

UC San Diego

UC San Diego Previously Published Works

Title

Weather impacts expressed sentiment

Permalink

<https://escholarship.org/uc/item/7kg7t66z>

Journal

PLOS ONE, 13(4)

ISSN

1932-6203

Authors

Baylis, Patrick
Obradovich, Nick
Kryvasheyev, Yury
et al.

Publication Date

2018

DOI

10.1371/journal.pone.0195750

Peer reviewed

RESEARCH ARTICLE

Weather impacts expressed sentiment

Patrick Baylis¹, Nick Obradovich²*, Yury Kryvasheyev³, Haohui Chen³, Lorenzo Coviello⁴, Esteban Moro^{2,5}, Manuel Cebrian², James H. Fowler⁶

1 Vancouver School of Economics, University of British Columbia, Vancouver, British Columbia, Canada, **2** Media Lab, Massachusetts Institute of Technology, Cambridge, MA, United States of America, **3** Data61, Commonwealth Scientific and Industrial Research Organisation (CSIRO), Melbourne, Australia, **4** Institute of Electrical and Electronics Engineers, New York, NY, United States of America, **5** Department of Mathematics and GIS, Universidad Carlos III de Madrid, Leganes, Spain, **6** Departments of Political Science and Medicine, UC San Diego, San Diego, CA, United States of America

✉ These authors contributed equally to this work.

* nobradov@mit.edu



OPEN ACCESS

Citation: Baylis P, Obradovich N, Kryvasheyev Y, Chen H, Coviello L, Moro E, et al. (2018) Weather impacts expressed sentiment. PLoS ONE 13(4): e0195750. <https://doi.org/10.1371/journal.pone.0195750>

Editor: Tobias Preis, University of Warwick, UNITED KINGDOM

Received: December 29, 2017

Accepted: March 28, 2018

Published: April 25, 2018

Copyright: © 2018 Baylis et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: We collected our Twitter data from the public domain in adherence with Twitter's Developer Agreement and we used aggregated Facebook data published previously. The data used in this study are restricted from public redistribution by both Facebook and Twitter. However, all intermediate estimates needed to evaluate the conclusions in the paper are present in the paper and/or the Supporting Information. The raw Twitter data may be procured through Twitter's GNIP service, and the official data sharing policy at Facebook is that they will work with researchers who want to replicate published findings.

Abstract

We conduct the largest ever investigation into the relationship between meteorological conditions and the sentiment of human expressions. To do this, we employ over three and a half billion social media posts from tens of millions of individuals from both Facebook and Twitter between 2009 and 2016. We find that cold temperatures, hot temperatures, precipitation, narrower daily temperature ranges, humidity, and cloud cover are all associated with worsened expressions of sentiment, even when excluding weather-related posts. We compare the magnitude of our estimates with the effect sizes associated with notable historical events occurring within our data.

Introduction

Mood and emotional state support human physical, psychological, and economic well-being. Positive emotions are associated with improved physiological factors such as cortisol levels and cardiovascular functioning [1] and amplify cognitive performance and mental flexibility [2]. They can also increase social connectedness and perceived social support [3] and may augment income and economic success [4]. Emotional states can also be transmitted through social networks [5,6], amplifying the broad-scale effects of altered individual emotions.

Prior work suggests that environmental factors—and ambient meteorological conditions in particular—may substantively impact emotional state. However, previous empirical investigations of this relationship have found conflicting results. Early studies found large associations between meteorological conditions and mood [7,8] but were limited by small sample sizes and limited generalizability. A number of studies in the most recent decade have found small to negligible associations [9–11], while others document associations that vary across individuals [12,13], associations observed at high levels of aggregation [14], or associations that are contingent on other factors [15]. A still more recent large-scale longitudinal analysis reports robust linkages between daily weather variation and reported well-being [16]. Whether, and if so, how meteorological variables shape human emotions remains an open question.

This pattern of divergent results is due in part to a lack of large-scale data on emotional states. To rectify this problem, we employ a correlate of emotional states: the sentiment of

Funding: This work was supported by Ministerio de Economía y Competitividad: FIS2013-47532-C3-3-P, FIS2016-78904-C3-3-P (<http://www.mineco.gob.es/>); and National Science Foundation DGE0707423, TG-SES130013, 0903551 (<https://www.nsf.gov/>). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

human lexical expressions on social media [17]. We report on associations between meteorological conditions and the expressed sentiment of tens of millions of United States residents across 3.5 billion posts on both Facebook and Twitter between 2009 and 2016. This work expands on Baylis (2015), who uses a billion Twitter posts over a shorter time-frame to estimate preferences for and valuations of different realizations of temperature in order to project the amenity cost of climate change [18]. In particular, the significantly longer sample period and the inclusion of the more representative, longer-form Facebook data in our analysis allows us to defensibly generalize our findings more broadly than was possible in Baylis (2015).

