

Experimental evidence of massive-scale emotional contagion through social networks

Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock

Significance

We show, via a massive ($N = 689,003$) experiment on Facebook, that emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness. We provide experimental evidence that emotional contagion occurs without direct interaction between people (exposure to a friend expressing an emotion is sufficient), and in the complete absence of nonverbal cues.

Abstract

Emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness. Emotional contagion is well established in laboratory experiments, with people transferring positive and negative emotions to others. Data from a large real-world social network, collected over a 20-y period suggests that longer-lasting moods (e.g., depression, happiness) can be transferred through networks [Fowler JH, Christakis NA (2008) *BMJ* 337:a2338], although the results are controversial. In an experiment with people who use Facebook, we test whether emotional contagion occurs outside of in-person interaction between individuals by reducing the amount of emotional content in the News Feed. When positive expressions were reduced, people produced fewer positive posts and more negative posts; when negative expressions were reduced, the opposite pattern occurred. These results indicate that emotions expressed by others on Facebook influence our own emotions, constituting experimental evidence for massive-scale contagion via social networks. This work also suggests that, in contrast to prevailing assumptions, in-person interaction and nonverbal cues are not strictly necessary for emotional contagion, and that the observation of others' positive experiences constitutes a positive experience for people.

- [computer-mediated communication](#)
- [social media](#)
- [big data](#)

Emotional states can be transferred to others via emotional contagion, leading them to experience the same emotions as those around them. Emotional contagion is well established in laboratory experiments (1), in which people transfer positive and negative moods and emotions to others. Similarly, data from a large, real-world social network collected over a 20-y period suggests that longer-lasting moods (e.g., depression, happiness) can be transferred through networks as well (2, 3).

The interpretation of this network effect as contagion of mood has come under scrutiny due to the study's correlational nature, including concerns over misspecification of contextual variables or failure to account for shared experiences (4, 5), raising important questions regarding contagion processes in networks. An experimental approach can address this scrutiny directly; however, methods used in controlled experiments have been criticized for examining emotions after social interactions. Interacting with a happy person is pleasant (and an unhappy person, unpleasant). As such, contagion may result from experiencing an interaction rather than exposure to a partner's emotion. Prior studies have also failed to address whether nonverbal cues are necessary for contagion to occur, or if verbal cues alone suffice. Evidence that positive and negative moods are correlated in networks (2, 3) suggests that this is possible, but the causal question of whether contagion processes occur for emotions in massive social networks remains elusive in the absence of experimental evidence. Further, others have suggested that in online social networks, exposure to the happiness of others may actually be depressing to us, producing an "alone together" social comparison effect (6).

Three studies have laid the groundwork for testing these processes via Facebook, the largest online social network. This research demonstrated that (i) emotional contagion occurs via text-based computer-mediated communication (7); (ii) contagion of psychological and physiological qualities has been suggested based on correlational data for social networks generally (7, 8); and (iii) people's emotional expressions on Facebook predict friends' emotional expressions, even days later (7) (although some shared experiences may in fact last several days). To date, however, there is no experimental evidence that emotions or moods are contagious in the absence of direct interaction between experiencer and target.

On Facebook, people frequently express emotions, which are later seen by their friends via Facebook's "News Feed" product (8). Because people's friends frequently produce much more content than one person can view, the News Feed filters posts, stories, and activities undertaken by friends. News Feed is the primary manner by which people see content that friends share. Which content is shown or omitted in the News Feed is determined via a ranking algorithm that Facebook continually develops and tests in the interest of showing viewers the content they will find most relevant and engaging. One such test is reported in this study: A test of whether posts with emotional content are more engaging.

The experiment manipulated the extent to which people ($N = 689,003$) were exposed to emotional expressions in their News Feed. This tested whether exposure to emotions led people to change their own posting behaviors, in particular whether exposure to emotional content led people to post content that was consistent with the exposure—thereby testing whether exposure to verbal affective expressions leads to similar verbal expressions, a form of emotional contagion. People who viewed Facebook in English were qualified for selection into the experiment. Two parallel experiments were conducted for positive and negative emotion: One in which exposure to friends' positive emotional content in their News Feed was reduced, and one in which exposure to negative emotional content in their News Feed was reduced. In these conditions, when a person loaded their News Feed, posts that contained emotional content of the relevant emotional valence, each emotional post had between a 10% and 90% chance (based on their User ID) of being omitted from their News Feed for that specific viewing. It is important to note that this content was always available by viewing a friend's content directly by going to that friend's "wall" or "timeline," rather than via the News Feed. Further, the omitted content may have appeared on prior or subsequent views of the News Feed. Finally, the experiment did not affect any direct messages sent from

one user to another.

