Daniel Drew Turner

for Coye Cheshire

INFO 203: Social Organizational Issues of Information

May 11, 2010

Crisis on Infinite Pages: Comments Gone Wild at Online Newspapers

I. INTRODUCTION

Over the last few years, newspaper web sites have been faced with a "crisis in comments." Major and minor publications are expending large amounts of time and resources (both in short supply) to debate how to keep online discussions forums from being overrun by flame wars, hateful speech, and other abuses. Though newspapers find significant intrinsic and extrinsic value in offering user commenting features on news stories and columns, barrels of ink have been spilled in major news outlets on articles with titles like "Is Internet Civility an Oxymoron?"[1], "If You Can't Manage Comments Well, Don't Offer Comments at All"[2], and "News Sites Rethink Anonymous User Comments"[3].

This is not just anecdotal. Over half the newsrooms that responded in a 2009 survey by the American Society of News Editors reported having to shut down at least one online forum in the previous year. This represents the most dramatic failure point of this form of online communication. But even short of that, ongoing issues with online forums and story-connected comments are creating a pressing problem for newspaper publishers.

---

1 http://online.wsj.com/article_email/SB10001424052748704246804575190632247184538-lMyQjAxMTAwMDEwODExNDgyWj.html

2 http://www.ojr.org/ojr/people/robert/201003/1836/

3 http://www.nytimes.com/2010/04/12/technology/12comments.html

As of 2006, approximately 80 percent of online sites for newspaper included some kind of comment system for news stories, or allowed users to submit user feedback of some kind (Beyers, 2006). In 2009 the American Society of Newspaper Editors, a nonprofit professional organization, surveyed over 1,000 newsrooms across the United States and received 267 responses. The survey found that 87.6 percent of the news outlets responding "invite online comments regarding specific stories" (Keyes, p 12).

But 38.9 percent also reported shutting down comment thread for a specific story within the previous year and 46.2 percent said they had to ban users within the same period. These are significant numbers, both at the macro (total number of troublemakers) and at the micro (negative effects to user experience) levels.

The reasons behind these shutdowns include spam; personal abuse; hateful language including racist, sexist or anti-gay comments; "flame wars" (two or more commenters filling comment threads with hostile and heated arguments); "trolls" (users whose purpose is only to start flame wars); sockpuppetry (one user assuming another identity to promote statements made by his or her original persona); scaring off potential sources[4]; astroturfing (assuming personas to promote a cause as though it had grassroots support (Klotz)).

"We've lost subscriptions over the comments," said Michael Freimann, the online editor of The Pantagraph in Bloomington, Ill, which saw "overwhelming" abuse of the system. "We've been bad-mouthed on the radio," he added. Freimann was quoted by Kurt Greenbaum in his online article "Reader comments online: have we lost control?"[5]. Soon after instituting an online comment system, the News & Observer or Raleigh, N.C. saw " the arrival of spam, profanity, harassment, and the need to spend time each day deleting inappropriate comments" (Gsell). An editorial for the Online Journalism Review outlined the editorial attitude of the Los Angeles Times towards reader comments as "shut up"[6]. The New York Times featured a story outlining how many news sites such as their own, washingtonpost.com, and others, were "rethinking" comments[7].

And in a particularly messy case, which goes far beyond the scope of this paper and into various legal and privacy issues, the Plain Dealer in Cleveland, Ohio, has been sued for $50

---

4 http://blog.washingtonpost.com/story-lab/2010/03/blowback_commenting_on_debt_co.html

5 http://tae.asne.org/StoryContent/tabid/65/id/123/Default.aspx

6 http://www.ojr.org/ojr/stories/070817niles/

7 http://www.nytimes.com/2010/04/12/technology/12comments.html

million by a local judge over the paper releasing her personal information. The newspaper released this information after a reporter, within company policy, looked up who owned the pseudonymous user identity that had been posting abusive comments – including one on the mental state of a relative of one of the paper's reporters. The user had also posted comments about cases before the local judge in question, which would be professionally improper for the judge to do[8]. Whatever the merits of the law suit, this is another example of how problematic anonymous and pseudonymous commenting has become for online newspapers.

To journalists, who enter a large problem or issue almost always through the lens of the individual[9], this problem may seem intractable in that though individuals are often the cause, it hinges on the behaviors of masses of users. This, however, makes this crisis a prime subject for the application of sociological research, studies on human social activity.

So far I have seen no evidence that in all the seminars, on- and off-line discussions, white papers, and the like, there has been any such analysis. Even The Crunchberry Project (www.crunchberry.org), a 2008 research and development project from the Medill School of Journalism's New Media Publishing class, tasked with improving online conversations around news, did not think to look across disciplines.

In this paper, I will survey sociological research on trust, cooperation, reputation, and identity – issues that, I believe, play key roles in motivating proper and poor actions in online forums. Perhaps I can tease out from this principles publishers can use to reduce comment abuse.

In his article "Terrorism Studies" for the Apr 26, 2010 *New Yorker*, Nicholas Lemann wrote, "The point of social science is to be careful, dispassionate, and analytical, to get beyond the lure of anecdote and see what the patterns really are." But, he warns, "the laboratory approach can't be made to scan neatly, because there isn't a logic that can be counted to apply in all cases."

Of course he was speaking of the motivations of terrorists. But the observation is applicable; the principles I find may help in the aggregate, but will not prevent the determined bad actor. The flamers will always be with us to some degree.

---

8 http://blog.cleveland.com/metro/2010/04/cuyahoga_county_judge_shirley.html

9 This is a sudden personal observation of the author, who worked as a journalist for over a decade. Most news stories, no matter how vast or minute, center usually on an individual interpretation or experience. I wonder how this affects us, as writers, in perceiving solutions.

II. PROBLEM SPACE

ON SCOPE AND APPROPRIATING EXISTING RESEARCH

The subjects of this paper are limited to online versions of established print newspapers with regional and national scope. The paper will not address online-only news outlets (such as talkingpointsmemo.com or drudgereport.com), web sites for broadcast news shows (such as cbsnews.com), blogs, or news aggregators. The reasons are both theoretical and practical.

First, it has been established names such as sfgate.com, washingtonpost.com, nytimes.com, and the like, that have driven the discussions of what is to be done with this crisis in commenting[10]. Second, these publications have a history of Letters to the Editor; this gives them the largest stake in and commitment to maintaining online comments. To shut them down would signal a serious problem (see above). Third, though a "hyperlocal" new site would offer a more controlled laboratory for observations, online interactions on such a small scale has not been studied well, and might not scale up or down. For example: research by Perrin and Vaisey has shown that local issues in a print newspaper are less likely to see the equivalent of an online "flame war" issue. Fourth, Chan-Olmsted and Park (2000) showed that broadcast TV stations' sites at the time, only 10 percent featured discussion forums – a much lower interactivity rate than web sites for newspapers.

