

Learning from Elitist Jerks: Creating High-Quality Knowledge Resources From Ongoing Conversations

Julia Bullard

School of Information, University of Texas at Austin, 1616 Guadalupe Street, Austin, TX 78701. E-mail: julia.a.bullard@gmail.com

James Howison

School of Information, University of Texas at Austin, 1616 Guadalupe Street, Austin, TX 78701. E-mail: jhowison@ischool.utexas.edu

Online-community management is commonly presented as the facilitation of conversation and contributions, especially converting readers to contributors. However, the goal of many discussion communities is to produce a high-quality knowledge resource, whether to improve external task performance or to increase reputation and site traffic. What do moderation practices look like when the community is focused on the creation of a useable knowledge resource rather than facilitating an inclusive conversation? Under what conditions is this style of moderation likely to be successful? We present a case study from online gaming—Elitist Jerks—in which aggressive moderation is used to transform a conversational medium into a high-quality knowledge resource, using the strategy of open censorship. We present a content analysis of moderator comments regarding censored messages. Our analysis revealed differences in types of contributor mistakes and the severity of moderator actions: infractions that interfered with both conversation and resource quality were punished harshly, whereas a set of infractions that supported conversation but undermined resource quality were more respectfully removed. We describe a set of conditions under which moderators should intervene in the conversion of conversation to knowledge resource rather than the conversion of lurkers to contributors.

Introduction

The common wisdom of community management is that the purpose of intervention is to protect and facilitate the generative practices of activity and conversation (Bishop, 2007; Kraut & Resnick, 2012; Preece & Shneiderman, 2009). A style of community management that prioritizes

conversation is consistent with using site activity as a measurement of success. It is also consistent with findings in online community scholarship that communities benefit from interpersonal relationships built on social presence and politeness (Wei, Crowston, Li, & Heckman, 2014) or that community activity is driven by a feeling of connectedness (Godin, 2008). Such a focus directs community managers to particular perceptions regarding community users, including the vilification of lurkers and prioritizing the conversion of visitors to “contributors” (Ebner, Holzinger, & Catarci, 2005). For example, moderators might provide contribution channels with low thresholds for effort, including introduction threads, the practice of thanking or echoing others’ posts, or providing space for off-topic threads (Kraut & Resnick, 2012; Millington, 2012).

However, discussion forums are not only spaces for conversation. Increasingly, these sites are treated and used as knowledge repositories, valued for complete, accessible, and accurate information (Lampe, Wash, Velasquez, & Ozkaya, 2010). Ease of navigating the “conversation archive”—the knowledge resource produced by prior discussion—is of particular concern for communities whose ongoing operation depends on reuse of previously shared content (Millen, 2000). Communities seek information technologies appropriate to their particular activities and aims, including the importance of conversation and expertise (Wenger, White, & Smith, 2009). For example, the impact of the visibility and accessibility of discussion site information through search engines has led existing communities to migrate to online discussion spaces from mailing lists (Vasilescu, Serebrenik, Devanbu, & Filkov, 2014). For inward-facing corporate communities, valuing a discussion forum as a knowledge resource might originate with a concern for external task performance, leading to a need for information that is of high quality and accessible (Yates & Wagner, 2010). For

Received July 1, 2013; revised April 25, 2014; accepted April 26, 2014

© 2015 ASIS&T • Published online in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/asi.23327

outward-facing communities, increasing information value improves the status of the site through metrics such as PageRank and ultimately leads to an increase in traffic and advertising revenue. Either way, this change in emphasis affects how a community values different kinds of users. For users interested in the informational content of a discussion board, the conversation itself is only the process that generates valued content. In this context, those coming to the site for information are a class of users who do not contribute directly to “activity” visible directly on the site but may benefit the community in other ways, such as increasing the status of the site by discussing and linking to it in external contexts and by applying knowledge gained to external practices (Takahashi, Fujimoto, & Yamasaki, 2003).

How then should a community manager approach the creation of a knowledge resource valued by information seekers? Facilitating a high-quality knowledge resource is a challenge, because there is an inherent tension between the needs of a conversation and the needs of a knowledge resource. A wide-ranging and stimulating conversation does not naturally result in a high-quality knowledge resource; topics are repeated with little additional information, and interpersonal exchanges might contain very little informational content while clouding search efforts. Although active conversations might contain useful knowledge, such knowledge is easily accessible only for those who continually participate. For those who visit the site on the occasion of a particular information need, or need a summarized introduction to the topic, conversational media lack clarity and accessibility.

For this reason, online communities have often undertaken additional work to generate canonical resources from their conversations, such as the FAQ genre that developed out of Usenet threads and efforts to maintain wikis capturing the community’s knowledge alongside discussion forums (Hansen, 2009; Wagner, 2006). These parallel practices—the conversation and the canonical resource—have a complex relationship. In one way, a high-quality knowledge resource enhances the conversational practice because veterans and newcomers have a common point of reference for ongoing discussion. However, being referred out of the conversation to review the FAQ, or being told to “RTFM,” can be an alienating experience for newcomers. Worse, as we discuss in detail here, the effort to create a separate but continually updated distillation of the conversation is substantial; communities struggle to find members motivated to create and maintain the parallel resource. Furthermore, an out-of-date knowledge resource may be worse than none at all, causing conflict regarding canonical, shared understanding. Indeed, many communities begin these projects and later abandon them, citing the deterrent of the “hassle” of keeping current two repositories and the desire of members to engage in conversation only. The additional effort and coordination required to create parallel knowledge resources seems incompatible with the genius of online communities: their convenience for participants. Online communities are successful in part because they reduce ancillary effort and

hassles to participation, such as the need to put aside time, commuting costs, or record-keeping overhead. How then might a community develop a relatively low-effort way of ensuring that the archive of their conversation is itself a useful knowledge resource?

