

Does Counting Emotion Words on Online Social Networks Provide a Window Into People's Subjective Experience of Emotion?

Citation for published version (APA):

Kross, E., Verduyn, P., Boyer, M., Drake, B., Gainsburg, I., Vickers, B., Ybarra, O., & Jonides, J. (2019). Does Counting Emotion Words on Online Social Networks Provide a Window Into People's Subjective Experience of Emotion? A Case Study on Facebook. *Emotion*, 19(1), 97-107. <https://doi.org/10.1037/emo0000416>

Document status and date:

Published: 01/02/2019

DOI:

[10.1037/emo0000416](https://doi.org/10.1037/emo0000416)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Does Counting Emotion Words on Online Social Networks Provide a Window Into People’s Subjective Experience of Emotion? A Case Study on Facebook

Ethan Kross
University of Michigan

Philippe Verduyn
KU Leuven and Maastricht University

Margaret Boyer, Brittany Drake, Izzy Gainsburg, Brian Vickers, Oscar Ybarra, and John Jonides
University of Michigan

Psychologists have long debated whether it is possible to assess how people subjectively feel without asking them. The recent proliferation of online social networks has recently added a fresh chapter to this discussion, with research now suggesting that it is possible to index people’s subjective experience of emotion by simply counting the number of emotion words contained in their online social network posts. Whether the conclusions that emerge from this work are valid, however, rests on a critical assumption: that people’s usage of emotion words in their posts accurately reflects how they feel. Although this assumption is widespread in psychological research, here we suggest that there are reasons to challenge it. We corroborate these assertions in 2 ways. First, using data from 4 experience-sampling studies of emotion in young adults, we show that people’s reports of how they feel throughout the day neither predict, nor are predicted by, their use of emotion words on Facebook. Second, using simulations we show that although significant relationships emerge between the use of emotion words on Facebook and self-reported affect with increasingly large numbers of observations, the relationship between these variables was in the opposite of the theoretically expected direction 50% of the time (i.e., 3 of 6 models that we performed simulations on). In contrast to counting emotion words, we show that judges’ ratings of the emotionality of participants’ Facebook posts consistently predicts how people feel across all analyses. These findings shed light on how to draw inferences about emotion using online social network data.

Keywords: big data, emotion, Facebook, LIWC

Supplemental materials: <http://dx.doi.org/10.1037/emo0000416.supp>

A growing number of studies suggest that it is possible to index how people feel by simply counting the number of positive and negative emotion words contained in their online social network posts (e.g., Coviello et al., 2014; Eichstaedt et al., 2015; Golder & Macy, 2011; Jones, Wojcik, Sweeting, & Silver, 2016; Kramer,

2010, 2012; Kramer, Guillory, & Hancock, 2014; W. Wang, Hernandez, Newman, He, & Bian, 2016). Guided by this approach, research has begun to use data from online social networks such as Facebook and Twitter to examine how people’s feelings change in response to different events (e.g., Bollen, Pepe, & Mao, 2009; Coviello et al., 2014; Jones et al., 2016), whether they are “contagious” over online social networks (e.g., Coviello et al., 2014; Kramer, 2012; Kramer et al., 2014), how they vary across time, space, and culture (e.g., Davalos, Merchant, Rose, Lessley, & Teredesai, 2015; Golder & Macy, 2011; Ritter, Preston, & Hernandez, 2014; W. Wang et al., 2016), and whether they predict consequential outcomes that span the gamut from county-level cardiovascular disease (Eichstaedt et al., 2015) to stock market performance (e.g., Bollen, Mao, & Zeng, 2011; Siganos, Vagenas-Nanos, & Verwijmeren, 2014) to happiness, distress, and psychopathology (Lin, Tov, & Qiu, 2014; Shen, Brdiczka, & Liu, 2015; Wojcik, Hovasapian, Graham, Motyl, & Ditto, 2015).

These findings are rapidly having a cumulative impact on several areas of research within psychology. They are being used as a foundation to develop and extend a wide variety of psychological theories in which human emotion plays an important role—for example, research on emotional contagion (e.g., Coviello et al.,

This article was published Online First April 5, 2018.

Ethan Kross, Department of Psychology, University of Michigan; Philippe Verduyn, Faculty of Psychology and Educational Sciences, KU Leuven, and Faculty of Psychology and Neuroscience, Maastricht University; Margaret Boyer, Brittany Drake, Izzy Gainsburg, Brian Vickers, Oscar Ybarra, and John Jonides, Department of Psychology, University of Michigan.

This research was supported by funds provided by the University of Michigan to Ethan Kross and a postdoctoral research fellowship to Philippe Verduyn from the Fund for Scientific Research-Flanders.

Correspondence concerning this article should be addressed to Ethan Kross, Department of Psychology, University of Michigan, 530 Church Street, Ann Arbor, MI 48104 or Philippe Verduyn, Department Work & Social Psychology, Maastricht University Universiteitssingel 40 (r. 4.753), 6229 Maastricht, the Netherlands. E-mail: ekross@umich.edu or philippe.verduyn@maastrichtuniversity.nl

2014; Kramer, 2012; Kramer et al., 2014), stress, happiness and well-being (e.g., Eichstaedt et al., 2015; Kramer et al., 2014; Ritter et al., 2014), emotional dynamics (Golder & Macy, 2011), and behavioral decision making (e.g., Bollen, Gonçalves, Ruan, & Mao, 2011; Siganos et al., 2014). But whether the conclusions that emerge from these studies are valid, rests on a critical assumption: that people's usage of positive and negative emotion words in their posts accurately reflects how they feel. Although this assumption is widespread in research using Facebook data, here we suggest that there are both conceptual and empirical reasons to challenge it, and we provide evidence to corroborate these claims.

Conceptual Concerns

Consider first the fact that simply counting the number of positive and negative emotion words contained in people's posts fails to take context into account (for a similar argument, see Pang & Lee, 2008; Panger, 2016; Tausczik & Pennebaker, 2010)—an idea that is central to many prominent theories of emotion (e.g., Barrett, 2006; Ellsworth & Scherer, 2003). To illustrate, reflect on the following statements: "I am so happy" and "You look so happy." A simple text analysis that counts the percentage of positive words would give these statements the same score—25% positive words—even though these statements convey obviously different messages. Simple word counting methods have no way of adjudicating between such statements.

Adding to this concern, recent studies indicate that people often manage their images on Facebook (e.g., Bazarova, Taft, Choi, & Cosley, 2012; Ellison, Heino, & Gibbs, 2006; Lin et al., 2014; Mehdizadeh, 2010; Qiu, Lin, Leung, & Tov, 2012; Siibak, 2009; Strano, 2008; Zhao, Grasmuck, & Martin, 2008). For example, in response to learning that a coworker earned a promotion, a person might post, "I am so happy for you," when they might, in fact, feel jealous. Thus, even if word counting methods were capable of capturing context, there is still reason to doubt whether the words people use in their online social networks accurately express how they feel.