But despite this (relatively) limited scope, I believe it is justified to refer to research on related online formats, such as Usenet postings, web forums, or bulletin boards.

In a study of online news media in China, Hong Kong, and Taiwan during the SARS crisis of 2002-2003, Alice Lee showed that online news sites can serve many of the functions of online bulletin boards (and, by extension, Usenet groups) in serving to build community and interconnection amongst readers, though to different degrees in the three countries. Even though Chinese sites did not offer the degree of user interactivity as other examples, Lee found that they "were also capable of providing social linkage and social amusement" and could "facilitate communication and emotion sharing among members of the community" (Lee, p. 19). This would support the idea that we can apply many findings of studies of Usenet and bulletin boards to news comments.

Diana Chung pointed to Rafaeli's definition of interactivity as "responsiveness": "Rafaeli

10 http://www.nytimes.com/2010/04/12/technology/12comments.html

studied computer-mediated groups and the communication exchanges among asynchronous multi-participant public discussion groups. … He defines interactivity as 'the extent to which messages in a sequence relate to each other and especially the extent to which later messages recount the relatedness of earlier messages'"(Chung, p. 5-6). (Rafaeli and Sudweeks, 1997, http://www.usc.edu/dept/annenberg/vol2/issue4/rafaeli.sudweeks.html). This also would seem to imply that person-to-person (p2p) interaction functions much the same online, whether it is in a Usenet group or a comment thread – again, justifying use of sociological and behavioral studies not focused specifically on the systems that are the focus of this paper.

WHY COMMENTS MATTER

In this section, I will show why newspaper cannot simply abandon online comment systems, for existential and practical reasons.

New technologies, and the way they allowed relatively unmediated communication of information, contributed greatly to the downfall of the Soviet Union, said Scott Shane: it was "a revolution driven by information" (Shane, p. 261), once a bureau chief in Moscow for the *Baltimore Sun*, was not arguing that email and other Internet-based forms are inherently anti-statist, but rather that they serve to advance a more democratic (in the sense as opposed to censorial) society.

Similarly, a quantitative study quoted by Christopher Kedzie, who worked for the Ford Foundation in Moscow, found that the "correlation between interconnectivity and democracy is positive" (Kedzie, p. 213). What can be called this democratizing influence is not limited to electronic media – television, radio, and the printed word have served. But a single post on the web has a much lower barrier to entry: it can be generated with a hand crank and a satellite phone, from an anonymous source, and appear attached to a front-page article on nytimes.com. This democratizing influence is an intrinsic value to newspaper, and in line with their stated mission; my own experiences working at publications and with other journalists offer anecdotal but strong support to this. And though it is beyond the scope of this paper, there is a good deal of research linking a well-informed populace to effective governance. (Note: all of these values may well be socially constructed, especially in the light of the fact that current ethical codes and expectations of "objectivity" only sprung up in the early twentieth century – but they are in place at the publications in question as I write this.)

However, despite this high intrinsic value, high-profile newspapers are still in crisis over comment systems and the rampant abuse and flame wars. Conferences are being held, articles

about articles are being written left and right, and almost half of the newsrooms in a 2009 survey reported having to shut down a forum in the previous year (ASNE survey).

(Please note: For the purposes of this paper, the terms “comments", “comment systems”, "comment threads", and the like will be considered to include systems that allow readers to attach their comments to a single online story, or post comments in a forum. All present the same problems for publishers and exhibit highly similar user behaviors.)

There are also more practical reasons for newspapers to maintain these systems, despite the headaches and heartaches they may bring.

First, and most simply, adding a comment system to a news site generates revenue. Though the history has been shaky[11], almost all newspaper sites have a revenue model that includes online ad displays; rates charged advertisers are higher on sites with more page views and unique visitors (so far, subscription and membership fees have not proved viable[12]). At the most immediate, most primitive level, a user adding a comment forces a reload of the site’s page. The more users posting comments, the more page views from the same size audience. And higher ad rates.

In the case of washingpost.com, the ombudsman Andrew Alexander wrote in his April 4, 2010 column that "The growth [in online comments] is critical to The Post's financial survival in the inevitable shift from print to online"[13].

Second, allowing comments can foster a sense of community among users. Kollock and Smith, in their study of Usenet behavior, noted that "the conversational 'floor' constitutes a commons: if access to the floor is allocated in an ordered way by speakers exchanging 'turns'" – not an inapt description of how comment systems force asynchronous communication (I will return later to their work on community definition, and how news-related online comments do not meet all their criteria). Ridings, Gefen, and Arinze define "virtual community" as "groups of people with common interests and practices that communicate regularly and for some duration in an organized way over the Internet through a common location or mechanism." (Ridings, Gefen, Arizne, p. 273). Though how regular "regularly" and how long "some duration" can be-

---

11 (http://www.stateofthemedia.org/2007/narrative_online_economics.asp?cat=3&media=4)

12 http://www.niemanlab.org/2009/04/paying-for-online-news-sorry-but-the-math-just-doesnt-work/

13 http://www.washingtonpost.com/wp-dyn/content/article/2010/04/02/AR2010040202324.html

come an academic question, it is not rare to see on any news site regular commenters; sites try to encourage this behavior through promotions, prizes, and other gifts. And Herring uncovered that the freedom to express views (as in: leaving a public comment) was one of the two main reasons people joined virtual communities.

The practical benefit when a newspaper can instill a sense of community (or "brand identity", as marketers would have it), in some users is that not only will these users drive more content but their presence can "seed" new users to become regulars. (This will become important also in the later discussion of registration.) And this increase in user participation, in turn, drives site revenue (as noted above).

Community can also act as an exchange: sites that offer users the value of a community space online can ask for more information, such as email addresses[14], sex, age, ZIP code, from users, and sell this aggregated anonymous data (along with online browsing habits ("In addition to personally identifiable information, we also collect certain non-personally identifiable information through technology and tools, including cookies, Web Beacons and log data. We aggregate and analyze this information in order to learn more about how our Web sites are used."[15]).

Offering interactive comments also serves to boost an online news outlet's credibility with readers. "Interactivity and transparency have vast implications for the elevation of credibility on Internet news sites," wrote Martha Stone, a 2001 Poynter Ethics Fellow and Co-director of an Online News Association Web Credibility Study.

And Ash, Hettinga, and Halpern found in a 2009 quantitative study that "If a news Web site wants readers to enjoy visiting the site, and continue to return, allowing comments is a good way to promote this sense of enjoyment." (Ash, Hettinga, Halpern, "Effects", p. 22). Oddly, that they also found that the presence of comments seemed to reduce the perception of quality of the journalism; they hope to study this further.

To conclude: Comment systems attached to online stories represent a powerful set of existential and practical values for news publications, especially ones with a long history of community engagement. This means these news outlets are strongly disinclined to simply abandon online comments.

