports of conflict on the site than the few moderators could investigate.

At this time, the Answerbag database contained approximately 3 million questions, 11 million answers, and drew in excess of 9 million unique visitors per month. The total number of registered users exceeded 200,000, with roughly 5,000 actively posting a combination of factual, conversational and sometimes unclassifiable content each day. The daily volume of flags ranged from approximately 50-150 per day, not including reports via e-mail, IM and other channels. With roughly 15,000 content transactions per day, less than 1% of the content submitted to Answerbag was flagged for review.

Prior to the redesign, roughly a third of user-reported flags were for obvious spam content posted by humans or bots, and quickly acted upon. Another third were submitted by users reporting miscategorized questions and suggesting a more appropriate category, which Answerbag encouraged by awarding points for each accepted flag, and was another low-effort review task for moderators. However, the remaining third focused on reports of conflict between users. These flags required significant moderator time to investigate, since they often occurred across multiple questions and threads. When accusations of coordinated user attacks were lodged, activity histories and IP logs had to be consulted. Since the value and long-term viability of the site relies on preserving a shared sense of appropriate behavior, these moderator efforts were viewed as a worthwhile expenditure of time, and they often uncovered evidence of worse transgressions than those that had been initially flagged. Similarly, some users who had content removed would find evidence of similar content on the site that was clearly against site policy, but which had not been reported. Strongly negative user reactions to the redesigned site, and the factionalism that followed, tended to amplify these issues, and resulted in more flags of all types. Moderators were swamped.

Several methods were attempted prior to the approach reported here, with little success. An approach focused on sudden pageview increases succeeded primarily in identifying popular topics; traffic spikes often indicate a comment thread where users went completely off topic and had an apparently enjoyable improvisational conversation, but in the vast majority of cases users interacted appropriately within these threads. Some pages where users had posted spam content and directed traffic to it via external links were also identified, but those were already being reported through the regular flag function. Another attempt posited that longtime members were most likely to engage in conflict, but creating a report listing the content around which high-ranking users were contributing succeeded only in duplicating the buzz or “hot topics” algorithm. Not surprisingly, experienced users interacted around a large volume of content, so the report was anything but a time saver for moderators trying to identify unreported inappropriate content.

A solution was needed to supplement user reports of inappropriate content with a way to detect unreported conflict instances, with very limited moderator resources, ideally with the side benefit of learning more about the roots of the conflicts, both of which motivated this study.

The research questions guiding this study are:

- **Which content-bearing terms are associated with instances of reported conflict on this site?**
- **What is the nature of the conflict?**
- **To what extent can those content-bearing terms reveal unreported conflicts elsewhere on the site?**

4. METHOD

This study is part of a long-term participant observation, and access to the Answerbag Q&A database backend, moderator reports and transaction logs are available. For the purposes of this study, the unit of analysis is a question posted to Answerbag, including all subsequent answers and answer comments, any of which may be reported as a potential instance of conflict. The number and length of answers and comments appended to a given question is highly variable, but since term frequency analysis depends on wordstock quantity, Q&A threads were not normalized for length. Additionally, users may copy and paste content from other sources, and include links, images and embedded video, and threads sometimes diverge from the original topic of the question or answer.

4.1 Operationalizing conflict

Following Shachaf [15], for the purposes of this study, conflict is operationalized as behavior that directly or indirectly threatens the viability of the site. This may include behavior that creates a negative environment for current and future users, diluting the quality of the content and the experience of the site, and diminishes user participation. Differences of opinion, from a single negative comment to a long thread of argument, usually do not meet this standard; indeed, conversations such as these are one of the attractions of social Q&A sites. However, when users attack, stalk or otherwise harass other users, that can not only dissuade continued participation of the participants, but of future visitors who may find the site via a search engine referral, browse through several screens of conflict and click away, never to return. Indicators of conflict in this study include:

- Personal attacks
- Reposting moderator-removed content
- Profanity (including euphemisms) directed at a user or group
- Posting links/screenshots of past conflicts
- Posting attacks on all posts by a particular user
- Vindictively or frivolously flagging content
- Creating multiple accounts to continue/escalate conflict
- Participating in/orchestrating coordinated attacks

Both user-reported flags and the unreported conflict instances discovered in this study were analyzed through these indicators. Elements such as vindictive flagging and creating multiple accounts were identified via transaction logs, and more interpretive indicators such as personal attacks were evaluated via content analysis. Content demonstrating at least one of the above indicators was coded as an instance of conflict.

