

## Federal Hot Water for Facebook over Emotional Manipulation Experiment

By [Violet Blue](#) for [Pulp Tech](#) | July 12, 2014 -- 13:00 GMT (06:00 PDT)

An emotional manipulation experiment Facebook conducted on 689,003 unknowing users in 2012, only recently brought to light, has now drawn scrutiny from Washington.

On Friday, [Senator Mark Warner](#) (D-Va.) [filed a formal complaint letter \(.PDF\)](#) addressed to the Federal Trade Commission calling the experiment's transparency, ethics and accountability into question.

According to reports, it is not clear whether Facebook users were adequately informed and given an opportunity to opt in or opt out.

I also have concerns about whether or not Facebook responsibly assessed the risks and benefits of conducting this behavioral experiment, as well as the ethical guidelines, if any, that were used to protect individuals.

Senator Warner cautioned the FTC, saying that experiments such as these are a "slippery slope" and expressed serious concerns about "future studies like this, without proper oversight or appropriate review," and their impact on consumers.

The Senator's letter regarding Facebook's "Emotional Contagion" experiment follows a media firestorm of public outrage that [began June 26](#), a formal FTC complaint [from the Electronic Privacy Information Centre](#), and in the UK, an investigation by the Information Commissioner's Office.

In its "emotional contagion experiment" Facebook tampered with the emotional well-being of 689,003 users to see how their emotions could be controlled; Facebook's hypothesis amounted to "let's see if we can plant unhappiness and make it spread."

According to Facebook's researchers, the unknowing users automatically consented to the emotional manipulation experiment because they were Facebook users. Consent was indicated, they said, when any user clicked "agree" to Facebook's 9,000-word Terms of Use agreement when signing up, and by continuing to use the site after any Terms updates.

"Who knows what other research they're doing?"

Everyone except Facebook agrees that what Facebook did was alarming, unethical and certainly carried some risk to its users' mental health.

When news of the study hit, Cornell University, [formerly supportive](#) in its association with the experiment, issued a [statement distancing itself](#) from involvement.

Facebook was then shown to have changed its Terms to include "research" [on its users] only [after the experiment had been done](#).

Those outside of Facebook involved with the experiment have been [accused of approval-laundering](#) by respected academics.

Susan Fiske, the professor of psychology at Princeton University who edited the study for Proceedings of the National Academy of Sciences of America, said "I'm still thinking about it and I'm a little creeped out, too."

Fiske [told The Atlantic](#) on June 28:

I had not seen before, personally, something in which the researchers had the cooperation of Facebook to manipulate people...

Who knows what other research they're doing?

Proceedings of the National Academy of Sciences of America issued [a formal statement to press](#) June 3 saying it had ethical concerns, "that the collection of the data by Facebook may have involved practices that were not fully consistent with the principles of obtaining informed consent and allowing participants to opt out."

The Linguistic Inquiry Word Count (LIWC2007) tool, used by Facebook to monitor and manipulate user emotions in the experiment, has been shown to be [wholly inaccurate and inappropriate](#) for Facebook's intended use.

A disgraceful time to "lean in"

Reactions from Facebook's executives have eroded the company's reputation further.

Scientist Adam Kramer, the experiment's primary author, responded June 29 by [posting a non-apology to Facebook](#) that utterly missed the point, saying they were sorry about the way they had described the experiment while attempting to re-frame the concept of user consent as if it was merely red tape.

In trademark Facebook spin, he blamed the public outrage on bad representation, as if news of emotional tampering in people's day-to-day lives was a trivial misunderstanding that only anxious people worried about.

On July 2, Facebook Chief Operating Officer Sheryl Sandberg sandbagged Kramer's broken PR strategy, again trying to re-cast the tsunami of public outrage as a miscommunication about the study's description.

Which would make sense if Facebook had actually communicated to the public, or its users, about the experiment in the first place.

The Washington Post [reported](#):

On Wednesday, Facebook's second-in-command, Sheryl Sandberg, expressed regret over how the company communicated its 2012 mood manipulation study of 700,000 unwitting users, but she did not apologize for conducting the controversial experiment.

