

Michelle P. Modic. Anarchy vs. Authority: Attitudes Towards Centralized Control in Social Tagging Systems. A Master's Paper for the M.S. in L.S degree. March, 2014. 39 pages. Advisor: Jane Greenberg

This study examines the perceptions of users towards centralized authority in social tagging systems. Research shows that scholars and individuals overseeing social tagging projects see a number of pros and cons towards allowing users more control. This study looks at another area of tagging, specifically what the users themselves think, by using content analysis of a message board thread on the topic of moderation in a particular social tagging system. The findings indicate that, like the scholars and individuals overseeing social tagging projects, their views are fairly mixed. Some want a system with strong moderation to weed out malicious users and fix mistakes. Others don't trust moderators not to serve their own ends, and want a more open and creative system that they ultimately see as more honest. The challenge for future social tagging system moderators is to please both groups without alienating either.

Headings:

Tags (Metadata).

Metadata.

Folksonomies.

Internet users.

ANARCHY VS. AUTHORITY:
ATTITUDES TOWARDS CENTRALIZED CONTROL IN SOCIAL TAGGING
SYSTEMS

by
Michelle P. Modic

A Master's paper submitted to the faculty
of the School of Information and Library Science
of the University of North Carolina at Chapel Hill
in partial fulfillment of the requirements
for the degree of Master of Science in
Library Science.

Chapel Hill, North Carolina

April 2014

Approved by

Jane Greenberg

Table of Contents

BACKGROUND STATEMENT	2
LITERATURE REVIEW	4
DATA COLLECTION METHODOLOGY.....	21
RESULTS	23
DISCUSSION.....	29
CONCLUSION	34

Background Statement

The tension between anarchy and authority underlines much of humanity's political thoughts and discourse. Are we more afraid of our fellow human beings, of the strong taking advantage of the weak? Or are we more afraid of authority, corruption, and possible authoritarianism? However, some areas lean more towards individualistic freedom and expression than others, and the internet is one such place (Zittrain). People who have few problems with authority in their daily lives can balk at it online. There's simply a pervasive aura of freedom about the internet, and many can become incredibly defensive when this is challenged in even the most insignificant ways. The internet may be our final frontier, the closest thing most people have to a "wild west."

The question for people and organizations becomes how to harness the power of the internet without imposing so much control that they alienate the majority of their potential users or customers. Those who implement social tagging systems or crowdsourced metadata projects inevitably stumble upon this quandary. How much should they moderate their participants? Too little moderation and useless, profane, and otherwise offensive tags could run amok. Too much moderation, however, and not only is creativity stifled but the user base largely run off. No project can survive without a dedicated user base.

Of course, not all users want the same things. Some want high levels of moderation and guidance, as it provides structure and eliminates the risk of having the entire system taken over by the useless or profane. Others see even the smallest exercise of control as unforgivable censorship, and campaign for the ability to tag things however they see fit. The latter group believes that the good intentions of the many will drown out the bad intentions of the few, and that the useful tags will rise to the top fair and square. The trick for project organizers therefore is to find a way to keep both groups happy or, failing that, to avoid driving either group away.

1.1 Purpose

This paper provides a snapshot into the opinions of the users of a social tagging system and examines their attitudes towards authority. The resulting dialogue and range of opinions can inform future projects on how best to assert their authority, and what reactions they can expect

Literature Review

While there are many papers that discuss social tagging, few touch on the attitudes towards centralized authority found in the participants. Much of the literature instead discusses the efficacy of crowdsourced and other social tagging projects- weighing the pros and cons and determining how best to handle the data once it has been generated. Some of the studies do focus on the benefits and drawbacks of allowing the crowd free reign over a tagging system, and by extension the positive and negative effects that greater moderation can have upon such a system. By looking at the thoughts and data contained within the literature, we can extrapolate the views of moderators and project heads on anarchy versus order in a social tagging system. This literature review will broadly cover the following categories: first general background about attitudes towards authority on the internet at large, followed by general surveys of social tagging, then some looks at the issues with social tagging systems, and finally an article about authority in social tagging environments.

1.2 The Internet and Authority

First, to set a backdrop for the stage upon which all social tagging projects play out, we look at the internet as a whole. The arguments of anarchy versus authority routinely play out upon the larger scale of the internet itself, with everything from

federal legislation to blog posts attempting to weigh in on the matter. One of the best concise encapsulations of all the different issues at hand comes in M. Christopher Riley's piece "Anarchy, State, or Utopia?: Checks and Balances in Internet Governance." In this piece, Riley runs over the dominant issues in the freedom or governance debate where the internet is concerned, and looks at the issues from three points of view: governments, businesses, and individual users. The differences between these points of view are enlightening, especially when they can so easily be transposed into the social tagging sphere. Riley identifies "five high-order Internet policy goals: freedom of speech and association; privacy; security, including cyber security; economic growth; and social order" (p. 12). While all five of these relate to social tagging system in their own ways, the first and last goals are both highly relevant and influential. It is, after all, the desire for free speech that underlines much of the anti-establishment attitudes in social tagging projects. Social order, on the other hand, is one of the primary reasons that users do want moderation: they want a way to delete or otherwise remove things which they find offensive or socially unacceptable.

Riley goes on to show the benefits and drawbacks of either extreme for the internet as a whole, which again directly apply to the microcosm of social tagging. A system that is too authoritarian, he points out, would be very effective and maintain a high level of social order. It would not, however, promote free speech and in fact would very likely actively curtail it. This would lead to a highly effective system that was sorely lacking in creativity. A user-centric system, by contrast, would promote an incredibly high level of free speech. Such a system, however, would be rife with

security concerns and individuals promoting themselves and their interests at the cost of others. Indeed, as Riley puts it, “many would frustrate states’ and other users’ interests in protecting cultural and social values” (p. 16). So it seems that the best strategy, for tagging projects as well as the internet at large, would be to find some sort of balance between the two.