4.2 Process

Step 1: Harvest wordstock. Content flagged for review as inappropriate by at least two users or moderators from November 2009 through February 2010 was collected. This yielded 2444 user-reported conflict instances connected to a question, answer or comment. Multiple user reports were sometimes connected to the same question; eliminating these duplicates yielded 1890 unique

questions associated with user-reported conflict. All words were extracted into a spreadsheet for analysis.

Step 2: Remove non-content-bearing terms, conflate term variants. Stopwords and other terms common in reported conflict instances (e.g. profanity) were removed, since these terms do not help distinguish individual Q&A in a search, and are therefore non-content bearing. A modified Porter stemmer [35] is used to conflate word form variants, primarily via suffix stripping.

Step 3: Relevance rank the remaining terms. Following classical IR, identify the remaining terms which are common in certain instances of reported conflict (term frequency), but otherwise rare across the sample (inverse document frequency). This yields a list of content-bearing terms most highly associated with reported conflicts during the timeframe of the sample. The top seven resulting terms were associated with 880 reported conflict instances.

Step 4: Search content-bearing terms across the entire site. Limiting the expanded search to the timeframe of the sample yields a set of Q&A containing terms associated with reported conflict, but which were not reported.

Step 5: Content analysis of resulting Q&A to identify instances of unreported conflict. The content analysis is conducted as an inductive constant comparison [36, 37, 38]. Transaction logs are consulted to ground instances of suspected conflict, for example by flag history, login times and multiple accounts associated with the same IP address. These top seven terms were associated with an additional 585 questions containing unreported conflicts during the timeframe of the sample.

5. RESULTS AND ANALYSIS

The seven content-bearing terms closely associated with user-reported conflict are listed in ascending order of the ratio of reported conflicts to observed occurrences across the 4-month sample. Unreported conflicts are not included in the percentages, as they were collected and measured outside the original sample.

1. Spam (33.1%)
2. Sockpuppet (37.5%)
3. Troll (40.1%)
4. Bagicide (43.5%)
5. Glitch (47.3%)
6. Rejigger (54.6%)
7. Fluther (59.7%)

A discussion of the content and transaction log analysis around each of these seven terms follows. While the numerical data represents the raw count of observed term occurrences and reported and unreported conflict instances during the period under review, the number of users involved in these conflicts could not be determined with certainty, primarily due to the likelihood that some participants in the conflict used multiple accounts. In rough terms, 75% of the conflict-related reports were associated with unique user accounts, while the remaining 25% involved multiple posts or conflicts by the same user, or were cases where coordinated activity by multiple accounts was suspected.

For the qualitative analyses, both the moderators and the author interpreted the data through their day-to-day familiarity with the site and its users, sometimes incorporating data from past and subsequent interactions outside the sample collected for this study.

5.1 Spam




Figure 1. Spam occurrences and surrounding conflicts

The spam flag was conceived by Answerbag designers as a way for users to identify inappropriate ads that slipped through the spam monitoring services and filters used to intercept bot-generated spam attempts. Some would post spam links on random pages throughout the site, others would post spam content on Answerbag, using it as a de facto hosting service, and link to it from elsewhere on the Web. While the latter variety of spamming drove additional traffic to the site, it directly detracted from the site's viability, and Answerbag users caught and reported this activity via the spam flag.

Defining spam as an unwelcome ad is the most common interpretation of the term, but some Answerbag users stretched this definition to include repetitive and/or nonsensical content, leading to conflict over appropriate site usage.

Why is it that when a dingbat comes on and starts spamming with stupid questions one after another, that instead of letting them fade off the main page people keep answering them?

Another user stretched the term even further, to include people who post floods of questions, or who create multiple accounts (sockpuppets, discussed below).

A spammer is someone who posts excessive amounts of the same questions, or questions that are very similar. Or someone who opens up multiple accounts.

However, debate over this term also appeared to bring the community together in some instances, which was a comparatively positive outcome. Users posted examples of appropriate and inappropriate content, compared notes on when to flag a piece of content as spam, and even invoked legal and historical parallels:

Kinda like the Supreme Court ruling on pornography, I know it when I see it.

