

Searching for Safety Online: Managing “Trolling” in a Feminist Forum

Susan Herring

School of Library and Information Science, Indiana University, Bloomington, Indiana, USA

Kirk Job-Sluder

School of Education, Indiana University, Bloomington, Indiana, USA

Rebecca Scheckler

Virginia Polytechnic Institute and State University, Blacksburg, Virginia, USA

Sasha Barab

School of Education, Indiana University, Bloomington, Indiana, USA

A common phenomenon in online discussion groups is the individual who baits and provokes other group members, often with the result of drawing them into fruitless argument and diverting attention from the stated purposes of the group. This study documents a case in which the members of an online community—a feminist web-based discussion forum—are targeted by a “troll” attempting to disrupt their discussion space. We analyze the strategies that make the troller successful and the targeted group largely ineffectual in responding to his attack, as a means to understand how such behavior might be minimized and managed in general. The analysis further suggests that feminist and other nonmainstream online forums are especially vulnerable, in that they must balance inclusive ideals against the need for protection and safety, a tension that can be exploited by disruptive elements to generate intragroup conflict.

Keywords CMC, conflict management, deception, disruptive behavior, feminism, trolling

Received 21 November 2001; accepted 1 April 2002.

We received feedback on this article from multiple sources. We are especially grateful to Blaise Cronin, Mary Gray, Geert Lovink, John Paolillo, Barry Wellman, and three anonymous reviewers for their insightful comments and suggestions.

Address correspondence to Prof. Susan Herring, SLIS, Indiana University, Bloomington, IN 47405, USA. E-mail: herring@indiana.edu

Online discussion forums allow for convenient and ongoing communication among groups of people separated in place and time. In the best of cases, such forums can evolve into communities whose members share information, experience a sense of belonging, and provide mutual support (Preece, 2000; Rheingold, 1993). Moreover, the relative anonymity of the Internet can make people feel safe talking about issues that might be considered sensitive, inappropriate or dangerous in face-to-face public conversation (Donath, 1999; cf. Kiesler et al., 1984). These properties make online forums especially attractive to individuals seeking support for suffering from disease or abuse, and to members of minority social and political groups such as homosexuals, racial minorities, and feminists. Such groups can be considered vulnerable populations, in that they tend to be stigmatized and discriminated against by mainstream society.¹

At the same time, online discussion forums provide a new arena for the enactment of power inequities such as those motivated by sexism, racism, and heterosexism. The relative anonymity of the Internet releases some of the inhibitions of a civil society, resulting in flaming, harassment, and hate speech online (Ess, 1996). Despite the illusion they can give of security and privacy (King, 1996), online forums can be accessed by individuals hostile to the purpose of the forums, actively seeking to disrupt and undermine them. Moreover, the asynchronous, distributed

nature of online discussion forums allows those motivated to disrupt discussions to have far-reaching effects. These practices, while clearly problematic, are nonetheless widespread and often tolerated, due in part to the pervasiveness on the Internet of civil libertarian values that consider abusive speech a manifestation of individual freedom of expression (Pfaffenberger, 1996).

As a consequence of these characteristics of discussion forums, nonmainstream groups must confront a number of tensions online. Primary among these is the need to balance safety for participants with openness to free expression and discussion. Even in a community of individuals who share common values and experiences, a diversity of viewpoints can be expected, and differing views must be tolerated and respected if a climate of support and trust is to be achieved and diversity encouraged. But must online communities tolerate viewpoints that are directly in conflict with the goals of the community itself? In the case of groups already on the fringes of society, this may include harassing or hostile speech that reproduces the discrimination they face in mainstream society. While some participants find that challenging prejudice online can be an empowering act of resistance, others find that it diverts energy and attention away from the goals of the group (Collins-Jarvis, 1997). If a decision is made to restrict participation in an online forum to those individuals who support its goals, a second tension arises between setting protective boundaries and avoiding “ghettoization” of the group (Hall, 1996). A common question underlying these tensions is: When—and where—is it legitimate to draw the line?

In this study, we document a case study in which the members of a vulnerable online community—a feminist web-based discussion forum—respond to an individual attempting to disrupt their discussion space. Inasmuch as the individual represents himself insincerely (in this case, as interested in discussing feminism), we characterize his behavior as *trolling*, and his messages that lure members of the community into fruitless argument as *trolls* (Donath, 1999). With the exception of Donath (1999), who defines the phenomenon and gives examples from Usenet newsgroups, trolling has received little scholarly attention to date. In this study, we move beyond Donath’s limited, albeit useful, observations to analyze the specific mechanisms used by a troller,² and the responses of the group. The troller in this case succeeded in disrupting the group for nearly 2 months, but the group failed to reach a consensus regarding how to deal with him, despite unanimous agreement that he was a problem. The analysis sheds empirical light on the mechanisms of online deception and disruptive behavior, and points up the challenges of dealing effectively with such behavior in large, distributed, online groups, where consensus is often difficult

to achieve (Dubrovsky et al., 1991; Sudweeks & Rafaeli, 1996).

The remainder of this article is organized into six sections. We first review the literature on trolling and on disruption of online feminist spaces. We then discuss the data and methods used in this study. The third section analyzes in detail the troller’s behavior, including his ideological manipulation of tensions inherent in the feminist movement. This is followed by an analysis of the community’s response in which members attempt to isolate the troller by challenging, shunning, and calling for the forum administrators to ban him. The fifth section discusses the forum’s relative lack of success in managing the troller, noting the challenges inherent in such situations. We conclude by proposing interventions to forestall trolling behavior, and identifying further avenues for research into disruptive behavior in online discussion groups.

BACKGROUND

Trolling

Trolling entails luring others into pointless and time-consuming discussions. The name derives from the practice used in fishing where a baited line is dragged behind a boat (*Oxford English Dictionary*, 1992),³ although some Internet discourse refers to the troll as a fictional monster waiting under the bridge to snare innocent bystanders. Trolling often starts with a message that is “intentionally incorrect but not overly controversial.” In this respect, trolling contrasts with flaming, which is “an electronic mail or Usenet news message intended to insult, provoke or rebuke, or the act of sending such a message” (*Free Online Dictionary of Computing*, 1998). Trolling further differs from flaming in that the goal of flame bait is to incite any and all readers, whereas the goal of a stereotypical troll is to draw in particularly naive or vulnerable readers. Catching inexperienced users or “newbies” is a commonly stated aim of trollers (Andrew, 1996; Donath, 1999). As one Internet user, Andrew,⁴ states on his web site dedicated to trolling, “The object of recreational trolling is to sit back and laugh at all those gullible idiots that will believe *anything*” (Andrew, 1996). In practice, however, trolling and flaming often merge, in that in both cases there is intent to disrupt the ongoing conversation, and both can lead to extended aggravated argument.

Donath characterizes trolling as “a game about identity deception” (Donath, 1999, p. 45) in which all the participants are not cognizant of the nature of the game.⁵ The troller tries to write something deceptive, but not blatantly so, in order to attract the maximum number of responses (Andrew, 1996; Donath, 1999). Andrew extols a

successful troll:

His troll ran for over a year, it is known to have generated in excess of 3,500 responses (an average of 1 response every 160 minutes for a whole year) and the greatest coup of all was when an innocentamerican[sic] student lost not only her internet account but was also expelled from high school for abuse of the computer systems. Somehow she had managed to get the blame for causing the troll.

In the context of Usenet, where trolling first arose, a highly successful troll is one that is cross-posted to, and responded to on, many different newsgroups, thereby disrupting multiple groups with a minimum expenditure of effort. Andrew (1996) distinguishes “career trollers”—individuals who deliberately set out to disrupt groups and/or make trouble—from others motivated simply by the desire to attract attention.

