

---

1-2016

## From Obedience to Contagion: Discourses of Power in Milgram, Zimbardo, and the Facebook Experiment

Timothy Recuber

*Princeton University*, trecuber@smith.edu

Follow this and additional works at: [https://scholarworks.smith.edu/soc\\_facpubs](https://scholarworks.smith.edu/soc_facpubs)



Part of the [Sociology Commons](#)

---

### Recommended Citation

Recuber, Timothy, "From Obedience to Contagion: Discourses of Power in Milgram, Zimbardo, and the Facebook Experiment" (2016). Sociology: Faculty Publications, Smith College, Northampton, MA. [https://scholarworks.smith.edu/soc\\_facpubs/7](https://scholarworks.smith.edu/soc_facpubs/7)

This Article has been accepted for inclusion in Sociology: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact [scholarworks@smith.edu](mailto:scholarworks@smith.edu)

# From obedience to contagion: Discourses of power in Milgram, Zimbardo, and the Facebook experiment

Research Ethics  
2016, Vol. 12(1) 44–54  
© The Author(s) 2015  
Reprints and permissions:  
sagepub.co.uk/journalsPermissions.nav  
DOI: 10.1177/1747016115579533  
rea.sagepub.com  


**Timothy Recuber**

Princeton Writing Program, Princeton University, USA

## Abstract

When the public outcry concerning the ‘Facebook experiment’ began, many commentators drew parallels to controversial social science experiments from a prior era. The infamous Milgram (1963) and Zimbardo (1973) experiments concerning the social psychology of obedience and aggression seemed in some ways obvious analogs to the Facebook experiment, at least inasmuch as all three violated norms about the treatment of human subjects in research. But besides that, what do they really have in common? In fact, a close reading of Milgram, Zimbardo, and the Facebook experiment reveals something about the way power—both as a subject of scholarly inquiry and as an element wielded by researchers—is conceptualized today. Although all three experiments were, in essence, measuring the researchers’ ability to induce an emotional or behavioral change in subjects, the Facebook experiment did much more than the others to hide such considerations and naturalize the exercise of power at work in that study. This paper thus argues that the invisibility of power in the discourse of the Facebook experiment demonstrates, in miniature, the more insidious elements of big data as a whole.

## Keywords

algorithms, big data, emotion, experiment, Facebook, power

---

## Corresponding author:

Timothy Recuber, Princeton Writing Program, Princeton University, E010 Baker Hall, Princeton, NJ 08544, USA.

Email: trecuber@princeton.edu

When controversy broke out surrounding the study of ‘emotional contagion’ undertaken by Adam Kramer, Jamie Guillory, and Jeffrey Hancock (2014)—quickly known by the shorthand ‘the Facebook experiment’—many commentators drew parallels to controversial social science experiments from a prior era. The infamous Milgram and Zimbardo experiments concerning the social psychology of obedience and aggression, respectively, seemed to be clear analogs to Facebook’s own ethically dubious research. For example, one *Popular Mechanics* writer claimed that ‘the parallels between Milgram and the Facebook experiment are obvious although the scale is much larger’ (Marche, 2014: para. 2). A psychology professor writing about the Facebook experiment for *The Huffington Post* used Zimbardo’s work as an example of a similar case where ‘egregious violations have occurred when social scientists and other researchers’ made ethical determinations about their own studies and ‘assumed the study would be minimal risk’ (Klitzman, 2014: para. 10). And a writer in *The Guardian* argued that the Facebook experiment had actually supplanted the Milgram study’s position as ‘the most famous psychological experiment ever’ (Chamorro-Premuzic, 2014: para. 1).

The popular associations between these studies, and their perceived ethical abuses, ought to be examined further. Besides the simple fact that Milgram’s (1963) study of obedience, Zimbardo’s Stanford Prison Experiment (Haney et al., 1973), and Kramer et al.’s (2014) Facebook experiment all generated immense controversy, what do they really have in common? Are ‘Milgram’ and ‘Zimbardo’ merely commonplace phrases for any kind of research with shaky ethics—‘fall guys for justifying the necessity of ethics review’, in the words of Martin Tolich (2014: 86)—or are there more substantive commonalities between these infamous experiments and the Facebook study? Does the fundamental difference in scale between these earlier studies and Facebook’s mammoth sample of users render such comparison irrelevant? Or can we learn something about the way power—both as a subject of scholarly inquiry and as an element wielded by researchers—is conceptualized today through this comparison?