---

14 Greenbaum, Kurt. "You've got revenue." *American Editor* 82.4 (2007): 24. *Academic Search Complete*. EBSCO. Web. 28 Apr. 2010

15 http://www.nytimes.com/ref/membercenter/help/privacy.html

WHY COMMENTS AREN'T WORKING

It's too easy to say, "The boy just ain't right." In this section I will break down what about the state of comment systems is not working for newspapers.

Newspapers have a deep investment in publishing readers' voices. The long history of Letters to the Editor has made this feature part of the definitional identity of a newspaper. The sense of being able to reach the writers and editors also is a key factor in user trust of an institution.

"Trust is, at its heart, an *interpersonal* phenomenon," wrote Daryl Koehn in the *Journal of Business Ethics*. Koehn went on to say that emailing a web site is like "putting a message in a bottle and throwing it out to sea"; he suggested that offering phone access would show a "willingness to engage directly" (Koehn, p. 17). In this context, I am combining Worchel's 1979 concepts of **situational trust** (people adjust their tendencies in response to situation cues, such as the quality and amount of communication between parties) and **learned trust** (the experience parties have gained from past situational trusts) to obtain a fairly everyday sense of the word, and to avoid the pitfall of requiring that trust be between human actors only – in this case, human users can certainly feel there is agency in the other party (the newspaper), or at least can see that there are humans behind it.

In the case of web sites for established newspapers – the scope of this paper – there can be a great deal of learned trust based on an individual paper's reputation (such as, "All the News that's Fit to Print") that can carry over to the web version, though a new set of situational cues can alter that base. And I would argue that users expect some form of bottom-up communication with the newspaper, whether in print or online; there the aspect of "quality and amount of communication" grows in importance.

But online comments are qualitatively and quantitatively different from Letters to the Editor, even if some aspect of the latter is meant to serve the same functions.

In a 2009 study in four Norwegian municipalities, Winsvold compared comments and forums in online newspapers to the Letters to the Editor in their print counterparts. He found the online comments to be of "low quality" and the commenters to have dubious motivations while they seemed have "participated for fun or because they liked a good fight" (Winsvold, p. 13-14).

In terms of quantity, Winsvold also found that "[a]nother major challenge to the position of online communication arenas was the huge volume of contributions resulting from their inclusiveness." (Winsvold, p. 14). I will later discuss this in terms of moderation, but the relevance here is that traditional methods of dealing with potentially damaging user comments – that is, individual, in-house editorial consideration simply does not scale, especially for the traditionally resource-tight news industry.

In the more granular data obtained by the 2009 ASNE survey, we can see the primary reasons for shutdowns of story-related discussion threads. The top reported reason was "Discriminatory comments involving race, ethnicity, gender or sexual orientation", followed closely by "Hurtful comments not discriminatory in nature" and "Obscenities, profanities, foul language." Legal, fairness, and accuracy issues were far behind, suggesting that on-topic signal (user-generated content) is far overwhelmed, in terms of what is problematic, by off-topic noise.

Paradoxically, since the likelihood of abuse and flaming rises as the number of commenters goes up ("fruitful cooperation has proven to be difficult to sustain as the size of the collaboration increases" (Wilkinson and Huberman)), this means that the more attractive a comment system is, the less attractive it will be. "Why do people choose to join a virtual community? The most frequently cited reason in the literature is to access information (Furlong, 1989; S. G. Jones, 1995; Wellman et al., 1996), which is also a reason for group membership cited often by social psychologists (Watson & Johnson, 1972)" (Ridings and Gefen). Too much abuse in the comments, the more noise and less signal, means a lower ability to access information. It's the Yogi Berra effect: nobody goes there any more – it's too crowded.

As I showed in Section "Why Comments Matter", these comment systems have high value in a range of contexts for newspapers. Yet, as the data from the ASNE survey indicate, many newsrooms have been forced to cut away the source of this value. And the content and frequency of industry press and seminars on the subject indicate that the newspaper industry is ready to pour scarce resources into this problem area (recently the washingtonpost.com and sfgate.com sites have stated they will undertake comment system overhauls; though they have not, as of this writing, finalized their plans, both will include registration requirement as well as more filtering and an "ignore user" feature).

With all this, it is obvious Something Must Be Done (to quote "Duck Soup").

III. TOOLS OF THE TRADE

In this section, I will give a brief overview of a few tools news sites are currently using with their comment systems, whether they were originally designed to combat the problems listed above or not. And I will attempt to address their relative strengths and weaknesses – though it is clear that none, nor any existing combinations, have so far sufficed to eliminate or appreciably control the crisis to date.

Policy

Most, if not all, sites with comment systems have Terms of Service (TOS) that delineate not only what a user can expect from the site (conditions of privacy, etc.) but also what the user agrees to by using the site's features. It is key to note that the user need not sign anything, or register an identity at the site, or even have read the TOS in order to be held to them – it's an effective "you post, you've agreed" situation (for example, see the TOS for the Sacramento Bee's web site: http://www.sacbee.com/terms-of-service/). It remains a conundrum as to whether one is held to agreeing to the site's TOS if one visits the page listing the TOS in order to evaluate them.

The TOS for most newspaper sites include injunctions against much of what I above identified as a problem in online comments on their sites. The washingtonpost.com site demands agreement from users not to post "inappropriate" remarks, including those that are hateful or racist, or those that advocate violence. And here is a representative sample from the nytimes.com TOS:

> "**3. USER GENERATED CONTENT: SUBMISSIONS INCLUDING COMMENTS, READER REVIEWS, TIMESPEOPLE AND MORE**
> 3.1 (a) You shall not upload to, or distribute or otherwise publish on to the Service any libelous, defamatory, obscene, pornographic, abusive, or otherwise illegal material.
> 3.1 (b) Be courteous. You agree that you will not threaten or verbally abuse other Members, use defamatory language, or deliberately disrupt discussions with repetitive messages, meaningless messages or "spam."
> 3.1 (c) Use respectful language. Like any community, the online conversation flourishes only when our Members feel welcome and safe. You agree not to use language that abuses or discriminates on the basis of race, religion, nationality,

> gender, sexual preference, age, region, disability, etc. Hate speech of any kind is grounds for immediate and permanent suspension of access to all or part of the Service.
> 3.1 (d) Debate, but don't attack. In a community full of opinions and preferences, people always disagree. NYTimes.com encourages active discussions and welcomes heated debate on the Service. But personal attacks are a direct violation of these Terms of Service and are grounds for immediate and permanent suspension of access to all or part of the Service."[16]

However, these policies have proven largely ineffective, even though all reserve the right, as in the nytimes.com TOS, to "terminate or suspend your access to all or part of the Service for any reason, including, without limitation, breach or assignment of these Terms of Service."

The reasons are twofold.