Empirical Concerns

Compounding these conceptual issues are empirical concerns about the two types of studies that have been used to validate word-counting algorithms as a tool to draw inferences about emotional experience from online social network data. On the one hand, a number of studies have examined the correlation between positive and negative word usage in people's online social network posts and their life satisfaction ratings (e.g., Kramer, 2010; Liu, Tov, Kosinski, Stillwell, & Qiu, 2015). Such work has yielded inconsistent findings, with some research revealing the expected links between these variables (Kramer, 2010), and other work failing to replicate these results (N. Wang, Kosinski, Stillwell, & Rust, 2014). Even more problematic, these studies use life-satisfaction scores as the criterion to validate the proposal that positive and negative emotional word usage indexes how people feel. Although life satisfaction represents an important dimension of subjective well-being, it is distinct from people's subjective experience of how positively or negatively they feel on a moment-to-moment basis (i.e., their self-report emotion ratings; Kahneman & Deaton, 2010; Kahneman & Krueger, 2006)—the construct that

is supposedly tapped by counting emotional word usage on Facebook (or Twitter).

Other studies have examined the correlation between positive and negative emotional word usage and judges' ratings of the emotionality of participants' posts (e.g., Schwartz et al., 2013). But as explained above, it is possible for people to express emotional states on their online social networks that do not reflect how they actually feel. In such cases, one wouldn't expect judges' ratings of the emotionality of participants' posts to track with participants' self-reported emotion. Moreover, even if judges' ratings of the emotionality of participants' posts did correlate with how people reported feeling, demonstrating a correlation between such ratings and the number of emotion words contained in people's online social network posts still doesn't provide evidence indicating that emotional word usage captures how people feel. As Figure 1 illustrates, it is possible for two variables to correlate with a third variable, but not with each other.

Overview of Research

So, do the positive and negative emotion words that people use in their social network posts provide a window into how they feel? We addressed this issue by pooling data from four experience-sampling studies of emotion in young adults. Each of these studies contained two types of data: people's self-reported ratings of how good or bad they felt throughout the day over several days, and their Facebook wall posts corresponding to the time period that these self-reported affect measurements were obtained. These data allowed us to examine whether people's self-reports of how they felt throughout the day predict, or are predicted by, their naturally occurring use of positive and negative emotion words in their Facebook posts.

The unique nature of these data also allowed us to address a secondary question of interest. Specifically, we suggested earlier that emotional word usage on Facebook may not accurately index people's feelings because counting the frequency of such words fails to take the context under which such words are used into account. On the one hand, Facebook posts that explicitly express positive or negative emotions (e.g., "I am so happy now") should provide a better predictor of people's self-reports of how they feel. On the other hand, if people use their Facebook posts to manage their public image (e.g., Bazarova et al., 2012; Ellison et al., 2006; Lin et al., 2014; Mehdizadeh, 2010; Qiu et al., 2012; Siibak, 2009; Strano, 2008; Zhao et al., 2008), then there is reason to doubt whether even such explicit statements will index how people feel. To adjudicate between these alternatives, we also asked judges to content-analyze participants' posts for the degree to which they expressed positive or negative emotion and then examined how these statements tracked with both their self-reported ratings of emotion and the number of emotional words contained in participants' posts.

Method

Studies

This article used data from four experience-sampling studies that obtained information on participants' moment-to-moment happiness and their naturally occurring Facebook posting behavior

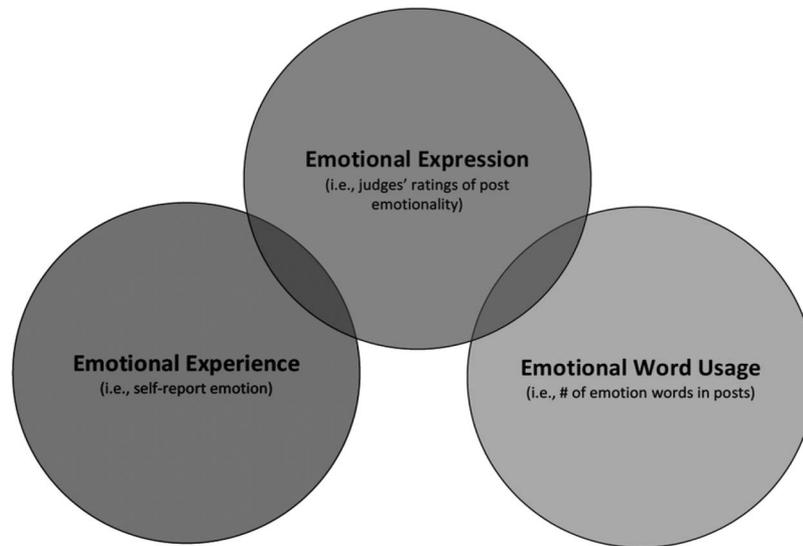


Figure 1. Diagram illustrating the transitive fallacy using the three variables that this article focuses on. Shaded regions of overlap between circles represent hypothetical significant correlations. The diagram demonstrates how it is possible for one variable (e.g., emotional expression) to correlate with two other variables (e.g., emotional experience and emotional word usage), without the latter variables correlating with each other. Assuming a correlation between the latter variables is an example of the transitive fallacy.

corresponding to the duration of the experience-sampling period of each study. Note that each of these studies focused on slightly different questions and included different measures and procedures. We focus here solely on those measures that were both identical across studies and necessary to address the article's goal of examining whether young adults' self-reports of how they feel throughout the day predict, or are predicted by, their naturally occurring use of positive and negative emotional words in their Facebook posts.

Results from two of these data sets have been previously published, although neither of them focuses on the questions of the current manuscript (Kross et al., 2013; Verduyn et al., 2015). The other two data sets have not been published previously. The University of Michigan Institutional Review Board approved all studies that were administered at Michigan; the social and societal ethics committee of KU Leuven approved the study that was administered at that site. Informed written consent was obtained from all participants prior to participation. Data used in the current analyses can be accessed at <https://www.openicpsr.org/openicpsr/project/101284/version/V1/view/>.

Participants

Across the four studies, we obtained data on 311¹ participants' moment-to-moment happiness and their naturally occurring Facebook posting behavior corresponding to the duration of the experience-sampling period of each study. The overwhelming majority of participants were students. Participants were compensated with course credit or money.

We performed two sets of analyses: (a) analyses examining whether the overall percentage of positive or negative words contained in people's Facebook posts correlates with their mean level of self-reported affect during the study (i.e., person-level

analysis), and (b) fine-grained analyses examining the relationship between the percentage of positive or negative words contained in participants' posts and their self-reported affect ratings preceding or following their posts (i.e., postlevel analysis).

The total sample size for the former analysis was $n = 185$ participants; the total sample size for the latter analysis was $n = 543$ observations derived from 158 participants (see Table 1 for demographics on these participants). We elaborate on why the number of participants for each of these analyses departs from the overall number of participants contained in our data file ($n = 311$) next.

Person-level analysis data exclusions. Following prior experience-sampling research (Felsman, Verduyn, Ayduk, & Kross, 2017; Verduyn et al., 2015) we excluded data from 17 participants who responded to <60% of the texts on a priori grounds.² The Facebook walls of 103 participants did not contain a single post during the study period. Consequently, there were no Facebook-posts to match against their self-reported affect scores. Participants who did not post did not differ in demographics from those who did post (all $ps > .21$). Because of technical problems, the timestamps for all of the wall posts of six participants in the Belgian sample were not saved and we could therefore not know whether their posts fell inside the study period. This resulted in a total dataset of 185 participants for all person-level analyses.

