

Facebook's flawed emotion experiment: Antisocial research on social network users

Research Ethics
2016, Vol. 12(1) 29–34
© The Author(s) 2015
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1747016115579535
rea.sagepub.com


David Shaw

University of Basel, Switzerland

Abstract

In June 2014, a paper reporting the results of a study into ‘emotional contagion’ on Facebook was published. This research has already attracted a great deal of criticism for problems surrounding informed consent. While most of this criticism is justified, other relevant consent issues have gone unremarked, and the study has several other ethical flaws which collectively indicate the need for better regulation of health and mood research using social networks.

Keywords

consent, harm, Facebook, IRB, REC, social networks

Introduction

In June 2014, a paper reporting the results of a study into ‘emotional contagion’ on Facebook was published (Kramer et al., 2014). This research has already attracted a great deal of criticism for problems surrounding informed consent (Bradshaw, 2014). While most of this criticism is justified, other relevant consent issues have gone unremarked, and the study has several other ethical flaws which collectively indicate the need for better regulation of health and mood research using social networks.

Using specialized software, Kramer, Guillory and Hancock manipulated the News Feeds of 689,003 Facebook users. For one week in 2012, half of these participants had some negative stories posted by their friends removed from the

Corresponding author:

David Shaw, Institute for Biomedical Ethics, University of Basel, Beroullistrasse 28, 4056 Basel, Switzerland.

Email: david.shaw@unibas.ch

Feed, and the other half had some positive stories removed. This means that for seven days over 300,000 Facebook users had either much more positive or much more negative News Feeds than usual. For all they knew, their friends were having either really good weeks or really bad weeks. The results showed that those in the “negative” group tended to then post negative stories themselves, and those in the “positive” group tended to post positive stories.

Consent

The authors claim that their research methods were “consistent with Facebook’s Data Use Policy (DUP), to which all users agree prior to creating an account on Facebook, constituting informed consent for this research.” However, none of the participants even knew that they were taking part in the study. This alone would tend to suggest that none of them actually gave informed consent. Of course, many Facebook users rarely read the full Terms and Conditions; perhaps the Data Use Policy clearly states that user data will be used for psychological research? In fact, the Policy simply states that “we may use the information we receive about you ... for internal operations, including troubleshooting, data analysis, testing, research and service improvement.” The clear implication here is that any data will be used for internal software and systems research, rather than for psychological experiments the results of which will be made public. Furthermore, even if users did realize that this wider sense of research was intended, the Policy says nothing about *selectively manipulating the data provided to users*, which would also be required for informed consent. In any case, informed consent normally means knowing exactly what is involved in the research. Facebook are clearly guilty of conducting human subjects research without valid consent. Facebook’s own response was “none of the data used was associated with a specific person’s Facebook account ... [there was] “no unnecessary collection of people’s data” (Hill, 2014), which focuses on confidentiality rather than consent and thus misses the point entirely.

Since the story initially broke, it has also emerged that Facebook’s Policy at the time of the study did not actually mention research, casting even more serious doubt on the claims of the researchers (Hern, 2014). It could even be argued that the posters of hidden stories were also participants in the research; they were not told their communications to their friends would be censored. The same Policy states “we may ... pick stories for your News Feed”, but this does not appear to cover hiding stories for the purposes of research.

Harm

Conducting research without consent is concerning, but in this case the results could be interpreted as suggesting that some participants may have been psychologically

harmful by the study. The main result is that exposure to negative posts made people feel worse; if valid, this means that hundreds of thousands of people were made less happy by the study, a fact that is hardly offset by the fact that a similar number were manipulated into having more positive emotions. Indeed, this latter group may have missed the opportunity to support friends who posted about difficulties or sadness, as these stories were hidden. And these friends in turn might be worse off for not receiving this support. As the authors state, “the well-documented connection between emotions and physical well-being suggests the importance of these findings for public health”. It is unfortunate that they did not realize the public health ramifications of their own study on its participants, some of whom were presumably children, as anyone over 13 can have a Facebook account and the researchers do not appear to have stratified by age. Of the USA’s 133,518,980 Facebook users, 12 million are children over 13, but 7.5 million children under the age of 13 also use Facebook every year in the US alone (Peckham, 2011). If we extrapolate these figures and apply them to the study, it would mean that 15% or more than 100,000 participants were children. Additionally, given that at least 5% of people are suffering from depression at any given time, we can assume that at least 30,000 people with depression were included in the study. It is certain that some people from these vulnerable populations were recruited into this study. Another problem is that, even though the study has now been publicized, the identities of the participants remain a mystery. This means that all English-speaking Facebook users still do not know whether or not they were manipulated. In this sense, even non-participants are affected by the consent failures of the study.

