

The Experience of 'Bad' Behavior in Online Social Spaces: A Survey of Online Users

John P. Davis

Social Computing Group

Microsoft Research

One Microsoft Way

Redmond, WA 98052-6399 USA

jpdavis@microsoft.com

+1 425 706-5860

ABSTRACT

Online antisocial or 'bad' behavior appears to be a serious and pervasive problem in a variety of online social settings. This paper presents results from an online survey designed to collect information about the frequency, context, and effects of aversive behavior, and the methods they employ to combat it. As expected, respondents perceived that bad online behavior occurs frequently and has a strong negative effect on online interactions. Most respondents reported that others' bad behavior had caused them to leave or avoid online social spaces. Participants reported that the methods they use to combat bad behavior are not very effective. Descriptions of bad behavior suggest that many are perpetrated by people users do not know, but a surprising number are perpetrated by acquaintances. Further, bad behavior often spans online domains, such that while the behavior may start in one place (e.g., a chat room), it may continue through other channels (e.g., email). Implications for designing more effective methods to thwart bad behavior using established social principles are discussed.

Keywords

Computer-mediated communication, bad behavior, deviant, aversive, online survey, social interaction

INTRODUCTION

Deviant or 'bad' behavior is pervasive in many computer-mediated social contexts. For example, chat rooms are often a venue where users make inappropriate sexual advances, swear at others or engage in other behaviors that most users deem inappropriate. Users in such environments are often relatively anonymous, and norms against such behavior do not exist or are weak. The price of bad behavior is high; users who are the targets of bad behavior may leave or avoid online social spaces. Effective responses, both at the individual and systems levels, are also costly and may not be very effective at curtailing bad behavior.

There have been several qualitative explorations and case studies of online spaces where deviant behavior has taken place and the methods used to deal with it (e.g., [12, 8]), primarily at the systems level. To date, however, there has been little systematic empirical research to explore the frequency with which bad behavior occurs and its effects, the experience of bad behavior from the perspective of the target of the behavior, the methods people employ to counter deviant behavior and the perceived efficacy of those methods. Clear understanding of individual users' current experiences with inappropriate behavior is key to developing more effective technological methods for reducing such behavior.

The current study explores users' experiences of bad behavior in several online settings through a survey that was deployed online. A primary goal of the survey was to collect information about the frequency, context, and effects of aversive behavior. The survey was also designed to gather information about users' perceptions of the efficacy of attempts to curtail bad behavior. In other words, the survey was designed to gather information about what methods users thought worked to prevent or diminish bad behavior and what methods did not. In addition, the survey was intended to gather detailed descriptions of the instances of what users' considered to be bad behavior, and their reactions to the behavior. From the information gathered through the survey, we hope to design better, more effective socially-based technological methods to prevent bad behavior before it occurs and correct it after it occurs.

BACKGROUND

The scope of the problem of aversive online behavior is massive. Tales of online verbal harassment in chat rooms and other group systems are common, for example [1]. Email spam is frequently identified as a pervasive annoyance that affects most users. The competitive nature of online games and the age of many gamers lead to a great deal of aversive behavior between players. Flaming and other forms of bad behavior also occur in newsgroups and other threaded discussion groups [8, 10]. Though anecdotal reports suggest that the problem of aversive behavior is nearly universal, little empirical data has been collected to support this notion. As mentioned earlier, bad behavior is costly to those who administer online contexts because

users who experience bad behavior will leave the space and never return.

What Constitutes Bad Behavior?

Whether or not a behavior is deemed 'bad' is determined by the target person's interpretation of the behavior. Something that is offensive to one person may be innocuous to someone else. Further, the context in which the behavior occurs is a strong determinant of whether it is judged to be bad. For example, in a game setting, certain behaviors on the part of users are considered to be 'part of the game,' while in other settings, identical behaviors would be deemed unacceptable. For the purposes of the study described here, 'bad behavior' was defined as *any aversive behavior users felt did not belong in a particular online environment*. In other words, respondents defined what behaviors they considered to be bad.