Figure 1 shows that the term had been associated with conflict prior to the redesign, rose during December along with most reported conflict, then quickly settled back to its previous levels, or slightly below. This may indicate that conflict over what constitutes spam may be relevant to community members under normal circumstances, but when a "force majeure" like a site redesign fragments the community, fewer individuals are concerned about the definition of spam.

5.2 Sockpuppet




Figure 2. Sockpuppet occurrences and surrounding conflicts

Answerbag policy did not forbid users from creating sock puppets, i.e. alternate accounts on the site. Some users felt it was appropriate to be able to post different content under different accounts, as long as the accounts didn't interact, uprate one another or otherwise game the system. Other users, often those who saw their answers downrated or ridiculed by the same few users, felt certain that sockpuppet accounts were being created by their enemies and used vindictively, which was the source of much of the debate and conflict surrounding this term.

What in the name of Thor is a sock puppet? Is it like a troll?

Q: Is it morally wrong to create a sockpuppet account for the purpose of making you look good and boosting your points?

A: People can do whatever the hell they want, and while I don't find it "morally wrong" per se, I find it to be a pretty sad and pathetic way to boost one's popularity and further a mythos about oneself. No, it just makes you a cheater.

When you accuse someone of being a sock puppet, are you not accusing that person of being a liar, inauthentic, duplicitous and cowardly? If you go by any other name other than that which is yours, you are a coward. What do you think?

Several Answerbag accounts, most notably one using the handle DreAnna, were used to "catfish" other users, usually for expressions of sympathy at their claimed life circumstances. This led to suspicion that some new users were not who they claimed to be.

Besides "DreAnna," do you think there are more ABers among us who are creating elaborate false identities and playing sick games and scams on people here?

After the redesign, many longtime users publicly declared that they would never return to Answerbag via a variety of means: altering their usernames, deleting content from their user profiles, or posting farewell questions. However, examining IP logs revealed that many of these users did return under new accounts, and when a new user posted content others felt only an experienced user would know, conflict would erupt and moderators had to ban several accounts on both sides. Figure 2 shows that in the final two months of the sample, there were more unreported than reported conflict instances related to sockpuppets, indicating that users may have realized that reporting a potential sockpuppet account might not be worthwhile, if it resulted in the reporting user being penalized as well.

5.3 Troll




Figure 3. Troll occurrences and surrounding conflicts

In Web parlance, the understanding of trolling lies somewhere between harmless provocation, willful satire and antisocial online behavior, conducted for the “lulz” of generating expressions of frustration from others. The latter can indicate or lead to more serious online attacks such as stalking, bullying or harassment, though these go beyond the generally understood limits of trolling. Donath [39] views trolling as a game of identity deception, where a troll attempts to legitimize their message by first creating the impression that they are members of the community, giving their trolling actions more attention and/or impact. Herring et al. [40] detail the struggle of a feminist online forum to balance anti-trolling policies with freedom of expression, to preserve the online community as a safe space. Figure 3 shows that the term troll was common in conflict instances throughout the period under study, and that Answerbag users showed evidence of struggling with defining the same balance of acceptable and unacceptable conduct.

Conflict surrounding the term troll generally began with a public accusation of improper behavior on the part of another user, labeling it trolling. This would generally result in a redoubling of the behavior by the accused user, with counteraccusations of censorship. Friends and followers of people on both sides would follow their activity notifications into the fray, and thereby create an instance of conflict. Following Donath [39], Answerbag users who conflated trolling with simple disagreement (i.e. anyone who disagrees with me is a troll) tended to include an expression of their legitimacy as a member of the site, in terms of their time as a member, number of contributions or more general “us vs. them” statements.

Alright, who is the immature troll who is going through ALL my questions and comments and down rating everything? What is your issue? Grow up please!

How do you feel about short uninformative/uninteresting answers? Do you consider that troll behavior? Because they only want to cause us annoyance and trouble. They are ruining the nice and friendly environment of AB.

Others both claimed site legitimacy and moral high ground by claiming to be above such petty concerns such as points, further escalating the conflict.

The points are good for creating trolling, vindictiveness, hatred and bigotry. The place would be much better without it. Give insignificant people a little tool of power like stupid points and look at what happens.

