

Reassessing the Facebook Experiment: Critical Thinking About the Validity of Big Data Research

Galen Panger
University of California, Berkeley
galen@ischool.berkeley.edu

October 19, 2015

Abstract:

The Facebook experiment of 2014 manipulated the contents of nearly 700,000 users' News Feeds to induce changes in their emotions. This experiment was widely criticized on ethical grounds regarding informed consent. This controversy, however, diverted attention from a more important concern the experiment was intended to address, which is the impact of Facebook use on well-being. In this paper, I explore the well-being concerns raised by prior research and argue that the experiment does not alleviate them, owing to poor research design. As the question of Facebook's impact on well-being is of great importance, both to Facebook and to society overall, there is a pressing need for more experimental research that is both sensitive to informed consent and carefully designed to yield reliable results. In turn, the lessons of this case have implications for general issues of validity that emerge in Big Data research, now in vogue at major scientific venues.

This is the postprint manuscript in revised form following peer review and acceptance for publication in Information, Communication & Society.

Introduction

In 2014, Facebook addressed the issue of well-being in a study entitled, ‘Experimental evidence of massive-scale emotional contagion through social networks’ (Kramer, Guillory, & Hancock, 2014). ‘The Facebook experiment’, as it is commonly known, sought to rebut concerns raised by academic research suggesting that Facebook has a negative impact on well-being (pp. 8788, 8790). Though this intention was overshadowed by an outpour of concern from scholars regarding the study’s ethics,¹ the study is a prominent entry into the well-being discussion. Coming from Facebook itself through one of the world’s most highly cited scientific journals, *Proceedings of the National Academy of Sciences*,² the study employs a ‘massive’ sample of nearly 700,000 Facebook users, making it perhaps the largest study of Facebook and well-being to date. As Facebook usage continues to grow (Rosenfeld, 2015), the mediating role Facebook plays in our social relationships continues to increase in importance. Has the Facebook experiment alleviated concerns that the company’s influence is not a positive one for society?

The study is also a prominent entry into a growing literature seeking to use large-scale social media data (e.g. status updates) to understand the human experience. Researchers have begun to use social media data to investigate everything from the happiness of nations (e.g. Dodds, Harris, Kloumann, Bliss, & Danforth, 2011; Kramer, 2010; Mislove, Lehmann, Ahn, Onnela, & Rosenquist, 2010) to workplace affect (De Choudhury & Counts, 2013) to human circadian rhythms (Golder & Macy, 2011) to public opinion (Mitchell & Hitlin, 2013), and a sizable sentiment analysis industry has emerged to facilitate research of this kind (Grimes, 2011).

¹ The controversy centred on the study’s lack of informed consent procedures. For reference, see Grimmelmann (2014) or refer to the archived front page of *Techmeme*, a technology news aggregator, from June 29, 2014 (<http://www.techmeme.com/140629/h1410>).

² See ‘About PNAS’, <http://www.pnas.org/site/aboutpnas/index.xhtml>

But critical questions remain about the validity of social media data for these uses and how the social media data generating process affects what these data can teach us about the human experience. The Facebook experiment's use of social media data to assess well-being crystallizes these questions.

Reflecting its growing popularity, scholars are devoting increasing attention to the methodological issues of social media data and other large-scale data sets known as 'Big Data' (e.g. boyd & Crawford, 2012; Crawford, 2013; Lazer, Kennedy, King, & Vespignani, 2014; Ruths & Pfeffer, 2014; Tufekci, 2014). Some issues, as Ruths and Pfeffer write in *Science*, 'are very basic (and long-studied) in the social sciences', and include things such as inappropriate extrapolation from non-random samples (2014, p. 1063). Other issues are more specific to the way Big Data are generated, such as how an Internet service's design affordances can introduce patterns in data that should not be mistaken as naturally arising (p. 1063). Researchers may also neglect shifts in language or user behaviour that can degrade the validity of a measure over time (Lazer et al., 2014). Because issues like these arise in the Facebook experiment, it provides an instructive case study of the general issues of validity that emerge in Big Data research.

This paper offers a detailed critique of the Facebook experiment to help researchers interpret its findings in light of the well-being concerns and to demonstrate how issues of validity in the use of large-scale social media data play out in a prominent case. I argue that researchers should exercise caution in interpreting the study's findings, for four major reasons: first, the study's experimental design may have low internal validity, due to simultaneous manipulation of positive and negative emotions in News Feed; second, the use of Facebook posts to assess subjects' emotional reactions means the study's finding of emotional contagion is indistinguishable from similar-patterned alternative explanations like mimicry and conformity;

third, the sentiment analysis package used to interpret the emotion in Facebook posts has unknown validity for the emotions of interest; and fourth, the social and design contexts influencing how Facebook posts are generated mean posts may have low validity for assessing emotional outcomes, particularly the negative outcomes of concern in prior research.

A note about terminology

In this paper, the term ‘well-being’ refers to the concept of ‘subjective well-being’ or ‘happiness’ popular in the economics and positive psychology literatures, as well as to a broader portfolio of related psychological outcomes such as self-esteem. In the literature, subjective well-being or happiness has two components: a person’s moment-to-moment emotional experiences (positive and negative) and her evaluation of her life as a whole, known as ‘life satisfaction’ (e.g. Diener, Oishi, & Lucas, 2009; Kahneman & Krueger, 2006). People with more positive and fewer negative emotional experiences, and higher life satisfaction, are considered happier. In turn, happiness is thought to promote a number of desirable outcomes, such as cardiovascular health, employment, stable romantic relationships, creativity and cognitive flexibility, and pro-social behaviour like donating money and volunteering (e.g. De Neve, Diener, Tay, & Xuereb, 2013; Diener et al., 2009; Kahneman & Krueger, 2006). Happiness exists in a dynamic relationship with many of these outcomes such that, for example, health, employment and supportive relationships further promote happiness (e.g. Layard, Clark, & Senik, 2012).

Because the Facebook experiment discusses emotional reactions to News Feed content, its results speak directly to emotional experience, the first component of happiness, but perhaps only indirectly to life satisfaction, the second (the two are moderately correlated, see Kahneman & Krueger, 2006). It also speaks only indirectly to the broader portfolio of well-being outcomes

addressed in research on Facebook and well-being, given that these other measures are often related to happiness but by varying degrees (e.g. Cohen, 2008). This lack of direct targeting of a larger range of well-being outcomes is another important limitation of the study; therefore, one recommendation for improving upon the Facebook experiment, noted below, is to include additional measures.

Facebook and well-being

Amid the ethical controversy surrounding the experiment, Facebook twice attempted to draw attention to the study's claims about well-being. In his sole public statement about the study, lead author Adam Kramer, a member of Facebook's data science team, wrote:

The reason we did this research is because we care about the emotional impact of Facebook and the people that use our product. We felt that it was important to investigate the common worry that seeing friends post positive content leads to people feeling negative or left out. ... And we found the exact opposite to what was then the conventional wisdom: Seeing a certain kind of emotion (positive) encourages it rather than suppresses it [*sic*]. (2014)

Mike Schroepfer, Facebook's Chief Technology Officer, later reiterated Kramer's statement when he announced new ethical rules for research at the company (2014).