What Causes Bad Behavior?

There are several possible causes of bad online behavior. Anonymity and deindividuation have each been shown to cause aversive behavior in 'the real world' of face-to-face social interactions [3], and they are even more pronounced in online settings [5]. True identities are difficult to pin down in online social spaces, and reputations do not carry the same weight online that they do in face-to-face interactions. Social presence, the feeling that others are in a social environment with you, is also diminished in online settings [9]. Users may not care if they hurt other users because they have little sense that others are 'real,' little expectation that their bad behavior has consequences for them, and little expectation that they will ever have to interact with the other person in the future.

Methods for Combating Bad Behavior

Prevention and remediation of aversive online behavior is difficult and expensive (see Suler & Phillips [1] for a useful taxonomy of interpersonal intervention strategies based on a case study of The Palace). High-level social methods for preventing and dealing with bad behavior involve carefully crafting the norms or standards for behavior in an online social space as the space develops. Rules and guidelines for behavior that users must accept before they are allowed to participate in a space or participation in an orientation session are examples of ways users can be inculcated with norms for acceptable behavior. Modeled behavior from other users serves the same purpose.

Additional social methods for preventing bad behavior can be built into an online social space. For instance, reputation systems are useful for this purpose. Such systems are largely intended to give users a method for checking into the social and behavioral history of someone they are interested in interacting with in some way (see Jensen, Davis & Farnham [4] for additional information about reputation systems). However, they may also help prevent deviant behavior by making users accountable for their actions because a good reputation is an important commodity in many contexts. User profiles, which often contain more personal and personalized information about

users, may also serve to decrease bad behavior. Access to personal information about other users may act to make the online space more social' and less anonymous (see Swinth, Farnham & Davis [11] for additional information about user profiles).

Other technical methods for detecting and dealing with the behavior exist, but they are expensive and difficult to design and implement. For example, MSN and AOL have both used some form of moderator to intervene on a user's behalf when the user clicks a button to report an incident involving bad behavior. Moderators work well to diminish bad behavior, but they are extremely expensive. Digitalcity.com has an application that detects foul language and bans perpetrators when they use profanity. Such applications can be useful for dealing with relatively simple instances of bad behavior (e.g., foul language), but they are not adept at dealing with more complex, subtle forms of aversive behavior. Other system level methods for dealing with bad behavior require real people to monitor users' behavior in some way and react to deviance. Because they involve live monitors to police a designated space, they are extremely expensive and administratively cumbersome. While they are more able to deal with the subtleties surrounding online behavior in sometimes complex contexts, they must also be adept at interpreting behavior and they must be carefully trained to avoid abuses of power.

Most of the responses to bad behavior take place at an individual level, rather than at a systems level. Users who experience bad behavior must deal with it in the immediate situation. They may confront the offender directly, tell an administrator about the behavior, flee the situation, simply ignore the behavior altogether or find another method. While a variety of responses to bad behavior are available to targeted users, they may not work particularly well.

Research Goals

This exploratory survey was designed to assess users' experiences with deviant or 'bad' behavior in online social settings. Specifically, the research goals were to gather descriptive information about the types of behaviors users perceived as deviant in different online contexts, and also to gauge the frequency with which such behaviors occur within those contexts. While researchers and online administrators have many ideas about what principles and interventions work to prevent or diminish bad behavior (e.g., [2]), an additional goal of the current research was to determine which methods *users* perceive as most effective. As described earlier, the ultimate goal is to design more effective systems for preventing and correcting deviant behavior in online social spaces.