In earlier Answerbag iterations, users could rate answers as useful (100%), somewhat useful (75%), or incorrect/not useful (50%). A user’s aggregate percentage would be available on their profile page, so other users could evaluate the answerer’s history along with the content of their answers. Discussions comparing the current Answerbag implementation to prior versions often led to conversations involving “old-timers” who lamented the loss of a quantitative metric of trolling. Broader conversations took place ascribing increased inappropriate behavior to the loss of the percentage downrating functionality. Through content analysis, several examples of negotiation about the meaning of the term troll were discovered, contextualized by both links to examples of other posted content as well as more abstract discussions about the meaning of the term in other contexts.

*How low does a % need to be for a person to be called a troll?
It’s my own arbitrary number but my number would be 75%.
I’d say anyone between 75% and 90% is a grouch, but perhaps not a troll.*

5.4 Bagicide




Figure 4. Bagicide occurrences and surrounding conflicts

Bagicide is a user-created portmanteau of Answerbag suicide. Though it generally refers to the act of abandoning or deactivating one’s account, it is most often used before the fact, when a user publicly announces their intention to leave the site and request that all their content be deleted. Sometimes users plan to return with new accounts, other times they express disappointment with their experiences on the site and claim they will not return.

The conflicts surrounding the use of this term relate to the sense of a user’s responsibility to the Answerbag community. Some view questions of this sort as unseemly attention seeking, and claim that those who propose it are fishing for responses along the lines of “Please don’t go!” Others invoke the social responsibility of maintaining a consistent identity, and note that deleting one’s content to the extent allowed by the site negatively affects existing discussions, and creates real harm.

It’s been asked before, but how do I do the “Bagicide” thing? I wish to start over as a beginner with no points.

What would/does committing Bagicide ultimately achieve? Does it award freedom and release from the points obsession, or is it counter-productive in that it may give the impression that points others awarded you are meaningless or not appreciated?

Drama. Puuuuke. Calls to mind the angsty teenager who commits/attempts suicide as a way of “punishing” their friends

and family for not treating them as they feel they should be treated. Narcissistic, selfish and melodramatic...

The question of whether Bagicide—or the threat of it—is appropriate site behavior is at issue when the term appears. Figure 4 shows that instances of conflict rose in December and January, when some users wished to remove all the content they had contributed as a form of protest in the wake of the redesign. It is notable that a high proportion of conflict instances were reported when this term was used, indicating that users felt that the threat of content removal was worth reporting to moderators. However, Answerbag users demonstrated the ability to have thoughtful and relatively harmonious discussions about Bagicide as well.

The best thing about Bagicide, IMHO, is that it sparks debate about ego and the importance of not allowing yourself to get caught up in points. Seems like a drastic step, though.

5.5 Glitch

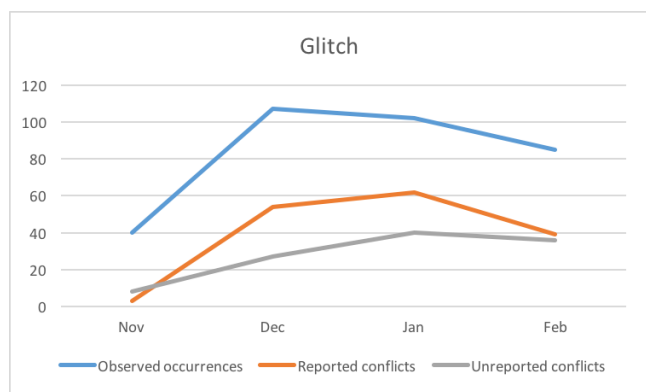


Figure 5. Glitch occurrences and surrounding conflicts

A glitch refers to a site malfunction. In the early days of Answerbag, prior to its sale to Demand Media, users would commonly report and help diagnose bugs and malfunctions, and email the site owner directly. That level of connection created a sense of shared endeavor and ownership, which site bugs arguably enhanced. But as the site grew, and gained the backing and technical resources of a large company, in the eyes of longtime users, malfunctions became less forgivable.