The incident cited above had disastrous results for a female high school student, an outcome claimed as success by the troller. Other effects included the disruption of discussion and waste of bandwidth for a year. For Donath, there is a further repercussion, the loss of trust that can occur in a discussion group disrupted by trolling. The seriousness of this effect is dependent upon the nature of the group. Groups that deal with emotionally charged and sensitive topics, for example, groups for victims of rape and sexual abuse, are more at risk than lighthearted ones. Some participants opt not to post to groups following incidents of loss of trust (Brail, 1996). In this regard, vulnerable and inexperienced computer-mediated communication (CMC) participants not only are more likely to be targeted by trolling, but also may be more adversely affected by it. New users tend disproportionately to be women, the young, and other nontraditional computer users (Mowbray, 2001). For this reason, perhaps, Andrew (1996) refers to the generic target of trolling as “she.”

Disruption of Online Feminist Spaces

When women gather online, and especially when they attempt to discuss feminism, they are not uncommonly the target of negative attention from individuals, mostly men, who feel threatened by or otherwise uncomfortable with feminism. The literature on disruption of online feminist spaces dates back to the early days of computer-mediated communication research. Balka (1993) traces the history of four feminist forums from the 1980s, all of which experienced some degree of male harassment. Ebben (1994) describes the evolution on Usenet of the soc.feminism newsgroup, which was started in response to an earlier incarnation of the newsgroup soc.women having been taken over by men, and which itself has subsequently been taken over by men posting antifeminist and misogynistic messages (Sutton, 1994). Collins-Jarvis (1997) documents the

crisis that befell Comserve’s Gender hotline when several males began bombarding the forum with antifeminist messages, causing female subscribers to flee the group, and the forum eventually to be shut down. In a similar vein, Reid (1994) reports an incident on a MUD for sexual abuse survivors, in which a male-presenting character with the name “Daddy” traumatized the community by shouting graphic descriptions of violent sexual acts to those present on the MUD.

In the interests of insuring a space in which women feel safe to participate, feminists online have sometimes taken the separatist route of excluding males from participation. Hall (1996) describes the practices of a women-only discussion group for lesbians and bisexuals, its reasons for not allowing men, and the challenges it faces in enforcing its women-only policy. The women-only policy of *Systers*, an online forum for women in computer science, is explained and justified by its founder, Anita Borg, in Camp (1996). Herring et al. (1995) suggest that women-only groups, regardless of whether they discuss feminism, are a reaction to patterns of male domination in mixed-gender discussion groups on the Internet. Women are discouraged from participating in computer-mediated communication (CMC) by men posting more, longer, and more aggressive messages (Herring, 1994, 1999; Herring et al., 1995; Kramarae & Taylor, 1993; Spender, 1995), and by complaints that women are dominating the conversation even when such is not the case (Herring et al., 1995). Women-only groups create environments in which women can speak and be heard on topics of interest to them. At the same time, such groups are controversial: They risk being exclusionary and thereby provoking further male resentment (Hall, 1996), and they can become “ghettoes” in which women’s online presence is marginalized relative to the Internet at large (cf. Herring, 1994).

As an alternative to excluding male participants, some women-centered groups respond to disruptive or harassing behaviors by implementing participation policies that make it more difficult for future disruption to occur. Thus the Gender hotline reopened with a moderator who now filters all messages received before posting them (Collins-Jarvis, 1997). The MUD for sexual abuse survivors described by Reid (1994) implemented a process of identity verification, and disabled the feature that allowed users to communicate simultaneously with everyone in the MUD. Other groups introduce a policy that allows disruptive participants to be banned from the group, as occurred in the present study.

Disruptive incidents that force group members to articulate explicit norms and rules may also have the unintended effect of strengthening an online group’s self-definition as a community. A well-known example is the virtual rape that took place on LambdaMOO, in which the characters of two women were taken over by a male-presenting

character, MrBungle, and made to commit violent sexual acts on themselves in a public forum. This incident was greeted with widespread outrage in LambdaMOO, although the group could not agree on how MrBungle should be dealt with, even after a public meeting was held in the MOO to discuss it. Ultimately, a single wizard took matters into his own hands and “killed” MrBungle’s character. As a result of these disruptive events, a system of self-governance was established on the MOO, complete with elected officials, effectively institutionalizing the community’s newly articulated value system (Dibbell, 1993).

The reactions of online groups to harassment and disruption can be situated theoretically with respect to two dialectics that run through the literature on women’s online discussion groups. The first is the tension between libertarian values on individual freedom of expression, on the one hand, and communitarian values on the good of the group, on the other. In the libertarian view, the Internet is a new frontier, free from rules. Although most see freedom of speech as a feature of democracy (Ess, 1996), some libertarian discussions go so far as to argue for anarchy of the Internet (Barlow, 1996). A communitarian view of freedom of speech, in contrast, recognizes that less empowered persons might require buffering so that their rights to speech are preserved, and for the good of the community as a whole (Ellsworth, 1989; Ess, 1996; Herring, 1996b, 1999; Reid, 1999).

The literature about online harassment underscores the tension between libertarian and communitarian values, in that harassment often arises in spaces known for their freedom, lack of censure, and experimental nature (Brail, 1996; Dibbell, 1993; Reid, 1999). Herring makes an explicit connection between online harassment and libertarian values in a study of gender harassment, noting that “This ‘rhetoric of harassment’ crucially invokes libertarian principles of freedom of expression, constructing women’s resistance as ‘censorship’” (1999, p. 151).

The second dialectic is found in the literature on feminist stances (Gur-Ze’ev, 1999; Hall, 1996; Kenway & Nixon, 1999). Hall (1996) explicitly contrasts bringing women into the extant culture, the liberal view, with the provision of separate women’s spaces, the radical view. The literature on online harassment provides ample evidence as to why women might want separate online spaces. In addition to the research on silencing just cited, numerous studies report the use of CMC to annoy, intimidate, and harass women online (Dibbell, 1993; Donath, 1999; Ebben & Kramarae, 1993; Ebben, 1994; Herring, 1994, 1996b, 1999, 2002; Herring et al., 1995; Shade, 1993; Sutton, 1994; We, 1993).

The present study describes one trolling incident targeted at a feminist group, and the tensions attempts to manage it provoked within the group between libertarian/communitarian and liberal/separatist values.

THE CASE STUDY

Context

The trolling incident occurred on a web-based discussion forum sponsored by a large-circulation feminist magazine published in the United States.⁶ The purpose of the discussion forum is to provide a space for dialogue advancing feminist concerns and issues. The forum has over 4000 members, of whom about 200 participate actively. In the discussion analyzed in this article, 41 individuals participated, 90% of them female and 10% of them male. Participants sometimes disagree on individual interpretations of feminist ideology and action, but generally share an agreement that women are politically disadvantaged compared to men, and that feminism is the best way to address this problem.

In early February 2000, this agreement was challenged from two different sources. Several gun rights advocates from another forum joined the feminist forum exclusively to advocate against gun control legislation, starting more than a dozen new threads to argue their point of view. During the same period, a new male participant, Kent,⁷ started posting messages that were intentionally antagonistic to the core values of the forum. In his introduction to the forum, Kent identified himself as a middle-aged man in a professional position that involved overseas travel. He claimed to have been previously removed from other feminist forums for his views, and he also claimed he would eventually be removed from this feminist forum. He described himself as openly hostile to feminism, and started attacking forum members in dozens of posts spread throughout the forum.

Over a period of 8 days alone, more than 80 posts were written to a thread discussing Kent’s participation in that thread. Partly as a result of this discussion, the forum administrators adopted a new policy for participating in the forum (see Appendix A). Kent was eventually banned from the forum as a result of the new policy.