Postlevel analysis data exclusions. Our study focused on the degree to which the emotional words contained in participants'

¹ This number reflects the subset of participants who participated in these studies who possessed the two types of data that were the focus of the current analyses (i.e., Facebook wall data and self-reported affect data).

² Including these participants did not alter any of the conclusions we report (for details, see the online supplemental material).

Table 1
Demographics

Study no.	Location	Duration	Person-level analyses							Postlevel analyses						
			<i>n</i>	Age	% Women	% White	% Asian	% Black	% Other	<i>n</i>	% Women	Age	% White	% Asian	% Black	% Other
1	USA	14 days	62	19.66 (2.37)	67.7	62.9	29.0	3.2	4.8	58	67.2	19.68 (2.44)	63.8	27.6	3.4	5.2
2	USA	6 days	53	20.70 (2.44)	73.6	52.8	37.7	5.7	3.8	48	75.0	20.77 (2.45)	50.0	39.6	6.3	4.2
3	USA	7 days	44	20.23 (1.84)	81.8	59.1	29.5	6.8	4.5	33	78.8	20.33 (1.98)	63.6	27.3	6.1	3
4	Belgium	6 days	26	21.50 (1.58)	100	100	0	0	0	19	100	21.89 (1.56)	100	0	0	0
Total	—	—	185	20.35 (2.25)	77.3	64.3	27.6	4.3	3.8	158	75.9	20.42 (2.35)	63.9	27.8	4.4	3.8

Facebook posts correlated with self-report affect ratings that were generated in close proximity to such posts. Thus, we excluded posts that were made outside of the experience sampling window (i.e., at night or during early morning). We also calculated the time lag separating each Facebook post (T), and the affect assessment that preceded ($T - 1$) and followed it ($T + 1$). Observations that were more than three standard deviations above the sample mean for the time lag separating each Facebook post and the affect assessment that (a) preceded it ($n = 9$), and (b) followed it ($n = 5$) were removed from all analyses (after removing these outliers, the time separating self-reported affect ratings and subsequent Facebook post: $M = 93.41$ min, $SD = 69.74$; time separating Facebook post and subsequent self-reported affect rating: $M = 115.12$ min, $SD = 78.55$). Removing these posts further reduced the sample by 27 participants resulting in a final dataset of 543 observations derived from 158 participants.³

Power. The multilevel analyses of the postlevel data (543 observations from 158 participants) constituted the primary focus of our initial inquiry and were well powered to detect small effects. We also performed subject level analyses ($n = 185$) to completely explicate the data set. These analyses, although not our primary focus, were well-powered to detect medium-sized effects (see online supplemental files for additional information on power). Critically, we considered the possibility that neither of the aforementioned analyses would be sufficient to observe remarkably small effects, a point that is noteworthy given that prior research using “big data” has demonstrated how results with extremely small effect sizes (e.g., $d = .002$) can surpass the $\alpha < .05$ level for drawing inferences about statistical significance when sample sizes are sufficiently large. Thus, as we describe in more detail below, we used bootstrapping techniques to perform simulations with up to 20,000 participants to demonstrate that any failure to detect “statistically significant” results in the theoretically expected direction was not a function of the sample size we used.

Experience Sampling Procedure

Participants were text-messaged five times per day between 10 am and midnight for 14 days (Study 1), 7 days (Study 3), or 6 days (Studies 2 and 4). Participants in Studies 1–3 received these messages on their own touchscreen smartphone; participants in Study 4 received these messages on a smartphone they borrowed from the experimenters for the duration of the study. Text-messages occurred at random times within 168-min windows per day. Each text-message contained a link to an online survey, which asked participants to answer a series of questions including, “How do you feel right now?” ranging from 0 (*very positive*) to 100 (*very negative*), which our analyses focused on. The scale was reversed

prior to analyses such that high scores reflect how positively participants felt. Participants always answered this affect question first.

Matching Facebook Wall Data and Experience Sampling Data

We matched each time-stamped participant-generated Facebook post to a corresponding experience-sampling interval. For example, if a participant completed an experience-sampling survey at 11:00 a.m. (i.e., Text 1) and 12:30 p.m. (i.e., Text 2), interval T_{1-2} would include any post made by the participant between those times. We also calculated the time lag separating each Facebook post (T), and the affect assessment that preceded ($T - 1$) and followed it ($T + 1$) to examine whether the amount of time separating each post and each affect assessment influenced the nature of the relationship between these variables. When multiple posts were made within one interval we used the average of their time-stamps for these calculations.

Facebook Wall Data Content Analyses

To prepare the Facebook wall data for coding we first identified all information by replacing all names with two-digit identification codes (e.g., Participant 01, Friend 01). We then content analyzed each time stamped post in two ways as described below.

Positive and Negative Emotion Word Frequency

We used the Linguistic Inquiry and Word Count (LIWC) to compute the percentage of positive and negative words posted (Pennebaker, Booth, & Francis, 2007). Posts were analyzed with the LIWC dictionary that corresponded to the official language(s) of the study populations: English for the U.S. samples and Dutch or French (one single post) for the Belgian sample.

For the person-level analyses the posts of each of the 185 participants were stacked in one file and the percentage of positive and negative emotion words was calculated for each participant. Together, the 185 participants generated 1,446 posts; 18.4 (42.7) percent of people did not post a single positive (negative) word.

For the postlevel analyses, the 858 posts made by the 158 participants within the experience sampling window were distributed across 543 intervals (some posts fell within the same interval). LIWC analysis was used to determine the percentage of positive

³ Including these participants did not alter any of the conclusions we report (for details, see the online supplemental material).

and negative words posted for each of the 543 intervals. If multiple posts were made within a single interval they were compiled into a single text file for LIWC analysis resulting in a single percentage of positive and negative words for each interval containing at least one post. Thirty-four percent of intervals did not contain any positive words; 75% of intervals did not contain any negative words.

Explicit Expressions of Emotion

After a researcher fluent in English and Dutch translated posts from the Belgian sample into English, two independent judges content analyzed posts for the degree to which they expressed how negative or positive the poster felt, ranging from -3 (*very negative expression*) to 3 (*very positive expression*). If a post contained an emotion word but did not express the participant's own emotional state, (i.e., "You look so happy"), coders were trained to score it zero. Interjudge reliability was excellent (overall sample: $\alpha = .88$, U.S. samples: $\alpha = .89$, Belgian sample: $\alpha = .77$). Therefore, final ratings were averaged across the two judges ($M = .17$, $SD = .73$). The majority of posts (79.6%) received a score of zero. Note that 99% of posts that were rated 0 did not convey any emotion; 1% of these posts conveyed an equal amount of positive and negative emotion.