Posts were determined to be positive or negative if they contained at least one positive or negative word, as defined by Linguistic Inquiry and Word Count software (LIWC2007) (9) word counting system, which correlates with self-reported and physiological measures of well-being, and has been used in prior research on emotional expression (7, 8, 10). LIWC was adapted to run on the Hadoop Map/Reduce system (11) and in the News Feed filtering system, such that no text was seen by the researchers. As such, it was consistent with Facebook's Data Use Policy, to which all users agree prior to creating an account on Facebook, constituting informed consent for this research. Both experiments had a control condition, in which a similar proportion of posts in their News Feed were omitted entirely at random (i.e., without respect to emotional content). Separate control conditions were necessary as 22.4% of posts contained negative words, whereas 46.8% of posts contained positive words. So for a person for whom 10% of posts containing positive content were omitted, an appropriate control would withhold 10% of 46.8% (i.e., 4.68%) of posts at random, compared with omitting only 2.24% of the News Feed in the negativity-reduced control.

The experiments took place for 1 wk (January 11–18, 2012). Participants were randomly selected based on their User ID, resulting in a total of ~155,000 participants per condition who posted at least one status update during the experimental period.

For each experiment, two dependent variables were examined pertaining to emotionality expressed in people's own status updates: the percentage of all words produced by a given person that was either positive or negative during the experimental period (as in ref. 7). In total, over 3 million posts were analyzed, containing over 122 million words, 4 million of which were positive (3.6%) and 1.8 million negative (1.6%).

If affective states are contagious via verbal expressions on Facebook (our operationalization of emotional contagion), people in the positivity-reduced condition should be less positive compared with their control, and people in the negativity-reduced condition should be less negative. As a secondary measure, we tested for cross-emotional contagion in which the opposite emotion should be inversely affected: People in the positivity-reduced condition should express increased negativity, whereas people in the negativity-reduced condition should express increased positivity. Emotional expression was modeled, on a per-person basis, as the percentage of words produced by that person during the experimental period that were either positive or negative. Positivity and negativity were evaluated separately given evidence that they are not simply opposite ends of the same spectrum (8, 10). Indeed, negative and positive word use scarcely correlated [$r = -0.04$, $t(620,587) = -38.01$, $P < 0.001$].

We examined these data by comparing each emotion condition to its control. After establishing that our experimental groups did not differ in emotional expression during the week before the experiment (all $t < 1.5$; all $P > 0.13$), we examined overall posting rate via a Poisson regression, using the percent of posts omitted as a regression weight. Omitting emotional content reduced the amount of words the person subsequently produced, both when positivity was reduced ($z = -4.78$, $P < 0.001$) and when negativity was reduced ($z = -7.219$, $P < 0.001$). This effect occurred both when negative words were omitted (99.7% as many words were produced) and when positive words were omitted (96.7%). An interaction was also observed, showing that the effect was stronger when positive words were omitted

($z = -77.9$, $P < 0.001$).

As such, direct examination of the frequency of positive and negative words would be inappropriate: It would be confounded with the change in overall words produced. To test our hypothesis regarding emotional contagion, we conducted weighted linear regressions, predicting the percentage of words that were positive or negative from a dummy code for condition (experimental versus control), weighted by the likelihood of that person having an emotional post omitted from their News Feed on a given viewing, such that people who had more content omitted were given higher weight in the regression. When positive posts were reduced in the News Feed, the percentage of positive words in people's status updates decreased by $B = -0.1\%$ compared with control [$t(310,044) = -5.63$, $P < 0.001$, Cohen's $d = 0.02$], whereas the percentage of words that were negative increased by $B = 0.04\%$ ($t = 2.71$, $P = 0.007$, $d = 0.001$). Conversely, when negative posts were reduced, the percent of words that were negative decreased by $B = -0.07\%$ [$t(310,541) = -5.51$, $P < 0.001$, $d = 0.02$] and the percentage of words that were positive, conversely, increased by $B = 0.06\%$ ($t = 2.19$, $P < 0.003$, $d = 0.008$).