The simplest way to answer these questions is through a close reading of all three studies. Although Milgram’s study of obedience and the Stanford Prison Experiment have been replicated many times with many slight variations in methods and results over the years, the initial articles based on each experiment (Milgram, 1963; Haney et al., 1973) were the ones that ignited public controversy, much like the publication of the Facebook experiment’s results in the *Proceedings of the National Academy of Sciences* ignited this recent furor. Despite these similar circumstances, a comparison of Milgram, Zimbardo, and the Facebook experiment does indeed reveal a change in the conception of power, attributable in part to the emergence of so-called ‘big data’ in social-psychological research. That change cuts to the heart of what makes big data, and the ubiquitous power of tech

companies like Facebook to manipulate it, such a problematic feature of contemporary life.

Through close reading of all three studies, this paper demonstrates that for all their ethical flaws, the Milgram and Zimbardo papers spent more time considering the emotional impact of their studies on their subjects, did more to justify the necessity of their potentially harmful research, and were more transparent about the contingencies created by the artificial manipulation of human subjects than the Facebook experiment paper. Although all three experiments were, in essence, measuring the researchers' ability to induce an emotional or behavioral change in the subjects, the Facebook experiment did much more than the others to hide such considerations and naturalize the exercise of power at work in that study. The invisibility of power in the discourse of the Facebook experiment thus demonstrates, in miniature, the more insidious elements of big data as a whole.

### **Considerations of potential harm to subjects**

In some obvious ways, the Milgram and Zimbardo experiments posed more clear and direct hazards to the tiny sample of volunteers on whom they were conducted than did the Facebook study, which diffused its smaller potential harm to hundreds of thousands of unknowing participants. In the immediate wake of the obedience study, many argued that the nature and extent of that harm went beyond what Milgram had acknowledged (see for example Baumrind, 1964). Similarly, the outcry against the prison experiment was intense enough that even Zimbardo himself later came to admit that it may have been unethical (Zimbardo et al., 2000). The outcomes and implications of these two classic studies continue to be debated today (see Haslam and Reicher, 2012), even as public discussion about the Facebook study and its place in this pantheon of troubling experimental research has just begun. But even with all of this debate, it is hard to know the precise nature of the harm done to the participants in these three experiments. However, one can easily pinpoint the ways that the researchers behind all three studies initially viewed this issue, based simply on what they wrote about it. In that way, comparing the perceptions of the possibility of harm to subjects lets us see how much the researchers considered the ramifications of their own power over their subjects.

As is by now well known, Milgram's initial experiment put 40 subjects in the highly uncomfortable position of having to decide between two 'deeply ingrained behavior dispositions, first, the disposition not to harm other people, and second, the tendency to obey those whom we perceive to be legitimate authorities' (Milgram, 1963: 378). Milgram found that in this situation, authority typically won out. In the respected academic setting of Yale University, in the presence of a seemingly credible authority figure administering the experiment, 26 of the study's

40 subjects proceeded to administer what they believed to be shocks of up to 450 volts to another subject of the experiment—though this person was actually a confederate of the experimenter and there were no real shocks delivered.

Though the majority of his participants acquiesced to authority, Milgram did take time to describe the visible signs of distress that such obedience caused those subjects. He admitted that the subjects ‘were frequently in a highly agitated and even angered state’ (1963: 376) and that the experiment generated ‘extraordinary tension’ (1963: 377) in those who administered the full range of simulated shocks. This tension was such that three subjects even engaged in ‘full-blown, uncontrollable seizures’, one of which was ‘so violently convulsive that it was necessary to call a halt to the experiment’ (1963: 375). Though it may seem laughably inadequate in hindsight, Milgram’s paper did let readers know that at the experiment’s end, ‘procedures were undertaken to assure that the subject would leave the laboratory in a state of well being’, including ‘a friendly reconciliation ... between the subject and the victim’ (1963: 374). Such debriefing is a key tenet of research involving deception of participants (Smith and Richardson, 1983), though Milgram’s particular debriefing of these subjects has been criticized (Perry, 2012).