Deploying the theory of emotional contagion, which implies that like emotions cause like emotions, co-authors Kramer, Guillory and Hancock (henceforth KGH) attempt to rebut concerns that Facebook, particularly positive posts in News Feed, makes people feel negative. If emotional contagion is the only phenomenon at work in News Feed, then positive posts should cause people to feel positive, not negative. Indeed, the results of the study do appear to show that positive posts cause people to feel positive and not negative and, conversely, that negative posts cause people to feel negative and not positive (p. 8789). This result suggests that the emotional

impact of News Feed has much to do with the people we connect to and the emotions they express, and that connecting to happier people makes us happier, in line with reasoning in other studies of emotional contagion (e.g. Fowler & Christakis, 2008). Further, if Facebook's design should somehow encourage people to post more of their positive feelings and fewer of their negative feelings, then this positive skew in Facebook posts should redound to everyone's benefit and News Feed should only improve happiness.

An 'alone together social comparison effect'?

In the Facebook experiment, KGH challenge a contrasting idea – that positive posts can make people feel *unhappy* – which they call the 'alone together social comparison effect' (2014, p. 8788), citing Sherry Turkle's book *Alone together* (2011) and Leon Festinger's foundational work on social comparison (1954). Both the book and the theory of social comparison have played significant roles in the scholarly debate about Facebook and in media coverage (e.g. Copeland, 2011; 'Get a life', 2013; Konnikova, 2013; Marche, 2012; Wortham, 2011). In *Alone together*, Turkle argues that technologies like Facebook enable a risk-averse way of communicating by offering us a chance to edit and curate nearly everything we disclose to others about ourselves and by making everything we disclose part of a permanent record. In this risk-averse mindset, we convey superficial portraits of our lives, expunged of messy but authentic details that complete who we are. At the same time, technologies like Facebook offer a convenience for staying in touch that reduces what we can expect from our friends (e.g. we now expect to learn about friends' major life events on Facebook, rather than from a personal phone call; see Chapters 10, 13–14). Turkle argues that the superficiality and lowered expectations of online friendships have prepared us for what she believes is a coming wave of robotic

companions (Chapters 1–7), with whom there is truly only a veneer of connection (p. 17) and with whom we will ultimately feel ‘utterly alone’ (p. 12).

The dynamics at the heart of *Alone together* and the Facebook experiment are closely related, but because the Facebook experiment investigates the impact of receiving different emotions in News Feed, while *Alone together* focuses on self-disclosure and convenience in relationships, the experiment does not test *Alone together* well. On the other hand, the Facebook experiment is better suited to test the concerns regarding social comparison, which centre on what happens to our well-being when we constantly receive, in News Feed, the selective self-disclosures Turkle details. An influential early study, ‘Misery has more company than people think’, proposed that people tend to underestimate the prevalence of negative emotions and overestimate the prevalence of positive emotions in others’ lives because of the way others selectively hide their negative emotions in social interactions. We naturally compare ourselves to this positively skewed image of our friends, the authors suggest, and feel more lonely and dissatisfied with our own lives as a result (Jordan et al., 2011). Although the study itself concerns social relations as a whole, the authors suggest Facebook exacerbates this positive skew in self-presentation because of ‘the complete control that users have over the public image they project to the world through their photo albums, status updates’, and so on (p. 133), echoing Turkle’s portrayal (2011). In contrast to the reasoning provided by emotional contagion, the reasoning offered by Jordan et al. suggests Facebook’s positive skew should worsen – not improve – well-being.

Social comparison provides a good basis for understanding why the ostensibly happy photos, posts and self-portrayals of our friends on Facebook might make us feel unhappy, particularly because the theory has been tested and extended by a range of research dating back

to Festinger's original work (1954). Social comparison has been demonstrated to affect our well-being in many studies (e.g. Aspinwall & Taylor, 1993; Brown, Novick, Lord, & Richards, 1992; Lin & Kulik, 2002; Lyubomirsky & Ross, 1997; Morse & Gergen, 1970; Salovey & Rodin, 1984; Wheeler & Miyake, 1992; Wood, 1989) and through a range of methods, including brain imaging (Dohmen, Falk, Fliessbach, Sunde, & Weber, 2010) and socioeconomic surveys (Hagerty, 2000). In fact, social comparison is often posited as the solution to what is known as the 'Easterlin Paradox', which finds that while individual happiness increases with income, many societies that get richer do not tend to get happier. This paradox may result because, as Layard, Clark and Senik write in the *UN World happiness report (2012)*, people in a growing economy only get richer in absolute terms; on average, they do not become richer compared to others. Surprisingly but consistently, evidence suggests that the happiness we receive from our own incomes depends a great deal on the incomes of others around us.

Since the 'Misery' study by Jordan et al., further research has demonstrated social comparison on Facebook and its potential to influence well-being (Fox & Moreland, 2015; Haferkamp & Krämer, 2011; Krasnova, Wenninger, Widjaja, & Buxmann, 2013; Lee, 2014; Panger, 2014; Tandoc, Ferrucci, & Duffy, 2015; Vogel, Rose, Roberts, & Eckles, 2014), as well as gathered additional evidence on the overall impact of Facebook use on well-being (Kross et al., 2013; Burke, 2011; Burke & Kraut, 2013; Hinsch & Sheldon, 2013; Przybylski, Murayama, DeHaan, & Gladwell, 2013; Sagioglou & Greitemeyer, 2014; Verduyn et al., 2015), much of it providing reason for concern. Krasnova et al., for example, suggest that the travel and leisure photos our friends post are a particular source of envy for us (2013), while Vogel et al. find that the number of Likes and comments on friends' posts can induce social comparison (2014), made worse perhaps by Facebook's promotion of posts with many Likes and comments in News Feed

(Backstrom, 2013). Other researchers might say Facebook facilitates a kind of social transparency (Stuart, Dabbish, Kiesler, Kinnaird, & Kang, 2012) that enables us to see depressing things like when friends are hanging out without us or when a recent ex has found someone new, while still others suggest we find it difficult to stop using Facebook but also do not find it particularly worthwhile, causing us to feel worse after an extended period of use (Sagioglou & Greitemeyer, 2014).

Facebook as a set of experiences

Perhaps the most nuanced evidence to date about Facebook's impact on well-being comes from Moira Burke, another data scientist at Facebook, who pairs back-end usage data from Facebook with longitudinal surveys of well-being. Burke's research suggests that, rather than look at Facebook as a single experience, we should treat it as a set of experiences each with its own potential impact on well-being. Burke distinguishes between 'directed communication' or conversing one-on-one with others, 'broadcasting' or sharing posts widely with friends, and 'passive consumption' or browsing the News Feed and looking through other people's profiles. For directed communication, Burke finds that receiving composed directed communication (e.g. written messages) is associated with a number of benefits for well-being, but receiving convenient, one-click directed communication (e.g. Likes) has no detectable effect on well-being. She also finds that broadcasting posts to friends has few detectable well-being effects on the broadcaster.³ In contrast, Burke finds that passive consumption is associated with a number of *negative* well-being effects, including lower perceived social support, lower bridging social

³ Burke finds that broadcasting is associated with lower life satisfaction (Burke, 2011) and marginally higher stress (Burke & Kraut, 2013).

capital (feeling part of a broader community), and marginally lower positive affect, higher depression and higher stress⁴ (Burke, 2011; Burke & Kraut, 2013).