METHOD

Participants

A nonrandom sample of 497 respondents was recruited to complete the survey through the Social Computing Group's Online Laboratory (<http://www.onlinelab.org>). Participants for another study conducted on the Online Laboratory had

been recruited to participate through a large online gaming web site, and many also participated in this survey. Most respondents were either advanced or expert users of the Internet (60.4%, n = 300), and most (51.7%, n = 257) reported using the Internet 21 or more hours per week. See Table 1.¹

Table 1. Percent of respondents' that reported participating in each online activity, Internet experience, and time spent online

<i>Type of Online Activity</i>	<i>n</i>	<i>%</i>
Send and receive email	386	77.7%
Chat or instant messaging	430	86.5%
Newsgroups or discussion groups	311	62.6%
Online multiplayer games	483	97.2%
None of the above	3	0.6%
Internet Experience		
Beginner	7	1.7%
Intermediate	57	11.5%
Advanced	152	30.6%
Expert	148	29.8%
Did not respond	133	26.8%
Hours Per Week Spent Online		
None	1	0.2%
1-5	3	0.6%
6-10	14	2.8%
11-15	32	6.4%
16-20	55	11.1%
21+	257	51.7%
Did not respond	135	27.2%

Survey Instrument

The survey instrument developed for this study contained approximately 40 questions that asked users about their general experiences with bad behavior and specifically about instances of such behavior in four online social contexts: games; chat rooms; email; and newsgroups. The specific questions were developed through conversations with a variety of online users, and items were iteratively pre-tested. While some of the questions were designed to capture information about the frequency with which different behaviors occurred, others were more exploratory in nature.

Filter questions allowed us to tailor the survey based on the user's individual responses so they did not see questions that did not apply to them (e.g., if a respondent had not experienced aversive behavior in a particular context, they did not see the set of questions associated with that section of the survey).

¹ Because this study is exploratory in nature and the sample was not randomly selected, one should be cautious about generalizing to the larger user population.

There were several response formats for the survey items. Some questions asked users to check all responses that applied to them (e.g., "Which of the following online systems have you used?"). Some questions asked respondents for a yes or no response (e.g., "Have you ever left an online environment or interaction because another user behaved inappropriately in some way?"). Additional questions asked respondents to select the answer that best applied to them (e.g., "How did you feel about the inappropriate behavior you described?"). In some instances, respondents were asked free-response questions to give them a chance to describe some aspect of the bad behavior they had experienced or, where appropriate, to further explain an answer they had given. Additional items asked respondents for demographic information, including a self-assessment of their level of expertise in the use of the Internet and the amount of time they spent online.

RESULTS

Frequency of Bad Behavior

To get a sense of how often users think deviant behavior occurs in online settings, we asked them to estimate its frequency. See Table 2. Overall, 66.4% (n = 243) of our respondents estimated that bad behavior occurs often or very often. Only 1.0% (n = 4) of users thought it never or seldom occurred. In addition, users were asked to report whether or not they had experienced bad behavior in each of the social instruments they said they had used. As Table 2 shows, a significant percentage of users in each domain said that they had experienced some form of bad behavior. These results mirror the commonly held intuition that aversive online behavior is quite pervasive.

Table 2. Reported frequency of bad behavior

	<i>n</i>	<i>%</i>
How frequently does bad online behavior occur?		
Never	2	0.5%
Seldom	2	0.5%
Occasionally	37	10.1%
Often	140	38.3%
Very often	103	28.1%
Not sure	82	22.4%
Have you experienced bad behavior (of those who reported that they used each)		
In email	131	33.9%
In chat or instant messaging	175	40.7%
In newsgroups or discussion groups	86	27.7%
In online multiplayer games	364	75.4%
Compared to face-to-face, how often does online bad behavior occur?		
Far less often	18	4.9%
Less often	15	4.1%
About as often	43	11.8%
More often	108	29.7%
Far more often	180	49.5%

Face-to-face (FTF) communication is often used as the benchmark or comparison against which computer-mediated communication modes are judged. We asked users to assess the frequency with which bad behavior occurs in online settings compared to FTF interactions. As Table 2 shows, 79.2% (n = 288) reported that bad behavior occurs more often or far more often online compared to FTF.