For Milgram the extreme tension his experiment generated was a surprising side effect, but the Zimbardo team set up a prison-like environment precisely to see if the normal tensions, degradations, and abuses of power in actual prisons could be recreated in an experimental setting. Yet the research team of Haney, Banks, and Zimbardo (1973) were surprised by the quickness and the extent to which their experiment was a success in its ability to create a toxic environment for its subjects. The experiment’s 21 subjects were divided into two groups—10 prisoners and 11 guards—and assigned to play these roles in a makeshift prison constructed in the basement of Stanford University’s psychology building. The results generated ‘great negativity of affect’ (1973: 9) in both prisoners and guards, and the planned two-week experiment had to be terminated after only six days because of ‘extremely pathological reactions’ and in five cases, ‘extreme emotional depression, crying, rage, and acute anxiety’ (1973: 10).

Thus, in the initial publications derived from these two controversial experiments, the authors were fairly frank in detailing even the unintended or unforeseen negative emotional repercussions of the experiment on its subjects. But by contrast, this was a topic that the Facebook team barely explored at all, despite the fact that their study was an explicit attempt to measure the emotional expression of its subjects. The Facebook study focused on its News Feed feature, which is ‘the constantly updating list of stories in the middle of your home page’ including ‘status updates, photos, videos, links, app activity and likes from people, Pages and groups that you follow on Facebook’ (Facebook, n.d.). Facebook co-founder and CEO Mark Zuckerberg called News Feed ‘one of the most important services that we’ve built’, because it ‘takes all the things that your friends are doing and puts them all

in one place' (quoted in Hachman, 2013). The Facebook experiment authors attempted to see whether subtle increases or decreases in the positive or negative content of users' News Feeds would result in the production of similarly positive or negative content by those users. For the researchers, this would constitute 'a form of emotional contagion' (Kramer et al., 2014: 8788), and indeed that is what they claimed to find in the results of this study performed on 689,003 unknowing Facebook users.

Yet there was not a single actual example of what such negative textual content consisted anywhere in the article, and only a passing mention of the way that positive and negative Facebook posts were operationalized. 'Posts were determined to be positive or negative if they contained at least one positive or negative word, as defined by Linguistic Inquiry and Word Count software' (Kramer et al., 2014: 8789) they wrote, but readers never actually got to see that list of positive or negative words. So later in the piece, when the authors revealed that 'for people who had positive content reduced in their News Feed, a larger percentage of words in people's status updates were negative and a smaller percentage were positive' (2014: 8789), the reader had no examples to draw from when trying to determine what this effect actually looked like in a typical subject's News Feed. Did subjects express mild irritation? Were they horribly depressed? The paper didn't show what these results looked like on the ground. In this way, the authors of the experiment rendered its micro-level effects invisible to readers.

Some of this was surely due to the vast scale of the research undertaken by the Facebook team, especially as compared to the small numbers of participants in the Milgram and Zimbardo studies. But nothing would have prevented Kramer, Guillory, and Hancock from providing more information about what the negative responses from users tended to look like, or from classifying the types of negative affect expressed, or from actually debriefing users who had been exposed to these experiments, had the researchers been reflexive at all about the potential for harm to these users, or the potential public backlash when their results were published. In this way, they were not so different from Milgram or Haney, Banks, and Zimbardo, inasmuch as those earlier researchers also failed to fully think through the personal harm to study participants that might have lingered even after they were debriefed.

## **Perceived benefits of the research**

Clearly, though, the Facebook researchers did have the capacity to envision negative affect stemming from this experiment, as they somewhat ironically made clear in their closing paragraph, where they made the case for the importance of their findings. They argued that despite the small effect size they measured, 'given the massive scale of social networks such as Facebook, even small effects can have

large aggregated consequences' (Kramer et al., 2014: 8790). They went on to claim that 'the well-documented connection between emotions and physical well-being suggests the importance of these findings for public health' (Kramer et al., 2014: 8790). This statement served to make the case that the research was worthwhile, with the implication being that Facebook might improve its users' lives with this information, but it also seemed to inadvertently remind readers of the study's large potential harm, given that they had admittedly caused at least a slight negative emotional reaction in over 300,000 subjects. In any case, besides this one hazy connection between emotions and 'physical well-being', there were no more concrete statements about the benefits to society of this experiment, and for the 300,000 negatively affected users, no potential personal benefit at all.