It's just what companies do, she said.

July 3, Facebook's Global Head of Policy Monika Bickert grabbed the shovel of callous irresponsibility from Sandberg and dug deeper, [saying](#) the nonconsensual mental health experiment was an example of Facebook's "innovation."

Bickert — somewhat presciently in light of Senator Warner's letter — also [told the audience at the Aspen Ideas Festival](#), "it's concerning when we see legislation that could possibly stifle that sort of creativity and that innovation."

Warner has a track record of heavy involvement in the Senate regarding consumer protections and cybercrime. This past week, the Senate Select Committee on Intelligence included his amendment to produce a comprehensive report of digital security threats and cybercrime in their bipartisan cybersecurity package.

Earlier this year, Senator Warner chaired a Senate Banking subcommittee hearing on the recent massive credit and debit card security breaches impacting major retailers like Target and Neiman Marcus and millions of American consumers.

In a [release accompanying his FTC letter](#) Warner said,

I don't know if Facebook's manipulation of users' news feeds was appropriate or not.

But I think many consumers were surprised to learn they had given permission by agreeing to Facebook's terms of service.

And I think the industry could benefit from a conversation about what are the appropriate rules of the road going forward.

If the cavalier attitudes expressed in Facebook's public statements about the matter are any indication, we'll need more than a conversation to prevent Facebook — and its ilk — from railroading academics, deceiving press and putting users at risk.

It's clear that Facebook is in a self-fulfilling bubble, much like Google, where it believes human beings are products; that they're only freaking out because they hate change or "don't get it" — and Facebook's ways of relating to the world have been severed in ways that facilitate a blatant disregard for the sanctity of other people's lives.

Because what's most foretelling of individual suffering, ultimately, is not the surveillance, the lying, or the messing with our heads, but the indifference of those in control.

## Sheryl Sandberg not sorry for Facebook mood manipulation study

By [Gail Sullivan](#) July 3, 2014

On Wednesday, Facebook's second-in-command, Sheryl Sandberg, expressed regret over how the company communicated its 2012 mood manipulation study of 700,000 unwitting users, but she did not apologize for conducting the controversial experiment. It's just what companies do, she said.

"This was part of ongoing research companies do to test different products, and that was what it was; it was poorly communicated," Sandberg, Facebook's chief operating officer, [told the Wall Street Journal](#) while travelling in New Delhi. "And for that communication we apologize. We never meant to upset you."

Sandberg's statement was the first public comment by a Facebook executive on the controversy since it erupted over the weekend, prompting anger from many Facebook users and [criticism from some academics](#) who said it was unethical to manipulate users' emotions without informed consent.

In the study, researchers at Facebook tweaked what hundreds of thousands of users saw in their news feeds, skewing content to be more positive or negative than normal in an attempt to manipulate their moods. Then

they checked users' status updates to see if the content affected what they wrote. They found that, yes, Facebook users' moods are affected by what they see in their news feeds. Users who saw more negative posts would write more negative things on their own walls, and likewise for positive posts.

Sandberg's apology is not likely to appease some, such as Robert Klitzman, a psychiatrist and ethics professor critical of the study, who said in a [column for CNN](#) that "the problem is not only how the study was described, but how it was conducted."

It seems that until now, Facebook data scientists have been pretty much free to do as they please. "There's no review process, per se," Andrew Ledvina, who worked at Facebook as a data scientist from 2012 to 2013, [told the Journal](#). "Anyone on that team could run a test," he said. "They're always trying to alter people's behavior." Ledvina told the Journal that tests were so frequent that some data scientists worried that the same users might be used in different studies, tainting the results.

Facebook has since implemented stricter guidelines, the Journal reported. Research other than routine product testing is reviewed by a panel of 50 internal experts in fields such as privacy and data security. Company research intended for publication in academic journals goes through a second round of review, again by in-house experts.