A second challenge for creating a high-quality knowledge resource out of a conversation is the incompatibility of these two communication genres. In facilitating the creation of a knowledge resource, community managers risk interfering with the flow of conversation. The unavailability of conflict between the two possible uses of an interaction—participation in the conversation and reuse of the record it leaves—is a common observation in other fields. For example, archivists constantly weigh the primary and secondary uses of curated collections and debate whether their aim should be to preserve the records exactly as they were created (see, e.g., Jenkinson, 1965) or to select and reorder the contents to facilitate imagined research uses (see, e.g., Schellenberg, 1956). Although archival practices do not risk disrupting the original process that is generating the resources, archival theorists constantly wrestle with representing documents, or conversations, completely and in their original state or selecting only ideal documents to reduce the effort of future users.

In online-community management, there are risks of any intervention in the flow of conversation. As found in e-democracy studies, moderating to improve the quality of shared information threatens the “open” nature of democratic discourse (Wright & Street, 2007). Moderator intervention, no matter how benevolent in intent, might appear to be censorship, undermining the openness of virtual communities (Schackman, 2010). Focusing on the usefulness of archives prioritizes content, but there is more to conversation than content alone; conversation is an experience enhanced by social lubricants, including linguistic “phatics” (greetings, thanks, encouragement) and “face work” (Goffman, 1959).

Given these challenges, how does a community build a valuable, searchable knowledge resource from a conversational platform? What does moderation look like when there is a clear need to facilitate a high-quality information resource from a conversational medium? We provide answers to these questions by examining the overall strategy and style of moderator intervention in a community that specifically aims to prioritize a knowledge resource over conversation. We identify the overall strategy and focus on the types of contributions that moderators remove from the record of the conversation and how different types of rejected contributions are treated. Finally, we step back to discuss the conditions under which we are likely to see such an approach to community management succeed.

Moderation Practices in the Elitist Jerks Community

We examined the elitistjerks.com forums (EJ). The EJ website is a community discussion space for competitive

World of Warcraft (WoW) players and the guild site of Elitist Jerks, a competitive guild in WoW. News about the guild, including recent achievements and recruiting, is published on the main page, and the administrators and some moderators of the discussion forum are Elitist Jerks members. EJ, as well as being a guild website, is a discussion space for WoW players interested in “theorycrafting” to support game play (see following).

The EJ forums are an appropriate exemplar for four reasons. First, they represent a naturalistic community solution for a community need, rather than a designed solution from information professionals, and thus represent the members’ priorities for managing content and conversation. Second, these forums are highly successful: the practices clearly work well for this community. Third, as a discussion board for the production of findings and strategy for the larger online community of WoW players, the site includes a clearer role for noncontributing readers than might occur in more general-purpose forums or those not so explicitly oriented to an outside goal, increasing the comparability with information resources developed for future users. Fourth, the presence of the Banhammer subforum (see following) provides uncommonly candid access to the rationale of moderators in applying censorship policies to specific incidents. The following sections will briefly describe the activity of the EJ community and its major mechanism for enforcing a high standard of contributions.

Theorycrafting

Massively multiplayer online games (MMOs) create opportunities for many different kinds of play, ranging from casual exploration to hardcore optimization. Games such as WoW appeal to casual gamers, but competitive players with ample time to invest have access to a different community and a different mode of play. Hardcore or power gamers are a small subset of players dedicated to optimal performance, whose play sometimes looks like work, cheating, or “psychologically pathological” (Taylor, 2006). Power gaming is characterized by “a focus on efficiency and instrumental orientation (particularly rational or goal-oriented), dynamic goal setting, a commitment to understanding the underlying game systems/structures, and technical and skill proficiency” (Taylor, 2006, p. 72). In WoW, guilds of power gamers compete to be the first in the world, in their region, or on their server to complete the most difficult content available.

The digital nature of video games means that, unlike analog board games, many of the rules are hidden from players, calculated by the software and knowable only through the results of interaction with the game. The “massive” nature of WoW promotes the emergence of productive, collaborative exploration of the underlying rules. In the practice of theorycrafting, players engage in rigorous experimentation and modeling of the game software’s underlying rules to support optimal in-game practices. Players run simulations and parse logs of in-game

data to determine causal relationships among game values, essentially reverse-engineering the organizing principles of the simulated world. For the hardcore players who participate in the construction of theorycraft knowledge, the goal is not only to defeat the most challenging game content before another guild, but to master the game itself by parsing out the underlying rules that determine the optimal tactics for any encounter to come.

Scholars in game studies and education are particularly interested in the emergence of theorycrafting as a practice that transforms WoW play for theorycraft participants and the wider community of players through the development of knowledge resources. Throughout the history of WoW, theorycraft has developed from a derogatory term to a small but necessary project, to the point that the current caliber of hardcore play is possible only because of the principles established in theorycrafting spaces (Paul, 2011). The production and circulation of game knowledge that begin on these forums have a significant effect on WoW play, even among players who do not participate in the site itself (Thomas, 2009). Strategies derived from theorycraft direct player activity in purchasing new in-game equipment, deciding on a pattern or priority of abilities during combat, and assessing the performance of peers. These choices can have critical impacts on a player’s in-game capital, their status within their player group, and even their ongoing association with a particularly demanding and competitive guild of players. For these reasons, players value the accuracy, currency, and reliability of theorycraft knowledge, as the cost of bad strategies is at least as great as the benefit of good strategies. As with leisure activities such as backpacking or collecting, game play taken seriously by its participants is a source of advanced information needs demanding high-quality information resources (Chang, 2009; Fulton, 2009; Stebbins, 1982). Although the WoW community of players provides EJ with an external source of motivated readers, these readers demand of EJ high quality, reliable, and accurate content. Just as there is a great deal of public interest in the conduct of climate scientists, because their results are significant for public policy (see, e.g., Edwards, 2010), the operation of EJ is important to the community and its participants.