1.3 Social Tagging Overviews

Moving into the realm of social tagging itself, one finds a number of articles discussing social tagging as a whole. One such article is “Survey on Social Tagging Techniques” by Gupta, Li, Yin, and Han. This article, as the title suggests, surveys various issues surrounding social tagging systems- including how and why such systems operate as they do. Its comprehensive overview of the different methods and motivations inherent in tagging provides a helpful vocabulary to use when examining both tags and those who use them. They begin by summarizing the different tagging systems employed by popular websites and social networks, before going on to explain that tags and folksonomies exist in large part to breathe life into the otherwise static nature of rigid externally-imposed taxonomies.

They go on to document different motivations users have for tagging, including social signaling, opinion expression, organization, and attracting attention. All of these motives can be seen influencing the users in my data pool, with organization tending to lean more towards authority and the others (especially opinion expression) leaning towards anarchy. The different kinds of tags are documented next, and these include: content-based, attribute, subjective, organizational, factual, personal, and self-referential. The debates over which of these are valid categories of tags, and

which should or should not be employed, becomes a major component of the debate in the message board thread. Those who wish to allow all kinds of tags naturally favor less or no moderation, while those who feel only certain categories should be allowed favor a more heavily moderated system.

This then moves the article to discuss an inherent divide in the way taggers operate: categorizers versus describers. According to the authors, categorizers use smaller and more formalized tagging vocabularies, while describers use broader and more varied vocabularies. This divide, perhaps, can be seen when examining discussions and debates surrounding social tagging systems. Finally, the authors briefly mention a linguistic classification of tags. Basically, most tags can be sorted into only eight categories: functional, functional collocation, origin collocation, function and origin, taxonomic, adjective, verb, and proper name. Of these, functional, taxonomic, and adjective tags are the most prevalent in the data. The authors move to discussing tag generation models and mathematics before moving into tag analysis and visualization.

To close their paper, they discuss a couple of prevalent social tagging problems. The first of these is spamming, which has been examined many times in the literature and is further discussed below. Their brief overview simply mentions that spamming can and does happen, and runs through a few rudimentary methods for identifying and dealing with spam attacks. The other problem they discuss is that of “canonicalization and ambiguities” as they call it. In short, this is the problem that arises when users use synonyms, varying forms of the same word, and words with ambiguous meaning. Possible solutions are again briefly discussed before the paper

moves to its conclusion. The vocabulary so helpfully outlined in this article can help us to discuss the issues and tendencies we see in the data, and the problems they outline set the stage for the literature to follow as well as demonstrating some of the organizational dangers of a system without enough moderation.

Another paper that gives an overview of social tagging, “Perspectives on Social Tagging” by Ding et al., takes a more mathematical approach. While this paper is chiefly concerned with applying its own metrics to existing tagging sites, it does begin with a helpful discussion of various approaches to tagging. They point out how taggers enjoy the social environments that social tagging sites place them in, and that tags often enable like-minded individuals to come together and form communities based on common interest. As they say, “tagging works because it strikes a balance between the individual and the social” (p. 2389). Both the individualistic and the community-oriented aspects of tagging are therefore important to the taggers themselves. Motivations can range from the selfish to the altruistic, and all along the spectrum in between. However, as one article they cite notes, not all motivations are positive. It point out that people can also use tags to promote their own personal interests or push their political views. This then begs the question of would greater moderation solve this problem, and if so should it? Users still grapple with this question, as the data will demonstrate.

Another important aspect of user-created tags is language. As the article says; unlike a formalized metadata scheme, user-created tag scheme or folksonomy “speaks the same language as users” (p. 2390). This enables greater browsing and findability among the target audience, because the words they use to search are then present on

the objects themselves. Of course, allowing tags to be completely user-generated can allow for troublesome ambiguities and redundancies. While a predefined tag list can solve this problem, the paper points out that one of the benefits of social tagging is its very controlled nature that allows natural language to be used in tagging as described above. However, according to the authors, the broader the context the less helpful natural language is. The more users a site has, the more varied their vocabularies will be, and the more standardized the tags will need to be in order to make sense to everyone. This issue of language, and the applications of natural language in small or large communities, will be an important one to keep in mind.

Finally, they conclude that tags have a valuable place in social networks. While the above issues are all relevant, their data shows that large enough communities will eventually ameliorate most of these issues on their own. For example, even absent any sort of controlled vocabulary a preferred tagging vocabulary will arise over time within the community. Items tagged according to convention will be more discoverable than those which are not, leading to a feedback loop that encourages staying within the societally-determined tagging framework. This suggests less need for moderation than one might expect. This sense of community can also lead to a local social culture, which can have a positive impact on tagging. All in all, hidden among the mathematics, this article demonstrates some of the upsides and downsides of a more anarchic approach to social tagging.

1.4 Opinions on Social Tagging

A less formal look into the minds and motivations of taggers comes from Rashmi's blog entry entitled "A social analysis of tagging." This blog post examines

the “why” of tagging, and how tagging helps people connect with one another. She begins by pointing out that browsing the internet can be a solitary experience, and that tagging can bring users into a social milieu without required any active socialization or conversation, which is a huge benefit for many. She also notes that tags can set trends, both within tagging itself and in the wider social milieu that those tags serve. Part of the fun of tagging, she muses, is the hope that she too might one day be a trendsetter. She goes on to say that tags lead to ad-hoc group creation, where these groups behave more like crowds than traditional groups.

This leads to a discussion of crowd behavior, based on the “wisdom of the crowds” in contrast to the historical discourse that painted crowds in an almost exclusively negative light. She points out that the four conditions that James Surowiecki said could lead to this “wisdom of the crowds” are all present in social tagging systems. The four conditions are: diversity of opinion, independence of members from one another, decentralization, and a good method for aggregating opinions. She does point out that the second and third conditions are somewhat dampened by the fact that taggers can see what tags others have applied and been influenced by them. In addition, the amount of decentralization varies by tagging system. But she believes that the four conditions are present enough to create the desired effect, in essence championing the crowds against a more centralized or authoritative system. She points out that thus far tags seem to her to be relatively free of the negative aspects of crowd behavior- namely mob mentality. How true her observation proves remains to be seen.