Effects of Bad Behavior on Respondents' Behavior

Survey Results

The finding that users perceive that bad behavior occurs often does not on its own suggest that such behavior is a real problem, however. It may be that while users see deviant behavior frequently, they do not view it as problematic. In fact, such behavior could be part of the allure of the Internet, as suggested by one respondent:

It amuses me. Sometimes it's like watching a soap opera

One of the strongest arguments that deviant behavior has a negative impact on social environments would be that it adversely affects the behavior of other users. To explore how inappropriate behavior affected the respondents' behavior, we asked them whether they had ever *avoided* visiting an online environment because they were concerned that someone might behave inappropriately and whether they had ever *left* an environment because someone had behaved inappropriately. A majority of respondents (56.1%, n = 279) reported that they had avoided visiting an online environment because they perceived a threat that other users might behave inappropriately. Further, 79.3% (n = 394) of respondents reported that they had left an online social environment because another user behaved inappropriately.

Qualitative Results from Descriptions

The survey findings are supported by the content of the events users described. Many incidents of bad behavior suggested that the behavior of the targets of bad behavior is strongly affected by the bad behavior directed at them:

Most message boards are manifestations of bad attitudes. I tend to stay away.

I don't post to message boards anymore as I don't want to be flamed.

One thing that really ticks me off is when you subscribe to something, then they give your email to other people. This has led me to not subscribe to anything for almost the last 2 years.

Unfortunately, it seems that users experience a great deal of bad behavior online and such behavior negatively impacts their own.

Responses to Bad Behavior and Their Efficacy

Survey Results

Respondents were asked to report how they thought bad behavior was responded to in general, and how effective they thought such responses are at decreasing the behavior.

See Table 3. Note that respondents seldom reported that the behavior was dealt with administratively. The response that users said occurred most often was that the perpetrator was attacked by other users (41.0%, n = 150). Importantly, most participants reported that responses to bad behavior are either somewhat or very ineffective at decreasing it (51.3%, n = 187).

Table 3. Responses and effectiveness of responses to online bad behavior (Note: ns and % based on number of respondents who answered these questions)

	n	%
<i>How is inappropriate behavior responded to most often?</i>		
Little or no response	79	21.6%
Perp "attacked" by other users	150	41.0%
Perp ignored by other users	91	24.9%
Perp dealt with administratively	20	5.5%
Other	26	7.1%
<i>In general, how effective are responses to bad behavior at decreasing it?</i>		
Very effective	14	3.8%
Somewhat effective	102	28.0%
Neither effective nor ineffective	61	16.8%
Somewhat ineffective	97	26.6%
Very ineffective	90	24.7%

Results from Each Online Social Space

Respondents were asked to report a particular incident of bad behavior in which they were the target in each of the online social environments. For each, they were asked to report how they responded to the bad behavior and the results of their response on the perpetrator of the behavior. See Tables 4, 5, 6, and 7. Note that the possible responses available to users differed according to the environment in which the bad behavior occurred, and that respondents could select as many responses as were applicable in the situation.

There are several important things to note from the results shown in the tables. The first is that the targets of the bad behavior often elect to *not* respond to the behavior (27.5% in online games, 37.7% in chat, 45.3% in newsgroups, and 56.1% in email). A second important result to note is that the most often cited outcome of responses to bad behavior was that the perpetrator escalated the bad behavior after the response. In email, escalation was less common, perhaps because the perpetrator is more identifiable than in other contexts. Reporting perpetrators to administrators appeared to be a last resort, perhaps because it takes more effort or because users do not consider administrative response to be effectual.

