

Conforming to the Flaming Norm in the Online Commenting Situation

Peter J. Moor
peter@petermoor.nl

ABSTRACT

A certain kind of online behavior, called flaming, consists of exhibiting hostility towards other people by insulting, swearing or using otherwise offensive language. An experiment has been conducted to test whether perceived norms have an effect on flaming behavior in the online commenting situation, a situation where people can comment on a certain stimulus. This has been done in a natural setting, where participants did not know about the experiment until they had commented on a text. Participants flamed more often when earlier commenters had done so, indicating that conformation to the flaming norm indeed occurred. The results could, however, not be fully explained by the perspective of the SIDE model used in this study.

1. INTRODUCTION

1.1 Flaming and the SIDE Model

Computer-mediated communication (CMC) is quite different from face-to-face (FTF) communication. Probably one of the most important differences is the (perceived) anonymity supported by computer environments such as the Internet. When online, people tend to act less inhibited. This is expressed in a higher level of self-disclosure (Joinson, 2001), which can be useful for computer administrated questionnaires (Weisband & Kiesler, 1996). Other forms of uninhibited behavior, however, are less appreciated. People may display great hostility online by insulting, swearing or using otherwise offensive language. This kind of behavior, usually referred to as “flaming”, has been thought to be due to the lack of social cues in computer environments (Collins, 1992). The inability to see the expression on a sender's face or hearing his voice when reading a typed message has been thought to affect people's perception of themselves and others. Submersion in the medium and reduced self-awareness might evoke a deindividuated state, which enables people to show impulsive behavior that is normally considered inappropriate, such as flaming (Siegel, Dubrovsky, Kiesler, & McGuire, 1986; Kiesler, Siegel, & McGuire, 1984).

An alternative explanation is offered by the social identity model of deindividuation effects, known as SIDE (Reicher, Spears, & Postmes, 1995). This model, based on the social identity theory (Tajfel & Turner, 1986; Tajfel & Turner, 1979) and the self-categorization theory (Turner, 1987), states that people in a deindividuated state do not lose their sense of individuality or self-awareness. Rather, their personal identities make room for social identities. Because of the anonymity, people are relatively indistinguishable and their memberships of online discussion groups are far more salient than their personal identities. A shift from a personal to a social identity, called “depersonalization”, is facilitated. Thus, rather than displaying impulsive and uninhibited behavior, people in CMC will conform to perceived group norms. This may explain why some research has found that flaming is rare (e.g. Coleman, Paternite, & Sherman, 1999; see Lea, O'Shea, Fung, & Spears, 1992 for a review), while other research suggests that it is very common (e.g. Alonzo & Aiken, 2004; Aiken & Waller, 2000). Some groups may maintain a much more hostile norm than other

groups, depending on the standards invoked by dominant group members.

The SIDE model has been found to be a better predictor of CMC behavior than the deindividuation theory in a review of 60 studies (Postmes & Spears, 1998). Among other kinds of behavior, flaming has been found to be influenced by social norms within a group (Postmes, Spears, & Lea, 2000). More research has confirmed that visual anonymity predicts group self-categorization, which in turn predicts group attraction and positive stereotyping of fellow group members (Lea, Spears, & De Groot, 2001). Also, people tend to conform to group norms more when they are anonymous (Postmes, Spears, Sakhel, & De Groot, 2001). According to Joinson (2003, p. 50), the SIDE approach to CMC has received little critiques.

One more thing needs to be discussed about flaming. Much research has been done on the topic, but very different definitions of the term have been used (Lea et al., 1992; Thompsen, 1996). While the original meaning of the word “flaming” was “to speak rabidly or incessantly on an uninteresting topic or with a patently ridiculous attitude” (Steele et al., 1983, cited in Thompsen, 1996), researchers often used the word to denote the expression of emotions during CMC. Most researchers only address the negative and hostile side of emotional behavior, using definitions consisting of words like profanity, hostility, insults and swearing. In this paper, flaming will be defined as “displaying hostility by insulting, swearing or using otherwise offensive language”. This definition seems to agree with most definitions used in previous research.