Of course, the benefit to subjects of participating in a study, in the form of increased self-knowledge, for instance, is just one among several factors used to justify social psychological experiments (Forsyth and Pope, 1984). And in fact, this element went largely unaddressed in the write-ups for all three studies. But in contrast to the vague statements about the larger social benefits of the Facebook study, the experiments undertaken by Milgram and Zimbardo were at least conceived with very noble intentions, which the authors made clear in their writing. Although at times Milgram (1963) admitted that 'obedience may be ennobling and educative' (p. 371), it is worth remembering that he was inspired by the Holocaust to investigate the role of obedience and the power of authorities. As he noted, the inhumanity in the gas chambers and death camps 'could only be carried out on a massive scale if a very larger number of people obeyed orders' (Milgram, 1963: 371). Similarly, Zimbardo's team saw their research as part of a humanitarian mission to expose the structural problems inherent in the American prison system. They referred to 'accounts of atrocities committed daily, man against man, in reaction to the penal system or in the name of it' (Haney et al., 1973: 2), and sought to correct the unfair notion that such violence occurs simply because prisoners are bad people. Instead, they aimed to 'separate the effects of the prison environment *per se* from those attributable to *a priori* dispositions of its inhabitants' (1973: 3). Thus they sought to expose the psychology of imprisonment in order to reform that institution, going as far as to conclude that 'the punishment of being imprisoned in a real prison does not "fit the crime" for most prisoners—indeed it far exceeds it' (1973: 17). Debates continue over whether either study actually ended up measuring what it claimed to measure, or adequately supported the hypotheses it put forward (see Gibson, 2013; Perry, 2012; Haslam and Reicher, 2012; Banuazizi and Movahedi, 1975). But the authors did justify their exercise of power over subjects in these lofty terms nonetheless, and research has shown that the public's perception of the social or scientific benefits of a study can powerfully influence its ethical judgments about that study (Schlenker and Forsyth, 1977).

It is worth reiterating here that the benefits of the research in the Facebook experiment were much less clearly or grandly articulated. The authors framed the positives of their research mainly in methodological and scholarly terms: although emotional contagion was well established ‘in laboratory experiments’, only one study (Fowler and Christakis, 2008) had attempted to show emotional contagion via social networks over a long time span, and the authors noted that the results of that study were controversial. So while extending the scholarly conversation is certainly a worthy goal, there were no concrete real world implications for social change mentioned as potential outcomes of the research.

## **The discourse of power**

More troubling than the comparison of these justifications, though, is an analysis of the language that each of these studies used to describe power—both the power to influence others that they were ostensibly studying, and the power of researchers over their subjects. Milgram’s (1963) language focused, as already mentioned, on concepts like ‘obedience’ and ‘authority’. For instance, his discussion of results began by marveling at ‘the sheer strength of obedient tendencies manifested in the situation’ (p. 376). By contrast, he mentioned terms like ‘power’ and ‘manipulation’ only one time each in the paper. But his language was quite reflexive when describing the experimental set-up: he considered the ‘appearance of authenticity’ (p. 373) and the ‘reality of the experimental situation’ (p. 375) for the subjects. He tried to put himself in his subjects’ shoes by listing the features of the experiment that explained such surprising levels of obedience, including the imprimatur of Yale University, and the fact that the subject perceived the experiment’s ‘victim’ to have ‘voluntarily submitted to the authority system of the experimenter’ (p. 377). In Milgram’s research, then, one did see a serious consideration of the researcher’s own power to create and manipulate a situation, and an attempt to think through the mental and affective mechanisms behind the subjects’ behavior, if only as a way to explain unexpected results.