Table 4. Responses to bad behavior in games (note that respondents could choose more than one option)

GAMES	<i>n</i>	%
<i>How did you respond to the bad behavior?</i>		
Chose not to respond	100	27.5%
Could not respond	21	5.8%
Responded directly to perp	214	58.8%
Informed administrator	37.6	37.6%
Left the game	109	29.9%
<i>What did the perpetrator do?</i>		
Escalated the bad behavior	183	50.3%
Repeated bad behavior	108	29.7%
Repeated bad behavior elsewhere	30	8.2%
Stopped behaving badly	56	15.4%
Administrator handled	62	17.0%
Don't know	94	25.8%

Table 5. Responses to bad behavior in email (note that respondents could choose more than one option)

EMAIL	<i>n</i>	%
<i>How did you respond to the bad behavior?</i>		
Chose not to respond	74	56.1%
Responded angrily (to perp)	13	9.8%
Responded reasonably (to perp)	25	18.9%
Informed administrator	33	25.0%
Changed email addresses	15	11.4%
<i>What did the perpetrator do?</i>		
Escalated the bad behavior	21	15.9%
Repeated bad behavior	26	19.7%
Repeated bad behavior elsewhere	19	14.4%
Stopped behaving badly	8	6.1%
Administrator handled	19	14.4%
Don't know	74	56.1%

Table 6. Responses to bad behavior in chat (note that respondents could choose more than one option)

CHAT	<i>n</i>	%
<i>How did you respond to the bad behavior?</i>		
Chose not to respond	66	37.7%
Could not respond	2	1.1%
Responded directly to perp	79	45.1%
Informed administrator	32	18.3%
Left the chat room	59	33.7%
<i>What did the perpetrator do?</i>		
Escalated the bad behavior	63	36.0%
Repeated bad behavior	31	17.7%
Stopped behaving badly	28	16.0%
Administrator handled	28	16.0%
Don't know	68	38.9%

Table 7. Responses to bad behavior in newsgroups (note that respondents could choose more than one option)

NEWSGROUP/DISCUSSIONS	<i>n</i>	%
<i>How did you respond to the bad behavior?</i>		
Chose not to respond	39	45.3%
Could not respond	2	2.3%
Responded directly to perp	10	11.6%
Posted angry response	10	11.6%
Posted reasonable response	24	27.9%
Posted humorous response	15	17.4%
Informed administrator	8	9.3%
Left the newsgroup	26	30.2%
<i>What did the perpetrator do?</i>		
Escalated the bad behavior	34	39.5%
Perp stopped behaving badly	16	18.6%
Administrator handled	8	9.3%
Don't know	38	44.2%

Why Does Bad Behavior Occur?

Respondents were also asked to select reasons they thought bad behavior occurred from a list of several possibilities. See Table 8. The most often selected reason was that online spaces are anonymous, an intuition that is supported by ample empirical research that has found that anonymity increases aversive behavior in FTF and CMC situations [3, 5].

Table 8. Reported influences and causes of bad behavior

	<i>n</i>	%
<i>What is the influence of bad online behavior?</i>		
Very negative	152	41.4%
Somewhat negative	155	42.2%
Neither positive nor negative	54	14.7%
Somewhat positive	5	1.4%
Very positive	1	0.3%
<i>Why does online bad behavior occur?</i>		
No punishment	197	39.6%
No fear of punishment	258	51.9%
Perp doesn't see effects on others	177	35.6%
Anonymity online	296	59.8%
Not enough methods to deal with it	161	32.4%
Too much freedom online	33	6.6%
Misbehavers crave attention	214	43.1%

Additional reasons for bad behavior that were often selected were that perpetrators did not fear punishment for their actions (selected by 51.9% of the respondents) and that perpetrators crave attention (selected by 43.1% of the respondents). The latter result suggests that the best method for dealing with some types of 'bad behaviors' is to ignore them. However, results from other questions demonstrate that although participants often reported that they did not

react to instances of bad behavior, they often felt that responses generally were ineffective. The idea of using ‘ostracism’ as a way to thwart bad behavior is discussed at the conclusion of this report.

Additional Qualitative Results

A surprising finding from the descriptions of the incidents respondents provided was the number of instances of bad behavior that spanned several online social interaction modes:

In a chat room, someone got overtly sexual with me and started harassing me. I then got many E-Mails from this person. I reported him and have never gone in a chat room since. This was about 5 years ago.