First, the sheer scale of incoming comments (over 320,000 a month at washingtonpost.com [17] is impossible to monitor by humans. The site's ombudsman wrote, "About 300 comments are deleted each day. But others slip through because The Post's staff of only a few monitors can't possibly scrutinize everything."[18] To date, there are no studies showing automated systems can effectively take up the slack; many problematic posts do not contain explicit keywords, or mask them by misspellings; it is difficult to judge when a commenter is enthusiastically and helpfully engaged and when harassing; most newspapers have the stance of preferring to err on the side of allowing speech.

Second, there is what Ford and Strauss call the "disposability" of online identity. As I will discuss more later, banned commenters can simply come back to comment systems, even ones that require "basic registration." If sites track IP addresses, it is not difficult for a user to spoof or log on from a different one.

As a result, TOS policies can give online publications justification for their efforts to police abuse, but these efforts are outmoded and insufficient.

---

16 http://www.nytimes.com/ref/membercenter/help/agree.html

17 http://www.washingtonpost.com/wp-dyn/content/article/2010/04/02/AR2010040202324_2.html

18 ibid.

All in Moderation

Moderators of online comments could be thought of as serving some, though not all, of the purposes "the editors" did in the context of Letters to the Editor in print versions. Online moderators should, at least in theory, vet submitted comments for problematic content and approve those that do not contain such while deleting those that do (and perhaps banning those commenters). Beyond theory, salon.com co-founder and former managing editor Scott Rosenberg wrote in 2010, "Show me a newspaper website without a comments host or moderation plan and I'll show you a nasty flamepit."[19]

There are two models of moderation systems. "In-house" moderators are hired (full- or part-time) staff of the publication in question. They are trained and supervised by the publication, and responsible to that organization. User moderators are just that: users (perhaps veterans) of the service or publication who have earned, in some way, the right and responsibility to oversee the content of the comment system.

With in-house moderators, the process may not be transparent. That is, comments under review are not publicly displayed before approval; the in-house moderation process is not obvious to the average user, or to the person who made the comment under review. This has the advantage of not allowing all content into the public space by default, so those who would intentionally start flame wars do not have power to do so by default. The nytimes.com site is an example of this kind of system. One drawback for the user is the significant delay in the comment going up, and the even longer delay in seeing responses. This violates Rafaeli's "interactivity" and can be for users more frustration than it's worth.

And this practice is resource-intensive in a resource-starved industry. A top paper such as the New York Times can get away with posting only select comments and making sure they are moderated "by the authors and editors working with a blog, supplemented by a team of paid moderators whom we train and supervise,"[20]. However, this is not the case for most publications. Saundra Keyes, a journalism professor at thee Reynolds School of Journalism, University of Nevada, Reno, and a columnist for ASNE.org., said, "Few newsrooms, of course, can devote the resources the Times devotes to achieving that goal through moderation"[21]. And as noted

---

19

http://www.salon.com/news/feature/2010/04/13/newspaper_online_comments_moderation_open2010

20 "long-time editor" for the Times, personal communication, May 9, 2010

21 http://www.igreenbaum.com/2009/04/19-qs-and-as-from-asnes-story-comment-webinar

above, washingtonpost.com has "only a few monitors." (For more on the nytimes.com moderation policies, see http://www.nytimes.com/ref/membercenter/faq/comments.html and the wonderfully titled blog post "The Top 10 Reasons We Deleted Your Comment" at http://cityroom.blogs.nytimes.com/2007/11/15/the-top-10-reasons-we-deleted-your-comment/.)

Though it is beyond the scope of this paper, I did do a casual survey of what online newsrooms offer in terms of in-house moderation. Most I saw, at the national and regional level, are along the lines of washingpost.com at best; a few have dedicated (if sometimes part-time) staffers monitoring comments at they come in, stepping in to delete comments and ban users that violate their TOS. Whether they review comments before they appear online (proactive) or repair damage after they appear online (reactive) varies from newsroom to newsroom, based perhaps on policy, or the technical capabilities of their software platform, or simply manpower assignment.

In contrast, user or distributed moderation systems allow any user to "score" any other user's comments. This is best seen on the site Slashdot.org, which I will talk more about in terms of reputation systems. Lampe and Resnick studied the distributed moderation system of that site and found that there was a general success in "floating" the best comments and "burying" the worst, though they recorded that "much of a conversation can pass before the best and the worst comments are identified" (Lampe & Resnick); this is due to the human time it takes for a significant amount of Slashdot.org users to read and evaluate each comment.

Other examples of this are digg.com and various discussion-board-like sites. For some stories, computerworld.com allows readers to give a "thumbs up/down" vote on a previous comment[22]; users can also click a "report this comment" link at the bottom of each comment. However, this kind of rating is not moderation: readers cannot filter out below a certain rating level, nor do users receive any ding or bonus to any reputation score.

Some sites do have simple buttons that block viewing of comments that contain profanity. Again, this is not moderation, as it does not attach any information to the commenter nor any metadata to the comment. Though this paradigm does provide my personal favorite example: the default view of comments at the humor magazine cracked.com includes a button marked "Show Profanity".

Overall, though, distributed moderation systems do not address at all a critical problem of comment abuse. In a user or distributed moderation system, all comments are posted to the

22 https://www.computerworld.com/comments/node/9175594

public – they have to be, for users to moderate them. Though they may eventually be buried, a skilled or willful troll can still pollute the discourse.

So, though existing moderation paradigms offer some limited promise and victories, they each have failed, as implemented, to stem the crisis.

Registration is not Identity

Many sites require some sort of **registration** for users to post a comment, or access other site features. For example:

> "NYTimes.com requires that you supply certain personally identifiable information, including a unique e-mail address and demographic information (zip code, age, sex, household income (optional), job industry and job title) to register. By using NYTimes.com, you are agreeing to our Terms of Service."[23]

However, this "basic registration" (which I define as requiring a valid email address but no other confirmable data – the other demographic information cannot be confirmed or disproven in the registration system) can easily be defeated. And in fact, usually is, even by well-intentioned users. It requires no confirmable identification information to sign up for a "valid email address" and sites like bugmenot.com do brisk business with not only the malicious but also the privacy-minded.

This results in a mild barrier to entry to some, as it adds a step and asks for what some might consider personal data while asking for "trust" (which, as mentioned above, can be seen as requiring human-to-human interaction, as is not the case in a registration system) that no ill will come from this. Yet it is no barrier to the willful comment abusers, or a very low one.

In short, I would contest that while registration for online comment systems does serve various purposes, current forms are not enough to prevent the crisis in commenting, and do not begin to consider why.

Reputation and Recommending

---

23 http://www.nytimes.com/ref/membercenter/help/privacy.html

But some sites can build on registration.

Reputation and recommendation systems have become commonplace on web sites that specialize in enabling b2p, b2b, and p2p sales of goods and services (for the purposes of this paper, I will treat the two as functionally interchangeable; recommendations and ratings combine to form a reputation). Ratings and reputation metrics can be attached to users themselves (think of buyer and seller ratings through the Feedback Forum on ebay.com, retailer rankings on Google Product Search) or to user submissions (such as "Was this review helpful to you?").