Although the negative findings are not the focus of her work, Burke's results confirm that browsing the News Feed may be detrimental for well-being. Her work also appears to confirm Turkle's belief that the most convenient ways to connect such as Liking and broadcasting do not convey benefits to well-being, while more effortful ways of connecting, like taking the time to compose a message to a specific person, do. Further, Burke's studies suggest a high ratio of passive consumption to directed communication and broadcasting. Over a month, Burke's median subject received about 100 total directed messages, comments and Wall posts and broadcasted fewer than 50 status updates, photos and other posts, but loaded News Feed nearly 800 times and viewed 140 profiles (Burke, Kraut, & Marlow, 2011). If passive consumption is the dominant use, and it reduces wellbeing, then Facebook may reduce well-being overall.

Reassessing the Facebook experiment

The Facebook experiment arrives in this context, offering emotional contagion as a 'contrast to theories that suggest viewing positive posts by friends on Facebook may somehow affect us negatively, for example, via social comparison' (KGH, 2014, p. 8790). Instead of finding that positive posts make people feel 'alone', 'left out' or 'depressed' (2014), KGH find that positive posts cause positive emotion. If true, these findings could substantially alleviate concern that Facebook represents a threat to well-being. Significant limitations in the

⁴ To estimate effects, Burke uses a linear multilevel model with a lagged dependent variable. Though observational, this model produces estimates that are less biased than cross-sectional models, 'in effect controlling for an individual's previous level of the outcome variable (e.g. stress) and all of the unmeasurable factors that contribute to it' (Burke & Kraut, 2013, p. 1424).

experiment's research design prevent the study from alleviating concern, however. Below, I raise and discuss four major limitations of the study and argue for caution in interpreting its findings.

Experimental design may have low internal validity

To test the hypothesis that emotional contagion rather than social comparison is at work in News Feed, lead author Kramer selected nearly 700,000 English-speaking users at random and assigned each of them to one of four groups (KGH, 2014). In the first group, he removed between 10% and 90% of the *positive* posts they would have seen in their News Feeds over one week, and in the second group he removed between 10% and 90% of *negative* posts they would have seen. The third and fourth groups were control groups where he removed equivalent numbers of posts at random.⁵ Kramer then took the *subsequent posts* produced by each group during that week-long period and analysed how positive or negative they became in their own expressions as a result.

In line with emotional contagion operating in reverse, the authors expected to see subjects who had positive posts removed become less positive and more negative in their own posts, while those who had negative posts removed were expected to become less negative and more positive, compared to control groups. Indeed, this is the precise pattern the authors find, though the effects are quite small. The estimated percentage changes in subjects' subsequent emotions are all 0.1% or less (not 10% or 1%, but 0.1%), with Cohen's *d* ranging from 0.001 to 0.02 (KGH, 2014).

⁵ KGH provide no further information about the average or distribution of the percentage of posts removed from News Feed. They maintained two control groups because positive and negative posts occur at different rates in News Feed and thus removing a given percentage of posts involves removing a different number in the positive experimental and control groups than the negative experimental and control groups.

A significant problem with this experimental design relates to its internal validity. Although KGH seem to believe that removing one kind of emotion from News Feed (e.g. positive posts) has no effect on the proportion of the other emotion in News Feed (e.g. negative posts), in fact removing one may shift the proportion of the other. This means the experiment's results unfortunately may reflect the causal effects of both emotions acting in News Feed at once, which presents a significant problem for the study's internal validity. Note that one alternative to KGH's claim that emotional contagion rather than social comparison may occur for people when they browse News Feed is that *both* emotional contagion and social comparison may occur, along with other possible emotional dynamics. If both emotional contagion and social comparison are at work, then positive posts in News Feed may at times cause positive emotions, and at other times cause negative emotions (e.g. sadness, envy). Similarly, it is possible that negative posts, too, cause a mix of positive emotions (e.g. compassion, *schadenfreude*) and negative emotions, depending on the context.

Because of the possibility that positive and negative posts may both cause a mix of positive and negative emotions, it was crucial that KGH shift the proportion of only one emotion in News Feed at a time while holding the other constant so that the independent effects of each emotion could be measured. Unfortunately, it is quite possible that they did not. To illustrate this for the removal of positive posts, note that KGH classify roughly 46% of posts in News Feed as positive and 22% as negative (p. 8789), leaving roughly 32% of posts classified as neutral or other. In a News Feed with 100 posts, for example, removing half of the positive posts leaves us with 77 total posts, now 23 of which are positive, 22 negative, and 32 neutral or other. Removing half of the positive posts does reduce their proportion in News Feed, from 46% to 30%, but unfortunately also increases the proportion of negative posts, from 22% to 29%. Removing about

six additional negative posts or replacing positive posts with neutral posts would hold the proportion of negative posts constant.⁶

If KGH had avoided shifting both emotions in News Feed at once, we might have observed subjects in the positive post removal group become both less positive *and* less negative owing to reduced positive emotional contagion and reduced social comparison. Instead, because negative posts may have been allowed to increase in proportion, we see subjects become less positive but *more* negative. Is this increase rather than decrease in subjects' negative posts a sign that social comparison does not happen? Unfortunately, although answering the question of

⁶ In this scenario, negative posts increase in proportion when any percentage (greater than zero) of positive posts is removed. This is demonstrated mathematically by solving the following equation for x , which represents the percentage of positive posts removed:

$$\frac{0.224}{(1-x)(0.468) + 0.224 + 0.308} - 0.224 > 0$$

The values 0.468, 0.224 and 0.308 stand for the exact percentages of positive, negative and neutral posts, respectively, rather than the rough percentages used in the example above. A similar exercise demonstrates that positive posts increase in proportion when any percentage greater than zero of negative posts is removed. Note that this scenario does not account for the possibility that some posts may be classified as containing both positive and negative emotion, and thus may be removed in either experimental condition. This possibility is not discussed in the Facebook experiment but was acknowledged in an email from Kramer, who did not know the percentage of posts that received this dual classification (personal communication, 23 June 2015). Because posts may be both positive and negative, removing one emotion from News Feed may actually *decrease* the proportion of the other, rather than increase it, if a large enough percentage of posts receive dual classification. However, because removing one emotion from News Feed resulted in subjects producing *more* of the other, I proceed under the assumption that the other likely also increased in proportion in News Feed. It is possible that, by chance, the percentage of posts receiving dual classification is exactly the percentage needed to hold one emotion constant in News Feed when the other is reduced, but this does not seem likely.

whether positive posts might ‘somehow affect us negatively’ was a key motivation (KGH, 2014, p. 8790; Kramer, 2014; Schroepfer, 2014), the experiment as designed is unable to answer this question. It is possible, for example, that any reduction in negative posts due to reduced social comparison as well as any increase in negative posts due to increased negative emotional contagion are both much larger than the <0.1% overall difference reported, but that they cancel out to this small result. Note that this cancelling out may occur for subjects’ positive posts in this same experimental group as well. When positive posts are reduced in News Feed while negative posts are allowed to increase, we see subjects’ subsequent positive posts decrease overall by 0.1%. However, we are unable to see if subjects’ positive posts actually increase – due to more negative posts in News Feed causing more positive posts – only to be cancelled out by a reduction in subjects’ positive posts due to less positive emotional contagion. Because of this internal validity problem alone, it is difficult to determine what the Facebook experiment contributes to our knowledge of the relationship between Facebook and well-being.