I used to receive all sorts of nasty, or "targeted" mail that was just awful. I noticed it most whenever I would visit a newsgroup or chat room. As a result I simply NEVER go into a chat room anymore.

While we may tend to think of bad behavior in terms of discrete incidents involving only a single online social domain, this is often not the case.

The way in which some of the instances of bad behavior were described was sometimes amusing. Following are a few scenarios that users described:

A teenage girl kept trying to convert me to Christianity! She sent me endless cloying emails and barf worthy URLs. The emails were all lengthy and forwarded beyond belief. Several were based on guilt and thinly veiled threats.

I was in a discussion for new mothers. Someone made negative comments about the name another chose for her child. The kid will get enough grief later.

A number of the descriptions talked about the qualities of the perpetrators as a factor that influenced their tendencies to behave badly, and their own views about online bad behavior:

Men have HUGE tempers when people do not agree with them. Have you ever noticed the colorful language that follows?

Basically it revolves around kids (<16 yrs old) getting on and acting stupid/inappropriate because they have the insulation of the internet to get away with it. Adults still do dumb things but they usually will at least discuss conflicts with out resorting to bad stuff.

My beliefs make it very hard to have anything be inappropriate behavior short of a guy with a machete or chainsaw massacring people in a K-mart.

Another finding was that to some extent, users are aware of the dangers of interacting in online social spaces, and they accept the risks involved. Several were circumspect about the trade-offs in such interactions

Unsolicited spam happens all the time. But we live in a free internet world and do not advocate restrictions. We have delete buttons and we have the ability to screen spam. That is sufficient on an individual basis.

The last unbridled wild west of the net. Usenet groups are just that. Unleashed. So you get what you expect.

Most of the incidents users reported contained descriptions of deviant behaviors that one might expect to occur in online settings. The incidents were largely perpetrated by people users did not know and were most often impersonal acts. One user summarized the most often mentioned instance of online bad behavior concisely:

SPAM, SPAM, SPAM

However, a surprising number of the incidents of bad behavior described by users came from people they knew in each of the social spaces we explored (i.e., email, online games, chat, and newsgroups). The following are typical examples:

Ex-girlfriend sent me an e-mail telling me that I shouldn't participate in a mailing list as she was there.

I got emailed from someone I had dated a couple of times, and in whom I was not interested, because he wanted to change my mind...

I received an accidental fwd from a secretary at my work. I responded stating that I wasn't the person she intended to email. Through it all she started emailing me in an attempt to start a "relationship." She also sent interoffice memos to me.

This last result has important implications for how we design better methods for correcting or limiting bad behavior, which are described in the conclusion to this paper.

CONCLUSIONS

Results from this study confirm that online bad behavior is a pervasive and real problem. Respondents reported that they experience such behavior frequently in several online social environments, and it adversely impacts their own behavior in several ways—they leave and avoid contexts where bad behavior occurs. The cost of bad behavior is therefore high. Typical users pay the price because they cannot enjoy the social aspects of the Internet fully, and administrators or owners of online social spaces lose valuable users that flee aversive interactions. One user summed up the situation in his experience with

newsgroups, but this description can easily be applied elsewhere:

There's no one incident. It normally it goes like this:

- 1) *someone posts a question, maybe a little stupid*
- 2) *someone picks on the question*
- 3) *flames start.*

Oldest story on the internet

Because many system level responses, like having live moderators to respond to bad behavior, are extremely expensive, it would be preferable to have individual users deal with bad behavior instead. However, the results also show that typical user responses directed at perpetrators of bad behavior are ineffectual.

Designing More Effective Methods

A broader approach to curtailing bad behavior must work to gather the responses of individuals and bring to bear the social influences that work well in traditional social spaces (e.g., [7]). More effective methods for correcting bad behavior could capitalize on simple social principles that are effective in other domains.