Eric Goldman, Director of the High Tech Law Institute at Santa Clara University (http://law.scu.edu/hightech/), has defined **reputational information** as "information about an actor's past performance that helps predicts the actor's future ability to perform or satisfy preferences" (E. Goldman, "Regulating Reputation Systems" talk at UC Berkeley iSchool, April 14, 2010). This functional sense of reputation is distinct from and unrelated to any shame or gossip connotations.

Two types of reputation systems come into play online, Goldman said.

First are **unmediated (or distributed) systems**, where users comment or rate directly, in a "word of mouth" or p2p mode. There are various formats for this online already. Slashdot.org has a mature system that assigns reputation scores both to users and individual comments. Registered users can rate the comments of other (registered or "Anonymous Coward") users, so visitors to the site can filter their view by minimum comment rating and see only the contributions that have been ranked to their standards. Users themselves are also scored for reputation. Starting with a "+1" score, users can increase their "karma" by posting comments that receive high scores from others, or by moderating comments. Or, in the case of ebay.com, "After a transaction is complete, the buyer and seller have the opportunity to rate each other (1, 0, or -1) and leave comments (such as 'good transaction,' 'nice person to do business with,' 'would highly recommend'). Participants have running totals of feedback points attached (visibly) to their screen names, which might be pseudonyms. Yahoo! Auction, Amazon, and other auction sites feature reputation systems like eBay's, with variations, including a rating scale of 1-5, several measures (such as friendliness, prompt response, quality product), and averaging instead of total feedback score" (Resnick et al., 2000).

Second are **mediated systems**; in these, reputation information from other users is aggregated, transformed, and published, after the model of bond ratings or credit scores.

I would suggest that there can also be a third type. In an **authoritarian reputation system**, publishers or managers of the web site could either publicly post their assessment of a user, a

user's comments, a product, etc., or save this information "behind the scenes" for future actions against the user, such as banning or inviting him or her to participate in the site's management or creation.

Though some ratings and reputation systems are trivial (for example: where users can rate a comment thumbs-up or thumbs-down, and all users can see who voted which way), but some rely on complex algebra and/or data not accessible to users, such as how one user has been rated over the course of years. Opacity may be intentional, in an attempt to prevent users from gaming the system (as Google has strongly guarded its PageRank algorithm), or simply "security through obscurity." I will discuss in below the relevant issues with and research on this.

More and more online newspapers seem to be adopting, or considering adopting, some sort of reputation system. The washingtonpost.com site, for example, which sees over 320,000 comments per month, will go to a "tiered" system in mid-2010; the site's ombudsman described it as "commenters being assigned to different 'tiers' based on their past behavior and other factors. Those with a track record of staying within the guidelines, and those providing their real names, will likely be considered 'trusted commenters.' Repeat violators or discourteous agitators will be grouped elsewhere or blocked outright."[24]

It's critical to remember that a necessary component of these systems is persistent identity. Whether this comes from a system-assigned ID or a user-chosen username, without having some way of identifying and naming the subject of a rating, and having that subject remain the same person/company, no recommender or reputation system can provide useful information to other users. And so reputation systems inherit the same issues that apply to identity.

24http://www.washingtonpost.com/wp-dyn/content/article/2010/04/02/AR2010040202324_2.html

IV.THE SOCIAL RESEARCH/ANALYSIS

In this section I will outline some basic concepts in research of the sociology of online communications. The idea is that by analyzing the strategies that people pursue in online communications, I will be able to uncover concepts that can be leveraged by existing tools and policies in order to more efficiently promote good community participation and diminish bad actors.

Social (Capital, Ties)/Coin of the Community

The study of what people exchange online and why has a strong trail in sociological theory. Perhaps this is a function of a consumerist society, but any interaction, even trivial speech, can be seen as a transaction in theory (note that it is called "exchanging pleasantries"), and these transactions can form a kind of economic basis for a virtual community. This basis is social capital, which is as a term has had some range of definitions, but will be considered in this paper the value of an individual's social networks and how much people within those networks are inclined to do things for each other.

Social capital in online contexts that value information (whether it be factual, opinion, or recommendations) can simply be the public demonstration of knowledge, and the sharing thereof. Ridings and Gefen (2004) synthesized previous studies (Binik, Cantor, Ochs, & Meana, 1997; Hiltz & Wellman, 1997; Rheingold, 1993a; Sproull & Faraj, 1997) to be able to say, "Knowledge and information are, in general, a valuable currency or social resource in virtual communities."

Similarly, Donath (1996) said, "Individual recognition is important in many newsgroup" (remember that above, I showed that we can apply Usenet-centric studies to this paper's scope) and, "On-line status is recognized and there is deferral to respected members" based on those members' previous display of knowledge." This display is "those members'" building of social capital.

Much of what has been written about exchange and social capital has been within the context of goods and services transactions, which can prove problematic to the public goods context of this paper (or even impossible to translate). But the idea of social capital seems to be a reliable measure in both contexts, motivating those participating in an eBay trade as well as a comment thread.

Also related to this context is the idea of **strong** and **weak ties**. Mark Granovetter devised the idea of relative measuring of interpersonal ties, dividing "strong" into basically "people you really trust" while "weak" are "merely acquaintances." While this alone seemed commonsense, Granovetter went on to show how weak ties were "indispensable to individuals' opportunities and to their integration into communities" (Granovetter, 1973).

The relevance here is that almost all relationships within the scope of this paper – that is, individuals interacting under usernames/aliases/online personae through comments form attachments that are almost the definition of weak ties. In fact, the relative anonymity of online comments, while leveling much of social, regional, sex, and age differences, may strengthen their relationships. In fact, "the reduction of social cues makes it far more difficult to develop the intimacy and confidence necessary to deepen relationships. Therefore, the Internet is more conducive for the development of weak ties rather than strong ties (e.g., Bargh & McKenna, 2004; Blanchard & Horan, 1998; Haythornthwaite, 2002)" (Best & Kruger, 2006, p. 397)

Ridings and Gefen (2004) also stated that what "makes virtual communities special in this regard as compared, for example, with traditional social groups is the magnitude and impact of "weak ties," i.e., relationships with acquaintances or strangers to obtain useful information through online networks (Constant, Sproull, & Kiesler, 1996). A virtual community can be an ideal place to ask relative strangers about information."

Best and Krueger also said that "some argue that these online social interactions meet the conditions necessary to facilitate the production of social capital (e.g., Ester & Vinken, 2003; Hill & Hughes, 1997; Rheingold, 1993)… the Internet offers opportunities for users to develop personal ties with others, even a shared sense of collective identity (Rheingold, 1993; Walthier, 1995)" and "individuals more actively pursuing and maintaining weak ties typically possess greater levels of social capital than those limiting their interactions to strong ties" (Best & Kruger, 2006).