Contagion indistinguishable from other sociobehavioural phenomena

KGH appear to be unaware of this potential internal validity problem when they call attention to a ‘cross-emotional encouragement effect’ (2014, p. 8790) in the process of addressing a second problem of their research design, which is the difficulty of distinguishing emotional contagion from similar-patterned sociobehavioural phenomena like mimicry and conformity in their data. Indeed, many studies using social media data may find it difficult to distinguish what people genuinely feel from what they say or do in the social context of social media. For example, we often mimic others in social situations to express our affinity for them (e.g. van Baaren, Janssen, Chartrand, & Dijksterhuis, 2009), and we often feel pressured to

conform (Asch, 1952; Berns et al., 2005; Zimmerman, 2014). Because both mimicry and conformity would produce the same ‘like causes like’ pattern as emotional contagion when KGH alter emotions in News Feed, they offer observationally equivalent explanations.

This inability to distinguish emotional contagion from other sociobehavioural phenomena is an important limitation because it implies the happiness expressed by others in News Feed may not actually transfer to us. Seemingly unaware that they may alter the proportions of both emotions at once in each experimental group, KGH argue that mimicry cannot account for the full pattern of results. Specifically, they argue that mimicry cannot account for the ‘cross-emotional encouragement effect’, wherein the removal of one emotion from News Feed results in subjects producing *more* posts of the opposite emotion (2014, p. 8790). The ‘cross-emotional encouragement effect’ for the removal of positive posts from News Feed, for example, is subjects’ subsequent production of more negative posts, which KGH argue cannot be the result of mimicry presumably because people cannot mimic what they do not see. However, I have shown that when KGH remove positive posts from News Feed, they may also allow negative posts to increase in proportion, which means the increase in subjects’ production of negative posts *can* be explained by mimicry, as well as conformity. This is an important limitation.⁷

⁷ Though KGH argue that mimicry cannot account for the ‘cross-emotional encouragement effect’, they do not explain what can or why the effect is happening absent an internal validity problem. Elsewhere, the authors argue for the independence of positive and negative emotion, stating that ‘positivity and negativity were evaluated separately given evidence that they are not simply opposite ends of the same spectrum. Indeed, negative and positive word use scarcely correlated’ (2014, p. 8789). If positive and negative emotion operate independently, then there is no *a priori* reason to believe less of one in News Feed should cause more of the other in subjects’ posts. Again, a plausible explanation for this is the internal validity problem. However, even if positive and negative emotion are

To overcome the above two limitations of the Facebook experiment in future related work, researchers should ensure they hold the proportion of one emotion constant while adjusting the other, and should employ other methods that better distinguish how subjects feel from what they say or do in social situations. I discuss one such method, experience sampling, below.

Sentiment analysis package has unknown validity for emotions of interest

In the Facebook experiment, KGH (2014) use Facebook posts as both cause and consequence; they first manipulate the proportions of positive and negative posts in News Feed to cause changes in subjects' emotions, and then assess those emotions using subjects' own subsequent Facebook posts. However, Facebook posts are composed of text, photos and other elements and at the time of the study do not include a mood assessment.⁸ KGH therefore use a sentiment analysis package known as Linguistic Inquiry and Word Count (LIWC) to assign Facebook posts an emotional value. LIWC, last updated in 2007, takes in text and outputs measures of four emotions (positive emotion, anxiety, anger and sadness), along with other linguistic measures (<http://www.liwc.net>). LIWC is a word count technique, which means it outputs the percentage of words in a given text that have been designated as markers of each emotion or word category. Using LIWC, KGH classified posts in News Feed as positive or negative when 'they contained at least one positive or negative word' (2014, p. 8789). To measure how positive or negative subjects subsequently became in their own posts, the authors

dependent, and less of one generally implies more of the other, then there is also no reason, and none is offered, why mimicry and conformity would not also show this dependent relationship.

⁸ Facebook posts do now have a simple optional mood indicator, however (see Constine, 2013).

took the percentage of each subject's words that were positive and negative and averaged those percentages across all subjects in each experimental group.

Though LIWC is widely used in academic research for sentiment analysis, it is surprisingly difficult to find evidence regarding its validity as a measure of emotion in Facebook posts or other social media. Traditionally, academics provide evidence of validity for the measures they use, meaning, among other things, they show those measures have a connection to some 'ground truth'. The Satisfaction With Life Scale, for example, is a self-reported measure of happiness that correlates well with a number of 'ground truth' indicators such as how often you smile and the level of stress hormones in your blood (Diener, Emmons, Larsen, & Griffin, 1985; Kahneman & Krueger, 2006). Computer scientists using supervised machine learning, too, normally provide evidence of validity (Davis & Goadrich, 2006).

When KGH claim that LIWC 'correlates with self-reported and physiological measures of well-being' (p. 8789), they cite three studies. Unfortunately, in the first, Guillory, Hancock and co-authors find that LIWC actually cannot tell the difference between groups sharing negative vs. neutral emotions in text chat ('there were not differences in use of negative emotion words between groups', p. 747 in Guillory et al., 2011). In the second, Golder and Macy (2011) explore what LIWC analysis of tweets might tell us about human circadian rhythms, but they provide no validation evidence and, in a supplemental, go into detail about their study's limitations. In the third, Kramer (2012) explores emotional contagion in social media without validating LIWC but mentions an earlier study (Kramer, 2010) in which he finds LIWC has a low but significant correlation with Satisfaction With Life ($r = .17$).⁹ While a significant result

⁹ Kramer (2010) correlates scores on the Satisfaction With Life scale with LIWC analysis of an average of 244 status updates per subject ($n = 1341$ subjects).

there suggests KGH are not completely off-base for using LIWC, a low correlation with Satisfaction With Life does not constitute validation for emotional experiences, which are the subject of the Facebook experiment and are conceptually distinct (e.g. Kahneman & Krueger, 2006). In addition, while Kramer (2010) combines and normalizes all of LIWC's emotion measures to derive the correlation with Satisfaction With Life, KGH (2014) compare LIWC's outputs for positive and negative emotion separately, without validating them as separate measures.

There are a number of reasons to be sceptical that LIWC is a good fit for social media or, especially, this particular study. A first concern is that while LIWC has a measure for sadness, which should provide coverage for the lonely and depressed feelings KGH (2014) and Jordan et al. (2011) believe are a result of social comparison, it does not have a measure for envy or jealousy, thought to be the other major result (e.g. Krasnova et al., 2013; Tandoc et al., 2015). Without further evidence, it is difficult to assume LIWC can detect envy or, thus, that the Facebook experiment can speak to envy. LIWC was also developed and validated for the analysis of long form writing rather than short Facebook posts, which may impact its validity, particularly when the authors note that a post may receive a classification on the basis of a single emotion word (p. 8789). Third, the 915 emotion words in the LIWC dictionary (<http://www.liwc.net/descriptiontable1.php>) were selected not from language people are likely to use on Facebook but rather from reference materials like psychological assessments (Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007). Finally, LIWC requires discarding everything about a post except its plain text, meaning it loses potentially valuable emotional signals, like emoticons or smiling in photographs (Abdullah, Murnane, Costa, & Choudhury, 2015; Paltoglou & Thelwall, 2012).

In a broader review of LIWC in studies of social media and other computer-mediated communication, I found that even when high-level evidence of validity is presented, the results are mixed. LIWC appears to be consistently too positive in its ratings, for example, labelling the conversation in social media around the H1N1 disease outbreak as positive overall (Gonçalves & Araújo, 2013). Similarly, in a study of instant messaging, LIWC judges that participants expressing sadness are still positive overall (Hancock, Landrigan, & Silver, 2007). LIWC may perform well in combination with other techniques (Kivran-Swaine & Naaman, 2011; Paltoglou & Thelwall, 2012) but, alone, appears far inferior to a number of machine learning techniques (Tepper, Benerjee., Banerjee, Tepper, & Sapiro, 2012). Further, there is scant evidence comparing LIWC's validity for positive vs. negative emotion, or showing its performance for specific emotions *underneath* the umbrellas of positive and negative emotion, such as sadness or envy. Uneven validity makes it difficult to compare emotional outcomes to one another as KGH do.