The Zimbardo group engaged in largely the same sort of post-hoc reflexivity. Somewhat unsurprisingly, given its subject matter, their paper mentioned ‘power’ much more frequently, describing the ‘uses of power’ of the guards or the ‘pathology of power’ in general throughout their research. Terms from Milgram appeared as well—‘obedience’, ‘authority’, and especially a concern with the ‘simulation’ of this prison-like environment. Unlike Milgram, who sought to measure mainly behavior and dealt with the intense and tortured emotional reactions his study generated as a surprising byproduct, the Zimbardo team was centrally interested in the ‘affective states of both guards and prisoners’ (Haney et al., 1973: 9) that arise in the prison situation, or ‘total institution’ (p. 1) as they famously conceptualized it. Their paper showed a keen awareness of the strange situation they had thrust upon



their subjects, though not always a sensitivity to it, as when they matter-of-factly described the way ‘the prisoners’ uniforms were designed not only to deindividuate the prisoners but also to be humiliating’ (p. 8). Like Milgram, the team behind the Stanford Prison Experiment clearly believed their findings were generalizable to real-life, un-simulated conditions, but they also puzzled over the extreme affective reactions engendered by their simulations. ‘The profound psychological effects we observed under the relatively minimal prison-like conditions which existed in our mock prison made the results even more significant’ (p. 11), they argued. Although their treatment of the subjects during the experiment was certainly cavalier and unethical, their write-up of the results did explicitly conceptualize the experimental prison as a site where their power as researchers was enacted upon both the guards and the prisoners alike.

The Facebook experiment took a different approach to the issue of power. Despite the fact that, as the Facebook experiment researchers noted, ‘even small effects can have large aggregated consequences’ (Kramer et al., 2014: 8790) on a site like Facebook and in a sample of over 600,000 people, the authors of that study did not describe themselves as wielding power or influence. They did not seem to see their study as an exercise of power over users, in the ways Milgram and Zimbardo did, though undoubtedly it was—albeit in a more diffuse and imperceptible way. Instead, theirs was a language of ‘contagion’ that mimicked the larger discourse of ‘virality’ on which so much of digital culture currently depends (see Sampson, 2012). This language erases the agency or individuality of consumers, and instead endows the viral phenomenon with ‘an autonomy of spreading that takes place in the crowd—a self-spreading tendency’ (Parikka, 2013: 133). The Facebook researchers set themselves up as simply tracking a viral movement of affect through this population—and yet they barely identified themselves as the bio-engineers of this virus, perhaps because of the ethical weight of extending the metaphor in that direction. The study, they contended, ‘manipulated the extent to which people were exposed to emotional expressions’ (Kramer et al., 2014: 8788). ‘Manipulation’ did not refer to subterfuge or false pretenses, as it did in both the Milgram and Zimbardo studies, but to variables—in this case, variable levels of ‘exposure’. ‘Power’ and its related terms went unmentioned in the Facebook experiment paper—only ‘affect’ remained in common with the Zimbardo study.

Most problematically, the language of ‘simulation’, ‘appearance’, and ‘situation’ fell out as well. The Facebook researchers understood themselves to be operating outside the ‘laboratory experiments’ (Kramer et al., 2014: 8788) that they mentioned in the opening of the paper, and as such they didn’t describe themselves as creating an unnatural situation, manipulating an environment, or simulating a state of social interactions and affective conditions. One might perhaps assume this was because Facebook is always algorithmically altering its users’ News Feeds in ways unknown to most users (Holmes, 2014; Hamilton et al., 2014), and thus a

further experimental manipulation was not seen as particularly different or less of a genuine experience. But therein lies the problem.

