

The Ethics of Online Experimentation with Unsuspecting Users

Raquel Benbunan-Fich

S/CIS Department, Zicklin School of Business,
Baruch College, CUNY
rbfich@baruch.cuny.edu

In the summer of 2014, the journal *Proceedings of the National Academy of Sciences* (PNAS) published a study about emotional contagion conducted by a Facebook researcher in collaboration with two academics.¹ The premise of the study was that emotional states can spread through social networks. The data was gathered for one week in January 2012 from 689,003 randomly selected users who interact with Facebook's English version.² Two parallel experiments were conducted; one for positive and the other for negative emotional content. In both cases, user exposure to friends' emotional content displayed via the News Feed was reduced. Both experiments had their own control condition, in which a similar proportion of posts in the News Feed were omitted regardless of their emotional content. The results indicate that users in the reduced-positive condition had a larger percentage of negative words in their own status updates and a smaller percentage of positive words. The opposite pattern was observed in the reduced-negative condition. Although the effect size for the manipulations was very small, the findings provided statistically significant evidence of emotional contagion via social networks.

The Facebook emotion experiment illustrates the potential and the risks of conducting large-scale research through social computing platforms. The upside of deploying experiments via social networks is the immediate access to a large sample of online users, the convenience of running simultaneous conditions, and the option to study human behavior from a new perspective.³ Only the providers of these online platforms can conduct this kind of research. The downside of this approach is the danger of overlooking fundamental safeguards for the protection of human subjects who unknowingly become participants in these experiments. The Facebook emotion experiment is just one example of a new wave of online experiments with unsuspecting users of social platforms. These experiments are novel because online companies modify the inner workings of key computer programs to change or curate information about the users themselves or about their friends (i.e. social network connections). As such, this is a deeper and more advanced form of behavioral experimentation compared to the traditional surface-level (or interface) testing aimed to improve the design of websites. When programming code is altered to

¹ Kramer, A.D.; Guillory, J.E. and Hancock, J.T. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, vol. 111, no. 24, 8788–8790.

² As of June 2014, Facebook the world's largest social networking site had over 1.3 billion active users. About half of them use the English version of the site.

³ Manjoo, F. "A bright side to Facebook's experiments with its users" *The New York Times*, July 2, 2014. http://www.nytimes.com/2014/07/03/technology/personaltech/the-bright-side-of-facebooks-social-experiments-on-users.html?_r=0

manipulate information about the users or about their connections, without forewarning, intentional deception takes place.

This essay analyzes several cases of company-sponsored online experiments with unsuspecting users and discusses the ethical aspects of such experimentation. Three aspects raise ethical concerns regarding this new type research: the absence of user consent to participate in research, the presence of intentional deception, and the complete lack of protection human subjects who partake in privately funded behavioral research. A careful consideration of these issues suggest the need to expand the ethical framework for human subject protection, and to incorporate an active oversight for private companies that conduct online experiments through social networking platforms.

Facebook experiments with its News Feeds

For the emotional contagion study, Facebook researchers developed a special program to automatically determine the valence of the emotional content selected for users' News Feed and implement the intended manipulation depending on the case. The News Feed is a feature that displays comments, videos, pictures and links posted by other people in the user's social network.⁴ Although this feature has been subject to evolutionary changes since its inception, Facebook describes it as a "digest" of relevant news updates based on each user's preferences and connections. While the content is personalized for each user by the News Feed program, the process to generate the display is common to all users; it shows updates from friends and pages with which users interact frequently. The assumption that the News Feed uses a similar algorithm for all users was violated. The addition of the extra program that analyzed the nature of emotional content and automatically decided which content to suppress depending on the condition introduce an additional layer of "editorial curation" that was never disclosed to users.