Procedures

Our analysis focused on Kent’s activities in the forum, as represented in a single thread of 111 messages posted between March 13 and March 21. This thread was selected because it contained the most explicit discussion by the group about how to respond to Kent’s behavior.

The data analysis used grounded theory methods (Strauss & Corbin, 1998; Glaser & Strauss, 1967) to develop a coding scheme for behaviors exhibited in the thread. After the first pass coding of the data, we reviewed our coding and identified patterns in terms of related themes.⁸ Once we came to an agreement about the main themes, and what they meant, we coded all member

posts accordingly and went back through the data to extract the examples used in this article.

The coding process was informed by the previous experiences of the authors in Internet discussion forums, and by the fact that one of the authors had participated in the forum for several months prior to this discussion. Our data analysis thus draws on observations and analyses of similar incidents in the past and on an intimate knowledge of the context of the forum.

ANALYSIS

The troller, Kent, was successful in disrupting communication on the forum for a period of almost 8 weeks. It is instructive to examine the strategies he used, and to consider why they were successful. In fact, Kent does not precisely fit the model of the disingenuous troller, but rather succeeds in disrupting the group by a combination of trolling and other means, including provocations attuned to the topic (feminism) and the audience (feminists; mostly women) of the group. Specifically, he exploits tensions within feminism itself with regard to freedom of expression and the legitimacy of separatist spaces, making it more difficult for the group to take effective action against him.

The Troller Provokes

Following Andrew (1996), we identify three definitional criteria for trolls:

1. Messages from a sender who appears outwardly sincere.
2. Messages designed to attract predictable responses or flames.
3. Messages that waste a group's time by provoking futile argument.

Kent's participation in the feminist forum meets all three criteria for trolling. Each of these behaviors is discussed in turn in what follows.

Outward Manifestations of Sincerity. In several messages, Kent presents himself as someone who is sincerely interested in debating the merits of feminism, and thus as a legitimate participant in the group. He appeals to feminists on the forum to provide "proof" against his anti-feminist claims:

Example 1

Kent 3-15-2000 04:48 PM

Every poster here has told me that I'm wrong and they are right about feminism. Do you see that? I at least offer proof. *I want to discuss*, not just drop a slogan and ride out throwing dismissive insults.

...

To prove or not to prove would obviously be the rightful subject of my entire time here. *If you disagree then you can say so.* You do. If you want to shut me up though, better be more convincing than just saying you disagree. [emphasis added]

He also claims to be sincerely unclear about why others have a problem with his postings, and appeals to them to explain it to him, implying (and in some cases, overtly stating) that he will modify the behaviors they find offensive:

Example 2

Kent 3-15-2000 01:38 PM

...

In summary what exactly is offensive about my posts? *If you can tell me I will either stop doing it or leave the board.* If however you refuse to tell me, and I've not been shy about asking SPECIFICALLY what standards I'm supposed to live by, then I will carry on doing it of course.

...

I'm aware of the trouble I'm causing you [web mistress]. I'm aware that you've been willing to go to that trouble (probably not on my account but on principle of course). I'm not ungrateful and *if you come up with some solution that involves me changing my behavior than please feel free to ask me.* [emphasis added]

In these messages, Kent attempts to present himself as rigorous and principled as regards the rules of debate, and potentially cooperative if others meet his conditions (providing proofs, answering his questions, etc.). Rather than provoking response through contentiousness, he appeals to others to respond in good faith. He attempts to appear outwardly sincere (criterion 1 of our definition of trolling) as a means of drawing others into responding.

Flame Bait. The posts by Kent that others find offensive are pejorative statements about feminism and feminists, including the members of the forum. Given the audience, these remarks are "designed to attract predictable comments or flames" (criterion 2). They include insults, name-calling, contentious presuppositions, and blatantly contentious assertions about feminism, as for example the following:

Example 3

Kent 3-13-2000 09:29 PM

...

Incidentally I take the silence over the gender wage gap hoax to mean that no feminist here even wants to TRY to defend their biggest lie: that men are paid more for the same work than women are.

[contentious presuppositions (the gender wage gap is a hoax; that men are paid more for the same work than women is a lie)]

Example 4

Kent 3-14-2000 01:51 PM

Gee, Simone, I dunno, maybe its because you're a bimbo who can't figure out the difference between an anecdote and

a statistic? If you want more money then get off your lazy ass and make some.

[name-calling; insults]

Example 5

Kent 3-14-2000 10:11 PM (quoted by another member)
Feminism is evil and bigoted and always has been. Just look at the bitches on this group. Frankly I don't see how you can bear to be near them.

[blatantly contentious assertion; name-calling; insults]

Still, although such posts are clearly intended to be offensive, it might be argued that Kent is still acting in good faith, according to his stated intent to debate the merits of feminism. Indeed, in one post he justifies the use of insults as necessary to his debate strategy:

Example 6

Kent 3-15-2000 01:38 PM

...

My personal feeling is that I don't want to insult people more than is necessary to expound my point of view. Those of you who think insults are not necessary should talk to those who have just said that, for example, use of the word "bigot" . . . should be considered an insult. Or talk to those who challenged me to get personal by dismissing all my points about the feminist leadership, history and activities as "just extremists" and said "point out hate on this board."

Now having said that personal challenges are a necessary part of my point of view [. . .]

Attempts to Provoke Futile Argument. However, further analysis makes clear that Kent is not, in fact, interested in the give-and-take of genuine debate, but rather in provoking futile argument (criterion 3). A fundamental uncooperativeness and perverseness is evident in Kent's covert rhetorical strategies, which include refusing to acknowledge others' points (or even that they have responded), willfully misinterpreting others' motives and views, and taunting others for not ignoring him.⁹

Example 7

Kent 3-15-2000 04:58 PM

[Context: Venus and others have made numerous arguments to which Kent has not responded]

...

Venus, if you ever get around to making any arguments I will reply to them in the same tone or better. Find me an example of where I haven't or didn't and I'll certainly apologise to whoever.

[Refusing to acknowledge others' points.]

Example 8

Kent 3-15-2000 08:10 PM

[Context: Simone pointed out that Kent had not responded to a challenge she made about the gender wage gap in response to Kent's previous challenge.]

Now Simone, isn't that better than biching and moaning? Sorry if you feel I've been ignoring you.

[Willfully misinterpreting another's point.]

Kent employs strategies of denial and distortion, even when others respond to his questions or attempt to correct his misrepresentations. As a consequence, he offers to modify his behavior if others will only "answer" his questions or provide "arguments" or "proof" ultimately appear insincere. He does not acknowledge anything that others post as answers, arguments, or proofs.

Finally, Kent manifests perversity in taking those who are critical of him to task for not doing what they say they will do, even though it is unfavorable to him:

Example 9

Kent 3-15-2000 10:21 PM

This conversation reminds me of the quote at the top of the board sometimes about high-heels. THINK. If you don't like reading my stuff than just DON'T ok?

Now is that so hard, you "strong women"?

Don't read it.

Don't reply to it.

Don't post stupid comments about it.

Don't make jokes about it.

Don't reply with pathetic insults.

Don't post about how your so NOT reading it.

Don't post asking others not to read it.

JUST CUT IT OUT FOT GOD'S SAKE

This post is paradoxical, in that Kent has been expending considerable effort to post over a period of several weeks. He clearly wants others to read and respond to him, yet here he exhorts the "women" in the group to ignore him. This allows him to taunt them for not being "strong" enough to do what they say they are going to do, and as such is a further put-down of feminism and the forum. At the same time, it sows confusion by introducing an appearance of arbitrariness into Kent's position. This strategy, together with the distortion and denial strategies described earlier, violates conventional rules of conversational cooperation (Grice, 1991[1968]). To the extent that Kent is engaging in such behavior intentionally, it suggests a motive on his part to "create chaos and confusion," one of the objectives of trolling (Andrew, 1996).