The point of this discussion is not to imply that studies using LIWC have no merit, but that without evidence showing to what extent LIWC captures the emotions of interest in the linguistic contexts of interest, it is impossible to evaluate findings, particularly in comparison to those of other psychological studies. This evidence is especially needed in the Facebook experiment, where LIWC's low validity would be multiplicative given that KGH use LIWC both to classify posts for removal from News Feed and to measure the emotion in subjects' subsequent expressions. Low validity or low correlations between LIWC's outputs and the emotions in posts, of course, would likely bias effect size estimates towards zero, which may be another major reason KGH find such small effects. But while KGH spend several sentences attempting to contextualize the study's very small effect sizes (e.g. 'given the massive scale of

social networks such as Facebook, even small effects can have large aggregated consequences' p. 8790), unfortunately, they spend no time discussing the role LIWC and other aspects of their overall research design may play in generating those estimates. In this light, the study's massive sample size looks less like a virtue and more like a way to compensate for low validity.

For the first step of classifying posts for removal from News Feed, barring further evidence demonstrating LIWC's validity, an improvement upon LIWC would be to employ supervised machine learning with Facebook posts. Because Facebook now allows users to label their own posts with a particular mood, as noted above, Facebook's researchers may already have a simple means of obtaining labelled data for training a new classifier. However, to minimize bias, researchers should make the mood indicator a required part of posting for a random sample of users, and private, as part of a study involving informed consent. In training a new classifier, researchers can also incorporate additional emotional signals or features of Facebook posts, like emoticons or facial expressions in photographs (e.g. Abdullah et al., 2015; Paltoglou & Thelwall, 2012). For the second step of measuring subjects' subsequent emotional responses to the contents of their News Feeds, I recommend replacing LIWC, sentiment analysis and Facebook posts altogether in favour of experience sampling, which I describe in the next section.

Facebook posts have unknown validity for assessing the emotional impact of News Feed

For researchers hoping to use Big Data to understand the human experience, the previous discussion highlights an important limitation: when the research requires a significant data transformation step such as sentiment analysis, that step may bias estimates, often towards zero, and perhaps for some outcomes of interest more than others. Even if some amount of bias is

unavoidable in the data transformation step, Big Data researchers should help readers anticipate the impact it may have on findings.

A similar level of scrutiny and discussion is needed for the underlying data itself. Indeed, a key question for the Facebook experiment is whether Facebook posts are a valid and appropriate measure of the emotional impact of News Feed. Very simple situational factors could influence validity, such as whether people generally spend time on News Feed before posting, or whether they often post first. If people tend to post first and *then* look at News Feed, then posts may not have high validity as a measure of News Feed's emotional impact.

A look at the design and social psychological contexts in which the data are generated is also needed. Simple design factors like Facebook's affordance for easily resharing others' posts,¹⁰ for example, could mean posts better represent emotionally congruent reactions than emotionally incongruent reactions. The resulting pattern in the data could then make emotional contagion (or, again, mimicry or conformity) appear to be the dominant phenomenon, whether or not this is true. Simple social psychological biases could also make some patterns in the data more likely than others. For example, researchers have found that people are more likely to share information with others on the Internet when they are emotionally aroused (Berger, 2011; Berger & Milkman, 2012). Because of this arousal bias, Facebook posts may have greater validity for high-arousal emotional consequences like excitement and outrage than for low-arousal outcomes like peacefulness or sadness. Differences in validity could therefore bias some estimates like sadness closer to zero than others like anger which, again, makes it difficult to compare emotional outcomes. It means, for example, that a reduction in sadness following the removal of positive posts in News Feed could be easily washed out by an increase in high-arousal anger or

¹⁰ Such a resharing button was likely present when the study was conducted (e.g. see video in Tow, 2011).

anxiety due to the simultaneous increase in negative posts in News Feed, which is discussed as a possibility above.

Perhaps the most important question for KGH is whether it is appropriate to employ Facebook posts to rebut emotional outcomes of Facebook use that people may be reluctant to share on Facebook. In their study of envy on Facebook, for example, Krasnova et al. (2013) find people are loath to admit to feeling envious, which the authors suggest is due to social stigma against envy. Pairing this reluctance to share certain feelings like envy with the ‘complete control’ Jordan et al. (2011) believe Facebook offers over self-disclosure suggests it may not be valid to use Facebook posts to determine the emotional consequences of News Feed.

As the works of Jordan et al. (2011) and Turkle (2011) highlight, issues of self-disclosure and self-presentation are a pillar of social media research precisely because of the unprecedented means social networks offer us to control what we say and how we present ourselves to others. On Facebook, we can edit what we say before we say it, upload just the photos that are flattering and untag those that are unflattering. Social media researchers often cite the work of sociologist Erving Goffman, who likens social interactions to theatrical performances where, in the presence of an audience, we take great pains to leave the right impression and to play the character we desire others to believe we are (1956). Within this theoretical perspective, it is a given that people may withhold feelings like envy or sadness that undermine the positive impression they wish to foster, and with Facebook they are more empowered to do so. At the same time that social networks like Facebook offer us greater means to control how we present ourselves to others, they may also make it more *necessary* to exercise this control because of the sheer number of people we must try to speak to at once – friends, parents, co-workers, classmates, acquaintances and so on – when we post. Among Goffman’s insights is that we present ourselves

differently to the different audiences in our lives and that embarrassments occur when we play the wrong character with the wrong audience. Thus, ‘seamen, whose home away from home is rigorously he-man’, Goffman explains colorfully, ‘tell stories of coming back home and inadvertently asking mother to “pass the f-cking butter”’ (1956, p. 8). However, whereas in typical offline situations we have a defined audience and can tailor our self-presentation accordingly, in social media the norm is to broadcast our posts to all audiences at once, leading to what social media researchers refer to as ‘context collapse’ (e.g. boyd, 2014, pp. 31–33, 47–53). Thus, to manage our self-presentation and avoid embarrassments under conditions of context collapse, ‘types of conversation, which cannot be shared among’ audiences, Goffman writes, ‘will be laid aside’ (1956, p. 85; see also boyd, 2014, p. 33). In other words, the response to context collapse is often to withhold.

Given these dynamics, then, it is perhaps unsurprising that self-censorship is a widespread phenomenon on Facebook. In a study involving the observation of a random sample of millions of Facebook users, Das and Kramer find that about 33% of posts are censored *at the last minute* (2013). That is, people begin to type a post (at least five characters) but then decide to delete it. If we censor fully a third of what we want to express at the last minute, it is not hard to imagine we censor quite a bit more before we reach for the keyboard. Confirming the potency of context collapse, the co-authors also find that Facebook users with a greater number of distinct friend communities or audiences censor themselves at the last minute more often (p. 125).

The prevalence of self-censorship on Facebook should raise significant doubts about the validity of Facebook posts, particularly given the likelihood that self-censorship is not randomly distributed among emotions or emotional responses to News Feed content. Unfortunately, however, KGH do not discuss self-censorship or any other social dynamic save mimicry which,

as I show above, they fail to dismiss. Together, the many unacknowledged social, psychological and design factors that affect why we post and why we do not are a fourth major limitation of the Facebook experiment.