A prior Facebook experiment also manipulated the content displayed in the News Feed for research purposes, albeit with more obvious curation and more startling results. In an experiment conducted during Election Day in 2010, Facebook researchers randomly divided 61 million American users into three conditions: social message (n=60 million), informational message (n=611,000), and a control group (n=613,000).⁵ Each of the two message-conditions was shown a different, but non-political, get-out-to-vote statement. The social message group was shown a statement at the top of their News Feed to encourage the user to vote, along with a link to local polling places, a clickable button with the words "I Voted," a counter indicating how many other Facebook users reported voting, and up to six small randomly selected profile pictures of the user's Facebook friends who had already clicked the I Voted button. The informational message

⁴ The NewsFeed algorithm automatically determines what to show for each user depending on the users' connections and activity on Facebook. The News Feed results are personalized for each user and updated continuously. <https://www.facebook.com/help/327131014036297/>

⁵ Bond, R.M.; Fariss, C.J.; Jones, J.J.; Kramer, A.D.; Marlow, C.; Settle, J.E. and Fowler, J.H. (2012) "A 61-million-person experiment in social influence and political mobilization." *Nature* 489, 295–298. http://cameronmarlow.com/media/massive_turnout.pdf

group was shown the message, poll information, counter and button, but none of the friends' pictures. The control group did not receive any message at the top of their News Feed. The results indicate that those who received the social message were more likely to vote than users who received the informational message and those in the control group. There was no difference in turnout between the informational and control condition.

The effects of the News Feed's manipulation on actual voting were validated "through examination of public voting records" from several states.⁶ According to the Supplementary Information of this study,⁷ several states provided to Facebook their publicly available voting records for research purposes,⁸ allowing Facebook to compile a list of over 6.3 million matched subjects with corresponding online user account and offline voter information. This matched sample was used to perform a statistical analysis on the relationship between online treatment conditions and real-world voting behavior.

Compared to the emotional contagion project, the political turnout experiment generated much less publicity and very little backlash. In both cases, Facebook manipulated the News Feed algorithm to either suppress emotional information or enhance get-out-to-vote messages displayed in the automatically-generated and user-personalized flow of updates. Both experiments crossed the line between online and offline behavior in different ways by analyzing how social influence and contagion can change people's voting behavior or emotional states. In the turnout experiment, Facebook's evidence indicates that social messages showing the faces of friends who voted "increased turnout directly by about 60,000 voters and indirectly through social contagion,"⁹ suggesting that "seeing faces of friends significantly contributed to the overall effect of the message on real-world voting."¹⁰

This political turnout experiment provided evidence of the connection between online information displayed via social networks and actual behavior, and how companies can effectively match online users with records of offline activity. Although this raises a number of privacy concerns, the researchers argue that they never "see" individual records because the combination is done with computerized programs. A similar argument was used by the researchers in the Facebook emotional contagion study when they claimed that the modification of the NewsFeed results was done with a custom-developed program. This essay is focused on unethical practices of online experimentation with social networking users, excluding cases of

⁶ Bond et al. (2012), p. 295

⁷ The Supplement for the Massive turnout project is available at http://fowler.ucsd.edu/massive_turnout_supplement.pdf.

⁸ "The cost of state records varied from \$0 to \$1500 per state... The resulting list of states included Arkansas, California, Connecticut, Florida, Kansas, Kentucky, Missouri, Nevada, New Jersey, New York, Oklahoma, Pennsylvania, and Rhode Island. These states account for about 40% of all registered voters in the U.S., and their records yielded 6,338,882 matched observations of voters and abstainers that we could use to compare to treatment categories from the experiment." Massive Turnout Supplement, p. 3-4.

⁹ Bond et al. (2012), p. 297.

¹⁰ Bond et al. (2012), p. 296.

potential privacy breaches. The ethical aspects of privacy violations in online experimentation deserve a separate treatment, which is outside the scope of this essay.