According to Thompsen (1996), most research seems to be focused on the act of flaming, while its perception may be just as important for its definition. Flaming can not be defined by the act only, because “a flame is not a flame until someone calls it a flame.” (p. 302). While this is probably a very good point, the present research is still focused only on the act of flaming. Many flames, especially when they are quite extreme, are probably considered flames by most people. We are interested in the occurrence of this behavior, not in its perception. We'll come back to the definition of flaming in the next section, when the online commenting situation has been explained.

1.2 The Online Commenting Situation

Most research on flaming has focused on situations where a group of people had to discuss a certain topic or reach agreement on a decision task. In the CMC situation, this was usually done either by e-mail or by some synchronous chat program. The major activity was discussing, which is by definition a long-term process.

A completely different situation on the Internet where flaming seems to occur is what in this text will be referred to as the online commenting situation. In the online commenting situation, people are asked to comment on certain content. This content can be anything like a news article, a video, a song, a text or a website. Sometimes commenting can be done anonymously, other times people are required to log in to an account so their nickname is visible with each comment. People can also comment on each other's comments like in discussions, but there is a stimulus that is meant to be the main subject of all comments. Besides, contrary to discussions, it is very common

that people write only one comment and then leave the webpage without the intent to return and read new comments. Unlike discussing, commenting is usually a short-term process.

One example is the widespread phenomenon of weblogs. A weblog (or 'blog') is a website where people publish their own texts. Some people have weblogs about the news, like columns in a newspaper. Other people use their weblogs as public diaries. The ability for readers to comment is quite common. Recently, some famous Dutch people closed their weblogs because they did not want to cope with the hateful and insulting reactions anymore (Van Stein Callenfels & Van Woerden, 2007; "Onvriendelijke reacties...", 2007).

Another example is YouTube.com, a website where people can upload their own videos. People are allowed to comment on videos, and sometimes these comments are clearly hostile. When a video is not appreciated, its producer may be called names or even asked to die from some unpleasant disease. Even when people tell what they do not like about the content of a video, this may be done with hostile terms like "this sucks" or "I hate this". One particular example is a video (titled "Crazy Frog Bros.") of two young boys lip-syncing and dancing to a song. Two extreme but actual comments on this video are "you guys are fags. go kill yourselves by strangling yourselves then post that video on youtube so you can get people to laugh at you" and (translated from Dutch) "mother of these 2 retards what is it like to give birth to 2 cancer tumors?". One might argue that such comments on two happy children are at least remarkable.

Flames like the ones mentioned could be used only to be funny without any intent to harm someone. According to Thompsen (1996), these might not even be flames when they are not perceived as such. Perhaps some visitors or commenters find such comments quite amusing, but other visitors or the creators of the video or text (who are directly addressed) can still feel offended. We argue that flames meant to be funny are still flames, because people can still perceive them as such and feel offended.

Although all research on flaming has focused on other kinds of CMC, websites like YouTube.com show that the phenomenon may be more common in the online commenting situation. Perhaps, asking people for their personal opinions is a bad idea in an environment where lack of social cues and anonymity may already stimulate people to exhibit uninhibited behavior.

1.3 Goal of the Present Research

The aim of the present research was to find out whether the SIDE model can be used to explain flaming behavior in the online commenting situation. If people conform to the perceived norm when writing comments, as the SIDE model predicts, flames in earlier written comments stimulate people to write flames in their own comments.

One might wonder why the SIDE model would be valid in the online commenting situation. In typical research on CMC, participants discuss topics. It can be argued that commenting is very different from discussing, because commenting is not a long-term communication process. However, early social identity research shows that social identities can be elicited even when people do not communicate at all (Tajfel, 1974). People who have been randomly divided into two groups and do not know or communicate with each other, tend to show in-group favoritism, indicating depersonalization. Apparently, perceived categorizations that make no sense at all can still cause depersonalization. Then, it seems reasonable to expect that categorizing oneself and others as commenters in the online commenting situation will suffice as well. Especially when

commenting people all share the same opinion that is opposed to the opinion expressed in a stimulus text, such a categorization can be expected to become salient and cause depersonalization effects, like conforming to perceived norms.

This research has focused on answering the following questions:

- Do people in the online commenting situation conform to the flaming norm set by earlier commenters?
- If such conformation occurs, can it be explained by the SIDE model?

2. METHOD

2.1 Overview

A natural experiment has been conducted, in which participants were made aware of participating in an experiment only after they had commented on a stimulus. To enhance ecological validity, participants were not instructed or forced to comment, so they had the choice to comment or not as if it were a natural commenting situation.