The upset over Facebook's mood study is "a glimpse into a wide-ranging practice," Kate Crawford, a visiting professor at the Massachusetts Institute of Technology's Center for Civic Media and a principal researcher at Microsoft Research told the Journal. Companies "really do see users as a willing experimental test bed" to be used at the companies' discretion.

Plenty of companies may do this sort of testing. But Facebook is different, John Gapper argued in the [Financial Times](#). Here's why, he said:

- "Facebook holds more intimate information about its users than other internet companies."
- Unlike testing products to see what appeals to users, which many companies do, with Facebook, "we are the product" being tested.
- "Facebook wields incredible power over the behavior of users. This is partly because of its size." He points to another [Facebook study](#) of 235 million users – noting that their sample size is four times the population of France.
- Facebook "focuses its judgments on personal material," unlike Google, which uses algorithms to analyze material across the Web. "An algorithm that selects from thousands of links about, say, Buckingham Palace feels like a service; one that weeds out the posts of friends and family feels like a moral guardian."
- "Facebook has demonstrated that it can alter behavior," he writes, citing studies that show users who see more status updates will write more themselves and another than encouraged users to become organ donors by allowing existing donors to display their that status.

## Everything We Know About Facebook's Secret Mood Manipulation Experiment

It was probably legal. But was it ethical?

ROBINSON MEYER JUN 28 2014, 2:51 PM ET

Facebook's News Feed—the main list of status updates, messages, and photos you see when you open Facebook on your computer or phone—is not a perfect mirror of the world.

But few users expect that Facebook would change their News Feed in order to manipulate their emotional state.

We now know that's exactly what happened two years ago. For one week in January 2012, data scientists skewed what almost 700,000 Facebook users saw when they logged into its service. Some people were shown content with a preponderance of happy and positive words; some were shown content analyzed as sadder than average. And when the week was over, these manipulated users were more likely to post either especially positive or negative words themselves.

This tinkering was just revealed as part of [a new study](#), published in the prestigious *Proceedings of the National Academy of Sciences*. Many previous studies have used Facebook data to examine “emotional contagion,” as this one did. This study is different because, while other studies have observed Facebook user data, this one set out to manipulate it.

The experiment is almost certainly legal. In the company's current terms of service, Facebook users relinquish the use of their data for “data analysis, testing, [and] research.” Is it ethical, though? Since news of the study

first emerged, I've seen and heard both privacy advocates and casual users express surprise at the audacity of the experiment.

We're tracking the ethical, legal, and philosophical response to this Facebook experiment here. We've also asked the authors of the study for comment. Author Jamie Guillory replied and referred us to a Facebook spokesman. Early Sunday morning, a Facebook spokesman sent this comment in an email:

This research was conducted for a single week in 2012 and none of the data used was associated with a specific person's Facebook account. We do research to improve our services and to make the content people see on Facebook as relevant and engaging as possible. A big part of this is understanding how people respond to different types of content, whether it's positive or negative in tone, news from friends, or information from pages they follow. We carefully consider what research we do and have a strong internal review process. There is no unnecessary collection of people's data in connection with these research initiatives and all data is stored securely.

And on Sunday afternoon, Adam D.I. Kramer, one of the study's authors and a Facebook employee, [commented on the experiment in a public Facebook post](#). "And at the end of the day, the actual impact on people in the experiment was the minimal amount to statistically detect it," he writes. "Having written and designed this experiment myself, I can tell you that our goal was never to upset anyone. [...] In hindsight, the research benefits of the paper may not have justified all of this anxiety."

Kramer [adds](#) that Facebook's internal review practices have "come a long way" since 2012, when the experiment was run.

### **What did the paper itself find?**

The study found that by manipulating the News Feeds displayed to 689,003 Facebook users, it could affect the content which those users posted to Facebook. More negative News Feeds led to more negative status messages, as more positive News Feeds led to positive statuses.

As far as the study was concerned, this meant that it had shown "that emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness." It touts that this emotional contagion can be achieved without "direct interaction between people" (because the unwitting subjects were only seeing each others' News Feeds).

