



MARIA KONNIKOVA

JULY 2, 2014

DID FACEBOOK HURT PEOPLE'S FEELINGS?

BY MARIA KONNIKOVA

It's no secret that social media can affect your mood



(<http://www.newyorker.com/online/blogs/elements/2013/09/the-real-reason-facebook-makes-us-unhappy.html>), making you experience certain feelings based on the information you see and the people you interact with. Those feelings are one of the reasons that people use sites like Facebook or Twitter to begin with. But what if you found out that what you felt was the result of a deliberate manipulation by the social network itself?

That's at the heart of a controversy about a new study (<http://www.pnas.org/content/111/24/8788.full>), in which researchers from Facebook manipulated the amount of positive and negative material in users' news feeds in an attempt to see whether social networks can lead to the same kinds of "emotional contagion"—a person "catching" other people's emotions from interacting with them—that exist in real life.

In the experiment, a Facebook data scientist named Adam Kramer and his colleagues tweaked the algorithm that determines which items appear in news feeds, so that, in some cases, users would see more posts with positive words, and in others, more with negative words. They then observed those users for a week to see if the changes affected what those people posted in turn. Would seeing others' happy or sad content make people post happier or sadder stuff themselves? Yes, the authors found. "The results show emotional contagion," Kramer and his co-authors write. "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."

This caught people's attention—but probably not quite in the way that the researchers had expected. Some Facebook users were appalled. If Facebook could alter a person's emotions with a minor tweak to an algorithm—not only positively but negatively—didn't that mean that the company had knowingly made thousands of its users sadder, without ever asking for their explicit consent?

If the research had been conducted at an academic institution, it would have had to clear all sorts of regulatory hurdles from a body known as an institutional review board, which would have typically made sure that the emotional manipulation wasn't going to cause undue harm. But researchers from Cornell who were involved in the study didn't help gather the data about users; they only looked at it once Facebook had collected it, making them, in this case, exempt (<http://mediarelations.cornell.edu/2014/06/30/media-statement-on-cornell-universitys-role-in-facebook-emotional-contagion-research/>) from the university's institutional review board protocols. As for the researchers at Facebook who had actually gathered the data, they held

that Facebook's data-use policies, to which users consent when they join the site, allowed them to implement and test the effects of these manipulations.

Beyond the uproar, though, there's another question: Did Facebook really manipulate users' emotions in the first place? The truth is that the best way to know what people are feeling is also the most straightforward: ask them. Psychological studies of emotion most often use a standard rating scale, called the Positive and Negative Affect Schedule, that has people rate how they're feeling on multiple dimensions on a sliding scale. That is often the only way to know if an emotional manipulation has worked: Did it significantly change someone's responses?

Facebook used a different approach. To determine whether a post was positive or negative, researchers scanned the post's language using a software program called Linguistic Inquiry and Word Count. If the program caught a positive or negative word, it flagged the post accordingly. Crucially, to protect privacy, Facebook adapted the program so that researchers never saw any of the actual text, just the flag. But that protection meant that the researchers couldn't individually interpret each post and its nuances. Context, punctuation, words that weren't in the L.I.W.C. program to begin with—none of this was captured or taken into account.

But understanding emotion is far more complex than looking for certain words—especially if those words are taken completely out of context, as they were in this study. “All they are looking at is the presence or absence of a single emotion word,” Elliot Berkman, a psychologist at the University of Oregon, whose research focusses on social and affective neuroscience, told me. To infer that someone is expressing or experiencing an emotion just because they use a particular word is “a huge leap,” he said. “I ate a Happy Meal,” for

instance, could be flagged in the same way as “I ate a Happy Meal and it made me throw up,” which could, in turn, received the same flag as “So happy!” or “So happy???”

The same positive or negative word can carry a range of meanings, many of which are divorced from actually experiencing an emotion. Writing “I’m sorry you’re sad,” for instance, doesn’t actually mean that you are sad yourself—only that you have become aware of someone else’s emotion and are acknowledging it. There is, too, a difference between emotion—“I’m feeling amazing”—and reaction—“Can you believe this picture?? Amazing!” Context is crucial.