And their empirical evidence, they claimed, did in fact "offer generalizable empirical evidence in support of the positive view of online relations; indicators of social capital positively relate to the level of interaction with people met on the Internet. …Although online social interactions likely do not produce strong connections that elicit intense loyalty, these results do suggest that they foster connections critical to expanding networks and producing residuals such as generalized trust" (Best & Krueger, 2006).

So, in a way, news stories online provide a qualitatively different and more social capital building environment than topic-centric Usenet newsgroups, as people of many social and atti-

tudinal stripes will visit and comment on a news story (overlooking the "edge cases" highly homophilic and/or partisan news and opinion sites, which are not in the scope of this paper).

To sum up, social capital is also strongly related to what Ridings and Gefen call in virtual communities "the social support that the community can provide. Social support is 'the degree to which a person ' s basic social needs are gratified through interaction with others (Thoits, 1982, p. 147). Social support may also be linked with individual motivation to join groups because of the sense of belonging and affiliation it entails (Watson & Johnson, 1972…. House (1981) offers a more specific definition of social support: a flow of emotional concern, instrumental aid, information, and/or appraisal (information relevant to self-evaluation) between people (p. 26)." I take this "appraisal" to be, in part, the fact that other members of the community "listen" to the user's comments.

Supporting that, is that "Herring [(1996)] found that the freedom to express views and to receive social support were the main reasons individuals joined and used virtual communities. Her study of two email distribution lists found that people participated to exchange opinions, beliefs, understandings, and judgments though a social interaction with others, but where the pure exchange of information took on a secondary role" (Ridings and Gefen).

From this, we can conclude: most active users in online comments are motivated to increase their own social capital and maintain a network of weak ties. Feedback on these efforts, or lack of efforts, should be quite evident to the user: as capital and networks grow stronger, the more responses the user should see to comments and requests. How to entice flamers to "play nice" is a question that has not been studied in this context – but it is nontrivial to discover that these metrics of social capital and weak networks, which can and often are measure quantitatively, can be used to identify and potentially isolate abusive commenters.

Trapping Cooperation

The concept of cooperation between individuals (often strangers in real life) online is another fertile area of study. Though the bulk of the work has been focused on transactional situations, such as bidding and sales on eBay.com or craigslist.com, there's been work on the question of what drives cooperation, which often comes at some cost, when there are no "real" goods, services, or money to be exchanged.

This is relevant to online commenting systems because one form of cooperation could be said to be not not cooperating. That is, not "just getting along", but instead violating the social or explicit (see "Policy") norms. As stated before when defining the problem space of this paper,

these types of violations include trolling, abuse, starting a flame war, spamming, hateful language, sockpuppetry, astroturfing, and so on.

What is important to keep in mind is that though the effects of social capital are real, as seen above, and in the aggregate build cooperation and reinforce socially normative behavior, we can't rely on them in individual cases.

In individual cases, such as a trade between strangers, there can be a "tension between individual and collective rationality," where "behavior that is reasonable and justifiable for the individual leads to a poorer outcome for all" (Kollock and Smith, 1994). You may want to drive the biggest, meanest car on the market, but this may not be the best for the rest of the world.

They relate this tension in terms of Hardin's Tragedy of the Commons: "Hardin described a group of herders having open access to a common parcel of land on which they could let their cows graze. It is in each herder's interest to put as many cows as possible onto the land, even if the commons is damaged as a result. The herder receives all the benefits from the additional cows and the damage to the commons is shared by the entire group. Yet if all herders make this individually reasonable decision the commons is destroyed and all will suffer" (Kollock and Smith, 1994). In the case of online comment systems, the "commons" is the environment of relatively frictionless information transfer, with a sense of community, and information relevant to the users who have chosen to enter this "common" (that is, chosen to read this particular story, in this particular publication, and then gone on to the comments section).

For there to be a solid basis for reciprocity (of trust, perhaps of civility, even), which is a requirement for extended cooperation (and preserving the "commons" in the sense above), there should be: "ongoing interaction, identity persistence, knowledge of previous interactions, and strong group boundaries" (Kollock, 1999). As even a quick glance at your average online news story comment thread shows, none of these conditions may be met, and it is rare to see all of them met.

From this, we can conclude: As in commercial transactions, parties who anticipate ongoing relationships and exchanges with others have a vested interest in maintaining the "commons"; the trick is unpacking how to either get bad actors to buy in to some or all of Kollock's conditions, or design a system that will necessarily restrain them into at least some of these conditions.

Identity and Groups

The term "identity" has in sociology, as in our personal lives, many meanings. I will limit existential questions for the scope of this paper.

As I showed in above, people find value in belonging to a community, even a virtual, or temporary one. By default, people will form into groups, whether in the real world or online, or even within comment threads attached to a news story that will move off the "front page" of the web site within a day, provided there is a low lag time between comments being made and appearing online (Amichai-Hamburger, 2005).

A side effect of group creation and identifying with that group is that "even a trivial allocation of people to a group is likely to create a situation of ingroup favoritism" and "group members perceived their own group performance as superior on a cognitive task as compared with that of the other group" (Amichai-Hamburger, 2005). It's easy to see how this "minimal group paradigm" can become supercharged, in a sense, in the world of online journalism, where sometimes the information comes quickly and with minimal context, sometimes in an "if it bleeds, it leads" format. And then there are political stories.

One thing we can take away from a seemingly unrelated study is that users might be more likely to sign up, and reveal more of their identity, to publications they can identify as ingroup in some way.

In a study published in 2009, O'Donnell et al. tested how much of an invasion of privacy closed-circuit TV (CCTV) surveillance would be perceived, depending on the subjects' perceived social relationship with those in charge of the CCTV cameras. O'Donnell's team found that "surveillance is perceived as more acceptable when it originates from a group with which one identifies or shares an identity"; substitute "share personal information online" and we can see the result, which has practical implications when linked to my previous discussion of how to promote a sense of community in readers.

The flip side of honest identity and identification with a group is the problem described above of the disposability of online identity. (It's worth noting that nytimes.com may ask for users' real names -- "Please fill in the name field with your real name or initials. We have found that people who use their names carry on more engaging, respectful conversations"[25] – but few ever do.)

---

25 http://www.nytimes.com/ref/membercenter/faq/comments.html

Donath wrote that "the cost of identity deception to the information-seeking reader is potentially high" (Donath, 1996); in the context of online commenting, this means that an honest reader of a comment thread can suffer considerable deception. She also notes that the deceiver may not be uncovered as such until he or she has posted many erroneous comments. This is one of the most important tussles in the whole online commenting crisis: detection can be a time-consuming process, requiring multiple abuses, but each abuse is of immediate damage, and can compound over time (a flame war can outlast its instigator). The war is asymmetric.