Participants were attracted to a webpage offering a text and the ability to comment. Four comments were already given, as if they were placed by earlier readers of the text. Only four comments were used, because a few comments were thought to be more attractive to read than a lot. For the same reason, the text itself was relatively short as well, although not too short. It had to be clear that this was an online commenting situation and that this text was the main stimulus.

Two conditions were used in this experiment, differing in the nature of the existing comments (flaming or non-flaming). The effect of the condition on flaming behavior in new comments was analyzed in an attempt to answer the first research question.

Names were given with the comments, to make the website look natural. The names of the commenters were still very anonymous, though, e.g. "Freddy". When a participant wanted to comment, a (nick)name was asked as well.

Flaming in the online commenting situation is a form of expressing disagreement, so the stimulus text expressed an opinion clearly opposite to the opinions of most participants. This way, a situation was created in which flaming may occur.

When participants commented, they were directed to a page informing them about the experimental setting and asking them to answer a couple of questions. These questions were based on assumptions inspired by the SIDE model, in an attempt to answer the second research question. To stimulate participants to cooperate, the number of questions was kept low (17 in total).

2.2 Participants

Participants were recruited from three websites about free Windows software (www.moor-software.com, 2pic.moor-software.com and ecm.moor-software.com). For a few weeks, visitors of these addresses were first directed to a page informing them that the author of the website had read a certain text and that he thought it might be interesting for his visitors. People were encouraged to follow a link to the text (which was located on a different website), but they could also click on a second link to ignore the message and proceed to the website they were looking for. When the text link was clicked, the text was opened in a new browser window.

2.3 The Stimulus Text and Conditions

The page with the text showed the name of the author, the text itself, four comments and the ability to comment. The page was designed to give a serious, business-like impression as if it were

a weblog. The page was located on a different address from the software websites (markg.freehostia.com) and any connection with the websites was avoided. Participants were led to believe that this text was written and published by somebody else than the author of the software websites, and that reading this text had only been suggested because of its content.

Participants were randomly allocated one of two conditions (which will be explained shortly). The allocated condition for each unique IP address was saved to prevent returning visitors from switching conditions.

The text itself pled for prohibiting the distribution of free software by law, arguing that the quality is usually low and large software companies are losing profit. This is certainly an opinion that most free software users (and probably, other people) disagree with. It may even be regarded as a quite ridiculous opinion. The author made clear that he was very serious, though. When people would suspect that the text had been written only to provoke disagreements, flaming might be considered an appropriate response. Because we consider flaming as negative and inappropriate behavior, the text itself must not be a clear invitation to commit flaming. Also, only polite language was used and no direct insults at people or groups were made. The text's title was "Why free software distribution should be prohibited" and the author was called "Mark G".

The existing comments were all disagreeing with the text. The difference between the two conditions was the content of these comments. In the first condition, the comments were written in a polite way. The author of the text was told why his argument did not make sense and the existence of free software is a good thing. In the second condition, the comments were flames. One may argue that there is a considerable difference between mild swearing (e.g. "That's just bullshit") and insults aimed directly at the author ("Are you trying to be funny or are you really this dumb? Go fuck yourself"). To make the difference between the two research conditions as clear as possible, all comments in the second condition contained insults that were clearly hostile.

From now on, the conditions will be referred to as the "non-flaming" and the "flaming" condition. The virtual writers of the existing comments are referred to as the "earlier commenters".

2.4 Coding the Comments

Because our definition of flaming still leaves a lot of room for personal opinion and interpretation, eight raters have rated each comment to be either flaming or not. The number eight has been picked quite arbitrarily. A small number of raters would have biased the results, but eight seemed sufficient.

The raters were all Dutch students, both male and female, aged between 20 and 24. Most of them knew about the experiment and research goal, but they did not know in which condition the comments had been given.

Each rater received a small briefing, the stimulus text, the definition of flaming and a list of comments to rate. Raters were not given any guidelines about flaming apart from its definition, because this would have biased their ratings in the direction of the guideline composer's interpretation of words like "insulting".

Comments were called flames if, and only if, the majority of the raters had rated it as such. There were no cases where four raters had rated a comment as flaming and four as non-flaming. If this had been the case, a ninth rater had probably been involved.