A practical way to improve upon Facebook posts and LIWC as a measure of the emotional consequences of News Feed would be to use a different method known as experience sampling (Csikszentmihalyi & Larson, 1987; Hektner, Schmidt, & Csikszentmihalyi, 2007), considered by some to be the ‘gold standard’ for assessing emotional experiences, which play a key role in happiness (e.g. Kahneman, Krueger, Schkade, Schwarz, & Stone, 2004). Experience sampling involves interrupting people at random times as they go about their lives to ask how they are feeling or what they are doing in the moment, and could be used by future researchers to gauge what users are feeling as they browse News Feed. A pop-up with a mood assessment including specific emotions of interest should suffice. Experience sampling, as part of a study involving informed consent, is more likely to encourage open and honest self-disclosure because it is private and confidential, insulated from the demands of the social context. Because it is private, experience sampling is also less subject to sociobehavioural phenomena like mimicry and conformity. Further, because it is solicited randomly in time, experience sampling is less likely to involve momentary selection biases like arousal bias. Of course, there are trade-offs involved in any study in which subjects participate voluntarily, including the fact that people who participate may be different (e.g. have more agreeable personalities) than those who do not and perhaps in a way that is relevant to outcomes. The interruptions specifically involved in experience sampling can also be intrusive, leading to non-response and attrition issues, though with a large sample size relatively few interruptions would be necessary.¹¹ However, the trade-

¹¹ For more on these trade-offs and how they are typically addressed, see Hektner et al. (2007, especially pages 50–58).

offs involved are likely worth the reduction in relevant biases achieved through the method's confidentiality, random solicitation and, arguably, use of self-report measures, which remain the best way to access the 'subjective' quality of subjective well-being (Hektner et al., 2007, p. 10; Scollon, Prieto, & Diener, 2009, p. 171). Given concern about special reluctance to report envy even in private (Krasnova et al., 2013), researchers may consider taking additional steps to encourage openness, such as using a list experiment (Antin & Shaw, 2012).

Data from experience sampling, along with a better classification and removal procedure for posts in News Feed, should reduce bias in estimated effect sizes, distinguish emotional outcomes from sociobehavioural phenomena like mimicry and conformity, and help researchers get closer to understanding the real momentary emotional consequences of News Feed. Because well-being is more than momentary emotions, however, researchers should also employ some of the validated well-being measures used in Burke's work on Facebook and well-being, such as satisfaction with life, social support and loneliness (2011), which assess some of the longer-term conclusions people draw about their lives.

Discussion and conclusion

Even if researchers make all of these improvements, the scholarly discussion about Facebook, News Feed and well-being will continue. The emotional content of posts may not be the chief vector for social comparison in News Feed and social comparison may not be the chief reason why features like News Feed may be detrimental to well-being. More experiments are needed that are both sensitive to informed consent and carefully designed to yield reliable results, for example, with a subject pool that has consented to have their usage of Facebook features limited and aspects of their Facebook experience modified, say for a month, so that

researchers can assess the well-being impact of those interventions via experience sampling¹² and validated scales. What impact does taking a month-long break from News Feed have on well-being? What happens when we remove Like and comment counts (e.g. Dewey, 2014)? What does a News Feed optimized for well-being look like? Of course, Facebook researchers and their academic collaborators are best positioned to improve upon the Facebook experiment and to answer many questions like these, given their privileged access to the data and inner workings of Facebook. But will those with critical perspectives gain access? If not, how does that affect the research landscape? Indeed, some believe the rise of Big Data is creating a ‘divided’ research community, between those with inside access to platforms and those without (Ruths & Pfeffer, 2014, p. 1063).

In the end, the Facebook experiment reminds us that we have only begun to grapple with the many complexities of research involving Big Data – not just ethical complexities, but also methodological complexities. Big numbers are impressive. The scale of data coming from services like Facebook is exciting, and data transformations involved in Big Data research, like sentiment analysis, seem magical. As many practitioners understand well, and as others have noted (e.g. Crawford, 2013), however, they are not magic, and data do not automatically become better suited to answer research questions just because they are abundant.

In some ways, the era of Big Data has turned classic biases of judgement upside down. While Nobel laureate Daniel Kahneman and his colleague Amos Tversky once worried about how ‘sample size neglect’ was causing researchers to over-extrapolate findings from small samples (Tversky & Kahneman, 1974), today we might worry about giving sample size too

¹² Mobile push notifications would actually enable researchers to solicit mood reports throughout the day, on or off Facebook, providing a more complete view over time of the emotional impact of experimental interventions.

much weight. ‘Domain knowledge neglect’ or ‘research design neglect’ may be the greater threats in the Big Data era. Pre-specified sentiment analysis techniques like LIWC may be convenient but may not adapt well to social media, especially as language and expression evolve with those media. Facebook posts, too, while abundant, may have significant biases owing to the fact that they occur in design and social contexts which place limits on what they can tell us about the true emotional consequences of News Feed. While the well-being discussion continues, let us take advantage of exciting new sources of data without engaging in magical thinking.

Acknowledgments

I thank my advisor, Steve Weber, colleagues in my department’s Doctoral Research and Theory Workshop, consultants in Berkeley’s D-Lab and my reviewers for their helpful feedback on this paper.

Notes on contributor

Galen Panger is a doctoral candidate in the School of Information at the University of California, Berkeley, and is co-director of the Berkeley Center for Technology, Society & Policy. His research interests include the use of social media data to understand well-being and public opinion, social media psychology and digital citizenship.

References

Abdullah, S., Murnane, E. L., Costa, J. M., & Choudhury, T. (2015, March 14–18). *Collective smile: Measuring societal happiness from geolocated images*. Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, Vancouver, BC, Canada, pp. 361–374.

- Antin, J., & Shaw, A. (2012, May 5–10). *Social desirability bias and self-reports of motivation: A study of Amazon Mechanical Turk in the US and India*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Austin, TX, USA, pp. 2925–2934.
- Asch, S. (1952). *Social psychology*. New York, NY: Prentice-Hall.
- Aspinwall, L. G., & Taylor, S. E. (1993). Effects of social comparison direction, threat, and self-esteem on affect, self-evaluation, and expected success. *Journal of Personality and Social Psychology*, *64*(5), 708–722.
- Backstrom, L. (2013). News Feed FYI: A window into News Feed. *Facebook for Business*. Retrieved from <https://www.facebook.com/business/news/News-Feed-FYI-A-Window-Into-News-Feed>
- Berger, J. (2011). Arousal increases social transmission of information. *Psychological Science*, *22*(7), 891–893.
- Berger, J., & Milkman, K. (2012). What makes online content viral? *Journal of Marketing Research*, *49*(2), 192–205.
- Berns, G., Chappelow, J., Zink, C., Pagnoni, G., Martin-Skurski, M., & Richards, J. (2005). Neurobiological correlates of social conformity and independence during mental rotation. *Biological Psychiatry*, *58*(3), 245–253.
- boyd, d. (2014). *It's complicated: The social lives of networked teens*. New Haven, CT: Yale University Press.
- boyd, d., & Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, *15*(5), 662–679.