As a central location of theorycrafting discussion and knowledge, EJ is “disproportionally influential” in the WoW community (Paul, 2011). EJ’s influence and reputation as a leading source of reliable information has been attributed to their strict posting rules (Thomas, 2009). Those threads that do meet the high standard of EJ posting rules consist of thousands of posts, are read millions of times, and contain published data sets of empirical tests of in-game mechanics, spreadsheets to calculate weights of in-game values, and analytical reasoning comparing various models.

The Banhammer Subforum

Visitors to the forum can read the contents of the EJ forums, but only registered users can post. Registration is

free, but members must provide identifying information about their WoW character, to prevent users from posting anonymously. New accounts cannot post messages (are “muted” in the parlance of the site) for 1 day, after which they cannot start new threads until they have established their “posting capability” by posting 10 “successful” replies to other threads. Subsections of the discussion board include a “sticky” post advising contributors to a topic thread to “read the last 5 pages to be sure the topic you had in mind hasn’t recently been covered” (Forums Rules, 2010). The Forum Rules document, to which new users must agree, includes a directive to “familiarize yourself with The Banhammer . . . it will give you some examples of what not to do.”

If a post is not successful, it is deleted from the thread in which it was posted and moved to the Banhammer subforum by a moderator, who adds an explanation of why the rule was broken. The formal rules are often invoked by number and title, although some subject lines instead indicate that multiple or obvious rules were broken. The rules of the EJ forums are as follows (we will explain them through example).

1. All posters are to make an effort to communicate clearly.
2. All opinions should be stated as succinctly as possible.
3. All discussion should be both polite and civil.
4. Whining in any form is forbidden.
5. Threads should be started if and only if there is some reasonable topic to discuss.
6. Do not post unless you have something new and worthwhile to say.
7. Do not beg for hand-holding.
8. All accounts must have a valid WoW profile.
9. Do not sign your post.
10. Do not respond to terrible posts.

Each thread in the Banhammer corresponds to one rejected post, and includes the text of the original post, the name of the contributor, the moderator’s commentary, and a label that indicates whether the moderator action taken over the message is a warning or an infraction. Infractions have point values that determine for how long a member is muted, with 2 points being 1 day and 10 points being a permanent ban, or “The Banhammer.” Points applied to a member are additive but decay over time.

By publishing deleted posts in the Banhammer subforum, the moderators of EJ practice “open censorship.” Wright (2009) advocates open censorship as a moderation strategy for online communities that require some intervention to improve discussion quality but wish to avoid the appearance of malicious or biased interference. Although open censorship has been a particularly important strategy for e-democracy forums, which are targets of “troll” attacks but have to maintain a reputation for fairness, in the EJ forum this strategy offers additional benefits. First, it offers new users clear examples of unwelcome behaviors both by making visible the deleted posts and by including moderator rationale for the deletion decisions, providing another means for potential posters to learn the values of the site and avoid

making similar mistakes. Second, this visibility benefits researchers such as ourselves. Just as this uncommon access to deleted comments and moderator rationale makes the community more knowable to its participants, open censorship provides more varied points of analysis to understand the community through research.

By examining how moderators deal with different kinds of censored posts, we can provide evidence about the community’s values. In this way, we test the idea that, in practice, EJ prioritizes the knowledge archive over conversation and uncover how their strategy works, to assess when this strategy will be appropriate for other online communities, allowing us all to learn from the Elitist Jerks.

Methods

To understand how moderators were using the Banhammer, we developed operationalizations for two concepts, perceived poster behavior and moderator behavior. We first drew on existing literature, the first author’s insider knowledge of the community, and a small sample of threads from a recent period of Banhammer posts (March–April, 2012; $n = 70$) to develop a content analytic scheme for both concepts. Having demonstrated that this scheme had interrater reliability, we then examined a larger sample of threads ($n = 200$), chosen to cover a period long enough to reveal the rhythms of the community’s collective life by analogy to recommendations from ethnographers (Fetterman, 1998). Finally, we conducted statistical analysis to examine how moderator behavior was associated with poster behavior, allowing us to demonstrate how the EJ moderation strategy works.

Content Analysis Scheme Development

A small initial sample ($n = 70$) of threads was taken from a recent period of Banhammer posts (March–April, 2012). The subject lines and references to specific rules were stripped from the data. Codes emerged from the content of the moderator comments and were informed by the first author’s insider knowledge regarding theorycrafting and WoW practices and periodic returns to the literature. Instances varied according to the explanation that the moderator provided for the action and the tone of the comments. For example, a straightforward post from a moderator might gently suggest an alternative mode of conversation (moderator behavior) while explaining the ill-effect being avoided (poster behavior):

Please use a PM [private message] to thank someone. It helps to avoid thread clutter.

Another type of post does not contain any advice for the contributor but uses his or her mistakes as an illustrative example framed as an entertaining diversion:

This week on EJ Geographic: We track the common Nerd-rage poster through the wild jungles of confusion, wrong information and finally into the valley of crybabies.