For the person-level analysis, the emotion expression data for each participant were averaged resulting in a single rating of emotion expressivity for each participant. For the postlevel analysis, emotion expression was calculated for each interval containing at least one post. If the interval contained a single post, emotion expression was based on the expression rating of this single post. If multiple posts were made within a single interval they were coded separately and then averaged resulting in a single rating of emotion expression for each interval containing at least one post.

Overview of Data Analyses

Subject-level analysis. We first used regression analyses to examine whether the percentage of positive and negative words participants posted during the study period predicted their mean level of affect reported over the course of the study period. Next, we examined whether participants' mean level of emotion expression (as determined by judges) predicted their mean level of affect. Analyses involving judges' ratings controlled for the total number of words participants posted ($M = 107.4$, $SD = 151.82$); controlling for word count was not necessary in the analyses involving the emotion word count measures because those scores consisted of a percentage of the overall word count. Controlling for study sample did not substantively alter any of the results we report. Thus, we do not discuss this variable further.

Postlevel Analyses. We used multilevel analyses to examine whether the use of positive and negative emotion words in participants' posts (or judges' ratings of the emotionality of participants' posts) predicted, or were predicted by, participants' self-reported emotional experience. The intercept was allowed to vary across participants in each analysis to take into account possible dependencies following from the nested data structure (posts nested within participants). Level 1 predictors were grand-mean centered (because the absolute values of the independent variables are

expected to be related to the outcome variables), and entered as fixed effects in the multilevel models (i.e., slopes were not allowed to vary across participants) as participants contributed few observations ($M = 3.44$, $SD = 3.05$) making subject-specific relationships unreliable. Note that when allowing slopes to vary across participants, the corresponding variances of these random effects were not significantly different from zero (all $ps > .56$) and conclusions regarding fixed effects remained identical. In addition, it is notable that the regression weights we report are unstandardized. As with the person-level analyses, controlling for study sample did not substantively alter any of the results we report. Therefore, this variable is not discussed further.

Finally, if we wanted to examine whether the nature of Facebook posts at T_{2-3} predicted T_3 affect but did not have data on T_2 affect, then we predicted T_3 affect from T_{1-3} Facebook posts instead. Excluding intervals that did not consist of consecutive texts occurring on the same day (rather than following the aforementioned analytical scheme) did not substantively alter any of the results.

Multilevel Data Analyses Proceeded in 4 Phases

Phase 1: Affect predicting Facebook posts. First, we examined whether the way a person reported feeling prior to generating their post(s) (Time " $T - 1$ ") predicted the number of emotion words contained in their post(s) (Time " T "). As with the person-level analyses, we controlled for the number of words posted ($M = 21.54$, $SD = 28.06$) when predicting judges' ratings of the emotionality of participants posts as this number was not taken into account when coding this dependent variable.

Phase 2: Facebook posts predicting affect. Next, we examined the reverse relationship—whether the number of emotion words contained in participants' posts (Time " T ") predicted how participants reported feeling immediately after they generated that post (Time " $T + 1$ "). Once again, we controlled for the number of words posted when predicting self-reported affect by judges' ratings of the emotionality in the preceding posts.

Note that we analyzed the data both of the aforementioned ways because it was not clear, a priori, whether people's self-reports of how they felt before ($T - 1$) or after ($T + 1$) generating a post would correlate more strongly with emotional word usage (T). Thus, by modeling both relationships we aimed to explore both possibilities (for a similar approach, see Brans, Koval, Verduyn, Lim, & Kuppens, 2013; Kross et al., 2013; Verduyn et al., 2015).

Phase 3: Moderation by time lag. We examined whether the time lag separating our assessment of self-reported affect and when participants made their posts moderated any of the aforementioned relationships to account for the possibility that measurements obtained in closer temporal proximity might yield stronger relationships. Specifically, the amount of time separating Facebook posts (T) and their preceding affect rating ($T - 1$) was examined as a possible moderator at level 1 of the model when predicting the number of positive or negative words in Facebook posts (T) by affect ($T - 1$). The amount of time separating Facebook posts (T) and their subsequent affect rating ($T + 1$) was also examined as a possible moderator when predicting affect ($T + 1$) by the number of positive or negative words in Facebook posts (T).

Phase 4: Simulations. Each of the analyses described above were initially repeated with sample sizes that were simulated at 500, 1000, and every 1,000 subjects until 10,000 subjects were reached using bootstrapping techniques. We chose 10,000 subjects as an initial maximum since these were exploratory analyses and we had no a priori expectation concerning what number of simulations would be necessary for significant results to stabilize. As we report below, all of the analyses stabilized at $\alpha < .05$ with 10,000 participants. Thus, we did not run larger simulations under the assumption that doing so would not provide additional useful information. Because we resampled from the same set of subjects, multilevel models with multiple responses per subject have more data points than the number of subjects (on average, 3.44 data points per subject; e.g., sampling 10,000 subjects in a multilevel model will contain 34,400 data points on average).

Model equations and supplementary analyses. Model equations and supplementary analyses using complementary data-analytic approaches led to identical conclusions to those reported below and are reported in the online supplemental material.

Results

Descriptive Statistics

Table 2 presents the mean and standard deviation of the primary variables examined in the person- and postlevel analyses. Table 3 reports the correlations between these variables.

Person-Level Analyses

We first examined the association between self-reported emotion, positive- and negative emotion word usage and judges' ratings of the emotionality of participants' posts at the person level. The results of the regression analyses are presented at the top of Table 4. They did not reveal significant linear relationships among any of the Facebook emotional measures we administered and self-reported affect.

Preliminary Postlevel Analyses

Judges' ratings of the emotionality of participants posts correlated marginally significantly with the percentage of positive words ($B = .004$, $SE = .002$, $t(541) = 1.94$, $p = .053$, confidence interval [CI] $-.001$, $.008$) and significantly with the percentage of negative words ($B = -.027$, $SE = .007$, $t(541) = -3.72$, $p <$

$.001$, CI $[-.041, -.013]$ contained in participants' posts. This was to be expected, as human coders must rely to some degree on emotional language to make their judgments. However, as noted in the introduction, judges' ratings of the emotionality of participants' posts do not directly represent the psychological variable that emotional word usage is believed to capture—how people actually feel.

Next, we examined the temporal stability of participants' self-reported affect ratings. Participants' ratings of how they felt throughout the day were to a certain degree stable. That is, how people reported feeling at one assessment (T1) predicted how they felt at the next assessment (T2; $B = .31$, $SE = .04$, $t(541) = 7.71$, $p < .0001$, CI $[.23, .39]$). As such, if the percentages of positive and negative words contained in participants' posts are markers of emotional experience, they should be related to self-reported affect ratings preceding or following their posting.

Main Postlevel Analyses

We first examined whether people's self-reports of how they felt at one point in time (Time "T - 1") predicted their subsequent use of positive or negative emotion words in their Facebook posts (Time "T"; see Table 4 for all multilevel results).

Multilevel analyses indicated that this was not the case. Next, we examined whether the reverse sequence of events was true—perhaps people's usage of positive or negative emotion words in their Facebook posts at Time "T" predicted how they subsequently report feeling at time "T + 1." Multilevel analyses indicated that this was also not the case.