Another case of online manipulation of emotions

In addition to Facebook, other online companies have pushed the boundaries of online experimentation. At around the same time as the publication of the Facebook emotion experiment, the online dating website OKCupid¹¹ admitted that their researchers also conduct experiments with unsuspecting users. The company blog described three experiments conducted at different points in time.¹² In the first, they removed the pictures from all the profiles in “love is blind” day (January 15th, 2013). During the picture-blackout period, more communication and information exchange took place among “blind” users, but those conversations stopped when the pictures were restored at 4pm that day. In the second experiment, OKCupid manipulated the interface to show profile pictures with or without profile text and replaced the rating scales for personality and looks with one to measure how “cool” the person in the profile was perceived. The results show that the coolness ratings are entirely driven by the profile picture, and the profile text has no influence on the ratings given by other users. In the third experiment, OKCupid altered the compatibility percentage automatically provided by their matching algorithm to suggest that people were a much better or worse match than their actual match score.¹³

The OKCupid mismatching experiment, dubbed “the power of suggestion,” took pairs of users deemed as bad matches by the matching algorithm (30% compatibility level) and changed the level to display a 90% match, suggesting that they were a good fit for each other. As expected, misled users sent more first messages to their potential matches when they thought they were compatible, and engaged in conversations (or exchanged four messages or more) with their partners. At the conclusion of this experiment, affected users were notified of the correct match percentage. Upon examining the number of conversations that took place between mismatched partners, OKCupid was concerned that people interact due to the power of suggestion (induced by the fictitious compatibility level) and that their matching algorithm was useless to predict real compatibility. In order to rule out this possibility, OKCupid tested additional combinations where compatibility matches of 30, 60 and 90 were either accurately displayed or changed for one of the other two percentages. The ideal combination measured by the odds of an initial message turning into a conversation, was found in the 90%-90% condition, where people were a good match according to the algorithm, and they could see the true compatibility percentage displayed

¹¹ OKCupid is a free dating site that matches users through mathematical algorithms based on answers to questions about their preferences and tastes. The size of its user base was 12 million users according to estimates reported by Bloomberg BusinessWeek in September of 2014 (<http://www.businessweek.com/articles/2014-09-04/mining-okcupids-data-reveals-how-we-date-now>)

¹² OKCupid Company Blog (<http://blog.okcupid.com/index.php/we-experiment-on-human-beings/>)

¹³ Wood, M. OKCupid plays with love in user experiments. New York Times, July 28, 2014
<http://www.nytimes.com/2014/07/29/technology/okcupid-publishes-findings-of-user-experiments.html>

for them. By deliberately mismatching people, OKCupid tested its matching algorithm with unsuspecting users to examine their reactions.

A New Form of Online Experimentation

The first two experiments conducted by OKCupid (picture-blackout and pictures with/without text) are instances of what is commonly known as “A/B testing.” In this type of testing, a fraction of unsuspecting users is diverted to a different version of a web page to compare their behavior with others who are using the standard site. For example, the first experiment at OKCupid tested the effects of the profiles with and without pictures, and the second tested the effects of text or no-text along with the profile picture. In essence, A/B testing is a split testing of two versions of a website, or a feature within a website, to evaluate which one performs better according to relevant metrics (such as how long users stay on the site, and which links or buttons are clicked). However, the third experiment conducted by OKCupid is not an instance of A/B testing because it intentionally misleads unsuspecting users about their compatibility percentages to see how they would react. This is not a web-site design change; it is an instance of algorithmic manipulation and misinformation.

A distinction must be drawn between traditional A/B testing and this form of experimentation. In A/B testing, interface design characteristics such as arrangement of buttons, layout, or explanatory text are manipulated to test their effects. Many online companies routinely perform this type of tests with unsuspecting users to assess the impact of web-site design changes. A new form of experimentation emerges when the programming *code* of an algorithm is altered to induce *deception*. This is a deep form of testing, which I propose to call Code/Deception or *C/D experimentation* to distinguish it from the surface level testing associated with A/B testing. The Facebook emotion and turnout projects and the OKCupid “power of suggestion” experiment involved instances of deception either through the manipulation of information (reduction of emotional content) or outright lies (about compatibility matches). What makes C/D experimentation different from A/B testing is that the information manipulated is about the user or about people in the users’ social network. In contrast, the manipulation in A/B testing is about the design of site itself (buttons, explanations, descriptive text, colors, etc.).