Anonymity granted in online environments and a lack of accountability (fear of punishment) have been identified as two of the primary causes of bad behavior. There are several system level methods that can help decrease bad behavior through increasing the 'social presence' [9] of users in online social spaces. Online profiles are one way to decrease the anonymity associated with online social spaces. Reputation systems are also useful in making users more accountable for their actions in many situations. Careful construction of an online social space to ensure that norms against bad behavior develop early can be very effective. Kelly, Sung and Farnham [6] describe a case study of how designing for development of positive norms and a well-constructed reputation system can be effectively used for conduct control in a large, active web-based community.

Other social means for curtailing bad behavior remain largely unexplored. For instance, ostracism, which is the act of ignoring or excluding someone from a group, is a phenomenon that has a powerful impact on the target of the ostracism [13]. In chat rooms, users can often click an 'ignore' or 'mute' button to prevent other users' bad behavior (e.g., foul or offensive language) from appearing in the interface. However, critical feedback to the offender—that the user has ostracized him or her—is missing. One can imagine that as more and more users ignore an offending user while sending him or her feedback to that effect, offense rates would diminish, at least for some offenders in some domains. Such social feedback would be particularly efficacious if it were to come from users in the offender's social network. The descriptions of bad behavior reported here suggest that the behavior often originates with someone in the user's social network (though often peripheral), which suggests that systems for dealing with bad behavior could leverage this powerful

quality. Feedback from someone we know is impactful in many social settings.

It has become essential to address the causes of bad behavior in online environments. Our future research plans include testing the effectiveness of systems to combat bad behavior that include social mechanisms.

REFERENCES

1. Aiken, M., and Waller, B. Flaming among first-time group support system users. *Information & Management*, 37, 95-100.
2. Bruckman, A., Curtis, P., Figallo, C., and Laurel, B. Approaches to managing deviant behavior in virtual communities, in *Proceedings of CHI '94* (Boston, MA, April 1984), ACM Press, 183-184.
3. Diener, E., Fraser, S., Beaman, A. L., & Kelem, R. T. (1976). Effects of deindividuating variables on stealing by Halloween trick-or-treaters. *Journal of Personality and Social Psychology*, 33, 178-183.
4. Jensen, C., Davis, J., and Farnham, S. Finding others online: Reputation systems for social spaces, in *Proceedings of CHI '02* (Minneapolis, MN, April 2002), ACM Press.
5. Jessup, L. M., Connolly, T., & Galegher, J. (1990). The effects of anonymity on GDSS group process with an idea-generating task. *MIS Quarterly*, 14, 312-321.
6. Kelly, S. U., Sung, C., and Farnham, S. D. Designing for improved social responsibility, user participation and content in on-line communities, in *Proceedings of CHI '02* (Minneapolis, MN, April 2002), ACM Press.
7. Kollock, P., and Smith, M. (1996) Managing the virtual commons: Cooperation and conflict in computer communities. In *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*, Herring, S., Editor. Amsterdam: John Benjamins.
8. Phillips, D. J. (1996). Defending the boundaries: Identifying and countering threats in a Usenet newsgroup. *Information Society*, 12, 39-62.
9. Short, J., Williams, E. & Christy, B. (1976). *The social psychology of telecommunications*. Chichester: Wiley.
10. Smith, C. B., McLaughlin, M. L., and Osborne, K. K. (1997). Conduct control on Usenet. *Journal of Computer Mediated Communication*, 2 (4).
11. Swinth, K., Farnham, S. D., and Davis, J. P. (2002). Sharing personal information in online community profiles. Unpublished manuscript, Social Computing Group, Microsoft. Research.
12. Suler, J. R., and Phillips, W. L. (1998). The bad boys of cyberspace: Deviant behavior in a multimedia chat community. *CyberPsychology & Behavior*, 1, 275-294.
13. Williams, K. D., Cheung, C. K. T., and Choi, W. (2000). Cyberostracism: Effects of being ignored over the Internet. *Journal of Personality and Social Psychology*, 79, 748-762.