It is important to note that the time lag separating peoples' self-reports of how they felt during the day and their use of emotion words on Facebook varied from 0 min to 374 min, which raised the question: Might the null associations between these variables be a result of the variable time gap separating their assessments?

To address this question, we reran each of the aforementioned analyses controlling for the amount of time separating our assessment of self-reported affect and Facebook emotional word usage. Controlling for this factor did not substantially alter any of the results reported above (all $ps > .38$). We also examined whether the links between self-reported affect and Facebook emotional word usage became stronger as the amount of time separating the assessment of these variables shrank. Moderation analyses indicated that this also was not the case—time between measurements did not interact with affect to predict emotional word usage (or

Table 2
Means and Standard Deviations of Primary Variables

Study no.	Person-level dataset					Postlevel dataset					
	Affect	% pos	% neg	Judges' ratings	no. posts	Affect @ T - 1	Affect @ T + 1	% pos	% neg	Judges' ratings	no. posts
1	62.77 (14.38)	7.57 (3.87)	1.87 (1.63)	.08 (.27)	904	63.10 (25.98)	63.24 (27.73)	9.70 (11.90)	1.97 (4.70)	.13 (.73)	498
2	65.02 (14.39)	9.91 (15.43)	2.27 (5.63)	.18 (.49)	284	64.67 (23.41)	67.28 (22.45)	10.55 (16.53)	1.34 (3.32)	.21 (.63)	199
3	65.12 (12.18)	10.53 (12.32)	1.60 (2.35)	.32 (.71)	183	64.38 (22.86)	63.59 (23.60)	16.14 (24.44)	1.31 (3.48)	.28 (.80)	119
4	61.30 (7.06)	2.84 (5.68)	2.31 (5.11)	.03 (.62)	75	55.63 (21.78)	56.04 (20.82)	2.82 (6.25)	1.17 (3.98)	.12 (.78)	42
Total	63.77 (13.06)	8.28 (10.88)	1.98 (3.84)	.16 (.52)	1446	63.23 (24.86)	63.73 (25.91)	10.47 (15.34)	1.71 (4.26)	.17 (.73)	858

Note. Means are reported with standard deviations in parentheses. Affect = mean affect computed for the entire study period; % pos = the percentage of positive words contained in participants' posts; % neg = the percentage of negative words contained in participants' posts; Judge's ratings = the degree to which participants expressed positive (or negative) emotion in their posts, as determined by judges; Affect @ T - 1 = self-reported affect corresponding to a T - 1 observation; Affect @ T + 1 = self-reported affect corresponding to a T + 1 observation.

Table 3
Correlations Between Primary Variables at the Person- and Postlevel

Variable	1	2	3	4	5	6
1. Self-reported affect (aggregate)		—	—	-.059	-.026	.056
2. Self-reported affect (T + 1)	—		—	—	—	—
3. Self-reported affect (T - 1)	—	.434**	—	—	—	—
4. % Positive words	—	.024	.036	—	-.104	.056
5. % Negative words	—	.015	.027	-.128**	—	-.104
6. Judges' ratings	—	.140**	.142**	.085*	-.153**	—

Note. Correlations at the person-level are above the diagonal; correlations at the postlevel are below the diagonal. Self-report affect and judges' ratings are coded such that higher scores indicate more positive affect.
* $p < .05$. ** $p < .01$.

the other way around) in any of the aforementioned analyses (all $ps > .45$).⁴

If simple counts of the emotion words contained in people's Facebook posts do not index how people feel, might judges' ratings of the emotionality of participants' Facebook posts do so? To address this question, we reran each of the above analyses using judges' ratings of the emotionality of participants' posts instead of participants' emotional word usage levels.

Multilevel analyses indicated that participants' self-report ratings of how they felt prior to posting on Facebook significantly predicted judges' ratings of the emotionality of their subsequent Facebook posts in the theoretically expected direction—that is, the more positive participants reported feeling at time "T - 1" the more positive judges rated their Facebook posts at Time "T." The reverse relationship was also significant: Judges' ratings of the positivity of participants' Facebook posts at Time "T" significantly predicted how positive participants subsequently reported feeling at Time "T + 1." Controlling for the amount of time separating participants' self-reported affect assessment and Facebook posts did not alter the significance of any of these results (judges' ratings were still significantly related to how people reported feeling, all $ps < .02$). This variable likewise did not moderate any of the above relationships (all $ps > .84$).

Simulations of Larger Sample Sizes

Although the sample sizes used in the aforementioned analyses were sufficient to observe small (i.e., postlevel analyses) to medium (i.e., subject-level analyses) sized effects, the sample sizes used in big data research are considerably larger. Might the links between self-reported affect and Facebook emotional word usage become significant with larger numbers of observations? Although we recognized that any nonzero correlation will eventually become statistically significant with a large enough sample, we nevertheless thought that there would be value in determining when this occurs. Moreover, we reasoned that it was possible that statistically significant effects might emerge in the opposite of the theoretically expected direction (i.e., more negative words predicting more positive affect; more positive words predicting less positive affect) if word counting methods do not provide a valid means of drawing inferences about emotion from online social network posts. To address these issues, we used bootstrapping techniques to repeat each of the above analyses with larger sample sizes.

We first performed simulations on the links between self-reported affect and each of the Facebook emotion measures we

assessed at the person-level (positive emotion word usage, negative emotion word usage, and judges' ratings of the emotionality of participants' posts). As Table 5 illustrates, no significant links emerged between self-reported affect and any of these measures in simulations that consisted of fewer than 2,000 simulated participants. However, samples sizes involving 3,000 or more participants generated significant associations in the theoretically expected direction between self-reported affect and judges' ratings of the emotionality of participants' posts. We also observed significant associations between self-reported affect and positive emotion word usage at sample sizes involving 2,000 or more participants. However, these associations were in the *opposite* of the theoretically predicted direction—using more positive words correlated with feeling worse. As the simulations grew larger (i.e., 9,000 or more participants), significant relationships emerged between negative emotion word usage and self-reported affect in the theoretically predicted direction.

Next, we examined these same relationships at the postlevel. We first examined whether people's self-reports of how they felt at Time "T - 1" predicted each Facebook emotion measure at Time "T." As the bottom panel of Table 5 illustrates, self-reported affect at T - 1 predicted both judges' ratings of the emotionality of participants posts and negative word usage at Time T in all samples sizes. However, although the results for judges' ratings were in the theoretically predicted direction, the results for negative emotion word usage were in the *opposite* of the theoretically predicted direction—feeling better predicted using more negative emotion words. We also observed a significant relationship between self-reported affect and subsequent positive emotional word usage at sample sizes involving 2,000 or more participants in the theoretically predicted direction. Note that the pattern of results observed on these postlevel analyses directly opposes the results of the simulation analyses at the person-level reported above, in which negative word usage correlated in the expected direction with self-reported affect but positive word usage did not.

Simulations that examined the reverse sequence of events (i.e., Facebook emotion measures assessed at Time "T" predicting self-reported affect at Time "T + 1") generated a largely similar pattern of results (see the middle panel of Table 5). Judge's ratings of Facebook posts at Time "T" predicted self-reported affect at Time "T + 1" in the theoretically expected direction in all simu-

⁴ Self-reported emotional intensity was also not related to how soon participants posted after getting a text message ($r = -.01$).