2.5 Questionnaire

Participants were asked for their age, gender and country. Not only to inform us about their demographics, but also to analyze whether these variables have any effects on flaming behavior. For example, according to Aiken and Waller (2000), males exhibit more flaming than females.

Participants were also asked whether they agreed with the opinion expressed in the text and whether they had read the comments that were already given. These questions were asked to check whether most participants disagreed with the text and all of them had read the given comments.

To investigate whether social identities were involved in the commenting situation, twelve statements (see Attachment A) were given. Participants were asked to specify on a 5-point Likert scale to what extent they agreed with these statements. The points were "strongly disagree", "disagree", "don't know", "agree" and "strongly agree".

The first three statements concerned the perceived offensive intent and effects of the text. If social identities of participants were elicited in the online commenting situation, participants might be expected not only to conform to a flaming norm. They might also be expected to express feeling offended by the text to the extent to which they perceive the other commenters to be offended. Therefore, a correlation between agreement to statements 2 and 3 would plead in favor of the SIDE model.

Questions 4 to 9 were asked about the participant's perception of both the author of the text and the writer of one of the earlier comments. A social identity would be expected to cause a positive perception of the fellow commenter and a negative perception of the author of the text.

Although social identities may work on a subconscious level, the last three statements addressed a group feeling quite directly. Participants may be aware of their social identities and might therefore (strongly) agree with these statements, especially if they felt offended by the text.

3. RESULTS

3.1 Participants

1420 people visited the page with the text, 707 in the flaming condition and 713 in the non-flaming condition. 77 people (5%) gave 78 comments on the text.

16 comments (11 from the flaming condition and 5 from the non-flaming condition) were not used for the analysis. Eight of the unusable comments (two of which were written by one person) were "empty" or contained only a random sequence of characters, two comments were written in a strange language with a different character set, and the remaining six comments were possibly not aimed at the text but at the author of the software websites. For example, one comment was "i dnt giv a fuck were is my free download cunt". Apparently, this commenter had clicked on the link to the text while he intended to continue to the software website he wanted to visit. Although this comment was quite clearly not aimed at the text, as the empty comments probably were, the purpose of some other comments was very unclear. One commenter who had named himself "cgg" had given the short comment "Bullshit" and given some doubtful (and incomplete) answers to the questions. Because it was very difficult to tell whether this person had commented on the text or was frustrated by being on the wrong page, this comment was also not used for the analysis. By the way, not all unusable comments were flames.

62 comments were used for the analysis, 31 from both conditions. 41 of the participants had also completed the questionnaire, 11 had answered some but not all questions and 10 had not answered any questions.

50 participants had given their age, with an average age of 41.62 years ($SD = 14.90$). The youngest participant was 15 and the oldest was 73.

48 participants had given their country. 18 were from the United States, 9 from the United Kingdom, 6 from the Netherlands, 4 from Sweden, 6 from other countries in Europe (including Russia and Turkey), and 4 from other countries (India, Indonesia and New Zealand).

Only 2 of 52 participants claimed that they disagreed with the text and only one participant stated not to know about her agreement. Remarkably, her comment included “I DISAGREE BECAUSE”.

4 of 52 participants claimed that they had not read the comments that were given. One of these comments, given in the flaming condition, contained “I have no fowl language to offer, I just disagree completely”. These four participants probably had a short look at the flaming comments and decided not to read them. Three of them were indeed in the flaming condition.

3.2 Flaming

The eight raters agreed on 43 of the 62 comments, and for 11 comments, only one rater disagreed. Fleiss’ Kappa was 0.66, indicating that the inter-rater reliability was acceptable.

12 comments were rated as being flames, and 50 comments as not being flames. 9 of the flames were given in the flaming condition, meaning that 29% of the comments in the flaming condition contained flaming, compared to only 10% in the non-flaming condition.

To give an impression, some of the flames were “Are you Bill gates retarded brother, or are you on Microsofts payroll?”, “Mark G. your a twat!!!” and “Pillock”. Some people, apparently disliking the idea of prohibition, even asked whether the author of the text would feel more comfortable in Nazi Germany or Soviet Russia. One commenter, called “Someone who isn’t a retard”, ended his long comment with “So screw you Hitler”.