The results show emotional contagion. As Fig. 1 illustrates, for people who had positive content reduced in their News Feed, a larger percentage of words in people's status updates were negative and a smaller percentage were positive. When negativity was reduced, the opposite pattern occurred. These results suggest that the emotions expressed by friends, via online social networks, influence our own moods, constituting, to our knowledge, the first experimental evidence for massive-scale emotional contagion via social networks (3, 7, 8), and providing support for previously contested claims that emotions spread via contagion through a network.

Fig. 1.

Mean number of positive (*Upper*) and negative (*Lower*) emotion words (percent) generated people, by condition. Bars represent standard errors.

These results highlight several features of emotional contagion. First, because News Feed content is not "directed" toward anyone, contagion could not be just the result of some specific interaction with a happy or sad partner. Although prior research examined whether an emotion can be contracted via a direct interaction (1, 7), we show that simply failing to "overhear" a friend's emotional expression via Facebook is enough to buffer one from its effects. Second, although nonverbal behavior is well established as one medium for contagion, these data suggest that contagion does not require nonverbal behavior (7, 8): Textual content alone appears to be a sufficient channel. This is not a simple case of mimicry, either; the cross-emotional encouragement effect (e.g., reducing negative posts led to an increase in positive posts) cannot be explained by mimicry alone, although mimicry may well have been part of the emotion-consistent effect. Further, we note the similarity of effect sizes when positivity and negativity were reduced. This absence of negativity bias suggests that our results cannot be attributed solely to the content of the post: If a person is sharing good news or bad news (thus explaining his/her emotional state), friends' response to the news (independent of the sharer's emotional state) should be stronger when bad news is shown rather than good (or as commonly noted, "if it bleeds, it leads;" ref. 12) if the results were being driven by reactions to news. In contrast, a response to a friend's emotion expression (rather than news) should be proportional to exposure. A post hoc test comparing effect sizes (comparing correlation coefficients using Fisher's method) showed no difference

despite our large sample size ($z = -0.36$, $P = 0.72$).

We also observed a withdrawal effect: People who were exposed to fewer emotional posts (of either valence) in their News Feed were less expressive overall on the following days, addressing the question about how emotional expression affects social engagement online. This observation, and the fact that people were more emotionally positive in response to positive emotion updates from their friends, stands in contrast to theories that suggest viewing positive posts by friends on Facebook may somehow affect us negatively, for example, via social comparison (6, 13). In fact, this is the result when people are exposed to less positive content, rather than more. This effect also showed no negativity bias in post hoc tests ($z = -0.09$, $P = 0.93$).

Although these data provide, to our knowledge, some of the first experimental evidence to support the controversial claims that emotions can spread throughout a network, the effect sizes from the manipulations are small (as small as $d = 0.001$). These effects nonetheless matter given that the manipulation of the independent variable (presence of emotion in the News Feed) was minimal whereas the dependent variable (people's emotional expressions) is difficult to influence given the range of daily experiences that influence mood (10). More importantly, given the massive scale of social networks such as Facebook, even small effects can have large aggregated consequences (14, 15): For example, the well-documented connection between emotions and physical well-being suggests the importance of these findings for public health. Online messages influence our experience of emotions, which may affect a variety of offline behaviors. And after all, an effect size of $d = 0.001$ at Facebook's scale is not negligible: In early 2013, this would have corresponded to hundreds of thousands of emotion expressions in status updates per day.

Acknowledgments

We thank the Facebook News Feed team, especially Daniel Schafer, for encouragement and support; the Facebook Core Data Science team, especially Cameron Marlow, Moira Burke, and Eytan Bakshy; plus Michael Macy and Mathew Aldridge for their feedback. Data processing systems, per-user aggregates, and anonymized results available upon request.

Footnotes

- Author contributions: A.D.I.K., J.E.G., and J.T.H. designed research; A.D.I.K. performed research; A.D.I.K. analyzed data; and A.D.I.K., J.E.G., and J.T.H. wrote the paper.
- The authors declare no conflict of interest.
- This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.