Table 4
Person- and Postlevel analyses' betas, Standard Errors, P Values, and 95% Confidence Intervals

Model	B	SE	p value	95% CI
Person-level analysis				
Judge's ratings → self-report affect	1.42	1.86	.45	-2.25, 5.08
% Positive words → self-report affect	-.07	.09	.43	-.25, .10
% Negative words → self-report affect	-.09	.25	.73	-.58, .41
Postlevel analysis (multilevel)				
Self-report affect (T minus 1) → judges' ratings (T)	.004	.001	.002**	.001, .006
Self-report affect (T - 1) → % positive words (T)	.02	.03	.40	-.03, .07
Self-report affect (T - 1) → % negative words (T)	.00	.01	.55	-.01, .02
Judges' ratings (T) → self-report affect (T + 1)	3.23	1.34	.02*	.60, 5.86
% Positive words (T) → self-report affect (T + 1)	.01	.06	.81	-.11, .14
% Negative words (T) → self-report affect (T + 1)	-.05	.23	.84	-.49, .40

Note. CI = confidence interval; T = time lag separating each Facebook post; T minus 1 = time lag separating each Facebook and the affect assessment that preceded; T + 1 = time lag separating each Facebook and the affect assessment that followed it. The scales for self-report affect and judges' ratings of the emotionality of participants' Facebook posts were in the same direction—shifting from low to high represents a shift from negative to positive. Each predictor was tested in a separate model. All regressions involving judges' ratings control for word count. We report regressions for all analyses to maintain consistency across subject-level and postlevel analyses. Note, however, that regressions performed at the subject-level that examine % positive/negative words → self-report affect generate equivalent *p* values as zero-order correlations performed with these variables; the regression examining the link between judges' ratings and self-report affect is equivalent to a partial correlation between these variables that controls for word count.

p* < .05. *p* < .01.

lations. Negative emotion word usage predicted affect in the opposite of the theoretically expected direction at “T + 1” at simulations involving 4,000 or more participants. Finally, we observed a marginal relationship between positive word usage and self-reported affect in the theoretically expected direction with samples sizes involving more than 8,000 or more participants.

Discussion

Although the prospect of being able to predict people's subjective experience of emotion by simply counting the number of positive and negative emotion words contained in their online social network posts is enticing, the current findings question the validity of this approach. They demonstrate that young adults' self-reports of how they feel throughout the day do not predict subsequent Facebook emotion word usage. Nor does their usage of emotion words on Facebook predict how they subsequently feel. Moreover, although simulation analyses revealed significant relationships between positive and negative emotion word usage and self-reports of how people felt at high numbers of simulations, the direction of the relationship between different word count measures and self-reported affect was inconsistent across analyses; in three out of the six postlevel analysis models we performed simulations on, the word counting index correlated significantly with self-reported affect in the opposite of the theoretically expected direction (i.e., negative word usage predicting feeling better and more positive words feeling worse). Taken together, these results suggest that simply counting young adults' usage of emotion words on Facebook does not accurately index how they feel.

In contrast, we observed consistent significant (albeit “small”) associations in the theoretically expected direction between people's self-reports of how they felt and judges' ratings of the emotionality of participants' posts across all models at the

postlevel. Specifically, the more judges rated participants' posts as conveying positive (or negative) emotion, the better (or worse) participants' subsequently reported feeling. Similarly, the better participants reported feeling, the more judges rated participants' subsequent posts as conveying positive emotion. Moreover, judges' ratings predicted self-reported affect in the theoretically expected direction across all simulations in which a significant relationship was revealed. These findings suggest that it is possible to draw accurate inferences about how people feel from online social network data. They also suggest that people do not self-present all the time when they post information on Facebook (Buck et al., 2010)—if that were the case, then we should not have been able to identify a link between people's posts and how they actually feel.

An important question raised by these findings concerns what people's usage of emotion words on Facebook reflects if they do not capture people's subjective experience of emotion. Although our research was not designed to address this question, it is possible that people's posts mirror their observation of others (i.e., you look so happy; look how happy Fred is at the award ceremony). Future research is needed to explore this and other possibilities.

The current findings may appear to be controversial in the face of the large and growing number of online social network studies that employ word-counting methods to draw inferences about emotion. However, in their now classic review of the use of text-analysis programs to draw inferences about psychological processes from “offline” data, Pennebaker and colleagues (2003) concluded the following:

Virtually every psychologically based text analysis approach has started from the assumption that we can detect peoples' emotional

Table 5
Simulation Sample Size (N), Mean P Values, 95% Confidence Interval (CI), and Proportion of Results From Each Simulation That Were in the Theoretically Predicted Direction (% Pred. Dir.)

n	Judge's ratings			Negative word count			Positive word count		
	p	95% CI	% pred. dir.	p	95% CI	% pred. dir.	p	95% CI	% pred. dir.
Person-Level Analyses									
.5k	.310	-.69, 3.57	91%	.528	-.26, .19	79%	.279	-.15, .01	5%
1k	.196	-.11, 2.99	96%	.454	-.21, .07	88%	.137	-.12, -.02	1%
2k	.069 [†]	.29, 2.53	99%	.302	-.17, .01	96%	.032*	-.11, -.03	0%
3k	.023*	.57, 2.23	100%	.231	-.16, -.01	98%	.010**	-.10, -.04	0%
4k	.010**	.70, 2.16	100%	.157	-.15, -.02	100%	.003**	-.10, -.04	0%
5k	.004**	.75, 2.15	100%	.119	-.14, -.03	100%	.001***	-.09, -.05	0%
6k	.002**	.80, 2.06	100%	.096 [†]	-.13, -.03	100%	.000***	-.09, -.05	0%
7k	.001***	.82, 1.94	100%	.070 [†]	-.13, -.04	100%	.000***	-.09, -.05	0%
8k	.000***	.91, 1.95	100%	.055 [†]	-.13, -.04	100%	.000***	-.09, -.05	0%
9k	.000***	.91, 1.88	100%	.040*	-.13, -.04	100%	.000***	-.09, -.05	0%
10k	.000***	.95, 1.91	100%	.030*	-.12, -.05	100%	.000***	-.09, -.06	0%
Post-Level Analysis: Facebook Emotion Measure (T) → Self-Report Affect (T+1)									
.5k	.001**	1.93, 4.63	100%	.408	-.05, .20	11%	.467	-.05, .09	70%
1k	.000***	2.36, 4.25	100%	.292	-.01, .16	4%	.407	-.03, .07	77%
2k	.000***	2.58, 3.93	100%	.139	.02, .13	1%	.309	-.01, .06	90%
3k	.000***	2.77, 3.78	100%	.073 [†]	.03, .12	0%	.262	-.01, .05	93%
4k	.000***	2.80, 3.77	100%	.037*	.04, .11	0%	.207	.00, .05	96%
5k	.000***	2.87, 3.67	100%	.022*	.04, .11	0%	.164	.00, .04	96%
6k	.000***	2.88, 3.67	100%	.013*	.04, .11	0%	.146	.00, .04	97%
7k	.000***	2.91, 3.64	100%	.006**	.05, .11	0%	.109	.00, .04	99%
8k	.000***	2.94, 3.60	100%	.004**	.05, .10	0%	.095 [†]	.00, .04	99%
9k	.000***	2.96, 3.58	100%	.002**	.05, .10	0%	.074 [†]	.01, .04	100%
10k	.000***	3.00, 3.58	100%	.001***	.05, .10	0%	.055 [†]	.01, .04	100%
Post-Level Analysis: Self-Report Affect (T-1) → Facebook Emotion Measure (T)									
.5k	.000***	.002, .005	100%	.047*	.01, .02	0%	.253	.00, .05	95%
1k	.000***	.003, .005	100%	.006**	.01, .02	0%	.132	.00, .04	99%
2k	.000***	.003, .004	100%	.000***	.01, .02	0%	.035*	.01, .03	100%
3k	.000***	.003, .004	100%	.000***	.01, .02	0%	.012*	.01, .03	100%
4k	.000***	.003, .004	100%	.000***	.01, .02	0%	.002**	.01, .03	100%
5k	.000***	.003, .004	100%	.000***	.01, .02	0%	.001***	.01, .03	100%
6k	.000***	.003, .004	100%	.000***	.01, .02	0%	.000***	.01, .03	100%
7k	.000***	.003, .004	100%	.000***	.01, .02	0%	.000***	.01, .03	100%
8k	.000***	.003, .004	100%	.000***	.01, .02	0%	.000***	.01, .03	100%
9k	.000***	.003, .004	100%	.000***	.01, .02	0%	.000***	.01, .03	100%
10k	.000***	.003, .004	100%	.000***	.01, .02	0%	.000***	.01, .03	100%