Most non-flaming comments displayed disagreement with the text (e.g. “I find some freeware good and it is good for some home computer owners, who need not spend for their use and requirement”) but some comments seemed to respond to the flaming of the earlier comments (e.g. “I think it is a shame that opinions are voiced in terms of vulgarity”).

Table 1. Number of non-flaming and flaming comments in both conditions

Condition	Non-flaming comments	Flaming comments	Total
Non-flaming	28	3	31
Flaming	22	9	31
Total	50	12	62

Participants in the flaming condition flamed more often than participants in the non-flaming condition (see Table 1). Significance was established both with the Chi-square test ($\chi^2(1, N = 62) = 3.72, p = 0.05$) and Fisher’s exact test (1-sided, $p = 0.05$).

3.3 Differences on the questionnaire

All statistical analyses involving the questionnaire were 2-sided.

No significant differences between flaming and non-flaming commenters were found for age, gender or country (USA vs. Europe).

Many significant correlations were found between agreement to the statements of the questionnaire (see Appendix A). Despite the impressive number of correlations, some logically predictable correlations were not as significant as

might be expected (e.g. between statements 7 and 8: $r(47) = -0.27, p = 0.06$).

As predicted, a correlation was found between statements 2 and 3 ($r(48) = 0.35, p = 0.01$). One might think that agreement to either statement was influenced by the experimental condition a participant was in, but this seems not to be the case. First of all, the directions of these effects are different. Participants in the flaming condition felt slightly less offended themselves ($t(49) = 0.44, p = 0.67$), but were more convinced that the earlier commenters felt offended ($t(48) = -4.73, p = 0.00$). Besides, the correlation between statements 2 and 3 is significant within both the non-flaming condition ($r(21) = 0.55, p = 0.01$) and the flaming condition ($r(25) = 0.43, p = 0.02$). Apparently, in either condition feeling offended was significantly correlated with believing that the earlier commenters felt offended.

Participants were not more positive about their fellow commenter than about the author of the text, indicated by their agreement to statements 4, 5, 7 and 8. Differences were very small and the directions of these differences varied between conditions and statements. Agreement to statement 7 was lower for participants who had flamed, indicating that they liked their fellow commenter less ($t(47) = 2.40, p = 0.02$).

Participants did find their fellow commenter’s point of view more reasonable than the author’s ($t(49) = 3.33, p = 0.00$). This difference was also found within either condition (non-flaming: $t(22) = -1.82, p = 0.08$; flaming: $t(26) = -2.84, p = 0.01$). Participants in the flaming condition significantly found both points of view less reasonable than participants in the non-flaming condition (author: $t(49) = 2.30, p = 0.03$; commenter: $t(48) = 2.14, p = 0.04$).

Correlations between agreement to statements 10, 11 and 12 were all significant with p-values under 0.01, suggesting that all three statements measured the same underlying construct, an identification with the earlier commenters. The average agreement to these statements, however, was not very high (3.04, 3.30 and 2.82, respectively). Participants in the non-flaming condition agreed more to these statements, but this was only weakly significant for statement 10 ($t(47) = 1.75, p = 0.09$) and far from significant for statements 11 and 12.

Agreement to statement 2 was correlated with agreement to statement 10 ($r(47) = 0.43, p = 0.00$), statement 11 ($r(48) = 0.36, p = 0.01$) and, less significantly, statement 12 ($r(48) = 0.24, p = 0.09$). Agreement to statement 3, however, was only correlated slightly significant with agreement to statement 11 ($r(47) = 0.27, p = 0.06$).

Some of the correlations already mentioned were even stronger when comparing flammers and non-flammers within the flaming condition. Participants who had flamed, felt more offended by the text ($t(25) = -1.63, p = 0.12$), perceived the other commenters as being offended more often ($t(25) = -2.83, p = 0.01$), and liked their fellow commenter less ($t(25) = 2.47, p = 0.02$). No significant differences were found in agreement to statements 10, 11 and 12.

Because the number of flammers in the non-flaming condition was only three, no comparisons between flammers and non-flammers in this condition have been made.

4. DISCUSSION

4.1 Conforming to the flaming norm

Our first research question was whether people in the online commenting situation conform to flaming norms set by earlier commenters. The results of this experiment suggest that this is indeed the case. Participants in the flaming condition were considerably more eager to flame themselves than people in the non-flaming condition.