There is thus considerable evidence that Kent is a troller—that is, someone who is intentionally misrepresenting himself as interested in debating about feminism, but whose actual motive is to provoke and disrupt. Kent provides explicit support for this interpretation by his avowal that he has come to the group in order to provoke its members to kick him off, and in his boasting mention that he has been kicked off several feminist groups previously, suggesting that he views the activity as entertainment or sport.

At the same time, he is both less and more than a typical troll. He does not obey the principle of minimal expenditure of effort, nor does he attempt to engage other groups in the interaction. Moreover, he does not hide his trolling intent (although he does not refer to it in those terms), and his intentions are ultimately clear enough to the majority of group members—that is, he does not really “fool” them, inconsistent with another objective of trolling. Ultimately, his attempts at appearing sincere are too riddled with hostility and sarcasm to be persuasive.

Ideological Manipulation. What Kent lacks in deceptiveness, he makes up for in ideological manipulation of his audience. He exploits the tension between freedom of expression and the mostly female group’s interest in maintaining a civil environment by presupposing that the former should outweigh the latter:

Example 10

Kent 3-15-2000 04:48 PM

...

So to re-cap. What you are saying is that I should be banned because I keep saying all feminists are bigots and liars and so forth? In fact I should be banned for my well-documented and supported opinion on the very topic which this board is set up to discuss?

In another post, he suggests that women are “refusing to debate” with him because they fear the threat he poses to feminism; this ties back to his general claim that feminists are intolerant of debate. Later, he explicitly delegitimizes the communitarian value system that underlies the calls to ban him by describing the forum as a “girlies support group” characterized by “catty/emotional infighting,” a feature he also attributes to “early feminism as a movement.”¹⁰ This line of argument touches a nerve with this audience, in as much as feminism struggles to balance openness with a recognition of the need for women-centered spaces. Taken seriously, it places the group members in a double bind: If they allow Kent to continue, he pollutes their online environment with anti-feminist harassment; if they ban him, they close off debate and risk being labeled censorious. This bind may partially explain why forum members were unable to reach a consensus on how to deal with Kent.

The Group Responds

All but a small minority of participants expressed the view that Kent was a problem and agreed that his behavior was intended to undermine the forum. Moreover, most ultimately agreed that his posts were in violation of the group’s norms and values. However, despite widespread agreement on the existence and nature of the problem, the group could not agree on a course of action. Rather, participants split between calling for Kent to be banned, and calling for the group as a whole to ignore him in the hopes

that he would lose interest and go away. In practice, neither suggestion was followed: Participants engaged with Kent by trying to reason with him, and, when that failed, by insulting him in an escalating conflict, thereby falling into the trap the troll had laid for them. At the same time, the conflict led group members to negotiate explicitly what was appropriate discourse for the forum, reinforcing the group’s identity and leading to clearer limits on disruptive behavior. These responses are discussed next.

Calls for Administrative Banning. Many participants in the discussion proposed banning Kent administratively from participation on the forum system. Some explicitly invoked the communitarian or “radical” notion of protecting the forum as a “safe space” for feminists:

Example 11

Danielle (Member) 03-16-2000 01:09 PM

I can’t believe we’re discussing whether or not to ban [Kent]. There’s no question in my mind. Free speech does not include a long list of behaviors. Police are obligated to investigate death threats. Threats against the President or US Government have to be investigated by the Secret Service. Libel and slander laws limit what can be printed. Not all printed materials, like pornography, are available to all citizens, like children. Filing a false police report is a crime. There’s no excuse for [Kent]’s behavior. There’s no political rationalization.

I have only read a couple of his posts because I don’t need that shit. I see that some of you have engaged him and that’s your decision. But I wish you wouldn’t. *One thing that I’ve always looked forward to in feminism is the creation of “safe” spaces.* [Kent] is not going away on his own. I have no qualms about advocating the use of the [] boards with this in mind. [emphasis added]

Perhaps out of a concern that this stance could be interpreted as isolationist, censorious, or admitting female weakness (interpretations repeatedly articulated by Kent), others who favored banning offered a more legalistic justification, maintaining that Kent’s posts were in violation of the rules of the forum, and that he should be removed on those grounds.

Example 12

mizz-t (Member) 03-14-2000 02:55 PM

[Quotes an insulting post from Kent (example 4)]

—call me a rat, a tattle-tale, a trouble maker, a whiny bitch, whatever you like, but isn’t this violating the rules of the board already?????????

I have HAD it with this guy!!

The proposal to ban Kent met with considerable support, but it also encountered two obstacles. First, the forum members did not know what technological or administrative procedures would have to be followed in order to ban someone. In fact, no one had ever been banned from the forum before, and thus no procedures had been formally

established. Second, an equally vocal group of participants opposed banning on philosophical grounds.

Calls to Ignore the Antagonist Voluntarily. This group of members recommended simply not responding to Kent, suggesting that he would disappear if he did not get the kinds of angry responses he was seeking. Voluntarily ignoring Kent would deny him his audience, while maintaining the forum's dedication to free speech. These justifications are present in Emily's call to ignore Kent:

Example 13

Emily (Member) 03-16-2000 05:34 PM

I really don't think we should be banning or censoring anyone. If it becomes ridiculously extreme, sure why not. But I don't really think it is necessary here. I think if everyone starts ignoring him he's bound to go away eventually.

This is a characteristically libertarian approach to trolling, in that it relies on denying the troller's desire to stir up trouble (shunning) rather than administrative sanctions (banning).

At the same time, Glenda raises a central problem with shunning. Although shunning is presented as a passive strategy (i.e., just do nothing), in fact it requires considerable self-control not to respond to offensive provocation. Glenda argues that ignoring Kent is an appealing (and more effective) solution, but ultimately too difficult to carry out:

Example 14

Glenda (Member) 03-15-2000 12:04 AM

On to Kent—I keep TRYING to keep my mouth shut and ignore him, I really do! Guess I've got more masochistic tendencies than I thought. :rollseyes: Actually, I'm getting better at not offering any opinions of my own—just raking him over the coals for his. Shunning is ssoooo much more effective than banning—but that won't stop any junior members from unwittingly stepping into his trap.

Experienced forum members might ignore Kent's provocations, but new forum members would be tempted to respond in kind. Glenda's distinction makes an important point: Effectively shunning a disruptive individual requires a group consensus to follow through on ignoring the individual. Despite widespread agreement that ignoring Kent was a good idea, many participants continued to argue with him, thereby undermining the group's attempt to shun him.

Refuting the Antagonist's Claims. Forum members occasionally attempted to refute Kent's claims by answering his questions, suggesting counterexamples, or pointing out logical flaws. In a lengthy post, Marjorie challenges Kent on numerous fronts:

Example 15

Marjorie 03-15-2000 02:50 PM

'What was offensive'???? Being called a bitch, for starters. And referring to 'those bitches who run the shelters.'

And "Maybe you should try being a man and facing sexual rejection hundreds of times from bitches like you."

And "Yes you miserable **** you get to CHOOSE. The man, poor bastard, has no choice. Do you comprehend the difference princess?" And "What you mean is the feminist fag-boy self-flagellation view of men's issues?"

And "Gee, Marilyn, I dunno, maybe its because you're a bimbo who can't figure out the difference between an anecdote and a statistic?" And all the bullshit on Julie.