[. . .] [Quotes Original Post]

Breaking past the the [sic] treeline into the valley of crybabies, we find our poster in the last stage of existence. Here they experience a euphoric rage known as “internet superhero.” They talk tough, claim to have found the secret paradise of a “better site” with “more organized information” and of course “better moderators.” They continue to exhibit traits we saw back in the jungles of confusion and wrong information by misspelling words as they unleash this rage. Sadly, at this stage all one can do here is put them down. For those left behind on these dangerous forums the most humane way to do this is with the banhammer. In the valley of crybabies no one cares. . . .

Similar posts contained moderator comments that shamed contributors for their mistakes, implying the correct behavior, while being extremely critical of the individual:

You’re in a spreadsheet thread and you’re asking a question that could be answered by using the spreadsheet. That’s either a special kind of lazy or a special kind of dumb.

We first sought to develop codes that reflected moderators perception of poster behavior (infraction type). The categories of identified mistakes were indicated by two coders, with interrater reliability measured by Cohen’s kappa of 0.67, typically considered substantial agreement (Landis & Koch, 1977). We used the code “Noise” for posts accused of diluting the information content of the forum by creating “clutter.” Posts coded as “Noise” were by those moderators characterized as adding to the reading load of users without adding new, relevant, or legible information. We used the code “Damaging” for posts accused of undermining the knowledge system by submitting content that was incorrect or did not meet the standard of effort set by the community. Codes within this class distinguished between posts that included contributions that relying on pseudo- or nonevidence such as anecdotes (“Bad Argumentation”) and those that requested information that the contributor could deduce from available information (“Begging”). The Damaging code also described posts from anonymous accounts not linked to valid WoW profiles and therefore could not be held accountable and those that were identified as trolling or whining and risked derailing a productive conversation.

We developed two separate operationalizations of moderator behavior, one based on the content of their messages (criticism type) and the other on the “infraction points” applied (punishment value).

The first operationalization was based on understanding the type of criticism used by the moderator, based on criticism taxonomies from psychology (Baron, 1988; Peterson & Smith, 2010; Raver, Jensen, Lee, & O’Reilly, 2012; criticism type). *Destructive criticism* was defined as the use of sarcasm, threats, or humor at the contributor’s expense or the attribution of error to internal causes, such as laziness. *Constructive criticism* was defined as genuine encouragement to correct mistakes and repost; acknowledgment of the contributor’s effort and good intent; and a lack of sarcasm, threats, or unkind humor. Moderator comments too brief to

meet either criterion—usually rote recitation of the relevant rule broken—were coded as *Null*. The categories of moderator attitude were indicated by two coders, with interrater reliability measured by a Cohen’s kappa of 0.45, typically considered moderate agreement (Landis & Koch, 1977).

The second operationalization of moderator behavior was based on “infraction points” applied to infractions by moderators (punishment value). Banhammer threads included a point value from 0 to 10. Based on frequency of use and the first author’s experience, we recoded this to three levels. Posts receiving zero infraction points were labeled as receiving a “warning,” those receiving 0–9 points were labeled as an “infraction,” and those receiving the full 10 points were labeled as “Banhammer,” community parlance for an infraction resulting in an immediate and permanent posting ban.

In combination, our two operationalizations of moderator behavior provide triangulation. Infraction points have real impacts on poster capabilities, so they best show the revealed judgment of the moderator about the poster’s behavior, whereas the type of criticism leveled in the content of the moderation message provides greater insight into the reasoning applied by the moderators and thus into community values.

Content Analysis

We chose the length of our sampling period based on rhythms of the community’s collective life, by analogy to recommendations from ethnographers (Fetterman, 1998). A sample of 200 Banhammer threads was taken from those posted between January, 2010, and October, 2012, a season and a half of activity for EJ as determined by the flow of events in WoW. A “season” for EJ users is defined as the time between WoW software expansions. Each expansion introduces new content in the form of new game areas and new encounters, and new rules in the form of new playable characters and adjusted versions of existing character classes. Expansions renew the interest of veteran players by providing variety and attract new players by evening the playing field.

In the time between the release of the Cataclysm expansion (December, 2010) and the release of the latest expansion, Mists of Pandaria (September, 2012), the frequency of posts on the Banhammer has reflected the stimulating effect of new content. Increased Banhammer activity occurred immediately after new end-game content was introduced in game (December, 2009; December, 2010; June, 2011; and November, 2011) and when game mechanics were released in anticipation of major content launches (October, 2010). Within the past season and a half, the combination of these effects produced two major spikes in activity, covering the periods of January–April, 2010, and October, 2010, to February, 2011. Together, these 8 months (24% of the total months) account for 62% percent of the activity.

To reflect the typical activity of the Banhammer, the sampling method was adjusted to undersample from these

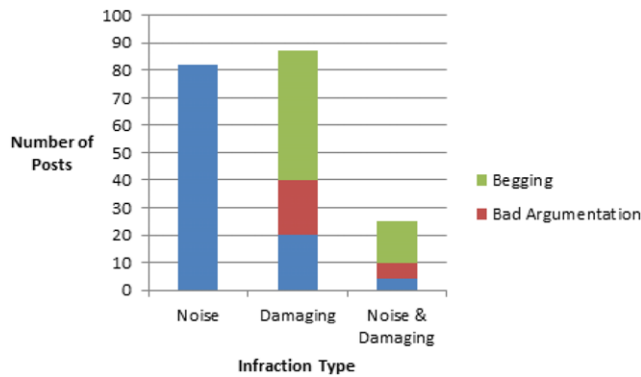


FIG. 1. Distribution of infraction types ($n = 197$). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

extremely prolific periods. 50% of the sample was taken from the other 26 months, in which activity is relatively stable, and 50% was taken from these two spikes. For the sample of 200 total posts, we took every 128th post from the postrelease months and every 77th post from the regular months. By choosing this time period and sampling in this manner, we ensured that we did not oversample special periods and gained insight into the long-term practices of the community.