The researchers add that never during the experiment could they read individual users' posts.

Two interesting things stuck out to me in the study.

The first? The effect the study documents is very small, as little as one-tenth of a percent of an observed change. That doesn't mean it's unimportant, though, as the authors add:

Given the massive scale of social networks such as Facebook, even small effects can have large aggregated consequences. [...] 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**.

The second was this line:

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 ( $z = -7.219, P < 0.001$ ).

In other words, when researchers reduced the appearance of *either* positive or negative sentiments in people's News Feeds—when the feeds just got generally less emotional—those people stopped writing so many words on Facebook.

Make people's feeds blander and they stop typing things into Facebook.

### Was the study well designed?

Perhaps not, says [John Grohol](#), the founder of psychology website Psych Central. Grohol believes the study's methods are hampered by the misuse of tools: Software better matched to analyze novels and essays, he says, is being applied toward the much shorter texts on social networks.

Let's look at two hypothetical examples of why this is important. Here are two sample tweets (or status updates) that are not uncommon:

- "I am not happy.
- "I am not having a great day."

An independent rater or judge would rate these two tweets as negative — they're clearly expressing a negative emotion. That would be +2 on the negative scale, and 0 on the positive scale.

But the LIWC 2007 tool doesn't see it that way. Instead, it would rate these two tweets as scoring +2 for positive (because of the words "great" and "happy") and +2 for negative (because of the word "not" in both texts).

"What the Facebook researchers clearly show," [writes Grohol](#), "is that they put too much faith in the tools they're using without understanding — and discussing — the tools' significant limitations."

### Did an **institutional review board (IRB)**—an independent ethics committee that vets research that involves humans—approve the experiment?

According to [a Cornell University press statement](#) on Monday, the experiment was conducted before an IRB was consulted. Cornell professor Jeffrey Hancock—an author of the study—began working on the results *after* Facebook had conducted the experiment. Hancock only had access to results, [says the release](#), so "Cornell University's Institutional Review Board concluded that he was not directly engaged in human research and that no review by the Cornell Human Research Protection Program was required."

In other words, the experiment had already been run, so its human subjects were beyond protecting. Assuming the researchers did not see users' confidential data, the results of the experiment could be examined without further endangering any subjects.

Both Cornell and Facebook have been reluctant to provide details about the process beyond their respective prepared statements. One of the study's authors told *The Atlantic* on Monday that he's been advised by the university not to speak to reporters.

By the time the study reached Susan Fiske, the Princeton University psychology professor who edited the study for publication, Cornell's IRB members had already determined it outside of their purview.

Fiske had earlier conveyed to *The Atlantic* that the experiment was IRB-approved.

“I was concerned,” Fiske told [The Atlantic](#) on Saturday, “until I queried the authors and they said their local institutional review board had approved it—and apparently on the grounds that Facebook apparently manipulates people's News Feeds all the time.”

On Sunday, other reports raised questions about how an IRB was consulted. In a [Facebook post](#) on Sunday, study author Adam Kramer referenced only “internal review practices.” And a [Forbes report](#) that day, citing an unnamed source, claimed that Facebook only used an internal review.

When [The Atlantic](#) asked Fiske to clarify Sunday, she said the researchers’ “revision letter said they had Cornell IRB approval as a ‘pre-existing dataset’ presumably from FB, who seems to have reviewed it as well in some unspecified way... Under IRB regulations, pre-existing dataset would have been approved previously and someone is just analyzing data already collected, often by someone else.”

The mention of a “pre-existing dataset” here matters because, as Fiske explained in a follow-up email, “presumably the data already existed when they applied to Cornell IRB.” (She also noted: “I am not second-guessing the decision.”) Cornell’s Monday statement [confirms](#) this presumption.

On Saturday, Fiske said that she didn’t want the “the originality of the research” to be lost, but called the experiment “an open ethical question.”

“It's ethically okay from the regulations perspective, but ethics are kind of social decisions. There's not an absolute answer. And so the level of outrage that appears to be happening suggests that maybe it shouldn't have been done...I'm still thinking about it and I'm a little creeped out, too.”