Design

In addition to concerns about consent and harm, the design of the study also raises other issues. First, it is unclear why such a massive sample size was chosen. It is widely recognized that it is unethical to recruit more participants than is required to meet a study’s objective. In this case, similar results could have been obtained using only a few thousand participants. As well as being scientifically questionable, this means that hundreds of thousands of people were subjected to emotional manipulation unnecessarily. Furthermore, doubts have also been raised about the validity of the methods, given that no real-world assessment of emotion was conducted, the text analysis tool was inappropriate for short sections of text, and the effect size was very small (Grohol, 2014). Given these flaws, the results can perhaps not be taken as evidence that significant psychological harm was inflicted on participants, but harm may nonetheless have occurred, and the issues surrounding consent and deception remain.

Review

The study also raises the issue of ethics review. Despite initial appearances that no ethics committee was involved, it has transpired that it was actually reviewed by a university IRB and determined to be exempt from review. Why? “Because the research was conducted independently by Facebook and Professor Hancock had access only to results – and not to any individual, identifiable data at any time – Cornell University’s Institutional Review Board concluded that he was not directly engaged in human research” (Cornell University, 2014). Most research ethics committees would immediately reject the suggestion that one member of a research team was not involved in human subjects research simply because he or she did not have access to identifiable data while the other researchers did. The only way in which a case could be made for review not being necessary would be to say that the data-gathering had already been conducted before review was sought – another situation where an IRB or REC would immediately refuse permission, as projects must be approved before data-gathering commences. It is quite clear that Cornell were mistaken here. (The editor of PNAS also supported this somewhat tenuous conclusion; Lee, 2014.) Furthermore, given that all users of English-language Facebook could have been participants, ethical approval should probably have been sought in every country where participants lived, in line with normal procedures for international human subjects research. Only English-language speakers were involved, but this still means that any Facebook users in Britain, Canada, Australia, New Zealand and everyone anywhere else in the world who had English set as their language option could have been included in the study, meaning that dozens if not hundreds of jurisdictions could have been involved. At a minimum, Facebook should have checked the regulations in each of these jurisdictions to see whether IRB or REC approval was required.

Saving Face(book)

In an attempt to forestall the rising swell of indignant reaction to his study, the lead researcher issued an apology and explanation at the end of June (Kramer, 2014). As is often the case with such attempts, it only makes things worse. Kramer begins by stating that “The reason we did this research is because we care about the emotional impact of Facebook and the people that use our product.” This claim is somewhat undermined by their failure to care about this during the conduct of the research. He goes on to state “my coauthors and I are very sorry for the way the paper described the research and any anxiety it caused. In hindsight, the research benefits of the paper may not have justified all of this anxiety.” It is unclear whether he means the anxiety it caused to participants or the anxiety caused to those learning about the study after the fact. Finally, he admits shortcomings in the review process, but claims that Facebook “have been working on improving our internal

review practices. The experiment in question was run in early 2012, and we have come a long way since then. Those review practices will also incorporate what we've learned from the reaction to this paper." Once again, we can see a failure to acknowledge the importance of comprehensive external review by an independent body.