Three of the 200 posts in the sample were found to be later reversed by the moderator and were removed from the set. Among the remaining, three were moderator identification of “spambots” and were not coded. Our data set for the findings reported here was thus 194 posts.

Findings

Coding the Banhammer posts by infraction type revealed an equal distribution of “noise” and “damaging” posts, as show in Figure 1. A small fraction (~13%) of posts were coded as both noise and damaging. In their commentary on these posts, moderators identified several separate violations. Among the posts coded as damaging or noise and damaging, the most common subcode was “begging” (~64%), with fewer being coded as “bad argumentation” (~23%), and the remainder divided between those posting anonymously or accused of whining or trolling.

Moderator comments were often so brief or rote as to be coded as neither constructive nor destructive criticism, as shown in Figure 2. Among the more substantial moderator comments, most (~36% of total) were destructive criticism. Constructive criticism was the least common mode of moderator comment (~14%).

Our analysis sought to understand the community’s priorities by assessing how different kinds of infractions were treated; those treated with destructive criticism or high infraction point totals would indicate the behaviors the moderators judged worst, whereas those treated with constructive criticism or lower infraction point totals would indicate the behaviors the moderators judged less harshly.

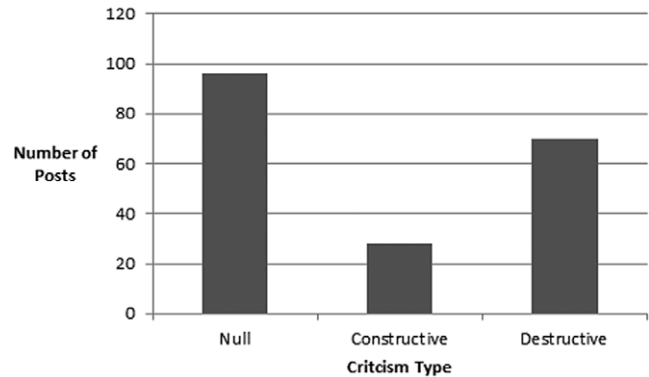


FIG. 2. Distribution of criticism types.

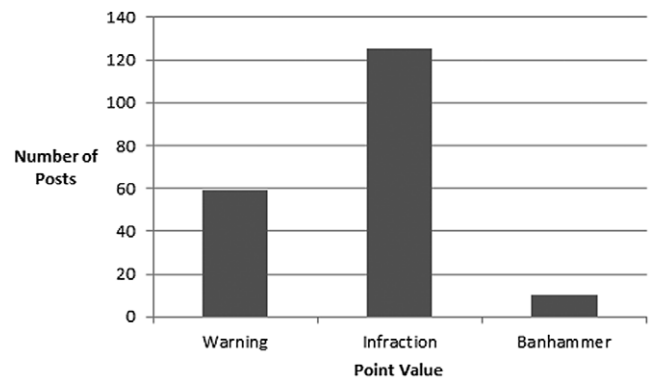


FIG. 3. Distribution of punishment values.

A chi-squared test of two independent proportions was used to determine the association between infraction type and criticism type. The association was statistically significant ($p < 0.05$), albeit with a small effect size (Cramer’s $V = 0.186$). We interpret this to mean that noise posts were the most likely to be treated with constructive criticism, whereas posts making mistakes or damaging alone were overrepresented among those treated with destructive criticism.

We also examined the relationship between infraction type and punishment value, our second operationalization of moderator behavior. Figure 3 shows the distribution of the point values. Most threads applied some point value (~64%), with a permanent ban (Banhammer) being the least common outcome (~5%).

A chi-squared test of two independent proportions was used to determine the association between infraction type and punishment value. The association was statistically significant ($p < 0.01$), with a small effect size (Cramer’s $V = 0.303$). Noise posts were most likely to receive a warning, whereas posts making damaging mistakes were most likely to receive an infraction value, and most (92%) of those with both mistakes to receive an infraction value. Damaging posts were highly overrepresented among the Banhammer values. This is consistent with the previous

finding regarding the attitude applied to damaging and noise posts, further illustrating a pattern in which damaging posts are treated by moderators as more deserving of punishment, whether in the form of harsh language or days banned from contributing.

Discussion

Among the types of contributions addressed in the EJ Banhammer, damaging is likely to be the most familiar to online community moderators. These posts are removed from their original contexts and vilified by moderators because they derail conversations and are unpalatable to valued, core contributors. The concerns and principles of the EJ community are specific to its aims; the need for hard data and scientific argument to make in-game decisions produces a set of values not likely shared with communities prioritizing emotional support or arguing matters of opinion. However, the idea that community managers should address posts that deviate from community values—such as a respect for the scientific method that produces generalizable conclusions—is not specific to this community. Contributions that violate core principles endanger the space created by core contributors, and their presence risks losing these valued members.

Trolling behavior is an extreme example of this type of unwanted contribution. By censoring trolling, moderators disrupt intentional attempts to derail the conversation. The other types of posts in this class, defined here as bad argumentation and begging, may not be intentional attacks on the flow of community discussion but similarly endanger the existing and valued conversation. Bad argumentation posts introduce conclusions lacking the evidence or argumentation other contributors require to formulate a response. Begging posts ask contributors to shift focus from the construction of generalizable conclusions to the production of specific answers unlikely to be of future use. No matter how alien the concerns of this community, this responsibility of community managers to protect the valued core conversation is recognizable.