C/D experimentation engenders additional risks for unsuspecting users, quite different from A/B testing. The failure to recognize the distinction may lead to the erroneous belief on the part of researchers and sponsors that this activity is covered by “terms of use” agreements. For example, the researchers involved in Facebook’s emotional contagion study argued that the acceptance of Facebook’s Data Use Policy, which is a condition for establishing a user account in Facebook, provided consent for their study. It should be noted that the Facebook data use policy in effect at the time of the emotional contagion experiment in January of 2012 did not mention the possibility of using information collected by Facebook for “research” purposes. To address this

gap, a few months later, in May of 2012, the policy was amended to reflect several changes including the addition of “research” in the list of potential “internal operation” uses.¹⁴

Implicit Consent via Terms of Service agreements

While Facebook’s terms of use did not include “research” as a possible use for the information collected, the user agreement in effect at OKCupid did incorporate the possibility of using data for research and analysis purposes. Typical user agreements include some language to indicate that the company will use the data for testing-troubleshooting and service improvements. The argument is that online companies automatically acquire *implicit consent* for research when a user accepts the Terms of Service (TOS).

User acceptance of the Terms of Service agreement by clicking on a checkbox is one of the requirements for account creation in most social networking and other commercial websites. The binding action is exercised with a click-through instead of a signature, and for this reason these contracts are known as click-through agreements. These agreements are complex and difficult to read and thus raise doubts on the validity of ‘informed consent.’¹⁵ Due to their length, most people fail to read the content of TOS agreements, and are unaware of their content. For example, the length of Facebook’s TOS at the time of the emotion study¹⁶ was about 6,700 words, and OKCupid’s current TOS are about 3,700 words. At a rate of 200 words per minute, it would have taken an average reader about 33.5 minutes to read Facebook’s TOS, and 18.5 minutes to read OKCupid’s. Yet, research shows that people spend an average of half a minute before clicking on the agreement box.¹⁷

Another study based on a content analysis of the TOS of 30 popular websites found that, due to language complexity and the use of legal terminology, users may not understand which rights they are granting when they post their creative content on these sites,¹⁸ even if they take the time to read the terms. A software solution in the form of a browser extension has been developed to help individuals understand in plain language the main provisions of TOS.¹⁹ The lack of reading or understanding of TOS applies to all conditions and restrictions that users “accept” when they click-through. Thus, users may not realize that they are implicitly consenting to participate in

¹⁴ Forbes "Facebook Added 'Research' To User Agreement 4 Months After Emotion Manipulation Study" <http://www.forbes.com/sites/kashmirhill/2014/06/30/facebook-only-got-permission-to-do-research-on-users-after-emotion-manipulation-study/>

¹⁵ Luger, E.; Moran, S. and Rodden, T. “Consent for all: revealing the hidden complexity of terms and conditions.” CHI '13 Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, April 27–May 2, 2013, Paris, France.

¹⁶ Facebook User Agreement dated Sept. 2011 is posted by the Coudrain Group at <http://thecoudrain.com/files/documents/Facebook-Data-Use-Policy.pdf>

¹⁷ Bakos, Y. Marotta-Wurgler, F. Trossen, D. R. Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts. *Journal of Legal Studies*, Vol. 43, No. 1, 2014.

¹⁸ Fiesler, C. and Bruckman, A. “Copyright terms in online creative communities.” CHI '14: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Apr 26 - May 01, 2014, Toronto, ON, Canada.