The p-value of this effect was 0.05, which is just on the edge of statistical significance. The unrounded p-value was even slightly higher (0.053 for Fisher's exact test, 1-sided). Since the number of participants was only 62, this p-value is probably sufficiently convincing.

One more thing needs to be said, however, about the six comments that were identified as being unusable because they were aimed at the author of the software websites rather than the author of the text. Although several commenters made their intention very clear, some wrote comments that were very difficult to interpret. Six comments, compared to the total number of 62 usable comments, have a considerable influence on the p-value that is found. For example, a premature analysis with the discussed "Bullshit" comment included yielded a p-value of 0.035. Only removing one comment shifted the p-value for almost two hundredth. In fact, most comments that have been removed because they were probably not aimed at the text seem to fit quite well in the norm conforming picture. So if we had decided not to remove these comments at all, we had probably found a p-value much smaller. Whether it would have been valid to use these comments could be disputed, but these commenters seem to have been influenced by the flaming norm as well as the commenters giving feedback to the text. What we are trying to make clear, is that subtle differences in our subjective and fallible decisions had a profound effect on the results of this experiment. A larger number of participants with less reluctant page visitors expressing their frustration would have yielded more robust results.

4.2 Explaining the results using the SIDE model

Our second research question was whether a conformation effect could be explained using the SIDE model. Although the SIDE model would predict the effect that we have found, so would other theories such as the social learning theory (Bandura, 1977). Crucial to the SIDE model is the activation of a social identity, a phenomenon we have tried to find by using a questionnaire.

Participants in our experiment felt more offended by the text when they thought that earlier commenters felt more offended, regardless of the commenters' flaming behavior. This correlation can be explained and even expected by the SIDE model, arguing that people find the text threatening to their social identities and as such to both themselves and the earlier commenters.

Participants did not like their fellow commenter better than they liked the author of the text, an effect that would certainly be predicted if social identities would have played a role. The author of the text would be expected to be perceived as an out-group member threatening what bonds the members of the in-group (i.e. their love for freeware or at least their disliking of prohibition). In fact, flammers even significantly liked their fellow commenter less. If flaming would have been a result of depersonalization, an effect in the opposite direction would be expected.

Agreement to statements about identifying with the earlier commenters was far from convincing. It was, however, correlated with feeling offended by the text. One might argue that a social identity is easier elicited when people feel threatened by an out-group member, which would make this correlation fit into the SIDE framework. It is remarkable, from this viewpoint, that a correlation between identifying with the earlier commenters and thinking that they were offended by the text, was not found.

Summarizing, it is very unclear whether depersonalization accounted for flaming behavior. Some results from the

questionnaire fit the SIDE model quite well, but some results certainly do not. Some effects that would be expected if social identities had interfered were not found, leaving us with the conclusion that other processes may have influenced flaming behavior rather than depersonalization.

Perhaps identifying with earlier commenters was countered by the negativity of their behavior. The only behavior of the earlier commenters known to the participants, was flaming. Participants in the non-flaming condition, however, did not identify with their fellow commenters more significantly.

4.3 Recommendations for future research

To convince oneself that flaming is far from absent in the online commenting situation, one only has to search for some websites offering this situation and read the comments. No experiment was needed to show that this negative kind of behavior is as common on such websites as it is in synchronous CMC, which has had much more attention from research on flaming. In a first exploration of this new context, the present research tried to use the SIDE model to explain flaming behavior. Some recommendations can be made for future research, based on the results of this experiment.

Flaming behavior was found more often when the apparent norm of earlier commenters consisted of flaming. The small number of participants and the additional difficulties introduced by people reluctantly directed to the text have probably affected the results considerably. More research is to be recommended, with more participants and a better way of recruiting them. A solution to both problems might be to involve an established and popular website (like YouTube.com) in experiments. This would also make the situation even more natural, yielding highly valid results. Perhaps the scope of such research could even be moved from flaming to other kinds of behavior. For example, if some creative form of art is shown online, where opinions are clearly a matter of taste, do people copy an opinion expressed by earlier commenters?

Not all of our results were expected from the SIDE model's perspective. More research might be conducted to investigate whether depersonalization plays any role in the online commenting situation. If future research confirms that people conform to perceived norms in the online commenting situation but fails to find evidence for depersonalization or related phenomena, other theories could be tested to find what processes are underlying this behavior.