For more, check [Atlantic editor Adrienne LaFrance’s full interview with Prof. Fiske](#).

### **From what we know now, were the experiment’s subjects able to provide *informed consent*?**

In its [ethical principles and code of conduct](#), the American Psychological Association (APA) defines *informed consent* like this:

When psychologists conduct research or provide assessment, therapy, counseling, or consulting services in person or via electronic transmission or other forms of communication, they obtain the informed consent of the individual or individuals using language that is reasonably understandable to that person or persons except when conducting such activities without consent is mandated by law or governmental regulation or as otherwise provided in this Ethics Code.

As mentioned above, the research seems to have been carried out under Facebook’s extensive terms of service. The [company’s current data use policy](#), which governs exactly how it may use users’ data, runs to more than 9,000 words and uses the word “research” twice. But as [Forbes writer Kashmir Hill reported Monday night](#), the data use policy in effect when the experiment was conducted never mentioned “research” at all—the word wasn’t inserted until May 2012.

Never mind whether the current data use policy constitutes “language that is reasonably understandable”: Under the January 2012 terms of service, did Facebook secure even shaky consent?

The APA has further guidelines for so-called “deceptive research” like this, where the real purpose of the research can’t be made available to participants during research. The last of these guidelines is:

Psychologists explain any deception that is an integral feature of the design and conduct of an experiment to participants as early as is feasible, preferably at the conclusion of their participation, but no later than at the conclusion of the data collection, and permit participants to withdraw their data.

At the end of the experiment, did Facebook tell the user-subjects that their News Feeds had been altered for the sake of research? If so, the study never mentions it.

James Grimmelman, a law professor at the University of Maryland, believes the study did not secure informed consent. And [he adds that Facebook fails even its own standards](#), which are lower than that of the academy:

A stronger reason is that even when Facebook manipulates our News Feeds to sell us things, it is supposed—legally and ethically—to meet certain minimal standards. Anything on Facebook that is actually an ad is labelled as such (even if not always clearly.) This study failed even that test, and for a particularly unappealing research goal: *We wanted to see if we could make you feel bad without you noticing. We succeeded.*

### **Did the U.S. government sponsor the research?**

~~Yes, according to a [press release from Cornell University](#) describing the study. Cornell's story identifies the [Army Research Office](#)—an agency within the U.S. Army that funds basic research in the military's interest—as one of the funders of the experiment.~~

Cornell has [now updated their](#) June 10 story to say that the research received no external funding.

### **Do these kind of News Feed tweaks happen at other times?**

At any one time, Facebook said last year, there were [on average 1,500 pieces of content](#) that could show up in your News Feed. The company uses an algorithm to determine what to display and what to hide.

It talks about this algorithm very rarely, but we know it's very powerful. Last year, the company changed News Feed to surface more news stories. [Websites like BuzzFeed and Upworthy proceeded to see record-busting](#) numbers of visitors.

So we know it happens. Consider Fiske's explanation of the research ethics here—the study was approved “on the grounds that Facebook apparently manipulates people's News Feeds all the time.” And consider also that *from this study alone* Facebook knows at least one knob to tweak to get users to post more words on Facebook.

## Facebook Experiments Had a Few Limits

by Reed Albergotti July 2, 2014

Few of Facebook's 1.3 billion users know much about the social network's Data Science team, but that's changed after reports have come out about a psychological study conducted on nearly 700,000 users. WSJ's Reed Albergotti joins Tanya Rivero on Lunch Break with an inside look at Facebook's research team.

Thousands of Facebook Inc. users received an unsettling message two years ago: They were being locked out of the social network because Facebook believed they were robots or using fake names. To get back in, the users had to prove they were real.

In fact, Facebook knew most of the users were legitimate. The message was a test designed to help improve Facebook's antifraud measures. In the end, no users lost access permanently.