Note. In bootstrapping it is possible that p values and CI do not cohere (i.e., a p value > .05 can be associated with a CI that does not contain the value zero). Statistical significance reported in the article is based on p values to align interpretation of these results with the criterion used for the rest of the analyses reported in the article. However, it is notable that using CIs as a criterion to infer significance does not alter any of the conclusions we report.
[†] p < .10. * p < .05. ** p < .01. *** p < .001.

states by studying the emotion words they use . . . [but] in reviewing the various word use studies, it is striking how weakly emotion words predict people's emotional state . . . taken together, it is our sense that emotion researchers should hesitate before embarking on studies that rely exclusively on the natural production of emotion words. (p. 571)

The current findings are consistent with this conclusion. They extend this idea to the digital domain, where emotion word counting methods are increasingly being used as a tool to draw inferences about people's subjective emotional states.

It is important to emphasize that the current findings do not question the existence of significant relationships between the numbers of emotions words contained in people's online social network posts and the various outcomes of interest reviewed earlier. Rather, they question the interpretation of such findings—that is, it is possible that such word counts predict these outcomes, but not for reasons that they are claimed to.

It is also important to recognize that the current work focused specifically on whether people's usage of emotion words on Facebook predicted their subjective emotional states—the variable that prior social media research that has used emotion word counting methods to draw inferences about. However, it is well established that emotions have physiological and behavioral components. Future research should examine how judges and word counting methods predict these aspects of an emotional experience as well.

Four caveats are in order before concluding. First, our analyses focused on data from young adults who represent a significant user base of Facebook—according to a recent report approximately 82% of Internet users who use Facebook are between the ages of 18 and 29 (Duggan, 2015). Thus, our findings are potentially relevant to a large segment of the population. Nevertheless, future research is needed to examine whether these findings generalize to

older Facebook users, Facebook users from other cultures, and noncollege students. Second, a number of recent papers have used people's usage of positive and negative words on Twitter to draw inferences about emotions in the same manner described here (Bollen, Mao, et al., 2011; Bollen et al., 2009; Eichstaedt et al., 2015; Golder & Macy, 2011; Ritter et al., 2014). Identifying whether the current findings extend to this platform as well is critical for assessing the generality of our conclusions. Third, it is possible that other aspects of people's writing on social media (e.g., their usage of punctuation and function words) may index people's subjective emotional states more accurately. Although addressing this issue was not the focus of this work, future research should explore this possibility. Finally, in the present study we collected a limited number of posts from participants. Consequently, it was not possible to reliably separate within- from between-subjects effects at the post level. Future studies should consider sampling participants over larger periods of time to address this issue.

Concluding Comment

Within a relatively short time span, online social networks have revolutionized the way human beings interact, providing psychological scientists with a new platform to ask and answer important, and at times provocative, questions about human emotion. Although the current research suggests that such interactions can provide important insights into people's emotional lives, they suggest that the dominant way that such insights are currently gleaned—by counting the number of emotion words people use in their posts—may not provide a valid means of doing so.