What's offensive is the fact that you repeatedly run back to the basest terms you can: bimbo, bitch, princess, baby, girl. Anything to let us know we're less than people.

And for all the prolific posting you do, you can't name a single instance of anyone here hating men, not caring about men, blaming men for everything, etc. Anyone who comes close is immediately reminded (by us) that THAT'S NOT WHAT FEMINISM IS ABOUT. All you can do is go on and on about how you know what we *really* mean, even when it's the opposite of what we say. Twist the words around, call us liars and bigots, and treat us like shit, all to "prove" your point. You call EVERYTHING we say "feminist propaganda lies." As determined by you, definition courtesy of you. Why the fuck should *anyone* bother?

But what is most insulting is the continual whining about how put-upon you are. How HOUNDED you are. Like you didn't come here and TELL US you were going to hound us. Like you didn't come here SPECIFICALLY to get this reaction. Like your previous experiences didn't clue you into the fact that when you treat people like shit, THEY WON'T WANT YOU AROUND. And you DO treat people like shit, constantly. Me? I get it because I point that fact out. You seem to have no ability whatsoever to refrain from your verbal abusiveness, and then we get to hear about how not only is it our fault, but you're doing it FOR us, out of the goodness of your heart.

Personally, I'm sick to death of YOUR lies (all feminists are bigots and liars, we blame men for everything, blah, blah, blah). And how lucky you are to just be handed this forum for saying anything you want, and how badly you abuse it and us, and then go running to your "oh, poor me" excuses while calling *us* victims, because you're sooo innocent of any wrongdoing. Grow up.

This post is characterized by an angry and aggressive tone, including heavy use of sarcasm such as "you're sooo innocent of any wrongdoing." Marjorie's intent appears to be to shame Kent by "telling it like it is."

Insulting the Antagonist. In other posts, forum members simply return insults for insults, effectively lowering themselves to Kent's level:

Example 16

Sharon 03-15-2000 04:54 PM

i don't think you could even give us an accurate summary of your ass Kent . . .

This example came after Kent's lengthy explanation of his goals on the forum (Example 6). Sharon directly insults Kent's ability to participate in the discussion, using vulgar language to do so. In other posts, forum members portray Kent as immature, suggesting he is an 8-year-old child pretending to be an adult.

Another strategy used by some forum members is the off-record insult. Donald made an inflammatory statement in the context of an ongoing discussion about Kent, yet without mentioning Kent's name:

Example 17

Donald 03-14-2000 10:23 PM

Tell me if this sounds like anyone you know:

"Batterers are very into making excuses and presenting themselves as victims. They really see other people . . . as abusing or attempting to control them. It's the way to rationalize, minimize or deny their own behavior."

Just curious.

Kent responded with anger a few posts later, clearly interpreting the comment as directed at him personally. In response, Donald invoked plausible deniability, taking the opportunity indirectly to insult Kent further, implying he was paranoid for becoming offended:

Example 18

Donald 03-15-2000 03:45 PM

I just want to point out that if people assume I'm talking about them when I say bad things, that's not my fault. My comments are like birds I set free on the wind, and if someone wants to catch them and hold on to them as their very own, that is their choice.

"Paranoia, paranoia. . . everybody's coming to get me . . ."

Donald distances himself from the insult by quoting another text rather than making his own statement. For the most part, however, members avoided insulting Kent as a person, instead criticizing his posting style and his disruptive effects on the forum.

Challenging Kent by refuting his claims and insulting him undermined attempts to shun him. As in the cases of gender harassment described by Herring (1999), insults and refutations were used by the troller as a springboard for further attacks.

Negotiating What Is Appropriate. In Example 6, Kent expressed a controversial view of what the appropriate norms of online debate should be. He wrote that personal attacks were appropriate and necessary to achieving the political goal of challenging feminism. This view, and the behavior that accompanied it, forced the group to define more concisely what they believed to be appropriate and inappropriate styles of participation in the forum.

According to some members, an appropriate challenge focuses on a person's ideas, while an inappropriate chal-

lenge focuses on the person expressing the ideas.

Example 19

Donald 03-14-2000 03:17 PM

Yeah, I think (hope) that just about everybody can see the difference between attacking someone's ideas and attacking someone personally. It may seem like a slight distinction from a semantic standpoint, but it is wholly significant.

"This argument is dumb" vs. "You are dumb."

"You sound like a bimbo" vs. "You are a bimbo."

"I disagree" vs. "I think you're a stupid bitch who needs to get the fuck up off her lazy fat ass and stop sitting around the house stuffing her face full of twinkies and shooting heroin and also giving pamphlets about the Devil to little children who happen to come by selling Girl Scout cookies."

We could also discuss the quantity of space taken up with a given argument, constructive vs. destructive arguments, being respectful vs. being a tool, etc.

A further point of discussion, introduced by the forum moderator, concerned the use of obscenities. Several participants in the discussion addressed the difference between a nonspecific use of obscenities, "What the fuck?," and obscenities directed toward a specific person such as "Fuck you." Nonspecific use of obscenities was considered to be emphatic, while obscenities directed at a specific person were considered hostile.

Most importantly, and with surprisingly little discussion, the group came to an agreement that personally insulting or offensive speech that persisted after warnings from the moderator would not be tolerated: After three such warnings, the offender would be banned from the forum. A new policy statement to this effect appeared on the forum Web site for the first time on March 15 (Appendix A). However, it would not be applied to Kent until two weeks later, at which time the forum moderator—like the wizard who "killed" MrBungle in the LambdaMOO case (Dibbell, 1993)—acted independently to ban him.

DISCUSSION

Why was this group not more effective in defending itself against the troller's attacks? We propose four explanations for this lack of success, the first ideological, the second psychological, the third experiential and the fourth relating to the nature of online forums.

As with the MrBungle case, forum members were caught between conflicting ideologies. Liberal and libertarian views advocate letting everyone participate, and combating problematic speech through debate. Communitarian views focus on maintaining safe space; together with radical views, they lead to the creation of separate environments such as those focused on women's concerns (Hall, 1996). Kent effectively exploited this tension—inherent in

the situation of any group that is vulnerable by virtue of being a target of discrimination or harassment—by pushing the bounds of harassing behavior, at the same time invoking principles of free speech and open debate. Moreover, by daring forum members to ban him—indeed, by making getting banned his goal—Kent guaranteed that he would “win” regardless of the outcome of the forum’s deliberations. This form of ideological manipulation was especially effective given that his audience was a feminist forum committed in principle to inclusiveness.

At the same time, the troller’s token displays of interest in feminist issues, and his token expressions of willingness to be convinced by evidence, psychologically manipulated members into continuing to engage with him, thereby prolonging an interaction that had seriously disruptive effects on the forum. Why do people respond to provocation, even when they recognize intellectually that angry responses are what is being sought? Grice (1991[1968]) observed that meaningful communication rests on a default assumption of mutual cooperation, leading communicators to assume that others are generally trying to be truthful, clear, consistent, etc., even when surface appearances suggest otherwise. Moreover, communicators are rationally motivated to protect one another’s social face, on the premise that harmony is more likely than conflict to produce desirable social outcomes for all involved (Brown & Levinson, 1987).¹¹ In contrast, a troller is fundamentally uncooperative: He seeks to confuse and deceive, rather than to be clear. The troller in the present study is also fundamentally unconcerned with maintaining others’ social face: On the contrary, like a flamer, he seeks to maximize face threats by means of insults and put-downs. Such behavior appears irrational to many online communicators, to whom it might never have occurred that anything useful could be gained by harassing and disrupting others. Accordingly, they persist in attempting to reason with the disruptive individual, to appeal to his better nature, or, failing that, to shame him. Their belief in the universality of the social contract may partially blind them to what we take to be the troller’s actual motivations: the desire to attract attention, including negative attention; and the desire to exercise control and feel superior by manipulating others to fall into a trap of the troller’s design.¹²

The target group in this case was also inexperienced with trolls, and thus responded more naively to the troller than a more experienced group might have done. Although the members of the forum sensed that Kent was trying to make trouble, they did not apply the word “trolling” to his behavior, nor did they appear to have available the mental concept of trolling by any other name. This is perhaps not surprising, given that trolling has traditionally been a Usenet phenomenon, and Usenet, although large, is still unknown to many Internet users. More surprising is that as feminists, the members of the forum did not relate Kent’s

behavior to a wider pattern of antifeminist harassment on- or offline. When early feminist Internet discussion groups such as WMST-L (the Women’s Studies List; Herring, 1996a) and those discussed by Balka (1993) were harassed by males wanting to “discuss feminism,” they quickly recognized the behaviors as online manifestations of the same disruptive harassment that feminists often encounter in the real world, and responded accordingly.