She goes on to espouse further benefits of tagging. Tagging, she says, is malleable. It refrains from imposing a rigid structure on the content and leads to self-expression. Indeed, tagging can lead to all kinds of ad-hoc collaboration and collective self-expression. Tagging also enables the fast transfer of information and ideas across the internet. As she puts it, “tagging helps in the spread of ideas, memes, trends and fashions.” This leads her to ask what role tags play in the ebb and flow of concepts, and note that said role is greater than zero. She goes on to detail two specific ways in which this can play out: a concept in search of a name, and a name in search of a concept. The former describes how a nameless concept can find itself named by social consensus by way of the tags applied to it. A popular term will emerge fairly quickly, and that term will be adopted as the accepted name for that concept. This demonstrates the power of naming that tags can exert. The second method discusses the way tags can help shape the conversation around a previously nebulous concept, until it starts to have more definite boundaries and definitions. This underlines the overall strengths and weaknesses of an anarchic and decentralized social tagging system, where tagging allows social coordination even when a concept is fuzzy, and yet because it never forces a decision, “you can reach a tagging frenzy even if the concept is ultimately rejected.” Her final conclusion regarding tagging systems is positive, seeing them as a valuable and useful tool. She notes, however, that when designing a tagging system one must understand both how it serves individuals and what sort of social formations it supports. This understanding is crucial to creating a functioning social tagging system, and this paper hopes to contribute in a material way to said understanding.

Tom Reamy's "Folksonomy folktales" sets out to look at what he believes are the misconceptions surrounding folksonomies, some of which have been demonstrated by the articles above, and to point out what he believes are the pros and cons of a decentralized and user-centric system. He first sets out to debunk the notion, seen in the Rashmi post above, that folksonomies are an example of the wisdom of the crowds. A key factor of the wisdom of the crowds, he says, is that no one is aware of anyone else's actions. This clearly is untrue for folksonomies and social tagging systems, where seeing what other users have done is a key component of how the system operates. He goes so far as to say that what actually happens in folksonomies isn't the wisdom of the crowds but the madness of the crowds, due to the prevalence of the bandwagon effect. He then goes on to debunk the myth that folksonomies are comprehensive classification schemes, saying that instead folksonomies are flat sets of keywords ranked by popularity.

He then goes on to discuss the supposed benefits of folksonomies, and the drawbacks inherent in these perceived advantages. He is particularly harsh on the notion that the decentralized nature of folksonomies is a good thing. He notes that folksonomy advocates seem to "devoutly wish" to avoid centralized authority, but says that the notion is itself flawed. As he puts it: "folksonomy sites do have a central authority, and it is the most oppressive and most dangerous type of central authority there is- the authority of the majority. Against the will of the people there is no recourse, no way of insuring the rights of the minority" (p. 8). This is certainly the most virulently anti-anarchic and pro-authority statement we have seen thus far in the literature. In Reamy's view, no matter what problems one might have with a

centralized authority, nothing is worse or more terrifying than the tyranny of the majority. He goes on to espouse the need for a central authority to which users might appeal, should they find themselves in the minority regarding a particular tag.

He concludes his essay by saying that he sees folksonomies as very useful, but only in certain situations and only with certain limitations. He believes that folksonomies need to be combined with taxonomies in order to impose a sense of order. He says that having the system suggest terms to the user will help keep the terms from running into the redundancy problems mentioned in other articles. He also suggests finding ways to incorporate folksonomies into faceted schemes which can be used to aid faceted search and navigation systems. In the end, folksonomies are not the revolutionary idea that others purport them to be, Reamy feels, but rather one useful possible tool among many. They can be useful when integrated into other systems, but left alone as free-roaming systems without an authority they lose their usefulness.

1.5 Issues with Social Tagging

While nowhere else in the literature are decentralized social tagging systems attacked quite so vehemently, many articles examine the problems and issues they raise as well as how to go about solving these problems. Guy and Tonkin, cited by Gupta et al. above, look into these problems in their article “Folksonomies: Tidying up Tags?” They believe that the number one problem people have with folksonomies is the inexact and often sloppy nature of the tags. These include the presence of compound words and personal tags, the lack of synonym and homonym control, and other similar issues. The resulting set of tags is both uncontrolled and chaotic. They

point out that not everyone sees this as a problem- that because tags exist to help the tagger it doesn't matter how or if they operate as a cohesive whole.

The authors see two key ways to improve the metadata created in folksonomies: educating users to add "better" tags and/or improving the systems to the same end (p. 6). The most common types of sloppy tags include: misspelt tags, compound word tags, tags that don't conform to singular vs. plural form conventions, personal tags, and single-use tags which only appear once in the database. There were also tags in various languages, as well as tags that had been transliterated from other languages and scripts in numerous and inconsistent ways. Educating the users would involve coming to some sort of consensus on these issues (choosing whether to use singular or plural nouns, etc.), and then following said consensus. On the internet, that's easier said than done. On the other side, improving the systems would primarily involve introducing more tag suggestion mechanisms. In these, the system recommends tags based on the ones a user has entered, in order to keep the vocabulary more consistent. The authors also suggest allowing for more discussion spaces where users can discuss tagging issues and possibly work out some of these issues themselves.

The discussion then moves to the notion of tidying up tags in and of itself. Namely: would cleaning up tags even be a good idea? As the authors point out, one of the main attractions of a folksonomy is its openness. Imposing strict vocabulary control would impede that. Additionally, not everyone thinks the same way. Having a number of different tags on any given resource allows for more people to find it, because it increases the number of different viewpoints and thought processes that

link to it. It's less an issue of everyone tagging within some ideal of "correctness," but rather if at least one person has tagged something the same way another individual user would. As the authors point out, dialects differ among class, educational, and age lines; which further fractures the user base. In addition, different internet communities and subcultures use different jargon, and having their terms also present will enable other members of that community to find a resource even if the tag might look nonsensical to an outside observer. If there are tags for everyone's different methods, then the tags have achieved their purpose of allowing for greater discoverability, even if the tags on any given resource might seem repetitive or sloppy when viewed in a vacuum. Getting rid of the idiosyncratic terms might help to create a more well-defined taxonomy, but is that the end goal of a folksonomy? In the end, this is the question the authors leave us with. Sloppy tags are one of the core faults that many find in folksonomies- but can this problem really be solved and if so should it? For an article that at first seems to be firmly on the side of centralized control, it ends uncertain on this issue and ends up providing a number of strong arguments in favor of decentralized anarchy.