The lack of informed consent for which so many commentators rightfully criticized the Facebook experiment is in this case a symptom of a larger problem with Facebook, and with the algorithmic manipulation of our online environments in general. Sophisticated sentiment analysis programs now listen in to everything we say online. But this represents ‘a particular kind of listening in which no one individual message is being heard. What takes place instead is an ongoing search for patterns. Thus, the rhetoric of listening readily slips into that of visibility, with its connotations of monitoring and oversight’ (Andrejevic, 2014: 54). In this slippage one sees a clear connection to the more overt monitoring of behavior and manipulation of environments explored in Milgram and Zimbardo, because these algorithms are also used to actively shift users’ behavior. As Rob Horning (2014) has argued:

To the degree that they have access to the devices we use to mediate our relation to everyday life, companies deploy algorithms based on correlations found in large data sets to shape our opportunities—our sense of what feels possible. Undesirable outcomes need not be forbidden and policed if instead they can simply be made improbable. We don’t need to be watched and brainwashed to make them docile; we just need to be situated within social dynamics whose range of outcomes have all been modeled as safe for the status quo. (para. 3)

Only in such a context, when the normal state of things consists of constant experimentation on unknowing populations, and constant real-time tweaking of environments, could researchers fail to see their manipulated News Feed as another kind of laboratory, created like any other lab to test and measure variables in a controlled environment.

Moreover, the larger, unseen algorithmic power of tech companies seeped into the very language with which power was described in the write-up of the Facebook experiment. The discourse of ‘contagion’ naturalizes power, paints it as something that moves organically across an equally susceptible mass of users, people, and subjects. The reality of the Facebook experiment could not be more divergent—researchers placed these hundreds of thousands of users in a convincingly simulated online environment that was markedly different from their normal environment, and this new environment generated new behaviors and different emotions than subjects would otherwise have felt. When one looks at the Facebook experiment in this light, it does have much in common with the Milgram and Zimbardo studies. But the fact that the Facebook experiment authors felt no need to seek the consent of their users beforehand, reveal this experimental simulation to them afterwards, or even reflect on these issues in their write-up of the experiment says volumes about the things giant technology firms feel entitled to do to users without their knowledge today.

## Conclusion

‘I wonder if Facebook KILLED anyone with their emotion manipulation stunt. At their scale and with depressed people out there, it’s possible’ (Weinstein quoted in Goel, 2014: para. 4). This quote from privacy activist Lauren Weinstein came early on in the *New York Times*’ coverage of public outcry over this experiment, and clearly the immediate effects of such emotional tinkering on subjects’ personal mental health are a legitimate concern. But once again, the fact that the Facebook experiment’s authors felt no need to even spill ink over such considerations points to some other issues as well. As Zeynep Tufekci (2014) argued, the negative reaction to the study ‘suggests that algorithmic manipulation generates discomfort exactly because it is opaque, powerful and possibly non-consensual’ (para. 69). This opacity blanketed the language of the paper itself, making the researchers’ own power invisible, and casting the ethics of the research process as seemingly unworthy of consideration. In the age of big data, where we are all routinely and unknowingly the subjects of similar algorithmic social media experiments, this opacity is indeed cause for alarm. Milgram at least revealed his deception to his subjects after the fact. Zimbardo at least debriefed his prisoners and guards at the study’s early conclusion. Today, such experimentation is a part of our everyday lives. But unlike experimental subjects of the past, in most cases we know not what we are being made to obey, nor by what unseen authority we have been deceived, nor to what contagions we have been exposed.

## Funding

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

## References

- Andrejevic M (2014) *Infoglut: How Too Much Information Is Changing the Way We Think and Know*. New York: Routledge.
- Banuazizi A and Movahedi S (1975) Interpersonal dynamics in a simulated prison: A methodological analysis. *American Psychologist* 30(2): 152–160.
- Baumrind D (1964) Some thoughts on ethics of research: After reading Milgram’s ‘Behavioral study of obedience’. *American Psychologist* 19(6): 421–423.
- Chamorro-Premuzic T (2014) We are Facebook’s digital guinea pigs: The web as a real-life experiment. *The Guardian*, 4 July. Available at: <http://www.theguardian.com/media-network/media-network-blog/2014/jul/04/facebook-emotion-social-psychology-experiment>.
- Facebook (n.d.) How News Feeds work. Available at: <https://www.facebook.com/help/327131014036297/>.
- Forsyth DR and Pope WR (1984) Ethical ideology and judgments of social psychological research: Multidimensional analysis. *Journal of Personality and Social Psychology* 46(6): 1365–1375.
- Fowler JH and Christakis NA (2008) Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *British Medical Journal* 337(a2338): 1–9.