The censorship of noise contributions, however, does not have a clear parallel with typical community management or moderation priorities. Unlike damaging posts, noise posts do not seem to violate the core principles of the conversation valued in this community. For example, the reason that saying “thank you” is identified as unwelcome is not because this community has a moral objection to gratitude; rather, the moderators ask that these contributors use a different channel for these messages, one that does not add to the reading load of members looking for substantive posts. Noise contributions are removed not because they risk offending the core contributors that drive activity and interest in the site but because they interfere with the act of reading and information seeking. This priority in moderator censorship is specific to online communities that recognize and value the needs of readers coming to the conversation for its informational content. As a community management

principle, censoring noise posts is a move away from treating discussion boards as vehicles for conversation toward a genre of knowledge resources.

Our concern in analyzing moderator activity in this community was to distinguish between those types of contributions that disrupt conversation and those that disrupt the knowledge resource. However, there are no clear distinctions among the types of damage done by unwanted posts. For example, bad argumentation posts disrupt the conversation by failing to provide the signals necessary to sustain scientific debate while simultaneously adding unsubstantiated or misleading information into the knowledge resource. There is often overlap across negative outcomes from unwanted posts. Interesting conclusions can be drawn from where these outcomes diverge. Our analysis is concerned with how moderators address posts that undermine one or both of the outcomes of a discussion board. We classify damaging posts as undermining both. By violating the community’s core principles, these posts disrupt the conversation. By contributing posts that, at best, do not serve the reader’s information needs, such as requests for specific answers, or, at worst, add unsubstantiated and misleading claims, these posts dilute the information content and undermine the accuracy and reliability of the knowledge resource. The severity of moderator responses to these contributions is consistent with this classification. We found that moderators were most likely to apply greater point values to damaging infractions, adding tangible punishment in the form of days muted from contributing, and to treat these contributors with harsh, destructive criticism.

Reserving the harshest commentary for those contributors who violate community principles is strategic behavior on the part of moderators. This strategy assumes that contributors who act counter to the values of core contributors are least likely to be reformed from bad to good contributors; rather, they can serve as clear examples to new members of unwanted behavior and an ongoing justification for moderator vigilance. Contributors who commit “noise” infractions, however, are not acting counter to the community’s values but rather have made an error in the proper phrasing and placement of content in this forum. These contributors more often receive warnings than punishments for their infractions, and moderators responding to their posts provide guidance and encouragement directing them to the correct way of participating in the forum. Future posts from these contributors might not be visible; rather, many of the “thank you” posts are instead routed to private channels to the other members, retaining their conversational nature but not cluttering the archive.

When Would the EJ Strategy be Appropriate?

The EJ forum is an example of an online community with clear motivations to curate a user-driven conversation into a quality information resource. Members of the community, and the larger theorycrafting and competitive *World of Warcraft* communities, accept the high threshold of quality

required to contribute to the forum and the harsh policies of its moderators, in part because it directly benefits their external goal to increase their performance in the game. What is not immediately clear is what elements of the EJ community are necessary for this knowledge resource-over-conversation strategy to succeed and which are incidental.

That EJ is a gaming community is incidental; this forum's relationship to *World of Warcraft* is only one type of possible connections between the circulation of knowledge in a discussion space and application to external performance and decision making. What is necessary is the external aim of the information. Contrasted against communities of interest or support, in which the circulation of information is an end in itself, communities with an external aim, such as communities of practice, would benefit from a moderation strategy that values the resulting information resource above the generative process of conversation. Along with this is the criterion that there be a substantial cost to bad information; the external aim must have some associated risk. In the case of EJ, bad information can form the basis of bad in-game strategies, leading individual members to perform badly, spend in-game capital on inferior gear, and risk their positions in highly competitive teams. In other communities of practice, unreliable information carries similar risks, from loss of face to loss of position over flawed decision making.

The last criterion necessary for a discussion community to benefit from this moderation strategy is the ease of creating a quality information resource in another medium. Information genres such as wikis allow communities to organize and share knowledge in a channel parallel to the conversation. For example, EJ briefly maintained a wiki, The Theorycrafting Think Tank, in parallel with the discussion forums, but the wiki was eventually abandoned, leaving only the discussion forum as the location of knowledge creation and organization. The EJ administrators explain, "It turns out that maintaining two separate copies of authoritative information is too much of a hassle for the majority of contributors" (Theorycrafting Think Tank, 2008). Wikis do offer affordances more appropriate to organizing and accessing information than a discussion forum, but it is the discussion itself that creates the information. For communities that create information and knowledge through conversation, it is more appropriate to curate the conversation into the knowledge resource than to create and maintain a parallel resource. For example, the math community uses a question-and-answer site to discuss research-level problems, accepting the limitations of that medium because it facilitates the discussion necessary for knowledge creation (Tausczik, Kittur, & Kraut, 2014). Especially when currency of information is particularly valuable, the parallel channel requires constant maintenance. Parallel resources such as wikis require additional expenditure of effort on the part of contributors; curating a discussion instead requires additional effort on the part of moderators.

The traditional approach to discussion board activity and moderator intervention is modeled in Figure 4. For a




FIG. 4. Conversation-centric discussion activity and moderator intervention.




FIG. 5. Knowledge resource-centric discussion activity and moderator intervention.

community in which conversation is the goal (e.g., an emotional support community), design decisions and moderator intervention prioritize the conversion of readers to contributors. Converting readers to contributors is necessary to offset the gradual exit of long-term members and to grow the community. Combined with moderator interventions that promote activity from existing contributors, this approach increases the quantity of posting content that in turn attracts readers, traffic, and new potential contributors.