One possible way of limiting the conflict, or slowing the production of news ones, Donath suggested, could be through "imposing high costs on deception" (Donath, 1996). Her example of a job seeker providing a padded résumé and methods of cutting through the deception, though, hinge upon noticing and differentiating between **assessment signals**, which are reliable because they come at a cost (like "big horns on a stag", and **conventional signals**, which can easily be faked (like wearing a T-shirt that says "Powerlifter" on it).

It remains to be seen if such discriminations can be made reliably in the context of online comments, and within the relatively brief time scale required to retain a sense of "interactivity." Donath admits that it is an interpretive and subjective process, and so possibly better suited to humans than machines (at least at this stage). However, as I showed above, human-powered comment moderation must make a Sophie's choice of either letting hundreds of thousands of potentially abusive comments through, or abandoning interactivity.

(This is not to say there are some levels of automation that work – there is a long history of the "kill file" or "bozo filter". But those only catch the most obvious spam and unsophisticated attempts. The proof is that these have been in place for years, yet the problem remains unabated.)

One more interesting study finding was that "given the right circumstances, online users easily forget about their privacy concerns and communicate even the most personal details without any compelling reason to do so. This holds true in particular when the online exchange is entertaining and appropriate benefits are offered in return for information revelation" (Berendt et al., 2005). This dovetails with what I've discussed above regarding community, social capital and exchanges, as well as the minimal group paradigm.

This would indicate that publications should, in addition to (of course) making their sites entertaining and informative, effectively link this with asking for stronger, or more binding, registration formats. What shape those would take, and at what point users would express resistance, would be a good topic for further research.

However, the above may not be an acceptable conclusion for those who place a higher value on anonymity. A dissenting voice was Rosenberg wrote that "anonymity isn't the problem":

> "The problem is that once an online discussion space gets off to a bad start it's very hard to change the tone. The early days of any online community are formative. The tone set by early participants provides cues for each new arrival. Your site will attract newcomers based on what they find already in place: people chatting amiably about their lives will draw others like themselves; similarly, people engaging in competitive displays of bile will entice other putdown artists to join the fun.
>
> So turning things around isn't easy. In fact, it's often smarter to just shut down a comments space that's gone bad, wait a while, and then reopen it when you've got a moderation plan ready and have hand-picked some early contributors to set the tone you want. If I were running a newspaper with a comments problem, that's how I'd proceed. Don't waste your time trying to force people to use their real names in hope that this will improve the tenor of your discussion area; build a discussion area that's so appealing from the start that it makes people want to use their real names.
>
> ... When you opened up comments, was it really about having a conversation with the readers? Then have that conversation! Get the editors and reporters in there mixing it up with the public. Sure, there will be problems and awkward moments; there will also be breakthroughs in understanding."[26]

26

http://www.salon.com/news/feature/2010/04/13/newspaper_online_comments_moderation_open2010

Reputation and Recommending

It's hard to find (there may not be one, and for a reason) a study that says people do not care about reputation; there may be no such study. It may be true that "people" in the aggregate cannot help but do so. Individuals may not at times, perhaps in the shortest term, and as a sacrifice for a more significant personal or group goal. "Individual recognition is important… Online status is recognized and there is deferral to respected members" (Donath, p. 13).

In addition, the strong majority of sociological research supports the idea that some form of reputation tracking and presentation can serve to civilize, in a way, online discourse. "Reputation systems seek to establish the shadow of the future to each transaction by creating an expectation that other people will look back on it" (Resnick, et al., 2000) means that the projection of reputation information removes interacting parties from the eternal and consequence-free "now" but instead places them into historical context where past transactions carry information to the present and the future. Perhaps you could call this the "you can run, but you can't hide" or "karmic" model.

But research indicates that the instillation of a reputation system is not enough to moderate the worst behavior of persistent user (of course the "troll-and-go", where a bad actor logs on, posts a disruptive comment, and disappears, remains problematic). The reputation system must feature some degree of transparency to users. This is at least to instill a sense of positive usability from potential participants: "users like and feel more confident about recommendations that they perceive as transparent" (Sinha & Swearingen, 2002).

Goldman echoed and expanded on this recommendation. He promoted the idea of "translucent algorithms": some, but not all, information should be exposed on how reputations are rated, what data come into play, and what the data are (within the bounds of protecting personal privacy). When the proper balance is struck, Goldman said, bad actors, bots, and companies would not have enough traction to game the system, yet the move away from opacity would greatly increase the user confidence in the system.

He added that he had found mediated reputation systems were "more valuable socially" than unmediated systems. The latter, he said, have "high transaction costs associated with finding reputable data sources" – that is, the maintainers of the system or the web site publishers would need to invest time and effort to assess on their own the credibility of the data coming in and the users supplying it (this is a prime case where sockpuppetry and astroturfing can greatly skew a system).

But still, the rogue bad actor, who has no plans for future interactions (the "troll-and-go") could ignore all these social forces. This actor has no stake in gaining reputation or social capital. This actor pops up with a new registration and plays the role of Donald Duck in "The Band Concert", acting as a force of disruptive chaos.

"Game-theory analysis demonstrates that there are inherent limitations to the effectiveness of reputation systems when participants are allowed to start over with new names," wrote Resnick (2000). He saw two alternatives: distrust newcomers to the online system in question, which would place a penalty that would discourage users and risk destroying online communities, or demand and enforce registration[27].

[27] Resnick favored the "once-in-a-lifetime" identification process (Friedman & Resnick, 2001) that would assign one persistent ID to one person. Ford & Strauss proposed a similar process with their "pseudonym parties" and "accountable pseudonyms".

V. CONCLUSIONS AND RECOMMENDATIONS

It is clear that the weight of sociological research either explicitly or implicitly suggest that to combat comment system abuse, web sites take measures towards enforceable registration, with some barrier to entry serving to reduce the disposability of identity (or identities). There is no recommendation of what this barrier would be in practice or in user perception.

I would strongly suggest a good topic for future quantitative study would be what level of strictness (personal information? verification?) would result in how unacceptable a barrier to entry for the average user, and how much would present enough of a cost to effectively discourage identity swapping by bad actors.

Enforceable registration is a prerequisite for a muscular reputation system, another strong recommendation from the body of sociological research. Without persistent identity there can be no reputation information (which is, remember, an indication of reliability in future actions).

On a more granular level, I conclude that there are tangible benefits for web sites that can closely bind their readership into identifying as part of a group that includes the publication. Public radio stations giving branded tote bags as pledge gifts do so not just to advertise, but also to match up that toter's identity with that of the station. Perhaps marketing expertise could be of use here: think of Ford families and Chevy families. This is a strong finding that publications should not treat their online arms are separate entities (as Time magazine did with its 1994 pathfinder.com web "portal"), but link the online and print experience as closely as possible for readers.

I also conclude that publishers should invest in research on what would be the assessment signals that can identify malicious users. This needs to be more than just keywords, as current filtering options can stop this (though not 100%) at the server level. It is beyond the scope of my research and knowledge of existing technology, but I do want to stress that investigations in this area should be seriously versed on this area (Donath's talk on "Signals, Truth, and Design"[28] would be a good starting point).