Figure 5 shows that the term “glitch” appeared in the Answerbag database prior to the redesign, but during the period studied here it greatly increased in frequency. Many instances were simply users collectively raging against the site maintainers, who were seen as incompetent and ungrateful. Conflicts occurred when the term became conflated with one’s support of the site administrators.

Why didn't you take the hours you put into this "upgrade," fix what we asked nicely to be fixed and made it a better site? All you have done is put all this time into a site that has even more glitches and everyone hates.

Okay Answerbag. Now I am pissed. Why do you glitch me like this? I just tried to ask a long complicated question and it just plain ol' disappeared. You've got to be kidding me. I followed all steps correctly.

Even seemingly positive questions including the term glitch often branched into factionalism and conflict. Responders argued that glitches would never be fixed unless members made a collective

effort to reduce the site’s traffic, and that the site owners had not earned such loyalty.

I have officially pronounced myself a glitch-bagger. I will not be happy until all the site glitches are fixed. But I will remain. Are you with me?

Some users attempted to respond with humor, posting questions that made light of the situation and recognized that users were all in this together.

Should there be The Glitch Awards? Categories like Least Frustrating Glitch or Likeliest Glitch to Result in the AB Headquarters Being Burned Down, etc.?

However, as these questions drew more and more responses, users who thought the community had been seriously damaged and disrespected by the site owners tended to “crash the party” on these lighter questions, accusing others of not taking the situation seriously enough, and creating conflict.

5.6 Rejigger

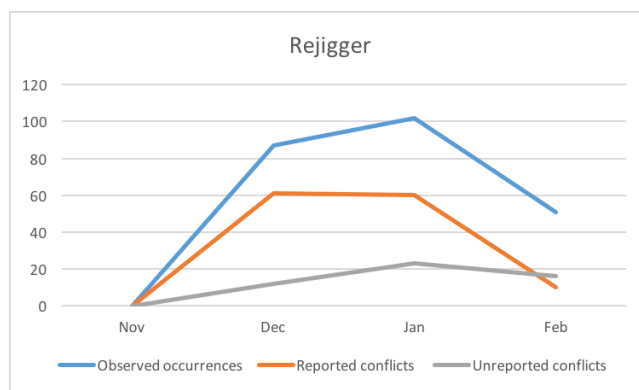


Figure 6. Rejigger occurrences and surrounding conflicts

Part of the 2009 redesign attempted to address the cumulative advantage problem, and establish more equity in points opportunity. The initial points structure allowed users to earn the right to give or take away as much as six points per answer by accumulating positive ratings of their content from other users, while new users could give or take away only one point per answer. Longtime users with large networks of similarly high-ranking friends also had the advantage of the notification function—their friends would receive a notification anytime they posted content, making their posts far more likely to receive large numbers of uprates from large numbers of other high-ranking users, often regardless of the relative quality of their posts. Also, some users had accumulated points by flagging content for removal or recategorization, and had discovered that it was far easier to earn points by clicking once on someone else’s content than generating their own. Moderators reported that certain users would submit hundreds of flags in a single session and rise in level while contributing no content.

To address this situation, the redesign included altering the rating structure to more closely resemble that of Facebook—all users regardless of level or experience could bestow one like to a given piece of content. Downrates were seen as the seeds of much site conflict, and eliminated entirely. Importantly, the points users had accumulated under the previous rating scheme would be retroactively converted to the new system using an algorithm that emphasized content contributions and de-emphasized content

flagging, but this was not disclosed to users. The new rating system was announced on the site blog by a Demand Media employee, who described the change as “rejiggering” the points structure. Despite assurances that longtime users would not be penalized under the new system, some were.

As Figure 6 shows, the term rejiggering and its variants had not appeared on the site before, but it immediately became code for an unfair rule change, or a too-breezy description of a serious situation. Some new users expressed support for the changes, while some experienced users felt betrayed. The most strident conflicts occurred between experienced users who accused each other of caring too much about points and leaderboard position, the other side took the position that administrators had disrespected the contributors by changing the rules with no consultation or notice.

Q. How long does it take to rejigger points, and how did they get jiggered to start with?

A. You say that nobody has lost a level or anything. As far as I am concerned that is a kick in the guts and a total insult. Especially since the people you put ahead of me have run off to other sites. Maybe it is a ploy to get them back. I was an Illuminati ranked Number 1 with 690+ thousand points and you have screwed me with your lies and promises.