We contrast this approach to the cycle of discussion board activity and moderator intervention observed in the EJ community, modeled in Figure 5. In these environments, moderator intervention does not promote the conversion of readers to contributors. Rather, moderator activity is likely to *decrease* the total number of posts (or posters) by selectively removing those that do not meet the high standard of quality set by the community. The only increase moderators prioritize is the likelihood of the site being useful to readers, by ensuring the topicality, accuracy, intelligibility, and reliability of posts. The natural loss of members over time—and the likely loss of additional members who react negatively to this style of moderator intervention—is offset by the existence of an external interest that drives new readers to the site and necessitates the high threshold of content quality.

To extend this model to other communities, it is necessary to identify circumstances in which information seeking might be at odds with other purposes of the information resource. For example, particular ways of curating archives for prioritizing specific research aims disrupts the original order of the collection (Henry, 1998; Schellenberg, 1956).

The 2013 decision of *Popular Science* (LaBarre, 2013) to remove commenting features from their online articles reflects a similar concern; carving out an open conversation space for explaining and discussing scientific findings and their implications ended up reducing the clarity of the finding themselves. The function of *Popular Science* as an information resource was threatened by the presence of conversation, so conversation was removed. Scholarly communication is a familiar example of an information curation project that separates seeking or use from creation; through journal databases and library reference services, user access is achieved without disrupting the resource itself, and the conversation within the community occurs through highly structured channels.

Conclusions

We have described a type of moderator intervention appropriate to online communities that seek to provide intelligible, accessible, and reliable information. The style of moderation found in the EJ community is an extreme example but represents a strategy generalizable across a wide variety of communities. We suggest the following core criteria under which online communities will benefit from moderator intervention that sacrifices conversation in favor of information resource quality: first, that there is an external driver for traffic and interest in the site, and, second, that the knowledge circulated in the site can impact external performance. In the case of EJ, the first criterion is satisfied by the enormous success of *World of Warcraft* in attracting and maintaining a large community of players. The second criterion is satisfied by the impact of theorycrafting on in-game play in *World of Warcraft*, such that accurate information can provide a strategic advantage in game, whereas inaccurate information can negatively affect the player's in-game capital and position in highly competitive groups.

Although many important kinds of online communities do not share these characteristics, including emotional support communities and communities struggling to attract participants, these criteria are shared with a number of different types of online communities: corporate knowledge-management systems, communities of practice, and scholarly communication. In the case of scholarly communication, this style of community management already exists. There are high barriers to contribution, clear external drivers to participate as a contributor, and valued external outcomes for knowledge gained. Scholarly communication goes beyond the discussion board medium we consider here, but the parallels illuminate the concerns of both venues, to create knowledge through conversation while prioritizing the assurance of the reliability, intelligibility, and accessibility of that knowledge. For those communities that provide external drivers to participate in and use a knowledge-management system, the quality, accessibility, and actionability of the information found should be the primary concern of administrators. For communities of practice that use online spaces to share and create knowledge among

engaged experts, moderator intervention that supports information quality will also be appropriate.

Conversation has been recognized as an important tool for knowledge management, particularly as a means to acquire knowledge from a community (Wagner, 2006); the moderation strategy introduced here is complementary, helping to refine and present the knowledge in conjunction with search engines. Most research in online communities has focused on attracting participants once a conversation is underway, but it is possible for moderators to enforce a high standard of quality that benefits readers with specific information needs. Future research should explore how and when a moderation strategy focused on obtaining a high-quality knowledge resource might undermine the conversation that is generating the knowledge, especially in communities with few outside drivers. Search engine and information retrieval researchers might wish to learn from the strategies to deemphasize results that look like noise or damaging posts as well as seeking to understand the impact of these strategies on the findability of knowledge. As more communities prioritize not only soliciting but also managing, organizing, or curating their members' knowledge, the EJ strategy of open censorship targeted at removing both damaging and noise posts to maintain the quality of the knowledge resource is an exemplary model upon which other communities might draw.

References

- Baron, R.A. (1988). Negative effects of destructive criticism: Impact on conflict, self-efficacy, and task performance. *The Journal of Applied Psychology*, 73(2), 199–207.
- Bishop, J. (2007). Increasing participation in online communities: A framework for human–computer interaction. *Computers in Human Behavior*, 23(4), 1881–1893. doi: 10.1016/j.chb.2005.11.004
- Chang, S.L. (2009). Information research in leisure: Implications from an empirical study of backpackers. *Library Trends*, 57(4), 711–728.
- Ebner, M., Holzinger, A., & Catarci, T. (2005). Lurking: An underestimated human-computer phenomenon. *MultiMedia*. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 12(4), 70–75. doi: 10.1109/MMUL.2005.74
- Edwards, P.N. (2010). *A Vast Machine: Computer Models, Climate Data, and the Politics of Global Warming*. Cambridge, MA: MIT Press.
- Fetterman, D. (1998). *Ethnography: Step-by-Step*. Thousand Oaks, CA: Sage Publications.
- Forums Rules. (2010). *Elitist Jerks*. Retrieved from https://web.archive.org/web/20120508013504/http://elitistjerks.com/t77/a54-forum_rules_updated_7_14_10/
- Fulton, C. (2009). Quid pro quo: Information sharing in leisure activities. *Library Trends*, 57(4), 753–768.
- Godin, S. (2008). *Tribes: We Need You To Lead Us* (p. 151). New York: Portfolio.
- Goffman, E. (1959). *The Presentation of Self in Everyday Life*. New York: Anchor Books.
- Hansen, D.L. (2009). Overhearing the crowd: An empirical examination of conversation reuse in a technical support community. In *Proceedings of the Fourth International Conference on Communities and Technologies* (pp. 155–164). New York: ACM Press. doi: 10.1145/1556460.1556484
- Henry, L. (1998). Schellenberg in cyberspace. *The American Archivist*, 61(2), 309–327. Retrieved from <http://archivists.metapress.com/content/f493110467x38701/fulltext.pdf>
- Jenkinson, H. (1965). *A Manual of Archive Administration* (2nd ed.). London: Lund Humphries.