REFERENCES

- Aiken, M., & Waller, B. (2000). Flaming among first-time group support system users. *Information and Management*, 37, 95-100.
- Alonzo, M., & Aiken, M. (2004). Flaming in electronic communication. *Decision Support Systems*, 36(3), 205-213.
- Bandura, A. (1977). *Social learning theory*. Englewood Cliffs, NJ: Prentice-Hall.
- Coleman, L. H., Paternite, C. E., & Sherman, R. C. (1999). A reexamination of deindividuation in synchronous computer-mediated communication. *Computers in Human Behavior*, 15, 51-65.
- Collins, M. (1992). Flaming: the relationship between social context cues and uninhibited verbal behavior in computer-mediated communication. Available on the WWW at: <http://www.emoderators.com/papers/flames.html>
- Joinson, A. N., (2001). Self-disclosure in computer-mediated communication: the role of self-awareness and visual anonymity. *European Journal of Social Psychology*, 31(2), 177-192.

- Joinson, A. N. (2003). *Understanding the Psychology of Internet Behaviour: Virtual Worlds, Real Lives*. Basingstoke: Palgrave Macmillan.
- Kiesler, S., Siegel, J., & McGuire, T. W. (1984). Social psychological aspects of computer-mediated communication. *American Psychologist*, 39, 1123-1134.
- Lea, M., O'Shea, T., Fung, P., & Spears, R. (1992). 'Flaming' in computer-mediated communication. In M. Lea (ed.), *Contexts of computer-mediated communication* (pp. 89-112). London: Harvester Wheatsheaf.
- Lea, M., Spears, R., & De Groot, D. (2001). Knowing me, knowing you: anonymity effects on social identity processes within groups. *Personality and Social Psychology Bulletin*, 27, 526-537.
- Onvriendelijke reacties maakt bloggen impopulair. [Unfriendly comments make blogging unpopular]. (2007, May 21). *Dagblad De Pers*, p. 6.
- Postmes, T., & Spears, R. (1998). Deindividuation and anti-normative behavior: a meta-analysis. *Psychological Bulletin*, 123, 238-259.
- Postmes, T., Spears, R., & Lea, M. (2000). The formation of group norms in computer-mediated communication. *Human Communication Research*, 26(3), 341-371.
- Postmes, T., Spears, R., Sakhel, K., & De Groot, D. (2001). Social influence in computer-mediated communication: The effects of anonymity on group behavior. *Personality and Social Psychology Bulletin*, 27, 1242-1254.
- Reicher, S. D., Spears, R., & Postmes, T. (1995). A social identity model of deindividuation phenomena. In W. Stroebe & M. Hewstone (eds.), *European Review of Social Psychology* (pp. 161-198), vol. 6. Chichester: Wiley.
- Siegel, J., Dubrovsky, V., Kiesler, S., & McGuire, T. W. (1986). Group processes in computer-mediated communication. *Organizational behavior and human decision processes*, 37, 157-187.
- Steele, G., Woods, D., Finkel, R., Crispin, M., Stallman, R., & Goodfellow, G. (1983). *The hacker's dictionary*. New York: Harper & Row.
- Tajfel, H. (1974). Social identity and intergroup behaviour. *Social Science Information*, 13(2), 65-93.
- Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (eds.), *The social psychology of intergroup relations* (pp. 33-48). Pacific Grove, CA: Brooks/Cole.
- Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behaviour. In S. Worchel & W. G. Austin (eds.), *Psychology of Intergroup Relations* (pp. 7-24). Chicago: Nelson-Hall.
- Thompson, P. A. (1996). What's fueling the flames in cyberspace: A social influence model. In L. Strate, R. Jacobson, & S. Gibson (eds.), *Communication and cyberspace: Social interaction in an electronic environment* (pp. 297-315). Cresskill, NJ: Hampton.
- Turner, J. C. (1987). A self-categorization theory. In J. C. Turner, M. A. Hogg, P. J. Oakes, S. D. Reicher & M. S. Wetherell, *Rediscovering the social group: A self-categorization theory* (pp. 42-67). Oxford: Blackwell.
- Van Stein Callenfels, H. P., & Van Woerden, I. (2007, May 22). Genoeg van het gescheld. [Enough of the swearing]. *Dag*, pp. 8-9.
- Weisband, S., & Kiesler, S. (1996). Self-disclosure on computer forms: meta-analysis and implications. Proceedings of CHI96. Available on the WWW at: http://www.acm.org/sigchi/chi96/papers/Weisband/sw_txt.htm