Did Umar Farouk Abdulmutallab attempt to “rejigger” the plane?

The so-called “rejiggering” is, IMO, somehow inequitable and awful...and certainly has not appropriately rewarded all the work that I, and others, put [in].

In the immediate aftermath of the redesign, one of the few working site functions was the ability for users to change their usernames. Doing so made the new name append to every piece of content the user had ever posted on the site, and some users made it an avenue of protest, generating user-specific conflict appearing across multiple questions, answers and comment, as revealed by transaction logs:

Rejiggering a Fat Baby’s Ass (Answerbag member profile)

5.7 Fluther




Figure 7. Fluther occurrences and surrounding conflicts

Fluther [41] is another social Q&A site. In the immediate aftermath of the redesign, when communication among users was limited by the site-wide malfunctioning of several new features, some users leveraged their offline connections to notify their friends to migrate to other Q&A sites, Fluther among them.

Figure 7 shows that the term Fluther had not appeared in the Answerbag database before, but its usage spiked in December 2009 and January 2010, and was the term most strongly associated with reported and unreported conflict in this study. Some users felt that it was disloyal to abandon Answerbag for another site, while others argued passionately that implementing a broken site and cutting off people’s communication had to be met with a strong collective response.

Some users attempted to start conversations about the plans of other users, but including the phrase “abandoned ship” in this example was viewed as an unwarranted attack, leading to conflict.

Looks like about 80 plus people have abandoned ship and are at Fluther. Are you staying on AB?

In a high percentage of these conflict instances, the cleavage of the community was apparent.

What do you think of all the people who went to Fluther? Apparently there is a Flutherite here who is saying that the flood of ABers to Fluther is rude and disruptive.

Many are using Fluther now and just watching AB fall the hell apart!

I am transitioning from AB to Fluther now. I am going all over AB cussing out and being a bitch to people talking sh-t about people leaving.

I was stunned, shocked and amazed. Not that so many of them chose to leave but that so many of them chose to leave after flinging insults at those of us who chose to stay.

Later in the period studied here, conflict ebbed around the term Fluther, but some users returned with reports and reconnaissance of life in the other community, leading to renewed friction.

Fluther is too much a ‘Nanny State’ for me. You have to appease the Spanish Inquisition to get a question passed. I got fed up with it. AB will improve I have no doubt.

While some occurrences of the term were in the context of feature comparison rather than conflict between users, this was the term most strongly associated with instances of conflict. Also, Figure 7 shows that while reported conflict instances including the Fluther term fell between January and February, unreported instances rose. This may indicate that those who used the term did not feel that moderator review was effective or worthwhile.

6. CONCLUSION

Identifying conflict and contested terms in an online community is like classic information retrieval in reverse. Instead of users seeking the right terms to identify relevant site documents, the site is seeking the right terms to identify relevant user contributions.

This study addressed the question of whether content-bearing terms could be used to identify instances of conflict on the Answerbag social Q&A site. Returning to the research questions:

Which content-bearing terms are associated with instances of reported conflict on this site? The seven content-bearing terms most strongly associated with instances of reported conflict on Answerbag during the period under study were spam, sockpuppet, troll, bagicide, glitch, rejigger and fluther.

What is the nature of the conflict? Though both the terms and the conflicts tended to be centered around the redesign of the site, the four-month timeframe of the sample also revealed that some

terms such as spam and troll were being discussed by the community both before and after the redesign, and that other terms rose and fell in frequency and conflict association as the community acclimated to a particular usage or interpretation, or simply moved on to other things. In other instances, apparent conflict led to positive discussion and negotiation about the meaning of contested terms, and appropriate site behavior.

To what extent can those content-bearing terms reveal unreported conflicts elsewhere on the site? Analyzing content-bearing word occurrences in reported conflict instances and searching them across the entire site yielded additional unreported conflict instances associated with those words in every case. In five of the seven cases, unreported conflicts outnumbered reported conflicts, with a general trend toward fewer reported and more unreported conflicts for the same terms over the duration of the sample.