Another article concerning itself with solving the problems of social tagging systems is "Combating Spam in Tagging Systems: And Evaluation" by Koutrika et al. Spam, according to them, is any malicious use of a tagging system to drive up hits for a particular page or resource, to push users away from a particular page or resource, or simply to badmouth someone or something by use of negative tags such as "evil." In their paper, the authors looked at the following components of the spam problem: how many malicious users a tag system can tolerate; whether anything can be done

about collusion between malicious users; what types of tagging systems are more prone to spam; whether limiting tags help combat the problem; whether encouraging more tags help combat the problem; the effectiveness of a moderator to combat spam; and whether some sort of system could be devised to weed out the “bad” taggers. However, as they point out, the notion of a malicious tag in and of itself can be difficult to quantify. Would any negative adjective count? Some seemingly negative tags might actually be highly useful or applicable for some users. They also note that no matter what system is employed to combat spam, malicious users can always collude to find a way around it.

After a detailed discussion of their methodology and the data derived from their research, the authors reached a number of conclusions. The first was that tagging systems are threatened not only by malicious users, but also by “lousy” ones who confuse the tagging system. This makes it more difficult to single out malicious users or malicious tags, when they can be hidden by the noise made by the “lousy” users. They also noted that, while practices such as limiting the number of tags per user did help to keep the amount of spam down, it also had a negative overall effect on the tagging system by hampering positive users. Any measures taken to prevent spam, therefore, must be careful not to constrain the “good” users of the tagging system.

Additionally, every anti-spam measure has its Achilles’ heel, and dedicated spammers will find it. Spammers will always try to exploit a tagging system, and sophisticated anti-spam measures might only lead to more sophisticated spammers. Finally, the true power in any tagging system lies with the users. If enough positive and productive users are tagging, their activity will often drown out and end up

shielding against the spammers. In other words, “the more tags generated by more responsible users, the better” (p. 33). All in all, this study thoroughly demonstrates the dangers malicious users present in any social tagging scheme. On the one hand, this would seem to support the notion that greater centralized authority is necessary. On the other hand, as they demonstrate, sufficiently motivated spammers will find a way to work around whatever authority system is in place. Centralized authority has to be careful not to destroy their own system by driving away or hampering the good users in an attempt to silence the bad ones. It comes back to the classic question of anarchy versus authority, and which leads to the better outcome. Ultimately, where spammers are concerned, the best strategy appears to lie somewhere in the middle.

1.6 Social Tagging and Authority

Finally, in his paper “Folksonomies and the New Order: Authority in the Digital Disorder,” Jens-Erik Mai explores the notion of authority and the role of the professional in the social tagging sphere. He situates folksonomies in the larger conversation about classification’s role in “creating order in the universe of knowledge” (p. 115). He mentions the foremost importance of naming in creating this order, and contrasts more formal traditional taxonomies with the pluralistic nature of folksonomies. Naming is a personal thing after all, and everyone names the world from their own point of view. Organizations and centralized authorities can impose a point of view, but absent that one ends up with as many different competing points of view as there are taggers in the system. This relates to the point made by Guy and Tonkin above, where every user understands the tags differently and having a more

diverse set of tags might make the system more usable to those with unique perspectives.

Mai goes on to point out that while folksonomies and social tagging systems are often likened to democracy, they are actually far more akin to libertarianism. Everyone's whims are allowed to flourish without any sort of centralized authority or voting process to determine which tags apply and which do not. The authority in these systems, therefore, rests in the users. However, that doesn't mean that professionals have no role in folksonomies. Rather, this shift to user-centric systems "requires (...) that we design systems that actually do facilitate the creation of order in the universe of knowledge in a responsible, democratic, and meaningful manner" (p. 119). Mai then describes the four progressive levels of social technologies: sharing, cooperation, collaborative production, and collective action. Folksonomies, he states, are at the sharing phase. They are still in a state of anarchy run rampant, with little cooperation or collaboration. Mai, like Guy and Tonkin, believes that have a space for users to discuss tagging and tagging issues would help foster group identity and lead to greater collaboration. He sees cooperation and collaboration as necessary tools for folksonomies to progress up the ladder from their current lawless state into a more useful tool in making sense of the "universe of knowledge."

He sums up the article by noting the contrast between the authoritarian nature of traditional models of knowledge organization and the more anarchic nature of folksonomies. He compiles a chart describing the major differences between the two. Traditional models value transparency, consistency, stability, and professionalism; while user-centric systems value inclusiveness, openness, collaboration, and

interpretation. The former has a difficult time understanding its users' needs, while the latter sometime struggles to get people involved. The clearest difference comes in the categories of "Naming" and "Authority." Traditional systems have objects named centrally by professionals, while social systems have objects named locally by users. The first takes its authority from external sources thereof, while the other generates its own authority from the user base and the users' trust. Both have their uses, and the job of information professionals going forward will be to determine the correct model for a given situation.

This examination of the literature demonstrates the variety of viewpoints on social tagging schemes, and the necessity or lack thereof of a centralized authority to exert control. Some see the crowds as posing the biggest problems for social tagging schemes, with authority and moderation as the best means to combat this. Others see too much authority as stifling the creativity that makes social tagging systems so vibrant, and looks to the crowds for wisdom and innovation. But the question is: how do the crowds themselves feel about this? Or, to put it another way, how do the users of social tagging systems feel about centralized authority in those systems? More research needs to be done to truly ascertain the needs, thoughts, and attitudes of the users regarding control and authority. The following case study attempts to shed some light upon this heretofore underrepresented viewpoint by examining the views of a particular subset of users.

1.7 Research Questions

1. How do users perceive authority and moderation in social tagging systems?

2. Is there more animosity directed at authority figures than at the masses or vice versa in the social tagging environment?
3. What are the strengths and weaknesses in each side's arguments in this environment?