The mistaken bioethics backlash

Despite the initial condemnation of Facebook from a wide variety of commentators, there was something of a backlash against criticism of Facebook from some ethicists in the following weeks. Michelle Meyer and colleagues argued in a *Nature* commentary that "the vitriolic criticism of this study could have a chilling effect on valuable research. Worse, it perpetuates the presumption that research is dangerous: (Meyer et al., 2014). These authors argued that what Facebook did in this experiment was no different from what it normally does on a daily basis with its millions of users, and so there was really nothing unethical about the research: "Some have said that Facebook 'purposefully messed with people's minds'. Maybe; but no more so than usual." However, this claim is contradicted by the authors' correct statement that the experiment "had the effect of concentrating the feeds with negative and positive content, respectively", which is obviously not something that happens routinely. In this sense, it was a clear deviation from normal Facebook practice. Furthermore, Meyer and colleagues argue that no consent was required, and that seeking consent might have affected the results of the study. However, Facebook claimed not only that consent was necessary, but that it was obtained, as discussed above, so these arguments somewhat miss the point. Furthermore, Meyer and colleagues also seem to assume that only US citizens were involved in the study, a US-centric view that overlooks the fact that Facebook is a global corporation and that the laws of dozens of several jurisdictions must also be considered, as noted above. Even if Meyer and colleagues are correct that review was not required under US legislation, it is misleading to claim that this is the only jurisdiction that matters. In addition, the authors seem to accept PNAS and Cornell's mistaken reasoning regarding ethics review. While Meyer and colleagues are perhaps correct that this experiment was not outrageously wrong, it was nonetheless unethical, and Facebook's attempts to justify the researchers' actions only made things worse.

Conclusion

This Facebook study was conducted without consent and without appropriate oversight, and may have harmed both participants and non-participants. Kramer's apology also puts the vast number of participants in context; a full 0.04% of

Facebook's users, or 1 in 2500, were unwitting subjects in this unethical research. Many of these people were almost certainly children, and many of the participants were probably suffering from depression. It is surprising and worrying that one of the world's most prominent companies should treat both the emotions of its users and research ethics so carelessly. Steps must be taken to ensure that international psychological and medical studies involving social network users are regulated to the same standard as other human subjects research.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- Bradshaw T (2014) Facebook hit by psychology test backlash. *Financial Times*, 29 June. Available at: <http://www.ft.com/cms/s/0/d2f493d6-ffa7-11e3-9a4a-00144feab7de.html> (accessed 2 July 2014).
- Cornell University Media Relations Office (2014) Media statement on Cornell University's role in Facebook 'emotional contagion' research. Available at: <http://mediarelations.cornell.edu/2014/06/30/media-statement-on-cornell-universitys-role-in-facebook-emotional-contagion-research/> (accessed 5 November 2014).
- Grohol JM (2014) Emotional contagion on Facebook? More like bad research methods. *World of Psychology* blog. Available at: <http://psychcentral.com/blog/archives/2014/06/23/emotional-contagion-on-facebook-more-like-bad-research-methods/> (accessed 2 July 2014).
- Hern A (2014) Facebook T&Cs introduced 'research' policy months after emotion study. *The Guardian*, 1 July. Available at: <http://www.theguardian.com/technology/2014/jul/01/facebook-data-policy-research-emotion-study> (accessed 2 July 2014).
- Hill K (2014) Facebook manipulated 689,003 users' emotions for science. *Forbes*. Available at: <http://www.forbes.com/sites/kashmirhill/2014/06/28/facebook-manipulated-689003-users-emotions-for-science/> (accessed 2 July 2014).
- Kramer A (2014) Facebook post. Available at: <https://www.facebook.com/akramer/posts/10152987150867796> (accessed 2 July 2014).
- Kramer A, Guillory JE and Hancock JT (2014) Experimental evidence of massive-scale emotional contagion through social networks. *PNAS*. Available at: <http://www.pnas.org/content/111/24/8788.full> (accessed 2 July 2014).
- Lee Z (2014) PNAS response re Facebook study on emotions. Available at: <https://twitter.com/ZLeeily/status/482937917921579008> (accessed 2 July 2014).
- Meyer M, Lantos J, London A et al. (2014) Misjudgements will drive social trials underground. *Nature* 511: 265.
- Peckham M (2011) Shocker: Millions of Facebook users underage. *Time*. Available at: <http://techland.time.com/2011/05/10/shocker-millions-of-facebook-users-underage/> (accessed 5 November 2014).