Why might the reactions of online feminist groups to disruption have changed in the intervening decade? First, the women in the early groups and those in the present study are different populations, individually and demographically. The early adopter, highly educated, upper-middle class ’70s-style feminist of the 1980s and early 1990s has given way to the younger and more demographically diverse feminists who populate Web forums, most of whom, consistent with the demographics of Web forum participants more generally, have come online within the past 5 years. In the absence of direct continuity between the earlier and the current groups, there is little opportunity for the experiences of the former to have informed the awareness of the latter. Second, tolerance for disruptive behavior on the Internet appears to have increased since the early 1990s. The notion that hostile or harassing speech online was protected under the First Amendment was a contested proposition in the early days, advanced primarily by men, and resisted by some women, who were accused of “censorship” for their views (Herring, 1996b, 1999). Since that time, the notion has attained near-hegemonic status, and is voiced by both women and men. Kent invoked free speech rhetoric in his abusive messages, and the members of the forum tolerated them for longer than they might have otherwise, we suspect, because of the prevailing view that abusive online communication is “free expression,” and hence legitimate. Such a view both encourages trolling behavior and makes it harder to resist.

An additional factor that abets disruptive activity is the difficulty of achieving consensus in online groups. Text-based CMC has been claimed to result in more frequent disagreements, greater polarization on controversial issues, and longer times to reach consensus than face-to-face interaction (Sproull & Kiesler, 1991; cf. Sudweeks & Rafaeli, 1996). These effects can be overcome to some extent by centralizing authority in the role of a moderator or group administrator, removing the requirement for absolute group consensus on each decision. The group analyzed in the present study operated in a decentralized manner that required consensus (as opposed to a majority vote) in order for a decision about the troll to be implemented. Since consensus could not be achieved, no decision was taken. The web mistress’s subsequent intervention effectively implemented a more centralized authority model, which the forum members appeared to welcome, since the forum did not provide individual members with the

technological means to ban or filter messages from other users. Thus both the social organization and the technological properties of the forum made it difficult for users to protect themselves from harassment by the troller, thereby inadvertently facilitating the troller's disruptive goals.

CONCLUSIONS

We conclude by suggesting several proactive interventions that might help to forestall individuals and groups from being trolled, yet not suppress debate.

The first is to make use of filters to block the harasser's messages. We believe that such technological enhancements can play a useful, if limited, role. A "killfile" capability permits individual users not to view posts from selected other individual users. Killfiles shift some of the filtering abilities of moderators from a centralized administrator to the individual participant, thereby preserving a decentralized structure and individual freedom of speech (Spertus, 1996). At the same time, killfiles do not exclude the posts from the view of other readers, nor from the archives of the forum. As in the case of the virtual rape previously cited, social damage can effectively be done to individuals without their reading the offending post (Dibbell, 1993). Moreover, since killfiles are reactive, users necessarily view some objectionable messages before they set a killfile, making it only a partial filter even for individual users.

Equally important is to educate users about trolling. Trollers stereotypically (although not exclusively) prey on inexperienced Internet users and populations that are vulnerable for other reasons. Forum administrators might warn users about the patterns that trollers follow. Simply naming the danger can heighten people's awareness of it. Because the danger is emotional rather than physical, we can imagine that warning about trolling might be similar to warning about phone pranks or sales scams, where awareness of the modus operandi is often sufficient to forestall the effect of the advantage-taking event. A number of Web sites give advice to individuals about how to avoid being trolled,¹³ including how to identify trolls (a message makes unbelievable claims, is cross-posted to a large number of newsgroups, etc.), how to respond to trolls (resist the temptation to respond; set a filter to block messages from the troller's email address), and where to report trollers (e.g., to their Internet Service Providers, which often have policies against trolling, flaming, and other forms of online harassment).

Perhaps while we are educating users, we might also inform them of the lack of anonymity of Internet communication, no matter how safe and secure a discussion site may appear. Users need to be aware of the practice of archiving Internet transcripts, of how easily messages can be disseminated to other Internet venues, and of the fact that

at least one system administrator always has access privileges to the contents of their servers, even when messages have been deleted. Greater awareness might lead users to reflect before responding hastily to provocative messages, since such messages could potentially come back to haunt them later.

The interventions just cited apply primarily to individuals. The present case also points to the need for online forums to articulate policies, guidelines for appropriate participation, and penalties for violating those guidelines, in advance of harassment episodes taking place. Public online spaces are likely to experience disruption from trolling and flaming unless policies and capabilities are implemented for excluding problem users. It is necessary in this regard to distinguish clearly between cooperative debate (however heated) and uncooperative provocation (however masked). Unambiguous and strong moderation from the start can avoid many problems (for an example, see Korenman & Wyatt, 1996). Some evidence suggests that groups vulnerable to harassment and trolling benefit especially from stricter centralized moderation (Herring, 2000).

We should also seek to understand the social psychology of trolling so as to be better able to predict its occurrence and limit its harmful effects. Why do trollers troll? In cases where stigmatized groups such as feminists or homosexuals are the target, we might speculate that the underlying motivation is hatred towards people who are perceived as different or threatening by the troller. When groups are targeted for their vulnerability alone, as in the case of disabled people or inexperienced users, a motivation suggests itself in terms of control and an enhanced sense of self-empowerment. However, not all disruptive online behavior is targeted at at-risk populations—trollers have disrupted discussion groups for sports car racers (male),¹⁴ video software artists (technologically sophisticated)¹⁵ and even Internet researchers (sophisticated about online social practices).¹⁶ A common denominator in the trolling cases that we are aware of is that trollers appear to enjoy the attention they receive, even—and perhaps especially—when it is unremittingly negative.¹⁷ This evidence suggests that ignoring the provocations of a troller might indeed be an effective way of thwarting him (or her),¹⁸ and that the definition of trolling needs to be expanded beyond the traditional stereotype to encompass a wider range of attention-seeking disruptive behaviors.

Given the proliferation of Internet communities and the numbers of people who stand to be impacted by trolling and disruptive online behaviors, more research is clearly needed in this area. Interviews with trollers, some of whom are openly proud of their trolling exploits, might be carried out to shed light on the motivations for trolling.¹⁹ Comparative studies are needed to determine if patterns of trolling differ when mainstream groups are the target rather than

minority groups, when trolling is instigated via a single message or sustained participation, and when the troller is multiple individuals rather than a single individual²⁰ or a software agent rather than a person.²¹ Studies might also be carried out to show how trolling is affected by the availability of technical tools that give participants greater control over the online environment.