Attachment A – Correlations between agreement to the statements of the questionnaire

Statement	1	2	3	4	5	6	7	8	9	10	11	12
1. I think the purpose of the text was to offend people.	--	0.50** <i>p</i> = 0.00	0.33* <i>p</i> = 0.02	-0.12 <i>p</i> = 0.40	0.12 <i>p</i> = 0.40	-0.27 <i>p</i> = 0.06	0.04 <i>p</i> = 0.79	0.04 <i>p</i> = 0.80	-0.09 <i>p</i> = 0.54	-0.04 <i>p</i> = 0.81	0.03 <i>p</i> = 0.85	-0.06 <i>p</i> = 0.66
2. I felt offended by the opinion expressed in the text.		--	0.35* <i>p</i> = 0.01	-0.06 <i>p</i> = 0.67	0.24 <i>p</i> = 0.10	-0.21 <i>p</i> = 0.15	0.08 <i>p</i> = 0.60	0.05 <i>p</i> = 0.71	0.22 <i>p</i> = 0.12	0.43** <i>p</i> = 0.00	0.36* <i>p</i> = 0.01	0.24 <i>p</i> = 0.09
3. I think the earlier commenters felt offended by the opinion expressed in the text.			--	0.06 <i>p</i> = 0.70	0.17 <i>p</i> = 0.25	-0.32* <i>p</i> = 0.02	-0.26 <i>p</i> = 0.08	0.06 <i>p</i> = 0.69	-0.07 <i>p</i> = 0.65	0.08 <i>p</i> = 0.61	0.27 <i>p</i> = 0.06	0.18 <i>p</i> = 0.21
4. I think Mark (the author of the text) is a nice person.				--	-0.36* <i>p</i> = 0.01	0.24 <i>p</i> = 0.09	-0.32* <i>p</i> = 0.02	0.41** <i>p</i> = 0.00	-0.23 <i>p</i> = 0.10	-0.13 <i>p</i> = 0.39	-0.08 <i>p</i> = 0.57	-0.21 <i>p</i> = 0.14
5. To me, Mark seems like an unpleasant person.					--	0.03 <i>p</i> = 0.82	-0.03 <i>p</i> = 0.83	-0.29* <i>p</i> = 0.05	0.26 <i>p</i> = 0.07	0.33* <i>p</i> = 0.02	0.19 <i>p</i> = 0.19	0.13 <i>p</i> = 0.36
6. Mark's point of view looks reasonable to me.						--	0.05 <i>p</i> = 0.75	0.08 <i>p</i> = 0.57	0.02 <i>p</i> = 0.88	-0.19 <i>p</i> = 0.20	-0.14 <i>p</i> = 0.32	-0.13 <i>p</i> = 0.38
7. I think Freddy (who wrote one of the comments) is a nice person.							--	-0.27 <i>p</i> = 0.06	0.37** <i>p</i> = 0.01	0.20 <i>p</i> = 0.17	0.03 <i>p</i> = 0.83	0.22 <i>p</i> = 0.13
8. To me, Freddy seems like an unpleasant person.								--	-0.51** <i>p</i> = 0.00	-0.35* <i>p</i> = 0.01	-0.29* <i>p</i> = 0.04	-0.15 <i>p</i> = 0.29
9. Freddy's point of view looks reasonable to me.									--	0.53** <i>p</i> = 0.00	0.39** <i>p</i> = 0.01	0.34* <i>p</i> = 0.02
10. I can identify myself with the earlier commenters.										--	0.71** <i>p</i> = 0.00	0.65** <i>p</i> = 0.00
11. I feel I have something in common with the earlier commenters.											--	0.56** <i>p</i> = 0.00
12. I perceive myself and the earlier commenters as members of a common group.												--

** . Correlation is significant at the 0.01 level (2-tailed).

* . Correlation is significant at the 0.05 level (2-tailed).

For all correlations, N was between 48 and 52.