- Brown, J. D., Novick, N. J., Lord, K. A., & Richards, J. M. (1992). When Gulliver travels: Social context, psychological closeness, and self-appraisals. *Journal of Personality and Social Psychology*, 62(5), 717–727.
- Burke, M. (2011). *Reading, writing, relationships: The impact of social network sites on relationships and well-being* (Unpublished doctoral dissertation). Carnegie Mellon University. Retrieved from <http://repository.cmu.edu/dissertations/185/>
- Burke, M., & Kraut, R. (2013, February 23–27). *Using Facebook after losing a job: Differential benefits of strong and weak ties*. Proceedings of the 2013 Conference on Computer Supported Cooperative Work, San Antonio, TX, USA, pp. 1419–1430.
- Burke, M., Kraut, R., & Marlow, C. (2011, May 7–12). *Social capital on Facebook: Differentiating uses and users*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada, pp. 571–580.
- Cohen, S. (2008). *Basic psychometrics for ISEL-12*. Retrieved from <http://www.psy.cmu.edu/~scohen/scales.html>
- Constine, J. (2013). Facebook tries letting you share emoticons of exactly what you're feeling, reading or eating. *TechCrunch*. Retrieved from <http://techcrunch.com/2013/01/30/facebookvisual-sharing/>
- Copeland, L. (2011). The anti-social network. *Slate*. Retrieved from http://www.slate.com/articles/double_x/doublex/2011/01/the_antisocial_network.html
- Crawford, K. (2013). Think again: Big data. *Foreign Policy*. Retrieved from <http://foreignpolicy.com/2013/05/10/think-again-big-data/>
- Csikszentmihalyi, M., & Larson, R. (1987). Validity and reliability of the experience-sampling method. *The Journal of Nervous and Mental Disease*, 175(9), 526–536.

- Das, S., & Kramer, A. (2013, July 8–11). *Self-censorship on Facebook*. Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media, Cambridge, MA, USA, pp. 120–127.
- Davis, J., & Goadrich, M. (2006, June 25–29). *The relationship between precision-recall and ROC curves*. Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA, USA, pp. 233–240.
- De Choudhury, M., & Counts, S. (2013, February 23–27). *Understanding affect in the workplace via social media*. Proceedings of the 2013 ACM Conference on Computer Supported Cooperative Work, San Antonio, TX, USA, pp. 303–316.
- De Neve, J. E., Diener, E., Tay, L., & Xuereb, C. (2013). The objective benefits of subjective well-being. In J. Helliwell, R. Layard, & J. Sachs (Eds.), *World happiness report* (pp. 54–79). New York, NY: Sustainable Development Solutions Network.
- Dewey, C. (2014). The (one) simple thing fueling your social media addiction. *Washington Post*. Retrieved from <http://www.washingtonpost.com/news/the-intersect/wp/2014/11/12/the-one-simple-thing-fueling-your-social-media-addiction/>
- Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The satisfaction with life scale. *Journal of Personality Assessment*, 49(1), 71–75.
- Diener, E., Oishi, S., & Lucas, R. E. (2009). Subjective well-being: The science of happiness and life satisfaction. In S. J. Lopez & C. R. Snyder (Eds.), *Oxford handbook of positive psychology* (pp. 185–194). New York, NY: Oxford University Press.
- Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A., & Danforth, C. M. (2011). Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PloS one*, 6(12), e26752.

- Dohmen, T., Falk, A., Fliessbach, K., Sunde, U., & Weber, B. (2010). Relative versus absolute income, joy of winning, and gender: Brain imaging evidence. *Journal of Public Economics*, 95(3–4), 279–285. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0047272710001854>
- Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2), 117–140. Retrieved from <http://hum.sagepub.com/content/7/2/117>
- Fowler, J. H., & Christakis, N. A. (2008). Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham heart study. *British Medical Journal*, 337(2), a2338–a2338.
- Fox, J., & Moreland, J. J. (2015). The dark side of social networking sites: An exploration of the relational and psychological stressors associated with Facebook use and affordances. *Computers in Human Behavior*, 45, 168–176.
- Get a life!. (2013). *Economist*. Retrieved from <http://www.economist.com/news/science-andtechnology/21583593-using-social-network-seems-make-people-more-miserable-get-life>
- Goffman, E. (1956). *The presentation of self in everyday life*. New York, NY: Doubleday.
- Golder, S., & Macy, M. (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051), 1878–1881.
- Gonçalves, P., & Araújo, M. (2013, October 7–8). *Comparing and combining sentiment analysis methods*. Proceedings of the First ACM Conference on Online Social Networks, Boston, MA, USA, pp. 27–38.

- Grimes, S. (2011). Text-analytics demand approaches \$1 billion. *InformationWeek*. Retrieved from [http://www.informationweek.com/software/information-management/text-analytics-demand-approaches-\\$1-billion/d/d-id/1097709?](http://www.informationweek.com/software/information-management/text-analytics-demand-approaches-$1-billion/d/d-id/1097709?)
- Grimmelmann, J. (2014). The Facebook emotional manipulation study: Sources. *The Laboratorium*. Retrieved from http://laboratorium.net/archive/2014/06/30/the_facebook_emotional_manipulation_study_source
- Guillory, J., Spiegel, J., Drislane, M., Weiss, B., Donner, W., & Hancock, J. (2011, May 7–12). *Upset now? Emotion contagion in distributed groups*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Vancouver, BC, Canada, pp. 745–748.
- Haferkamp, N., & Krämer, N. C. (2011). Social comparison 2.0: Examining the effects of online profiles on social-networking sites. *Cyberpsychology, Behavior, and Social Networking*, *14*(5), 309–314.
- Hagerty, M. (2000). Social comparisons of income in one's community: Evidence from national surveys of income and happiness. *Journal of Personality and Social Psychology*, *78*(4), 764–771.
- Hancock, J., Landrigan, C., & Silver, C. (2007, April 28–May 3). *Expressing emotion in text-based communication*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, pp. 929–932.
- Hektner, J. M., Schmidt, J. A., & Csikszentmihalyi, M. (2007). *Experience sampling method: Measuring the quality of everyday life*. Thousand Oaks, CA: Sage.

- Hinsch, C., & Sheldon, K. M. (2013). The impact of frequent social Internet consumption: Increased procrastination and lower life satisfaction. *Journal of Consumer Behaviour*, 12(6), 496–505.
- Jordan, A. H., Monin, B., Dweck, C. S., Lovett, B. J., John, O. P., & Gross, J. J. (2011). Misery has more company than people think: Underestimating the prevalence of others' negative emotions. *Personality and Social Psychology Bulletin*, 37(1), 120–135.
- Kahneman, D., & Krueger, A. (2006). Developments in the measurement of subjective well-being. *Journal of Economic Perspectives*, 20(1), 3–24.
- Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, 306(5702), 1776–1780.
- Kivran-Swaine, F., & Naaman, M. (2011, March 19–23). *Network properties and social sharing of emotions in social awareness streams*. Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work, Hangzhou, China, pp. 379–382.
- Konnikova, M. (2013). How Facebook makes us unhappy. *New Yorker*. Retrieved from <http://www.newyorker.com/tech/elements/how-facebook-makes-us-unhappy>
- Kramer, A. (2010, April 10–15). *An unobtrusive behavioral model of gross national happiness*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA, pp. 287–290.
- Kramer, A. (2012, May 5–10). *The spread of emotion via Facebook*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Austin, TX, USA, pp. 767–770.
- Kramer, A. (2014). Okay so. A lot of people have asked me ... [Facebook post]. Retrieved from <https://www.facebook.com/akramer/posts/10152987150867796>