Ultimately, we see Internet trolling as a manifestation of a much broader phenomenon whereby individuals take pleasure in disrupting the social order out of anger, perversity or contempt. It would be interesting to compare online disruption with disruptive behavior in face-to-face groups in playgrounds, classes, meetings, support groups, and social events. It may be that the textual nature of computer-mediated disruption makes its strategies more apparent and leads to insights that would otherwise elude us in spoken interactions, but which can then be applied to offline as well as online situations. At the same time, the Internet poses unique opportunities for abusers and concomitant challenges for those wishing to avoid abuse. Among the challenges we see are the need to set clear boundaries on acceptable speech in order to ensure that civility, safety, and freedom can coexist.

NOTES

1. We exclude from this category political minorities such as white supremacists whose members come mostly or entirely from the dominant (white, heterosexual, able-bodied, male) group. Feminists are included on the grounds that the label carries a stigma outside of feminist contexts, and self-declared or suspected feminists are frequently targets of abuse. This is the case even within academia, where feminism has achieved a degree of legitimacy in terms of publication, curricular, and hiring practices. The vulnerability of feminists is, of course, related to the fact that most feminists are women. Women could also be considered a vulnerable group according to a broad interpretation of the definition given here.

2. In discourse about this phenomenon on the Internet, the terms “troll” and “troller” are used interchangeably to refer to a person who sends disingenuous, provocative messages (“trolls”). To avoid confusion, we use “troller” to refer to the person, and “troll” to refer to the message, consistent with Donath (1999).

3. The fishing metaphor in fact gives rise to three different meanings for the word:

- (a) To troll for flames (most common usage, as in the present discussion).
- (b) To troll for (cyber)sex (e.g., <http://www.leathernroses.com/cyber/trolling.htm>).
- (c) To dredge for data, e.g., search for some specific information on the Web (e.g., <http://opop.nols.com/write/HTTP.html>).

The second and third meanings predate the Internet. Trolling in the first sense has been practiced and discussed on the Internet since at least the early 1990s; two well-known early trollers are Usenet personalities James “Kibo” Parry and Ted Frank (<http://www.killfile.org/~tskirvin/faqs/legends/legends1.html>).

4. As with many people attempting to disrupt online communication, Andrew conceals his offline identity. A search of his web site reveals no mention of his last name, location, or other identifying characteristics.

5. For a game-like approach to trolling, see Perry’s (n.d.) criteria for earning the “Trolling Merit Badge,” which involves successfully trolling seven kinds of targets: newbies, idiots, political newsgroups, the opposite sex, “oldbies,” famous persons, and kibo (cf. note 3).

6. Access to the discussion forum requires users to register as members by providing their names and e-mail addresses, but membership is free and open to anyone.

7. We have replaced all forum participant names with pseudonyms.

8. Examples of initial coding categories include “taunt,” “blame,” and “contentious assertion.” Examples of themes include “outward manifestations of sincerity,” “ideological manipulation,” and “refuting claims” (see Analysis).

9. All typographical errors in the examples were produced by the original authors.

10. We note in passing that Kent makes a number of remarks that demean not just feminists, but women in general. He evokes misogynistic stereotypes through the use of expressions such as “girlies support group,” “catty,” and (earlier) “hysterical” to describe a woman who was critical of him.

11. Individuals deviate from these ideals, sometimes deliberately, but cooperation and politeness remain default expectations in everyday communication. See Herring (1994, 1996b), however, for discussion of gender differences in the value placed on politeness.

12. This behavior is rational, in as much as it provides the troller with the outcomes he desires, albeit arguably for psychologically unhealthy reasons.

13. See, e.g., Gilmore (2000), Sackman (1995), <http://www.agoraphobiazone.net/trolling.html>.

14. Barry Wellman (e-mail communication, Mar 30 2002 15:19:40).

15. Mieszkowski (2002).

16. An extended trolling episode took place on the Association of Internet Researchers List (AIR-L) between fall 2001 and spring 2002.

17. Group responses to a troller are not always entirely negative. In both the AIR-L case and the cases reported by Mieszkowski (2002), some members defended the troller. What motivates people to do this is a question requiring further research.

18. The troller described by Mieszkowski (2002) self-presents as a woman—actually several women—under the pseudonym Netochka Nezvanova.

19. A challenge associated with interviewing trollers is that they may not represent themselves in a straightforward way, as Mieszkowski (2002) observed.

20. See footnote 18.

21. Serdar Argic, a notorious troll from the early 1990s who posted thousands of messages to multiple discussion groups about the Turkish–Armenian conflict in World War I, is believed to have been a software agent (bot) (Rob Furr, message posted to air-l@aoir.org, 21 Mar 2002 14:45:09; for other theories, see <http://www.kkc.net/eyenet/1994/net0728.htm>).

REFERENCES

Andrew. 1996. *The troller’s FAQ*. Retrieved 1/10/01 (<http://www.altairiv.demon.co.uk/faq/posts/trollfaq.html>)