Data Collection Methodology

The following research draws from the examination of an individual message board thread on NeoGAF.com of 358 posts debating the level of moderation and authority imposed over the new tagging system on Steam, using the content analysis research method. This thread was chosen, rather than one of the many on Steam's official message boards, because posters on a third-party site are more likely to be honest and less likely to be censored. The thread, titled "Valve curbs abusive Steam Tags after games tagged-'not a game,' 'hipster garbage'", discusses the new tagging system implemented on the Steam website and the sudden introduction of moderation into the tagging system by Valve, the company behind Steam. Originally users could tag games however they wanted, leading to a rash of snarky and humor-based tags in addition to more traditional ones. While many of the tags were in good fun, some were insulting and downright offensive. Additionally, many depended on in-jokes and gaming knowledge that might be confusing to those not in the know. Valve began to crack down on these tags, deleting both the offensive tags and the funny but seemingly useless ones. The thread touches on such topics as the appropriate amount of moderation for the tagging system, how tags should be used, which tags were the funniest, and whether or not the new tagging system was going to work out at all.

The researcher went through each post in the thread and took note of when various authority-related opinions were expressed. The numbers count the number of

opinions expressed, rather than the number of posts, simply because some posts touched on a number of very different things and would otherwise be hard to categorize. So the researcher took each commonly-expressed opinion that was found relevant, plus a few less common ones that had been expected to show up more often, and counted each occurrence thereof. Once the counts were done, the opinions were sorted into generally pro-authority or anti-authority sentiments, and those opinions that either didn't fit or which turned out to be irrelevant were tossed out. Then the researcher compared the total amounts of opinions expressed in both categories.

Results

The following results cover the views of the thread as a whole, before examining each side's views in closer detail. It finishes with a broad numerical comparison of authority views versus anarchy views. What follows is a brief explanation for every category that appears on the chart, in order to explain what each short phrase means. While many of the categories are self-explanatory, the following might need more of an explanation. "Tags are funny" refers to posts that simply commented on how amusing some of the joke tags were, and sometimes expressed support for those tags remaining intact in the system. "Negative tag actually useful" occurs when a poster proclaims that an otherwise negative tag like "hipster garbage" is actually useful to them despite its insulting nature, because they tend to like games that others would tag that way. "Don't only delete negative tags" comes from the belief that tags expressing negative opinions were being deleted while positive opinions were allowed to stay, and that this was unfair and misrepresented the feelings of the community.

"General cynicism" refers to negative or cynical views expressed about Valve and/or the game developers, believing that they would use the tagging system for their own ends at the expense of the users. "Tags should protect consumers" reflects the belief of some users that tags are needed to protect consumers from buying buggy or incomplete games without their knowledge. "Censorship/'free speech'" refers to

times when a poster deliberately invokes either concept by name. “Only owners of games” means that only owners of any particular game should be able to tag it, instead of allowing any and everybody to tag every game. This comes up because that’s the system Steam has in place for reviews- only owners can review. Since one’s steam account contains one’s purchase history, they can easily enforce this.

The next four categories all refers to “devs,” which is shorthand for game developers. These four views reflect whether or not tagging should be left up to the developers, and whether or not developers should be able to delete tags from their games at will. “Pro-” and “anti-quality judgments” reflect the debate between users as to whether tags should contain quality judgments of the games themselves or if those judgments should be left to reviews with the tags serving neutral, categorization purposes. Lastly, “complaining about trolls” refers to the belief that any open system can and will be abused by malicious users, or “trolls.” Users with this viewpoint usually blame the trolls for their inability to have a more open or un-moderated system.

1.8 Count of All Relevant Opinions

First, the following table shows how often each of the following authority-related opinions regarding the moderation of Steam’s social tagging system were expressed in the thread. Note, again, that each post in the thread can show up more than once if it espouses multiple different opinions. On the other hand, posts on other topics are not included at all.