Analyzing unreported conflicts with content and transaction log analysis revealed users' changing views of appropriate behavior in this online community. By identifying and searching content-bearing terms in this manner, moderators who would normally have to review every user-reported conflict could instead analyze the reported content, extract content-bearing terms, and review those items first, whether reported or not. In this study, the results suggest that by searching these content-bearing terms, moderators could identify conflict instances with 33.1% to 59.7% confidence, allowing them to prioritize reviewing these situations, where community is both created and fragmented.

Beyond managing conflict, this method also shows promise as a way to identify conversations and negotiations around contested terms, perhaps as boundary objects. How community members define appropriate behavior should be of interest to members, moderators and administrators alike in every online community, and future research might attempt to trace a word's introduction, through the negotiation of meaning, and perhaps identify patterns common across online communities.

6.1 Epilogue

On December 15, 2015, Answerbag users who tried to access the site were greeted with a terse message:

Answerbag is no longer available. Thank you for your patronage.

Demand Media, the site owner, did not provide any further explanation. Several days later, the message was replaced by a notice that the domain name was for sale. At this writing, none of the content from the twelve years of the community's existence can be accessed.

However, the viability of the community remains. As they have done before, users connected with one another through alternate means in the wake of the shutdown, and migrated to several new Q&A sites. One of these, Answermug [42], is run by a former Answerbag member.

7. ACKNOWLEDGMENTS

I wish to thank Joel Downs for creating Answerbag and for providing access to its data. I would also like to thank the members of the erstwhile Answerbag community for their contributions and participation. It was a good run.

8. REFERENCES

- [1] Carlin, G. 1972. Seven words you can never say on television. *Class Clown* (sound recording). Little David/Atlantic.
- [2] Yahoo Answers. 2016. <https://answers.yahoo.com/>
- [3] Quora. 2016. <https://www.quora.com>
- [4] Wilson, G. L. and Kelling, J. Q. 1982. Broken windows: The police and neighborhood safety. *The Atlantic*. <http://www.theatlantic.com/magazine/archive/1982/03/broken-windows/4465/>
- [5] Gazan, R. 2011. Redesign as an act of violence: Disrupted interaction patterns and the fragmenting of a social Q&A community. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vancouver, Canada, May 07-12, 2011). CHI '11. ACM, New York, NY, 2847-56.
- [6] Preece, J. and Maloney-Krichmar, D. 2005. Online communities: Design, theory, and practice. *Journal of Computer-Mediated Communication* 10, 4.
- [7] Lampe, C., Wash, R., Velasquez, A., and Ozkaya, E. 2010. Motivations to participate in online communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, GA, April 10-15, 2010) CHI '10 ACM, New York, NY, 1927-36.
- [8] Ruggiero, T. E. 2009. Uses and gratifications theory in the 21st century. *Mass Communication & Society* 3, 1, 3-37.
- [9] Allen, N. J. and Meyer, J. P. 1990. The measurement and antecedents of affective, continuance and normative commitment to the organization. *Journal of Occupational Psychology* 63, 1, 1-18.
- [10] Jin, J., Li, Y., Zhong, X., and Zhai, L. 2015. Why users contribute knowledge to online communities: An empirical study of an online social Q&A community." *Information and Management* 52, 7, 840-49.
- [11] Gazan, R. 2011. Social Q&A. *Journal of the American Society for Information Science and Technology* 62, 12, 2301-2312.
- [12] Raban, D. R. and Harper, F. M. 2008. Motivations for answering questions online. In D. Caspi and T. Samuel-Azran (Eds.), *New Media and Innovative Technologies*, Tel Aviv: Tzivonim Publishing, 73-97.
- [13] Bowler, L., Monahan, J., Jeng, W., Oh, J. S., and He, D. 2015. The quality and helpfulness of answers to eating disorder questions in Yahoo! Answers: Teens speak out. *Proceedings of the ASIS&T 2015 Annual Meeting*. Silver Spring, MD: Association for Information Science and Technology.
- [14] Worrall, A. and Oh, S. 2013. The place of health information and socio-emotional support in social questioning and answering. *Information Research* 18, 3. <http://www.informationr.net/ir/18-3/paper587.html>.
- [15] Shachaf, P. 2010. Social reference: Toward a unifying theory. *Library & Information Science Research* 32, 1, 66-76.
- [16] Shah, C., Oh, S., and Oh, J. S. 2009. Research agenda for social Q&A. *Library & Information Science Research* 31, 4, 205-9.