- Kraut, R.E., & Resnick, P. (2012). Building Successful Online Communities: Evidence-Based Social Design (p. 312). Cambridge, MA: MIT Press.
- LaBarre, S. (2013). Why we're shutting off our comments. PopularScience.com. Retrieved from <http://www.popsci.com/science/article/2013-09/why-were-shutting-our-comments>
- Lampe, C., Wash, R., Velasquez, A., & Ozkaya, E. (2010). Motivations to participate in online communities. In Proceedings of the 28th International Conference on Human Factors in Computing Systems—CHI '10 (pp. 1927–1936). New York: ACM Press. doi: 10.1145/1753326.1753616
- Landis, J., & Koch, G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159–174. Retrieved from <http://europemc.org/abstract/MED/843571>
- Millen, D. (2000). Community portals and collective goods: Conversation archives as an information resource. In Proceedings of the 33rd Hawaii International Conference on System Sciences, 00(c) (pp. 1–9). Los Alamitos, CA: IEEE Press. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=926713
- Millington, R. (2012). *Buzzing Communities: How To Build Bigger, Better, and More Active Online Communities*. London: FeverBee.
- Paul, C.A. (2011). Optimizing play: How theorycrafting changes gameplay and design. *Game Studies*, 11(2). Retrieved from <http://gamestudies.org/1102/articles/paul>
- Peterson, K.M., & Smith, D.A. (2010). To what does perceived criticism refer? Constructive, destructive, and general criticism. *Journal of Family Psychology*, 24(1), 97–100. doi: 10.1037/a0017950
- Preece, J., & Shneiderman, B. (2009). The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on Human-Computer Interaction*, 1(1), 13–32. Retrieved from <http://www.cs.umd.edu/~ben/papers/Jennifer2009Reader.pdf>
- Raver, J.L., Jensen, J.M., Lee, J., & O'Reilly, J. (2012). Destructive criticism revisited: Appraisals, task outcomes, and the moderating role of competitiveness. *Applied Psychology*, 61(2), 177–203. doi: 10.1111/j.1464-0597.2011.00462.x
- Schackman, D. (2010). Commons or gated community? A theoretical explanation of virtual community and the example of Craigslist. *Journal of Community Informatics*, 6(2). Retrieved from <http://ci-journal.net/index.php/ciej/article/viewArticle/457>
- Schellenberg, T. (1956). *Modern archives: Principles and techniques*. Chicago: The Society of American Archivists. Retrieved from <http://files.archivists.org/pubs/free/ModernArchives-Schellenberg.pdf>
- Stebbins, R.A. (1982). Serious leisure: A conceptual statement. *Pacific Sociological Review*, 25(2), 251–272.
- Takahashi, M., Fujimoto, M., & Yamasaki, N. (2003). The active lurker: Influence of an in-house online community on its outside environment. In Proceedings of the 2003 International ACM SIGGROUP Conference on Supporting Group Work (pp. 1–10). New York: ACM Press. Retrieved from <http://dl.acm.org/citation.cfm?id=958162>
- Tausczik, Y., Kittur, A., & Kraut, R. (2014). Collaborative problem solving: A study of Mathoverflow. In CSCW '14 Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (pp. 355–367). New York: ACM Press. Retrieved from <http://www.cs.cmu.edu/~ylataus/files/TausczikKitturKraut2014.pdf>
- Taylor, T.L. (2006). *Play Between Worlds*. Cambridge, MA: MIT Press.
- Theorycrafting Think Tank. (2008). *Elitist Jerks*. Retrieved from https://web.archive.org/web/20130817075125/http://elitistjerks.com/t47/t21302-theorycrafting_think_tank/
- Thomas, D. (2009). Scalable learning: From simple to complex in *World of Warcraft*. On the Horizon, 17(1), 35–46. doi: 10.1108/10748120910936135
- Vasilescu, B., Serebrenik, A., Devanbu, P., & Filkov, V. (2014). How social Q&A sites are changing knowledge sharing in open source software communities. In CSCW '14 Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (pp. 342–354). New York: ACM Press. doi: 10.1145/2531602.2531659
- Wagner, C. (2006). Breaking the knowledge acquisition bottleneck through conversational knowledge management. *Information Resources Management Journal*, 19(1), 70–83. doi: 10.4018/irmj.2006010104
- Wei, K., Crowston, K., Li, N., & Heckman, B. (2014). Understanding group maintenance behavior in free/libre open source software projects: The case of Fire and Gaim. *Information & Management*, 51(3), 297–309. doi: 10.1016/j.im.2014.02.001
- Wenger, E., White, N., & Smith, J.D. (2009). *Digital Habitats: Stewarding Technology for Communities*. Portland, OR: CPsquare.
- Wright, S. (2009). The role of the moderator: Problems and possibilities for government-run online discussion forums. In T. Davies & S.P. Gangadharan (Eds.), *Online Deliberation: Design, Research, and Practice* (pp. 233–242). Stanford, CT: CSLI Publications.
- Wright, S., & Street, J. (2007). Democracy, deliberation and design: The case of online discussion forums. *New Media & Society*, 9(5), 849–869. doi: 10.1177/1461444807081230
- Yates, D., & Wagner, C. (2010). Factors affecting shapers of organizational wikis. *Journal of the American Society for Information Science*, 61(3), 543–554. doi: 10.1002/asi.21266