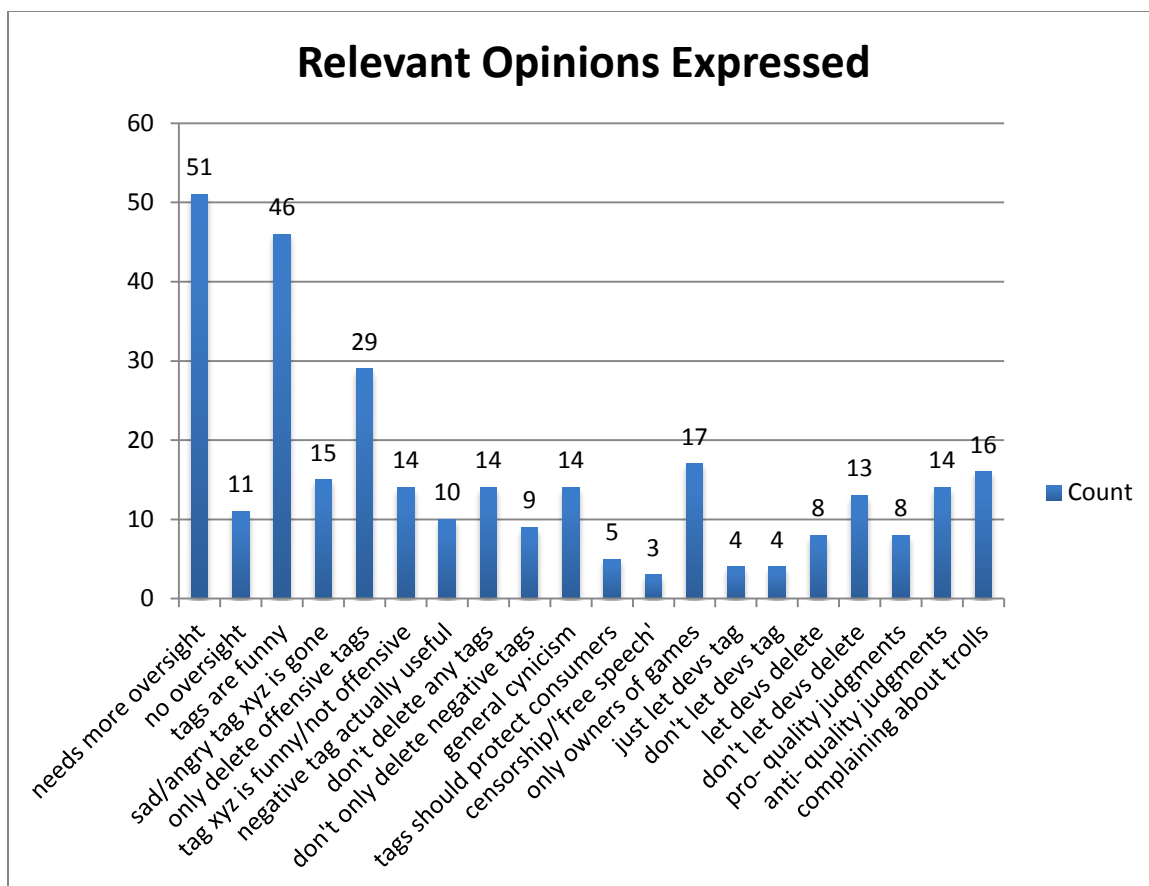


Figure 1: Number of Occurrences for Each Opinion

The most commonly-expressed opinion in the thread is the feeling that the system “needs more oversight.” This demonstrates the users’ desire for at least some level of centralized authority in order to help keep the system free of useless or offensive tags. Many even expressed surprise that the system hadn’t been more heavily-moderated from the start. General amusement at the funnier tags was next-most common, which is only surprising in that it wasn’t the most common opinion. Such users were generally lighthearted, but many also expressed pro-anarchy views or at the very least decried any attempt to remove the funnier tags. Given that, it naturally follows that the next most-common viewpoint was “only delete offensive tags.” Such users acknowledged that some level of authority needed to exist, but wanted it stripped

down to the bare minimum possible. Many of these posters additionally said that only one or two tags out of hundreds fell into the “offensive” category, making them far more in favor of anarchy than authority.

1.9 Count of Pro-Authority Opinions

The following chart examines the different views which have been classed as generally pro-authority or pro-moderation:

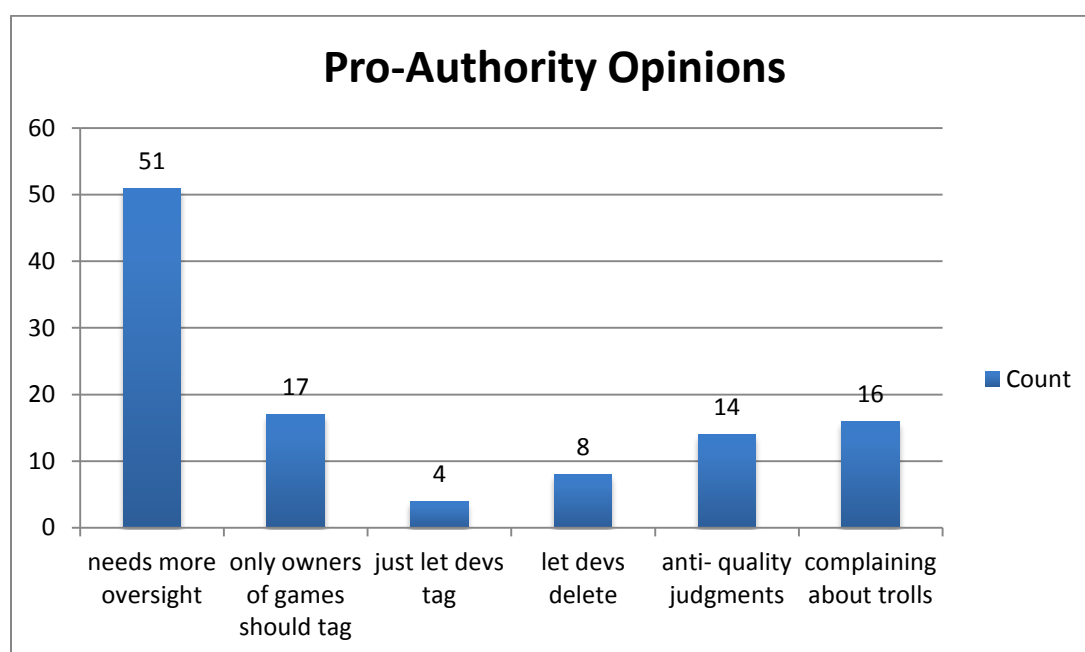


Figure 2: Number of Occurrences for Pro-Authority Opinions

In addition for the need for oversight, the most common views here wish to limit tagging in one way or another- by who can tag or how they should tag. When taken in conjunction with the “complaining about trolls” category, it seems that the dominant viewpoint here is that users cannot be trusted to successfully tag games themselves, and need to be policed and/or limited in some way in order to maintain the integrity of the tagging system.

1.10 Count of Anti-Authority Opinions

This chart shows the number of times each relevant anti-authority or anti-moderation opinion was expressed:

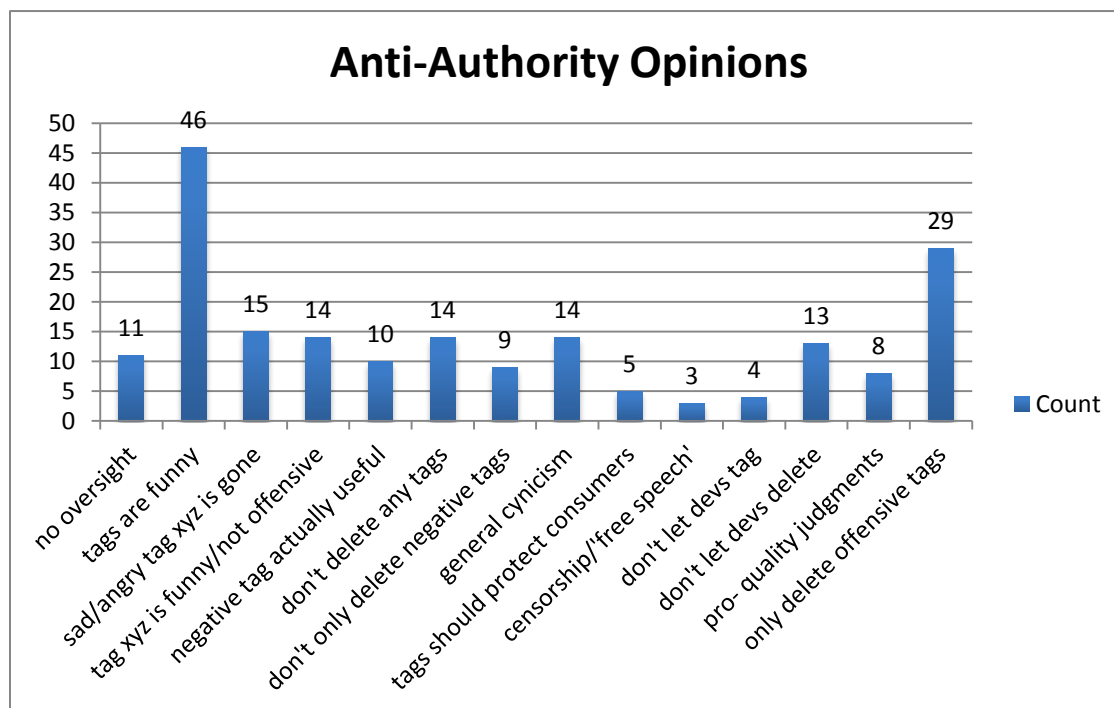


Figure 3: Number of Occurrences for Anti-Authority Opinions

Here, once you get past the top two opinions, many of the prevalent viewpoints seem to share a similar cynicism and negativity towards authority. “Don’t let devs delete” arises from the viewpoint that developers will delete any tag they deem even vaguely negative, giving all games positive or neutral tags only. Various defenses of tags, and anger over their removal, also fit into this category. These users don’t trust Valve to be able to judge what is or is not offensive and thereby worthy of removal. General cynicism, of course, only feeds into this. Finally, a decent-sized chunk of users believe that no tags should be deleted, and a similar amount believes in no

oversight at all. While these are minority viewpoints in the thread at large, they arise often enough to be notable.

1.11 Overall Comparison




Figure 4: Occurrence of Opinion Types Comparison

This chart shows that, despite “needs more oversight” being the most prevalent opinion in the thread, anti-authority viewpoints actually triumph overall. Even if the 46 occurrences of “tags are funny” were removed, which can be seen as neutral expressions of amusement, anti-authority opinions still prevail 138-110. Either way, each side has a significant number of proponents among the user community. The predominance of anti-authority opinions, however, should be taken into account.

Discussion

The crowds, therefore, seem to be as divided between anarchy and authority as the literature surrounding social tagging. Both sides have a number of reasons to back up their beliefs, and a number of users seem to fall somewhere in between. Those in favor of authority fear that malicious users will poison the system and that the lack of authority control will lead to redundant and useless tags. For them, the only way to have a functional tag system is to exert at least some level of moderation over the users' actions. Those in favor of an open and unmoderated system cite creativity and honesty as the benefits, and believe that the best tags will rise to the top and drown out the malicious ones. All in all, this debate demonstrates that what users really want is a fun and functional tagging system that reflect their viewpoints. Where they disagree is on which group (other users or moderators) poses the greater threat to that, and how to go about working around said threat.

Many of the pro-moderation posters were surprised that the system wasn't moderated to begin with. User "Wickerbasket," in post #20, sums up this viewpoint by saying "[t]hey definitely need to moderate it if they're going to allow user created tags; I'm surprised they even allowed that in the first place, there's so much room for abuse." A similar sentiment is expressed in post #90 by "GravityMan," who appreciates the attempt to involve the community but says "they need to be a bit

smarter about that.” This feeling that the internet community at large can’t be trusted to tag games without resorting to humor and insults pervades the entire thread. Many of the posters simply don’t trust their fellow users to utilize the tagging system the way they should. But malicious users aren’t the only issue, mistakes and inconsistency can be just as bad. In post #59, “Crub” argues for more quality control on the tags, saying that the site should “unify the tags so no two tags are the same (like ‘Point & Click’ and ‘Point and Click’).” Without some sort of moderation to unify similar tags, the entire system would end up much less useful than it could be.

The posters in the thread also go out of their way to suggest further moderating features such as creating a list of pre-set tags (post #113, post #342), setting up an approval system for new tags (post #114), only allowing owners of games to tag (post #183, among others), a central location to nominate tags before they can be used (post #246), or even hiring professionals to do the tagging instead of entrusting it to the masses (post #281). These suggestions all share the same goal of minimizing the effect of malicious users on the system, as well as helping to eliminate misspellings and redundancies. As “PBalfredo” points out in post #183, such measures would be necessary to clean up the tagging system “[e]ven if we lived in a paradise world were [sic] there were no trolls.” Between trolls and inconsistencies, the tag system cannot function on its own according to these posters. It needs a heavier hand to make sure that the tags are legitimate and that errors and redundancies are eliminated from the system. As “Nocturno999” concisely states in post #213, “[t]he system can be very useful with some strict moderation.”

On the other side, several posters express their faith in the community to sort out the tagging system on their own. These posters felt that, while trolls and others might well populate the system with negative or useless tags, such tags would be drowned out over time and rendered irrelevant. Thus, they believe that the tagging system will eventually regulate itself. In response to a post complaining that the system will always be overrun by trolls and therefore needs moderation, user “Honey Bunny” in post #310 says: “Will it always? I’d have thought in most cases the regular players/fans of the game would outnumber the trolls, so the accurate tags would have risen to the top, given time.” This belief in the ability of the users to balance out the trolls was earlier espoused by user “glaurung” in post #101, when he said “I do believe that the massive Steam user base will be able to balance this nonsense out right quick. Self regulating for the win.” Both see the trolls as too numerically insignificant to make a difference when compared to the vast numbers of Steam users and taggers who would tag games correctly and without malice. In their view, regulation isn’t needed to combat trolls. Rather, the community simply needs to ignore them and go about its business, and they will eventually be marginalized.