- Gibson S (2013) Milgram's obedience experiments: A rhetorical analysis. *British Journal of Social Psychology* 52: 290–309.
- Goel V (2014) Facebook tinkers with users' emotions in news feed experiment, stirring outcry. *New York Times*, 29 June. Available at: [http://www.nytimes.com/2014/06/30/technology/facebook-tinkers-with-users-emotions-in-news-feed-experiment-stirring-outcry.html?\\_r=0](http://www.nytimes.com/2014/06/30/technology/facebook-tinkers-with-users-emotions-in-news-feed-experiment-stirring-outcry.html?_r=0).
- Hachman M (2013) Facebook updates news feed with dedicated feeds, bigger photos. *Readwrite*, 7 March. Available at: <http://readwrite.com/2013/03/07/facebook-updates-news-feed-with-dedicated-feeds-more-photos>.
- Hamilton K, Sandvig C, Karahalios K and Eslami M (2014) A path to understanding the effects of algorithm awareness. *CHI Extended Abstracts on Human Factors in Computing Systems*. Available at: <http://social.cs.uiuc.edu/papers/pdfs/paper188.pdf>.
- Haney C, Banks C and Zimbardo P (1973) A study of prisoners and guards in a simulated prison. *Naval Research Reviews*: 1–16.
- Haslam SA and Reicher SD (2012) Contesting the 'nature' of conformity: What Milgram and Zimbardo's studies really show. *PLoS Biology* 10(11): 1–4.
- Holmes D (2014) These researchers are studying Facebook's news feed to make algorithms more accountable. *Pando Daily*, 23 July. Available at: <http://pando.com/2014/07/23/these-researchers-are-studying-facebooks-news-feed-to-make-algorithms-more-accountable/>.
- Horning R (2014) No life stories. *The New Inquiry*, 10 July. Available at: <http://thenewinquiry.com/essays/no-life-stories/>.
- Klitzman R (2014) Why Facebook should follow ethical standards—like everybody else. *Huffington Post*, 7 July. Available at: [http://www.huffingtonpost.com/robert-klitzman-md/why-facebook-needs-to-fol\\_b\\_5557862.html](http://www.huffingtonpost.com/robert-klitzman-md/why-facebook-needs-to-fol_b_5557862.html).
- Kramer ADI, Guillory JE and Hancock JT (2014) Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 24: 8788–8790.
- Marche S (2014) We can no longer ignore Facebook's vast power. *Popular Mechanics*, 1 July. Available at: <http://www.popularmechanics.com/technology/gadgets/news/we-can-no-longer-ignore-facebooks-vast-power-16944058>.
- Milgram S (1963) Behavioral study of obedience. *Journal of Abnormal and Social Psychology* 67(4): 371–378.
- Parikka J (2013) Review: *Virality: Contagion theory in the age of networks* by Tony D. Sampson. *Theory, Culture, & Society* 30(3): 131–136.
- Perry G (2012) *Behind the Shock Machine: The Untold Story of the Notorious Milgram Psychology Experiments*. Melbourne: Scribe Publications.
- Sampson TD (2012) *Virality: Contagion Theory in the Age of Networks*. Minneapolis: University of Minnesota Press.
- Schlenker BR and Forsyth DR (1977) On the ethics of psychological research. *Journal of Experimental Social Psychology* 13(4): 369–396.
- Smith SS and Richardson D (1983) Amelioration of deception and harm in psychological research: The important role of debriefing. *Journal of Personality and Social Psychology* 44(5): 1075–1108.
- Tolich M (2014) What can Milgram and Zimbardo teach ethics committees and qualitative researchers about minimizing harm? *Research Ethics* 10(2): 86–96.
- Tufekci Z (2014) Big data, surveillance, and computational politics. *First Monday* 19(7). Available at: <http://firstmonday.org/ojs/index.php/fm/article/view/4901/4097>.
- Zimbardo PG, Maslach C and Haney G (2000) Reflections on the Stanford Prison Experiment: Genesis, transformations, consequences. In Blass T (ed.) *Obedience to Authority: Current Perspectives on the Milgram Paradigm*. Mahwah, NJ: Lawrence Erlbaum.