¹⁹ The project is called “Terms of Service; Didn’t Read” and it features a categorization of TOS based on user ratings. (See <http://tosdr.org>).

company-sponsored research without additional notice. Strictly speaking, since OKCupid did anticipate, and explicitly listed, research as one of the potential uses, it could claim that it had obtained implicit consent for experimentation. In contrast, since Facebook failed to include research purposes in the list of anticipated uses of information collected it cannot assume that it had implicit consent for the emotion contagion study. In either case, implicit consent for research is not the same as *informed* consent for a specific study.

Informed Consent and Protection of Human Subjects

In academia, researchers often struggle with the amount of information to disclose in consent forms. On the one hand, explicit consent forms must include a description of the study, and outline the risks and benefits so that participants can opt-in with awareness of participation requirements and possible consequences. On the other hand, full disclosure of research details bears the risk of priming participants and altering their otherwise “natural” behavior to conform to experimenters’ expectations (i.e. the experimenter-participant artifact problem). To balance the ethical duty to inform prospective participants with the need to leave their behavior unaffected, consent forms usually include only general descriptions of research objectives, risks, and benefits. The main objective of the consent procedure is to *inform* prospective participants of the nature of the study and their chance to participate in it, if they so desire. Hence the term informed consent.

The informed consent requirement is part of the U.S. government’s set of regulations for the protection of human subjects, known as the “Common Rule.”²⁰ To ensure that all US federally funded research complies with existing regulations, research protocols must be reviewed and approved by an independent review board, before it is conducted. The review of research procedures is meant to ensure that risks to participants are minimized and commensurate with anticipated benefits. It must be noted that the Common Rule only applies to U.S. federally funded behavioral and biomedical research conducted at academic and other institutions.²¹ Private companies like Facebook or OKCupid are not legally required to follow the regulations in the Common Rule. This legal loophole creates inequality in the protection of human subjects that participate in company-sponsored research. The absence of a legal regulatory framework creates opportunities for unethical research practices with minimal adverse consequences.

²⁰ US Department of Health and Human Services, Federal Rules for the Protection of Human Subjects (Common Rule), <http://www.hhs.gov/ohrp/humansubjects/commonrule/>

²¹ The focus of this essay is behavioral research. Biomedical research conducted by pharmaceutical companies and other contract research organizations must also comply with the Common Rule. These companies conduct research to file applications seeking approval of new drugs or medical devices from the U.S. Food and Drug Administration (FDA). As such, the research procedures should have been previously reviewed and approved to ensure compliance with the protection of human subjects.

The core ethical principles used as the guideline for the Common Rule were initially outlined in the Belmont report.²² According to the report, the three basic principles for ethical research conducted with human subjects at institutions that receive federal funds are: (1) respect for people, (2) beneficence, and (3) justice. *Respect for people* includes two related principles: individuals should be treated as autonomous agents, and individuals with diminished autonomy (children, prisoners, etc.) are entitled to protection. *Beneficence* refers to the obligation to treat individuals in an ethical manner by protecting them from harm, and making efforts to ensure their wellbeing. Justice is concerned with the equitable distribution of the benefits and burdens of research. The C/D experiments described above did not comply with any of these three principles. First, since participants were uninformed and prevented from opting out, they were not treated as autonomous agents or afforded respect. Second, the manipulation of emotions or potential dating matches carries the risk of causing psychological harm to the users exposed to these tactics, without any efforts to ensure their well-being, both of which are inconsistent with the principle of beneficence. Third, an undue burden was imposed on those who were randomly selected for experimental conditions that were clearly intended to benefit the sponsoring companies (by allowing them to test the power of emotional contagion or the power of suggestion), and therefore an injustice was committed against these unsuspecting users. On the basis of this analysis, it can be concluded that these experiments did not follow ethical guidelines for the treatment of research subjects. The unethical treatment of participants stems from the lack of explicit consent and from the use of deception without safeguards.