Additionally, many posters don’t necessarily see the troll tags as a problem if said tags are funny rather than openly hateful. User “jabuseika” sums up this belief, which can be seen throughout the thread, in post #48 on page 1: “I mean, if it’s not hateful, it’s not hurting anyone, it’s not insulting the product, let it be.” This belief that only the most offensive tags should be removed, and the rest left to sort themselves out, was one of the pervasive views of the thread. The above numbers and charts bear this out, since “only delete negative tags” was the third-most common opinion. The “tags

are funny” opinion, which is closely related in that users didn’t want to see the funny tags deleted even if they were seemingly useless, was the second-most common.

However, some of the distaste towards authority tends towards the hyperbolic. The idea of developers being able to remove tags was deemed “scary” in post #162, a sentiment that many echoed. Another commented, “Why would Valve want to censor the truth?” (post #85). These can be seen as reflections of the distrust towards authority, especially corporate authority, which is often endemic to the internet. While these views might seem laughably over-dramatic, they need to be taken into account along with the rest. These views speak to a sizable population, as evidenced by how many of the commonly-expressed opinions deal with distrust in one form or another. This relates to one of the common themes of the thread- that developers occasionally sell games that are broken or otherwise worthless, and they want to protect fellow consumers against purchasing such games. Their distrust, in this case, has at least some valid backing. Those who distrust authority can also be valuable taggers with valuable insights to contribute, and the people in charge of social tagging systems need to look into ways to help allay these fears as much as possible.

In the end, both sides have a lot of traits in common. Both want a fully functional, yet creative and open, tagging system. Both want the system engineered in such a way as to minimize what they see as its greatest threats. Those who see trolls and lack of quality of control as the system’s biggest threats look to a more moderated approach in order to unify the system and expel negative and insulting tags. Those who see corporate interests and the squashing of creativity as the tagging system’s biggest threats look to maximize creative and minimize the role of authority.

They want the community to have the power to call out developers or other authority figures if and when necessary. Both sides, then, operate out of places of fear. The question simply becomes: what does each user fear most? That ultimately determines which side they fall in with. Tagging system moderators and developers need to know and understand all these fears, and try to find ways to allay each one without exacerbating another.

Conclusion

This paper set out to examine how users perceive authority and moderation in social tagging systems, look at whether more animosity is directed at authority figures than at the masses or vice versa, and determine the strengths and weaknesses in each side's arguments. It is hoped that these results can inform future projects on how best to assert their authority, and what reactions they can expect. The amount of moderation in any given social tagging system boils down to how much its creators trust or distrust their users. In the end, the results demonstrate that the feelings of those users depend on their relative levels of trust and distrust in both the moderators and their fellow users. The greater the distrust towards the users, the more strictly moderated the system. By contrast, a system that trusts its users is usually far more open. Such a system tends to attract users with a distrust of authority, who see the more tightly-moderated systems as overly-controlling and evidence of bad motives on the part of those in control.

Those seeking to initiate their own social tagging projects need to understand all of these fears and motivations, and how they interact. Those who distrust their fellow users need to see that there are controls in place to keep malicious users from running amok and to ensure a decent amount of order in the system. Those who distrust authority need to have free reign to tag as they wish within reason, and to feel that their creativity is being appreciated rather than hampered. Hopefully, this study

has shed at least a small amount of light into how the users of such systems feel.

While the literature is full of thought from the project creators and their peers, not much voice has thus far been given to those actually doing the tagging. Their thoughts and motivations are incredibly important for the success of any social tagging endeavor, and therefore must be taken into account.

These results should not be extrapolated too far, however, given the limitations of this study. While the chosen thread is fairly representative of users' opinions towards moderation and authority, it still represents a very small sample of the population. It also only includes those whose opinions, one way or the other, were strong enough to motivate them to post about it. Finally, this thread is on a gaming website discussing another gaming website. Its participants are therefore gamers, and their thoughts and opinions may or may not be representative of the general population.

Future studies should look into the viewpoints of other user communities on a much broader scale. Diversifying the population will allow both for the inclusion of more varied viewpoints and for a more accurate illustration of how users as a whole feel about authority in social tagging systems. Greater numbers will, of course, allow the results of these proposed future studies to be much more statistically relevant.

References

- Ding, Y., Jacob, E. K., Zhang, Z., Foo, S., Yan, E., George, N. L., & Guo, L. (2009). Perspectives on social tagging. *Journal of the American Society for Information Science and Technology*, 60(12), 2388–2401. doi:10.1002/asi.21190
- Gupta, M., Li, R., Yin, Z., & Han, J. (2010). Survey on social tagging techniques. *ACM SIGKDD Explorations Newsletter*, 12(1), 58-72. doi: 10.1145/1882471.1882480
- Guy, M. & Tonkin, E. (2006). Folksonomies: tidying up tags?. *D-Lib magazine*, 12(1), 7-20.
- Koutrika, G., Effendi, F. A., Gyöngyi, Z., Heymann, P., & Garcia-Molina, H. (2008). Combating spam in tagging systems: an evaluation. *ACM Trans. Web*, 2(4), 34 pages. doi:10.1145/1409220.1409225
- Mai, J.-E. (2011). Folksonomies and the new order: authority in the digital disorder. *Knowledge Organization*, 38(2), 114-122.
- Rashmi. (2006, January 18). A social analysis of tagging. Retrieved from <http://rashmishinha.com/2006/01/18/a-social-analysis-of-tagging/>
- Reamy, T. (2009, October). Folksonomy folktales. *KMWorld*, 18(9), 6+. Retrieved from http://go.galegroup.com.libproxy.lib.unc.edu/ps/i.do?id=GALE%7CA209239486&v=2.1&u=unc_main&it=r&p=AONE&sw=w&asid=8eeaed9370681d08a6455a16a512bdf2
- Riley, M. C. (2013). Anarchy, state, or utopia?: checks and balances in internet governance. *Internet Computing, IEEE*, 17(3), 10-17. doi: 10.1109/MIC.2013.24

Valve curbs abusive Steam Tags after games tagged-“not a game,” “hipster garbage.”

(2014). Retrieved February 23, 2014, from

<http://www.neogaf.com/forum/showthread.php?t=769706>

Zittrain, J. (2000-10-13). Balancing control and anarchy on the Internet. *The Chronicle of higher education*, 47(7), B20.