- Balka, Ellen. 1993. Women's access to on-line discussions about feminism. *Electronic Journal of Communication* 3 (1). Retrieved 11/11/01 (<http://www.cios.org/www/ejc/v3n193.htm>)
- Barlow, John Perry. 1996. *A declaration of independence of cyberspace*. Retrieved 1/18/01 (http://www.eff.org/pub/Censorship/Internet/censorship_bills/barlow_0296.declarationv)
- Brail, Stephanie. 1996. The price of admission: Harassment and free speech in the wild, wild west. In *Wired_women: Gender and new realities in cyberspace*, eds. L. Cherny and E. R. Weise, pp. 141–157. Seattle, WA: Seal Press.
- Brown, Penelope, and Levinson, Stephen. 1987. *Politeness*. Cambridge: Cambridge University Press.
- Camp, L. Jean. 1996. We are geeks, and we are not guys: The systems mailing list. In *Wired_women: Gender and New Realities in Cyberspace*, eds. L. Cherny and E. R. Weise, pp. 114–125. Seattle: Seal Press.
- Collins-Jarvis, Lori. 1997. *Discriminatory messages and gendered power relations in on-line discussion groups*. Paper presented at the 1997 annual meeting of the National Communication Association, Chicago.
- Dibbell, Julian. 1993. A rape in cyberspace: How an evil clown, a Haitian trickster spirit, two wizards, and a cast of dozens turned a database into a society. *Village Voice* 23 December.
- Donath, Judith S. 1999. Identity and deception in the virtual community. In *Communities in cyberspace*, eds. M. A. Smith and P. Kollock, pp. 29–59. London: Routledge.
- Dubrovsky, Vitaly J., Kiesler, Sara, and Sethna, Beheruz N. 1991. The equalization phenomenon: Status effects in computer mediated and face-to-face decision making groups. *Human Computer Interaction* 6:119–146.
- Ebben, Maureen M. 1994. *Women on the net: An exploratory study of gender dynamics on the soc.women computer network*. Unpublished doctoral dissertation, University of Illinois, Urbana-Champaign.
- Ebben, Maureen, and Kramarae, Cheris. 1993. Women and information technologies: Creating a cyberspace of our own. In *Women, information technology, and scholarship*, eds. H. J. Taylor, C. Kramarae, and M. Ebben, pp. 15–27. Urbana, IL: Women, Information Technology, and Scholarship Colloquium, Center for Advanced Study, University of Illinois at Urbana-Champaign.
- Ellsworth, Elizabeth. 1989. Why doesn't this feel empowering? Working through the repressive myths of critical pedagogy. *Harvard Educational Review* 59(3):271–297.
- Ess, Charles. 1996. The political computer: Democracy, CMC, and Habermas. In *Philosophical perspectives on computer-mediated communication*, ed. C. Ess, pp. 197–230. Albany: State University of New York Press.
- Free online dictionary of computing*. 1998. Retrieved 11/11/01 (<http://wombat.doc.ic.ac.uk/foldoc/contents.html>)
- Gilmore, Glynn. 2000. *Ignore trolls on newsgroups*. Retrieved 7/2/02 (<http://www.angelfire.com/biz4/glynnegilmore/page11.html>)
- Glaser, Barney, and Strauss, Anselm L. 1967. *The discovery of grounded theory: Strategies for qualitative research*. Chicago: Aldine.
- Grice, H. Paul. (1991 [1968]). Logic and conversation. In *Pragmatics: A reader*, ed. S. David, pp. 305–315. New York: Oxford University Press.
- Gur-Ze'ev, I. 1999. Cyberfeminism and education in the era of the exile of the spirit. *Educational Theory* 49(4):437–455.
- Hall, Kira. 1996. Cyberfeminism. In *Computer-mediated communication: Linguistic, social and cross-cultural perspectives*, ed. S. C. Herring, pp. 147–170. Amsterdam: John Benjamins.
- Herring, Susan C. 1994. *Gender differences in computer-mediated communication: Bringing familiar baggage to the new frontier*. Retrieved 11/11/01 (<http://www.cpsr.org/cpsr/gender/herring.txt>)
- Herring, Susan C. 1996a. Two variants of an electronic message schema. In *Computer-mediated communication: Linguistic, social and cross-cultural perspectives*, ed. S. C. Herring, pp. 81–108. Amsterdam/Philadelphia: John Benjamins.
- Herring, Susan C. 1996b. Posting in a different voice. In *Philosophical perspectives on computer-mediated communication*, ed. C. Ess, pp. 113–145. Albany: State University of New York Press.
- Herring, Susan C. 1999. The rhetorical dynamics of gender harassment on-line. *The Information Society* 15:151–167.
- Herring, Susan C. 2000. Gender differences in CMC: Findings and implications. *Computer Professionals for Social Responsibility Newsletter*, Winter 2000. Retrieved 11/11/01 (<http://www.cpsr.org/publications/newsletters/issues/2000/Winter2000/index.html>)
- Herring, Susan C. 2002. Gender[sic] violence: Recognizing and resisting abuse in online environments. *Asian Women* 14(summer): 187–212.
- Herring, Susan, Johnson, Deborah A., and DiBenedetto, Tamra. 1995. "This discussion is going too far!": Male resistance to female participation on the Internet. In *Gender articulated: Language and the socially constructed self*, eds. K. Hall and M. Bucholtz, pp. 67–96. New York: Routledge.
- Kenway, J., and Nixon, H. 1999. Cyberfeminisms, cyberliteracies, and the educational cyberspheres. *Educational Theory* 49(4):457–474.
- Kiesler, Sara, Siegel, Jane, and McGuire, Timothy W. 1984. Social psychological aspects of computer-mediated communication. *American Psychologist* 39:1123–1134.
- King, Storm. 1996. Researching Internet communities: Proposed ethical guidelines for the reporting of results. *The Information Society* 12:119–127.
- Korenman, Joan, and Wyatt, Nancy. 1996. Group dynamics in an e-mail forum. In *Computer-mediated communication: Linguistic, social and cross-cultural perspectives*, ed. S. C. Herring, pp. 225–242. Amsterdam: John Benjamins.
- Kramarae, Cheris, and Taylor, H. Jeannie. 1993. Women and men on electronic networks: A conversation or a monologue? In *Women, information technology, and scholarship*, eds. H. J. Taylor, C. Kramarae, and M. Ebben, pp. 52–61. Urbana, IL: Women, Information Technology, and Scholarship Colloquium, Center for Advanced Study, University of Illinois at Urbana-Champaign.
- Mieszkowski, Katharine. 2002. The most feared woman on the Internet. *Salon.com*, March 1, 2002. Retrieved 5/17/02 (<http://www.salon.com/tech/feature/2002/03/01/netochka/print.html>)
- Mowbray, Melinda. 2001. Reducing demographic bias. In *Online communities: Commerce, community action, and the virtual university*, eds. C. Werry and M. Mowbray, pp. 97–125. Upper Saddle River, NJ: Prentice Hall.
- Oxford English Dictionary*. 1992. Retrieved 9/28/99 (<http://etext.lib.virginia.edu/oed.html>)
- Perry, John. n.d. Net Scouts Home Page. Retrieved 7/2/02 (<http://www.perry.com/bizarre/scout.html>)
- Pfaffenberger, Brian. 1996. "If I want it, it's OK:" Usenet and the (outer) limits of free speech. *The Information Society* 12:365–386.

- Preece, Jennifer. 2000. *Online communities: Designing usability, supporting sociability*. Chichester, UK: John Wiley.
- Reid, Elizabeth. 1994. *Cultural formations in text-based virtual realities*. Unpublished master's thesis, University of Melbourne, Australia. Retrieved 6/15/01 (<http://home.earthlink.net/~aluluei/cult-form.htm>)
- Reid, Elizabeth. 1999. Hierarchy and power: Social control in cyberspace. In *Communities in cyberspace*, eds. M. A. Smith and P. Kollock, pp. 107–133. London: Routledge.
- Rheingold, Howard. 1993. *The Virtual community: Homesteading on the electronic frontier*. Reading, MA: Addison-Wesley. Retrieved 6/15/01 (<http://www.rheingold.com/vc/book>)
- Sackman. 1995. *Usenet fish tales—How to avoid being trolled*. Retrieved 7/2/02 (<http://netconference.about.com/gi/dynamic/offsite.htm?site=http%3A%2F%2Fwww.intersurf.com%2F%7Eaevinc%2Faev2trol.htm>)
- Shade, Leslie Regan. 1993. Gender issues in computer networking. Retrieved 11/11/01 (http://cpsr.org/cpsr/gender/leslie_regan_shade.txt)
- Spender, Dale. 1995. *Nattering on the net: Women, power, and cyberspace*. North Melbourne: Spinifex.
- Spertus, Ellen. 1996. *Social and technical means for fighting on-line harassment*. Retrieved 7/3/02 (<http://www.ai.mit.edu/people/ellens/Gender/glc>)
- Sproull, Lee, and Kiesler, Sara. 1991. *Connections: New ways of working in the networked organization*. Cambridge, MA: MIT Press.
- Strauss, Anselm, and Corbin, Juliet. 1998. *Basics of qualitative research: Techniques and procedures for developing grounded theory*. Newbury Park, CA: Sage.
- Sudweeks, Fay, and Rafaeli, Sheizaf. 1996. How do you get a hundred strangers to agree? Computer mediated communication and collaboration. In *Computer networking and scholarly communication in the twenty-first-century university*, eds. T. Harrison and T. Stephens, pp. 115–136. Albany, NY: SUNY Press.
- Sutton, Laurel. 1994. Using Usenet: gender, power, and silence in electronic discourse. In *Proceedings of the 20th Annual Meeting of the Berkeley Linguistics Society*, eds. S. Gahl, A. Dolbey, and C. Johnson, pp. 506–520. Berkeley, CA: Berkeley Linguistics Society.
- We, Gladys. 1993. *Cross-gender communication in cyberspace*. Retrieved 11/11/01 (ftp://cpsr.org/cpsr/gender/we_cross_gender)

APPENDIX A

Rules:

[Name] magazine's policy is to allow free debate on our boards, as long as users follow basic rules of human decency. Personally attacking, flaming or threatening another [name] board member is strictly forbidden. Threatening violence against a group of people (like Jews, or homosexuals, or feminists, for example) will also not be tolerated. If you violate these rules, you will be banned from the [name] boards.

So behave!