Uninformed and Deceived Research Participants

Deception through information manipulation takes place when a transmitting party intentionally introduces some form of misrepresentation to influence the behavior of the receiver.²³ Misrepresentation occurs either through concealment (hiding or selectively disclosing only part of the information) or through falsification (communicating false information as if it was true).²⁴ Intentional deception to stage experimental manipulations must be distinguished from the absence of full disclosure of research design information to prospective participants. Withholding the details of a study in the consent form is a common research practice and it is not considered an instance of deception.²⁵ In contrast, research practices that include “withholding of information to obtain participation, concealment and staged manipulations in field settings” are

²² US Department of Health and Human Services, The Belmont Report, <http://www.hhs.gov/ohrp/humansubjects/guidance/belmont.html>

²³ Johnson, P.E.; Grazioli, S.; Jamal, K.; and Berryman, R.G. Detecting deception: Adversarial problem solving in a low base-rate world. *Cognitive Science*, 25, 3 (2001), 355–392.

²⁴ Ekman, P. "Self-Deception and Detection of Misinformation." In Lockard, J.S. and Paulhus, D.L (Eds.) *Self-Deception: An adaptive mechanism*. Prentice Hall, NJ. 1988.

²⁵ Kimmel, A.J.; Smith, N.C.; Klein, J.G. Ethical Decision Making and Research Deception in the Behavioral Sciences: An Application of Social Contract Theory *Ethics and Behavior*, 21, 3, (2011), 222–251

instances of intentional deception.²⁶ Both Facebook and OKCupid's emotion experiments employed intentional deception with the potential to harm emotionally vulnerable users.

The nature of deception in both emotion-related experiments is different. In OKCupid's, deception occurred by giving users intentional and explicit misinformation about their compatibility matches (i.e. by lying about the results of the matching program). People sign up to OKCupid with the expectation that their matches will be accurately reported so that they can derive benefits from the site (dating, friendship, companionship, etc.), and were never told that the compatibility matches could be altered for research purposes. In Facebook's emotional contagion experiment, deception occurred through misrepresentation of the NewsFeed performance by the additional manipulation introduced to suit research purposes. This is an instance of violation of user expectations or assumptions that some consider deception while others do not.²⁷ According to Ralph Hertwig of the Max Planck Institute for Human Development in Berlin, "an unannounced change to the digital code controlling what gets posted on Facebook users' news feeds may be an 'implicit violation' of the site's contract with users who expect something else entirely."²⁸ Users who view the News Feed as an unadulterated collection of random or best updates from their connections, as extracted by the News Feed program, were deceived by this additional emotion manipulation.

Different ethical theories suggest alternative solutions for the dilemma of whether intentional deception in research is morally permissible. From the perspective of deontological theory, which maintains that the ethical decision is the one that reflects the adherence to moral rules or duties, deception in research is not an option. Doing so would conflict with the duty of being honest with participants. In the deontological view, consequences are not the most important consideration. In contrast, from the perspective of utilitarian theories, the morally right action is the one that maximizes the net expected utility for all parties affected by a decision or action. The focus of this theoretical perspective is on the consequences. Therefore, deception in research is acceptable if it produces net positive value when comparing the benefits (i.e. advancement of science) to the costs (i.e. potential harm to participants). The normative prescription of this approach is difficult to implement due to the practical challenges of determining who is affected and who is favored, and the range of potential consequences.²⁹

The use of deception in research is typically justified from a utilitarian standpoint with the notion that society at large might benefit from the outcomes of this type of research. This argument is only applicable if the final intention of researchers is the dissemination of their results for the advancement of science, which is not necessarily the objective that private companies like

²⁶ Kimmel et al. (2011), p. 226.

²⁷ Hertwig, R. and Ortmann, A. "Deception in Experiments: Revisiting the Arguments in Its Defense." *Ethics & Behavior*, 18(1), 2008, 59–92.