- Kramer, A., Guillory, J., & Hancock, J. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111(24), 8788–8790. Retrieved from <http://www.pnas.org/content/111/24/8788.full.pdf+html>
- Krasnova, H., Wenninger, H., Widjaja, T., & Buxmann, P. (2013, February 27–March 1). *Envy on Facebook: A hidden threat to users' life satisfaction?* Proceedings of the 11th International Conference on Wirtschaftsinformatik, Leipzig, Germany.
- Kross, E., Verduyn, P., Demiralp, E., Park, J., Seungjae Lee, D., Lin, N. ... , Ybarra, O. (2013). Facebook use predicts declines in subjective well-being in young adults. *PLoS ONE*, 8(8), e69841. Retrieved from <http://dx.plos.org/10.1371/journal.pone.0069841>
- Layard, R., Clark, A., & Senik, C. (2012). The causes of happiness and misery. In J. Helliwell, R. Layard, & J. Sachs (Eds.), *UN world happiness report* (pp. 58–89). New York, NY: Earth Institute.
- Lazer, D. M., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google flu: Traps in big data analysis. *Science*, 343(6176), 1203–1205.
- Lee, S. Y. (2014). How do people compare themselves with others on social network sites?: The case of Facebook. *Computers in Human Behavior*, 32, 253–260.
- Lin, L. F., & Kulik, J. A. (2002). Social comparison and women's body satisfaction. *Basic and Applied Social Psychology*, 24(2), 115–123.
- Lyubomirsky, S., & Ross, L. (1997). Hedonic consequences of social comparison: A contrast of happy and unhappy people. *Journal of Personality and Social Psychology*, 73(6), 1141–1157.

- Marche, S. (2012). Is Facebook making us lonely? *The Atlantic*. Retrieved from <http://www.theatlantic.com/magazine/archive/2012/05/is-facebook-making-us-lonely/308930/>
- Mislove, A., Lehmann, S., Ahn, Y. Y., Onnela, J. P., & Rosenquist, J. N. (2010). Pulse of the nation: U.S. mood throughout the day inferred from Twitter. Retrieved from <http://www.ccs.neu.edu/home/amislove/twittermood/>
- Mitchell, A., & Hitlin, P. (2013). *Twitter reaction to events often at odds with overall public opinion*. Pew Research Center. Retrieved from <http://www.pewresearch.org/2013/03/04/twitter-reactionto-events-often-at-odds-with-overall-public-opinion/>
- Morse, S., & Gergen, K. J. (1970). Social comparison, self-consistency, and the concept of self. *Journal of Personality and Social Psychology*, *16*(1), 148–156.
- Paltoglou, G., & Thelwall, M. (2012). Twitter, MySpace, Digg: Unsupervised sentiment analysis in social media. *ACM Transactions on Intelligent Systems and Technology*, *3*(4), 1–19.
- Panger, G. (2014, April 26–May 1). *Social comparison in social media: A look at Facebook and Twitter*. CHI '14 Extended Abstracts on Human Factors in Computing Systems, Toronto, ON, Canada, pp. 2095–2100.
- Pennebaker, J., Chung, C., Ireland, M., Gonzales, A., & Booth, R. (2007). The development and psychometric properties of LIWC2007. Retrieved from <http://www.liwc.net/LIWC2007LanguageManual.pdf>
- Przybylski, A. K., Murayama, K., DeHaan, C. R., & Gladwell, V. (2013). Motivational, emotional, and behavioral correlates of fear of missing out. *Computers in Human Behavior*, *29*(4), 1841–1848.

- Rosenfeld, E. (2015). Facebook user growth crushes estimates. *CNBC*. Retrieved from <http://www.cnn.com/id/102610670>
- Ruths, D., & Pfeffer, J. (2014). Social media for large studies of behavior. *Science*, *346*(6213), 1063–1064.
- Sagioglou, C., & Greitemeyer, T. (2014). Facebook's emotional consequences: Why Facebook causes a decrease in mood and why people still use it. *Computers in Human Behavior*, *35*, 359–363.
- Salovey, P., & Rodin, J. (1984). Some antecedents and consequences of social-comparison jealousy. *Journal of Personality and Social Psychology*, *47*(4), 780–792.
- Schroepfer, M. (2014). Research at Facebook. *Facebook Newsroom*. Retrieved from <http://newsroom.fb.com/news/2014/10/research-at-facebook/>
- Scollon, C. N., Prieto, C. K., & Diener, E. (2009). Experience sampling: Promises and pitfalls, strength and weaknesses. In *Assessing well-being* (pp. 157–180). New York, NY: Springer.
- Stuart, H. C., Dabbish, L., Kiesler, S., Kinnaird, P., & Kang, R. (2012, February 11–15). *Social transparency in networked information exchange: A framework and research question*. Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, Seattle, Washington, USA, pp. 451–460.
- Taheri, A. A., Tepper, M., Benerjee, A., & Sapiro, G. (2012, October 29–November 2). *If you are happy and you know it ... tweet*. Proceedings of the 21st ACM International Conference on Information and Knowledge Management, Maui, HI, USA, pp. 1602–1606.

- Tandoc, E. C., Ferrucci, P., & Duffy, M. (2015). Facebook use, envy, and depression among college students: Is Facebooking depressing?. *Computers in Human Behavior*, *43*, 139–146.
- Tow, S. (2011). Interesting News, Any Time You Visit. *Facebook Newsroom*. Retrieved from [https:// www.facebook.com/notes/facebook/interesting-news-any-time-you-visit/10150286921207131](https://www.facebook.com/notes/facebook/interesting-news-any-time-you-visit/10150286921207131)
- Tufekci, Z. (2014, June 2–4). *Big questions for social media Big Data: Representativeness, validity and other methodological pitfalls*. ICWSM '14: Proceedings of the 8th International AAAI Conference on Weblogs and Social Media, Ann Arbor, MI, USA.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. New York, NY: Basic Books.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.
- van Baaren, R., Janssen, L., Chartrand, T., & Dijksterhuis, A. (2009). Where is the love? The social aspects of mimicry. *Philosophical Transactions of the Royal Society B*, *364*(1528), 2381–2389.
- Verduyn, P., Lee, D., Park, J., Shablack, H., Orvell, A., Bayer, J., Ybarra, O., Jonides, J., & Kross, E. (2015). Passive Facebook usage undermines affective well-being: Experimental and longitudinal evidence. *Journal of Experimental Psychology: General*, *144*(2), 480–488.
- Vogel, E. A., Rose, J. P., Roberts, L. R., & Eckles, K. (2014). Social comparison, social media, and self-esteem. *Psychology of Popular Media Culture*, *3*(4), 206–222.

Wheeler, L., & Miyake, K. (1992). Social comparison in everyday life. *Journal of Personality and Social Psychology*, 62(5), 760–773.

Wood, J. V. (1989). Theory and research concerning social comparisons of personal attributes. *Psychological Bulletin*, 106(2), 231–248.

Wortham, J. (2011). Feel like a wallflower? Maybe it's your Facebook wall. *New York Times*. Retrieved from http://www.nytimes.com/2011/04/10/business/10ping.html?_r=0

Zimmerman, J. (2014). Facebook's study got it wrong: seeing our 'friends' happy can make us sad. *Guardian*. Retrieved from <http://www.theguardian.com/commentisfree/2014/jul/02/facebook-study-emotions-friends-totally-wrong>