Even if people change their word choice in response to what they see on Facebook, it’s hard to know why they’re doing it. “They show that changing the information you’re exposed to changes that process of word choice, but we don’t know anything at all about what that process means,” Ethan Kross, a psychologist at the University of Michigan, who has done extensive work on emotion on Facebook (<http://www.newyorker.com/online/blogs/elements/2013/09/the-real-reason-facebook-makes-us-unhappy.html>), said. “Is this actually contagion and is it leading to changes in emotional states? We don’t know.”

The social and public nature of Facebook makes people’s posts even more difficult to interpret. Let’s even assume that we see a legitimately happy (or sad) post, and that the study algorithm is refined further to catch nuance and context. How do we know that we’re encountering real happiness and not someone who is attempting to seem happier than he or she actually is? One effect that has been repeatedly documented is something called social comparison and social signaling (<http://www.newyorker.com/online/blogs/currency/2014/02/why-are-we-still-on-facebook.html>): people compare themselves to others who are similar to them, and then send out the types of signals that

they believe mark them as the person they see themselves as being. If you want people to think of you as someone who eats five-star dinners and takes trips to Aspen, you're likelier to describe your sad-looking takeout as something more exciting—or maybe leave that detail off your news feed entirely.

In that context, one way to view the increase in positive posts when people see positive content isn't that they've become happier. Instead, it's that they're trying to one-up the positive posts themselves—or at least show that they're on the same page. And when positive posts decrease, people might not actually feel less happy—they might just feel less of a need to self-aggrandize in order to keep up with their friends. If Facebook's study demonstrates anything, it might simply be our tendency to mimic others in our social group. It's not an emotional response so much as an online version of monkey see, monkey do.

One of the study's less-remarked-upon findings could support this thesis. It found that when people saw fewer posts with emotional content, they became less likely to post emotional content themselves. To rule out the possibility that the results simply showed people mimicking one another, Berkman suggested a third experiment. The researchers could choose neutral words to show some users more often—for instance, about taking a walk or going to lunch—and see whether those users then start posting about the same activity. (Berkman's bet is that they would.)

Mimicry is a well-documented phenomenon—and one that isn't all that interesting in its own right. What's potentially more intriguing about the Facebook study is that it sheds light on how the site could be used to study other questions. For the emotion study, Facebook looked at over half a million people during the course of a week; that's less than one per cent of Facebook's user base. With such unparalleled

access to subjects, psychologists could begin to study all kinds of phenomena. “Ethics aside, you can basically do the most powerful psych study ever with this kind of access,” Berkman said.

The key is to find questions about which the benefits of knowledge are so overwhelmingly apparent that they make the risks seem worthwhile—and where Facebook, with its huge amounts of data, facilitates an approach that couldn’t be achieved with a smaller, laboratory-based sample. “First, identify what this approach can do that other methods really can’t,” Berkman said. “And then that’s what you use it for.” In the case of the emotion study, Berkman was not sure whether all this data was needed to answer the questions that Facebook sought to address. “Can we justify a social network study of one million people?” Berkman asked. “Sure, if we are looking for a rare disease or a potentially dangerous mental state. Less so if you’re looking at a relatively obscure and well-established effect.”

One potential usage, which has already been tested on a social network, is to mine large sets of posts for signs of mental illness or depression. At this year’s Society for Personality and Social Psychology conference, a team from Microsoft discussed a study (http://www.munmund.net/pubs/chi_13.pdf) in which they had been using posts on Twitter to predict the risk of postpartum depression in pregnant women and new mothers. In that case, Berkman pointed out, the benefits likely outweigh the risks. While the cost of a false positive—telling someone that she is at greater risk of depression when she’s not—is high, the cost of missing the signs of mental illness is potentially even bigger. Microsoft’s research, in his mind, is an example of a good use of social networks for research. “Let’s let the data tell us something we can’t normally tell,” Berkman said, “and that you may not even know about yourself.”

Above: Facebook founder and C.E.O. Mark Zuckerberg. Photograph by Justin Sullivan/Getty.



Maria Konnikova is a contributor to [newyorker.com](http://www.newyorker.com), where she writes a weekly blog (<http://www.newyorker.com/news/maria-konnikova>) focussing on psychology and science.