²⁸ Quoted in <https://www.sciencenews.org/blog/scicurious/main-result-facebook-emotion-study-less-trust-facebook>

²⁹ Kimmel et al. (2011).

Facebook and OKCupid pursue with their internal research. In most instances, advancing knowledge about online user behavior is a circumstantial by-product of research conducted for the company's own benefit. In essence, those who stand to benefit the most are the companies pursuing the research, while those who bear the potential risks are the uninformed accidental participants. Some argue that deception is a last-resort research tool, only to be used when no other alternatives are available and more than minimal harm to subjects is not likely. Thus, research that is expected to cause psychological discomfort or severe emotional distress cannot be conducted with deception because the inherent risk to cause disproportionate harm to participants would offset any potential benefits. Independent oversight of research protocols is absolutely necessary when there is a power imbalance between those who conduct research for their own benefit, and participants who bear all the risks.

Facebook's social contagion experiments and the OKCupid's power of suggestion study manipulated emotions in different ways. With their respective manipulations came the potential to cause damage to emotionally vulnerable individuals, either by contagion of negative emotions or by despair of finding unsuitable dating matches. The researchers of the Facebook study recognize the link between online experiences and offline emotions and behavior: "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."³⁰ Given the potential carryover effects of online experiences to offline behavior, these two emotional manipulation studies were likely to cause psychological harm.

Complaints and Calls for Action

The attention garnered by the publication of Facebook's emotional contagion experiment prompted a flurry of social media comments³¹ and also more formal letters³² to the Federal Trade Commission urging them to investigate Facebook's research practices. The Electronic Privacy Information Center (EPIC) filed a complaint with the FTC on the basis that "the company purposefully messed with people's minds."³³ Although there is some history of complaints filed against Facebook's privacy practices, this is the first instance in which Facebook's research practices are called into question on non-privacy grounds. In October of 2014, Facebook announced revised guidelines concerning research, including internal review of research by a panel of senior employees in different areas.³⁴ Using members of the company rather than a panel of outside reviewers may not provide independence and objectivity in the evaluation of research projects. Thus, this attempt at self-regulation is insufficient to protect human subject participants.

³⁰ Bond et al. (2012), p. 298.

³¹ See hashtag #facebookemotionstudy results at: <https://twitter.com/hashtag/facebookemotionstudy?src=hash>

³² <http://james.grimmelman.net/files/legal/facebook/FTC.pdf>

³³ <https://epic.org/privacy/ftc/facebook/Facebook-Study-Complaint.pdf>

³⁴ <http://newsroom.fb.com/news/2014/10/research-at-facebook/>

The possibility of unethical behavior in company-sponsored online research is not restricted to Facebook. Unfortunately, it took a voluntary publication in a highly visible academic journal (PNAS), and a deliberate self-disclosure by OKCupid in its company blog to find out the details of these studies. There could be many more similar cases that have not surfaced or have not garnered enough public attention. There is also the risk that after public debate and policy changes companies might stop their collaborations with academics to avoid the public dissemination of results of their internal research. These conditions create momentum to make a specific call to guarantee the ethical treatment of online users who participate in company-sponsored research.

This call for action echoes recent ones issued by the same journal that published the emotion study^{35 36} and other outlets³⁷, and goes a step further by articulating a plan of work in three main areas. First, the FTC should expand the regulatory framework of human subject protection to cover online users who partake in all types of behavioral research, not only those covered directly by federal funding. Second, the regulations should incorporate special safeguards for deception and for the mitigation of negative effects from online to offline behavior. Third, designate an entity – a User Review Board –to regulate online behavioral research particularly studies that belong in the C/D category of experimentation. Each one is elaborated below.