- [17] Savolainen, R. 2012. The structure of argument patterns on a social Q&A site. *Journal of the American Society for Information Science and Technology* 63, 12, 2536–48.
- [18] Toulmin, S. E. 1958. The uses of argument. Cambridge University Press.
- [19] Cleverdon, C. W. 1960. ASLIB Cranfield research project on the comparative efficiency of indexing systems. *ASLIB Proceedings* 12, 421-431.
- [20] Cleverdon, C. W. 1967. The Cranfield tests on index language devices. *ASLIB Proceedings* 19, 6, 173-194.
- [21] Wilson, P. 1973. Situational relevance. *Information Storage and Retrieval* 9, 8, 457–71.
- [22] Swanson, D. R. 1986. Subjective versus objective relevance in bibliographic retrieval systems. *The Library Quarterly* 56, 4) 389-398.
- [23] Cosijn, E. and Ingwersen, P. 2000. Dimensions of relevance. *Information Processing & Management* 36, 4, 533-550.
- [24] Harper, F. M., Moy, D., and Konstan, J. A. 2009. Facts or friends?: Distinguishing informational and conversational questions in social Q&A sites. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, April 4-9, 2009) ACM, New York, NY, 759–68.
- [25] Liu, Z. and Jansen, B. J. 2015. Subjective versus objective questions: Perception of question subjectivity in social Q&A. *Lecture Notes in Computer Science 9021*. Berlin: Springer International, 131–140.
- [26] Mendes Rodrigues, E. and Milic-Frayling, N. 2009. Socializing or knowledge sharing?: Characterizing social intent in community question answering. In *Proceedings of the ACM International Conference on Information and Knowledge Management*. CIKM'09. (Hong Kong, China, November 2-6, 2009). ACM, New York, NY, 1127-36.
- [27] Burel, G., Mulholland, P., He, Y., and Alani, H. 2015. Modelling question selection behaviour in online communities. In *Proceedings of the ACM International World Wide Web Conference, WWW'15*, (Florence, Italy, May 18-22, 2015). ACM, New York, NY, 357–58.
- [28] Yamamoto, T., Nakamura, S., and Tanaka, K. 2011. Extracting adjective facets from community Q&A corpus. In *Proceedings of the ACM International Conference on Information and Knowledge Management*, ACM, New York NY, 2021–24.
- [29] Star, S. and Griesemer, J. 1989. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science* 19, 3, 387–420.
- [30] Bowker, G., and Star, S. L. 1999. *Sorting things out: Classification and its consequences*. Cambridge, MA: MIT Press.
- [31] Star, S. 2010. This is not a boundary object: Reflections of the origin of a concept. *Science, Technology, & Human Values* 35, 601-617.
- [32] Rehm, S.-V. and Goel, L. 2015. The emergence of boundary clusters in inter-organizational innovation. *Information and Organization* 25, 1, 27-51.
- [33] Chen, J., Xu, H., and Whinston, A. B. 2014. Moderated online communities and quality of user-generated content. *Journal of Management Information Systems* 28, 2, 237-268.
- [34] Andrews, D. C. 2002. Audience-specific online community design. *Communications of the ACM* 45, 4) 64.
- [35] Porter, M. F. 1980. An algorithm for suffix stripping. *Program* 14, 3, 130–137.
- [36] Glaser, B. G. 1965. The constant comparative method of qualitative analysis. *Social Problems* 12, 4, 436-445.
- [37] Goetz, J. P. and LeCompte, M. D. 1981. Ethnographic research and the problem of data reduction. *Anthropology and Education Quarterly* 12, 5, 51-70.
- [38] Lincoln, Y. S. and Guba, E. G. 1985. *Naturalistic Inquiry*. Sage.
- [39] Donath, J. 1999. Identity and deception in the virtual community.” In P. Kollock & M. Smith (Eds.), *Communities in Cyberspace*. London: Routledge, 29-59.
- [40] Herring, S., Job-Sluder, K., Scheckler, R., and Barab, S. A. 2002. Searching for safety online: Managing “trolling” in a feminist forum. *Information Society* 18, 5, 371-384.
- [41] Fluther. 2016. <http://www.fluther.com>
- [42] Answermug. 2016. <http://answermug.com>