The experiment was the work of Facebook's Data Science team, a group of about three dozen researchers with unique access to one of the world's richest data troves: the movements, musings and emotions of Facebook's 1.3 billion users.

The little-known group was thrust into the spotlight this week by reports about a 2012 experiment in which the news feeds of nearly 700,000 Facebook users were manipulated to show more positive or negative posts. The study found that users who saw more positive content were more likely to write positive posts, and vice versa.

Facebook Chief Operating Officer Sheryl Sandberg said Wednesday during a trip to India that the study was "part of ongoing research companies do to test different products" and was "poorly communicated."

The company said that after the feedback on the study, "We are taking a very hard look at this process to make more improvements."

Until recently, the Data Science group operated with few boundaries, according to a former member of the team and outside researchers. At a university, researchers likely would have been required to obtain consent from participants in such a study. But Facebook relied on users' agreement to its Terms of Service, which at the time said data could be used to improve Facebook's products. Those terms now say that user data may be used for research.

"There's no review process, per se," said Andrew Ledvina, a Facebook data scientist from February 2012 to July 2013. "Anyone on that team could run a test," Mr. Ledvina said. "They're always trying to alter peoples' behavior."

He recalled a minor experiment in which he and a product manager ran a test without telling anyone else at the company. Tests were run so often, he said, that some data scientists worried that the same users, who were anonymous, might be used in more than one experiment, tainting the results.

Facebook said that since the study on emotions, it has implemented stricter guidelines on Data Science team research. Since at least the beginning of this year, research beyond the routine product testing is reviewed by a panel drawn from a group of 50 internal experts in fields such as privacy and data security. Facebook declined to name them.

Company research intended to be published in academic journals receives additional review from in-house experts on academic research. Some of those experts are also on the Data Science team, Facebook said, declining to name the members of that panel.

A spokesman said Facebook is considering additional changes.

Since its creation in 2007, Facebook's Data Science group has run hundreds of tests. One published study deconstructed how families communicate, another delved into the causes of loneliness. One test looked at how social behaviors spread through networks. In 2010, the group measured how "political mobilization messages" sent to 61 million people caused people in social networks to vote in the 2010 congressional elections.

Many of Facebook's data scientists hold doctoral degrees from major universities in fields including computer science, artificial intelligence and computational biology. Some worked in academic research before joining Facebook.

Adam Kramer, the lead author of the study about emotions, said in a 2012 interview on Facebook's website that he joined the company partly because it is "the largest field study in the history of the world." Mr. Kramer, who has a doctorate in social psychology from the University of Oregon, said that in academia he would have had to get papers published and then hope that someone noticed. At Facebook, "I just message someone on the right team and my research has an impact within weeks, if not days."

Much of Facebook's research is less controversial than the emotions study, testing features that will prompt users to spend more time on the network and click on more ads. Other Internet companies, including Yahoo Inc., Microsoft Corp., Twitter Inc., and Google Inc., conduct research on their users and their data.

The recent ruckus is "a glimpse into a wide-ranging practice," said Kate Crawford, a visiting professor at the Massachusetts Institute of Technology's Center for Civic Media and a principal researcher at Microsoft Research. Companies "really do see users as a willing experimental test bed" to be used at the companies' discretion.

Facebook's team has drawn particular interest because it occasionally publishes its work in academic journals that touch on users' personal lives, including the study about positive and negative posts.

"Facebook deserves a lot of credit for pushing as much research into the public domain as they do," said Clifford Lampe, an associate professor at the University of Michigan's School of Information who has worked

on about 10 studies with Facebook researchers. If Facebook stopped publishing studies, he said, "It would be a real loss for science."

Dr. Lampe said he has been in touch with members of the Data Science team since the controversy erupted, "They've been listening to the arguments and they take them very seriously," he said.

Mr. Ledvina, the former Facebook data scientist, said some researchers debated the merits of a study similar to the one that accused users of being robots but there was no formal review, and none of the users in the study were notified that it was an experiment.

"I'm sure some people got very angry somewhere," he said. "Internally, you get a little desensitized to it."