The first recommendation is to expand subject protection to all kinds of research. According to the Common Rule 45 C.F.R. §46.102(d), research is defined as "a systematic investigation, including research development, testing and evaluation, designed to develop or contribute to generalizable knowledge."³⁸ Given the large scale of social network platforms, any systematic investigation with users should be considered research because it has the potential to directly (through publications or press releases) or indirectly (through innovations or best practices) contribute to the advancement of knowledge in user behavior. As such, internal research conducted in social networking platforms in the private sector must be subject to ethical regulations "to ensure that users' rights and interests are adequately safeguarded."³⁹ To this end, instead of "forcing a 20th century regulatory regime onto 21st century technologies,"⁴⁰ it is necessary to develop a practically applicable ethical framework consistent with the new realities of online user experimentation.

³⁵ Fiske, S.T. and Hauser, R.M. "Protecting human research participants in the age of big data." Proceedings of the National Academy of Sciences (PNAS), 111, 38, 2014, 13675–13676.

³⁶ Kahn, J.P.; Vayena, E. and Mastroianni, A.C. "Opinion: Learning as we go: Lessons from the publication of Facebook's social-computing research." Proceedings of the National Academy of Sciences (PNAS), 111, 38, 2014, 13677–13679.

³⁷ Goel, V. As Data Overflows Online, Researchers Grapple With Ethics. New York Times, Aug. 12, 2014. <http://www.nytimes.com/2014/08/13/technology/the-boon-of-online-data-puts-social-science-in-a-quandary.html>

³⁸ <http://www.hhs.gov/ohrp/humansubjects/guidance/45cfr46.html#46.102>

³⁹ Kahn et al. (2014), p. 13677.

⁴⁰ Kahn et al. (2014), p. 13679.

The second recommendation is to recognize and regulate the use of deception in privately-sponsored behavioral research. Although the use of deception in research has been a contentious topic, it is sometimes the only mechanism to create situations of interest that can be systematically studied. Kimmel and colleagues suggest a specific set of principles to follow when research involves deception. These principles highlight the need to forewarn and debrief subjects who consent to participate in research that involves deception.⁴¹ Following ethical practices in research that involves deception is of the utmost importance because participants are exposed to more than the minimal risk they would face in everyday life.⁴² Since participants and researchers are at opposite ends of the deception-use dilemma, and there is a power imbalance between them, an objective third party should analyze the potential benefits and costs of deception.

Third, an independent review board under the purview of the Federal Trade Commission (FTC) or another relevant entity should provide oversight to online research through social networking platforms. This is akin to the idea of Consumer Subject Review Boards (CSRB). Prof. Ryan Calo offers this idea as a thought experiment, not as a universal remedy. He argues that “[t]he accelerating asymmetries between firms and consumers must be domesticated, and the tools we have today feel ill suited. We need to look at alternatives. No stone, particular one as old and solid as research ethics, should go unturned.”⁴³ In this case, my proposed variation in the form of a User Review Board (URB) should be sensitive to the different types of online research (such as A/B testing vs. C/D experimentation). In particular, the board should examine the circumstances that allow the use of deception and the risks of carryover effects between online experiences and offline behavior, and ensure that the appropriate safeguards (forewarning, debriefing, etc.) are implemented along with the rest of the research procedures.

Conclusion

The analysis presented in this essay underscores the need to develop an ethical framework and an oversight body to regulate company-sponsored online research. With respect to the ethical framework, the provisions of the Common Rule must be revised and extended to any research that involves human subjects regardless of the nature of the funding (federal vs. private) or the type of institution that conducts such research. With respect to deception, there is a need to develop explicit rules to guide when and how this approach is acceptable in online research and what are the minimum safeguards for participants. With respect to oversight, an independent regulatory entity (or User Review Board) must be in charge of proactively reviewing sensitive research and balancing private interests with public benefits obtained from user participation. Users who participate in company-sponsored online research deserve an equal level of protection as human subjects who participate in federally-funded or academic research.

⁴¹ Kimmel et al. (2011)

⁴² Fiske and Hauser (2014)

⁴³ <http://www.stanfordlawreview.org/online/privacy-and-big-data/consumer-subject-review-boards>