Open commenting, especially in the field of the public good of journalism, is a public value. I hope this paper can serve to help resolve the crisis endangering it.

---

28 http://josephhall.org/nqb2/index.php/2006/09/27/donathdls

SOURCES

Amichai-Hamburger, Y. (2005). Internet Minimal Group Paradigm. *CyberPsychology & Behavior*, *8*(2), 140-142. Retrieved from Academic Search Complete database.

Ash, E. , Hettinga, K. and Halpern, D. (2009-08-05). Effects of a trend: The influence of user comments on readers' perceptions of online newspapers. *Paper presented at the annual meeting of the Association for Education in Journalism and Mass Communication, Sheraton Boston, Boston, MA*. Retrieved from http://www.allacademic.com/meta/p375864_index.html

Berendt, B., Günther, O., & Spiekermann, S. (2005). PRIVACY IN E-COMMERCE: Stated Preferences vs. Actual Behavior. *Communications of the ACM*, *48*(4), 101-106. Retrieved from Academic Search Complete database.)

Best, Samuel J., & Krueger, Brian S. (2006). Online interactions and social capital distinguishing between new and existing ties. *Social Science Computer Review*, *24*(4), 395-410. doi:10.1177/0894439306286855.

Beyers, H. (2006). What constitutes a good online news site? A comparative analysis of American and European awards. *Communications, 31*, 215-240.

Chan-Olmsted, S. & J. Park. (2000). From on-air to online world: examining the content and structures of broadcast TV stations' web sites. *Journalism and Mass Communication Quarterly*, *77*(2) 321-339.

Cheng, A., Friedman, E. (2005). Sybilproof reputation mechanisms. *Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems, August 22-22, 2005*.

Chung, D. S. (2004). Into interactivity? How news websites use interactive features. *Paper presented at the annual meeting of the International Communication Association, New Orleans Sheraton, New Orleans, LA Online.* Retrieved from http://www.allacademic.com/meta/p113336_index.html

Donath, Judith S. (1996). *Identity and Deception in Cyberspace*. Retrieved from http://smg.media.mit.edu/people/Judith/Identity/IdentityDeception.html

Ford, B., Strauss, J. (April 01, 2008). An offline foundation for online accountable pseudonyms. *Proceedings of the 1st workshop on Social network systems*, p.31-36.

Friedman, E., Resnick, P. (2001). The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy 10*(2), 173-199.

Granovetter, M. S. (1973). The strength of weak ties. *American Journal of Sociology*, *78* (6), 1360 – 1380.

Gsell, L. (2009). Comments Anonymous. *American Journalism Review*, *31*(1), 16-17. Retrieved from Academic Search Complete database.

Halvais, A. (2009). Do dugg diggers Digg diligently? I*nformation, Communication & Society 12*(3). 444–459.

Herring, S. C. (1996). Two variants of an electronic message schema. In S. C. Herring (Ed.), *Computer-mediated communication: Linguistic, social and cross-cultural perspectives* (pp. 81-106). Philadelphia: John Benjamins.

Iannaccone, L. (1994). Why strict churches are strong. *American Journal of Sociology*, *99*(5), 1180-1211. Retrieved from Academic Search Complete database.
Kedzie, Christopher R., "A Brave New World or a New World Order?", in S. Kiesler (Ed), Culture of the internet, (p. 213). Mahwah, NJ.

Keyes, Saundra. (2009, Winter). Fiery forums. *The American Editor*. 10-12.

Klotz, Robert J. (2007). Internet campaigning for grassroots and astroturf support. Social *Science Computer Review, Vol. 25, No. 1*, 3-12.

Koehn, Daryl. (2003, Mar). Business ethics in the global knowledge economy. *Journal of Business Ethics*, *Vol. 43* (No 1/2), 3-19. http://www.jstor.org/stable/25074972

Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace, In: P. Kollock/M. Smith (*eds.*), *Communities in cyberspace*, London.

Kollock, P. and Smith, M. (1995). Managing the Virtual Commons: Cooperation and Conflict in Computer Communities. (S. Herring ed.) *Computer-Mediated Communication*. Amsterdam: John Benjamins.

Lampe , C. and Resnick, P. (April 22-24, 2004). Slash(dot) and burn: distributed moderation in a large online conversation space. *Proceedings of the SIGCHI conference on Human factors in computing systems*, 543-550.

Lee, A. Y. (2004)**.** Online news media as interactive community bulletin boards: SARS coverage in the Greater China Regions**.** *Paper presented at the annual meeting of the International Communication Association, New Orleans Sheraton, New Orleans, LA*.

O'Donnell, A., Jetten, J., & Ryan, M. (2010). Who is watching over you? The role of shared identity in perceptions of surveillance. *European Journal of Social Psychology*, *40*(1), 135-147. Retrieved from Academic Search Complete database.

Perrin, A.J. and Vaisey, S. (2008). Parallel public spheres: distance and discourse in letters to the editor. *American Journal of Sociology*, *114*, 781–810.
Resnick, P., Kuwabara, K., Zeckhauser, R., Friedmanm E. (2000). Reputation systems. *Communications of the ACM*, *43* (12), 45-48.

Resnick, P., Sami, R. (October 19-20, 2007). The influence limiter: provably manipulation-resistant recommender systems. *Proceedings of the 2007 ACM conference on Recommender systems*.

Ridings, C., Gefen, D., & Arinze, B. (2002). Some antecedents and effects of trust in virtual communities. *Journal of Strategic Information Systems, 11*(3-4), 271-295.

Ridings, C., Gefen, D. (2004). Virtual community attraction: why people hang out online. *Journal of Computer-Mediated Communication*, *10* (1). Article 4. Retrieved from http://jcmc.indiana.edu/vol10/issue1/ridings_gefen.html

Shane, Scott. (1994). *Dismantling utopia: How information ended the Soviet Union*. Chicago : I.R. Dee.

Sinha, R., Swearingen, K. (2002). The role of transparency in recommender systems. *CHI '02 extended abstracts on Human factors in computing systems*, 830–831.

Solove, Daniel J. (2007). The future of reputation: Gossip, rumor, and privacy on the Internet. New Haven: Yale University Press.

Stone, Martha. (2002). *Interactivity and Transparency*. Retrieved from http://www.poynter.org/content/content_view.asp?id=4703

Wilkinson, D., and Huberman, B. (2007). Assessing the value of cooperation in wikipedia. *First Monday 12*, 4.

Wilson, T. (2005). FUDBusters. *Network Computing*, *16*(3), 22. Retrieved from Academic Search Complete database.

Winsvold, M. (2009). Arguing into the digital void?. *Javnost-The Public*, *16*(3), 39-54. Retrieved from Academic Search Complete database.