References

- Back, M. D., Stopfer, J. M., Vazire, S., Gaddis, S., Schmukle, S. C., Egloff, B., & Gosling, S. D. (2010). Facebook profiles reflect actual personality, not self-idealization. *Psychological Science, 21*, 372–374. <http://dx.doi.org/10.1177/0956797609360756>
- Barrett, L. F. (2006). Solving the emotion paradox: Categorization and the experience of emotion. *Personality and Social Psychology Review, 10*, 20–46. http://dx.doi.org/10.1207/s15327957pspr1001_2
- Bazarova, N. N., Taft, J. G., Choi, Y. H., & Cosley, D. (2012). Managing impressions and relationships on Facebook: Self-presentational and relational concerns revealed through the analysis of language style. *Journal of Language and Social Psychology, 32*, 121–141.
- Bollen, J., Gonçalves, B., Ruan, G., & Mao, H. (2011). Happiness is assortative in online social networks. *Artificial Life, 17*, 237–251. http://dx.doi.org/10.1162/artl_a_00034
- Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of Computational Science, 2*, 1–8. <http://dx.doi.org/10.1016/j.jocs.2010.12.007>
- Bollen, J., Pepe, A., & Mao, H. (2009, July). *Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena*. Paper presented at the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM), Barcelona, Spain.
- Brans, K., Koval, P., Verduyn, P., Lim, Y. L., & Kuppens, P. (2013). The regulation of negative and positive affect in daily life. *Emotion, 13*, 926–939. <http://dx.doi.org/10.1037/a0032400>
- Coviello, L., Sohn, Y., Kramer, A. D., Marlow, C., Franceschetti, M., Christakis, N. A., & Fowler, J. H. (2014). Detecting emotional contagion in massive social networks. *PLoS ONE, 9*(3), e90315. <http://dx.doi.org/10.1371/journal.pone.0090315>
- Davalos, S., Merchant, A., Rose, G. M., Lessley, B. J., & Teredesai, A. M. (2015). 'The good old days': An examination of nostalgia in Facebook posts. *International Journal of Human-Computer Studies, 83*, 83–93. <http://dx.doi.org/10.1016/j.ijhcs.2015.05.009>
- Duggan, M. (2015). *Mobile messaging and social media: 2015*. Washington, DC: Pew Research Center.
- Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., . . . Seligman, M. E. (2015). Psychological language on Twitter predicts county-level heart disease mortality. *Psychological Science, 26*, 159–169. <http://dx.doi.org/10.1177/0956797614557867>
- Ellison, N., Heino, R., & Gibbs, J. (2006). Managing impressions online: Self-presentation processes in the online dating environment. *Journal of Computer-Mediated Communication, 11*, 415–441. <http://dx.doi.org/10.1111/j.1083-6101.2006.00020.x>
- Ellsworth, P. C., & Scherer, K. R. (2003). Appraisal processes in emotion. In R. J. Davidson, K. R. Scherer, and H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 572–595). Oxford, UK: Oxford University Press.
- Felsman, P., Verduyn, P., Ayduk, O., & Kross, E. (2017). Being present: Focusing on the present predicts improvements in life satisfaction but not happiness. *Emotion, 17*, 1047–1051. <http://dx.doi.org/10.1037/emo0000333>
- Golder, S. A., & Macy, M. W. (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science, 333*, 1878–1881. <http://dx.doi.org/10.1126/science.1202775>
- Jones, N. M., Wojcik, S. P., Sweeting, J., & Silver, R. C. (2016). Tweeting negative emotion: An investigation of Twitter data in the aftermath of violence on college campuses. *Psychological Methods, 21*, 526–541. <http://dx.doi.org/10.1037/met0000099>
- Kahneman, D., & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences of the United States of America, 107*, 16489–16493. <http://dx.doi.org/10.1073/pnas.1011492107>
- Kahneman, D., & Krueger, A. B. (2006). Developments in the measurement of subjective well-being. *The Journal of Economic Perspectives, 20*, 3–24. <http://dx.doi.org/10.1257/089533006776526030>
- Kramer, A. D. (2010, April). *An unobtrusive behavioral model of gross national happiness*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.
- Kramer, A. D. (2012). *The spread of emotion via Facebook*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.
- Kramer, A. D., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences, USA of the United States of America, 111*, 8788–8790. <http://dx.doi.org/10.1073/pnas.1320040111>
- Kross, E., Verduyn, P., Demiralp, E., Park, J., Lee, D. S., Lin, N., . . . Ybarra, O. (2013). Facebook use predicts declines in subjective well-being in young adults. *PLoS ONE, 8*(8), e69841. <http://dx.doi.org/10.1371/journal.pone.0069841>
- Lin, H., Tov, W., & Qiu, L. (2014). Emotional disclosure on social networking sites: The role of network structure and psychological needs. *Computers in Human Behavior, 41*, 342–350. <http://dx.doi.org/10.1016/j.chb.2014.09.045>
- Liu, P., Tov, W., Kosinski, M., Stillwell, D. J., & Qiu, L. (2015). Do Facebook status updates reflect subjective well-being? *Cyberpsychology, Behavior and Social Networking, 18*, 373–379. <http://dx.doi.org/10.1089/cyber.2015.0022>
- Mehdizadeh, S. (2010). Self-presentation 2.0: Narcissism and self-esteem on Facebook. *Cyberpsychology, Behavior, and Social Networking, 13*, 357–364. <http://dx.doi.org/10.1089/cyber.2009.0257>
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval, 2*(1–2), 1–135.

- Panger, G. (2016). Reassessing the Facebook experiment: Critical thinking about the validity of Big Data research. *Information Communication & Society, 19*, 1108–1126.
- Pennebaker, J. W., Booth, R. J., & Francis, M. E. (2007). Linguistic Inquiry and Word Count (LIWC2007): A computer-based text analysis program [Computer software]. Austin, TX: LIWC. net.
- Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. G. (2003). Psychological aspects of natural language. use: Our words, our selves. *Annual Review of Psychology, 54*, 547–577. <http://dx.doi.org/10.1146/annurev.psych.54.101601.145041>
- Qiu, L., Lin, H., Leung, A. K., & Tov, W. (2012). Putting their best foot forward: Emotional disclosure on Facebook. *Cyberpsychology, Behavior and Social Networking, 15*, 569–572. <http://dx.doi.org/10.1089/cyber.2012.0200>
- Ritter, R. S., Preston, J. L., & Hernandez, I. (2014). Happy tweets: Christians are happier, more socially connected, and less analytical than atheists on Twitter. *Social Psychological & Personality Science, 5*, 243–249.
- Schwartz, H. A., Eichstaedt, J., Dziurzynski, L., Blanco, E., Kern, M. L., Ramones, S., . . . Ungar, L. (2013). Choosing the right words: Characterizing and reducing error of the word count approach. In *Proceedings of SEM-2013: Second Joint Conference on Lexical and Computational Semantics* (pp. 296–305). Stroudsburg, PA: Association for Computational Linguistics.
- Shen, J., Brdiczka, O., & Liu, J. (2015). A study of Facebook behavior: What does it tell about your Neuroticism and Extraversion? *Computers in Human Behavior, 45*, 32–38. <http://dx.doi.org/10.1016/j.chb.2014.11.067>
- Siganos, A., Vagenas-Nanos, E., & Verwijmeren, P. (2014). Facebook's daily sentiment and international stock markets. *Journal of Economic Behavior & Organization, 107*, 730–743. <http://dx.doi.org/10.1016/j.jebo.2014.06.004>
- Siibak, A. (2009). Constructing the self through the photo selection-visual impression management on social networking websites. *Cyberpsychology: Journal of Psychosocial Research of Cyberspace, 3*, 1.
- Strano, M. M. (2008). User descriptions and interpretations of self-presentation through Facebook profile images. *Cyberpsychology: Journal of Psychosocial Research of Cyberspace, 2*, 5.
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology, 29*, 24–54. <http://dx.doi.org/10.1177/0261927X09351676>
- Verduyn, P., Lee, D. S., Park, J., Shablack, H., Orvell, A., Bayer, J., . . . Kross, E. (2015). Passive Facebook usage undermines affective well-being: Experimental and longitudinal evidence. *Journal of Experimental Psychology: General, 144*, 480–488. <http://dx.doi.org/10.1037/xge0000057>
- Wang, N., Kosinski, M., Stillwell, D., & Rust, J. (2014). Can well-being be measured using Facebook status updates? Validation of Facebook's Gross National Happiness Index. *Social Indicators Research, 115*, 483–491. <http://dx.doi.org/10.1007/s11205-012-9996-9>
- Wang, W., Hernandez, I., Newman, D. A., He, J., & Bian, J. (2016). Twitter analysis: Studying US weekly trends in work stress and emotion. *Applied Psychology: An International Review, 65*, 355–378. <http://dx.doi.org/10.1111/apps.12065>
- Wojcik, S. P., Hovasapian, A., Graham, J., Motyl, M., & Ditto, P. H. (2015). Conservatives report, but liberals display, greater happiness. *Science, 347*, 1243–1246. <http://dx.doi.org/10.1126/science.1260817>
- Zhao, S., Grasmuck, S., & Martin, J. (2008). Identity construction on Facebook: Digital empowerment in anchored relationships. *Computers in Human Behavior, 24*, 1816–1836. <http://dx.doi.org/10.1016/j.chb.2008.02.012>

Received June 15, 2017

Revision received November 20, 2017

Accepted December 20, 2017 ■