In this manuscript, we examine three questions. First, do weather conditions associate with changes in the sentiment of human expressions? Second, are these associations robust to excluding discussion about the weather itself? Third, how do the magnitudes of these weather-sentiment associations compare to the effect sizes of other events that alter expressed sentiment?

Materials and methods

Data collection procedure

Our social media data consist of 3.5 billion posts in total, with 2.4 billion from Facebook and with 1.1 billion from Twitter. By using data from both social media platforms, we take advantage of the relative strengths of each as a data source on sentiment: Facebook data is more likely to be representative and to consist of text expressions revealing the user's underlying emotional state, while Twitter data provides for comparison of results across social media contexts. A central benefit of using the Facebook data is that the Facebook population is more likely to reflect the population at large. By the end of our sampling period, nearly 80% of online adults used Facebook, as compared to between 15 and 30% for Twitter, LinkedIn, Pinterest, and Instagram [19]. Surveyed adults also indicated that they use Facebook more frequently than any other social media platform. Of note, the authors complied with all relevant platform terms of service.

To measure expression of sentiment on Facebook, we use data from a previous work [5]. These data are based on "status updates" which are text-based messages that a user's contacts may view on their own Facebook News Feed. Our Facebook dataset starts on January 1st, 2009 and ends on March 31st, 2012, with 1,176 days in total. The Facebook data contain all users on the platform—both public-facing and private accounts—that chose English as their language, selected United States as their country, and could be matched to our selection of metropolitan areas by their IP-based geographic location at time of posting. The Facebook data we use here are described in detail elsewhere [5].

Our Twitter data consist of posts, or "tweets", that are short messages limited to 140 characters and are publicly viewable by default. Our Twitter data cover the period from November 30th, 2013 to June 30th, 2016, with 938 days in total. We collected tweets using Twitter's public Streaming API, placing a bounding box filter over the United States to gather our sample of precisely geo-located tweets. We then assigned tweets falling within a metropolitan area's boundaries to that specific area. This procedure allows for a high level of certainty that the tweet originated within a specific metropolitan area. We exclude re-tweets from our analysis and only consider direct user-generated content.

We employ gridded meteorological data from the PRISM Climate Group for our daily maximum temperature, temperature range, and precipitation measures [20]. We also employ cloud cover—a measure of sun exposure—and relative humidity data from the National Centers for Environmental Prediction (NCEP) Reanalysis II project [21] as these measures have been central to previous studies of the relationship between weather and emotional states (as daily

NCEP data run only to June 30th 2016 at the time of writing, we use that date as the end date of our analysis) [7,8]. We match daily meteorological variables in a location to the posts of users in that particular location on that day.

Measures of expressed sentiment

To determine whether a social media post uses words that express positive or negative sentiment, we rely on the Linguistic Inquiry Word Count (LIWC) sentiment analysis tool (see Table B in [S1 File](#)) [22]. LIWC is a highly validated, dictionary-based, sentiment classification tool that is commonly used to assess sentiment in social media posts [5,6,23,24] (of note, our results are similar under the use of alternative sentiment classifiers). In our analysis, we follow the psychological literature and treat positive and negative sentiment as separate constructs [25].

Dependent variable aggregation

We aggregate both our Facebook and Twitter data to the 75 most populated metropolitan areas in the US. For each post, we calculate whether the post contains either a positive or a negative LIWC term. We then average each user's posts on each day to produce a positive rate and a negative rate for each user. Then, for each city on each day, we average the values of users in that city to calculate our city-level dependent variables. In order to ensure our results aren't driven by ecological inference, we present user-level analyses in Fig A, B, and C and Tables D and F in [S1 File](#).

Analysis

City-level. To investigate if weather alters expressed sentiment at the city-level, we combine our aggregated city-level positive and negative sentiment scores with our daily meteorological data. We empirically model this relationship as:

$$Y_{jmt} = f(tmax_{jmt}) + g(precip_{jmt}) + h(\mu) + \gamma_t + \nu_{jm} + \epsilon_{jmt} \quad (1)$$

In this time-series cross-sectional model, j indexes cities, m indexes unique year-months, and t indexes calendar days. Our dependent variable Y_{jmt} represents our city-level measure of positive or negative expressed sentiment, respectively.

Our independent variables of interest are maximum temperatures ($tmax_{jmt}$) and total precipitation ($precip_{jmt}$). We also examine the marginal effects of temperature range, percentage cloud cover, and relative humidity, represented via $h(\mu)$. We empirically estimate our relationships of interest using indicator variables for each 5°C maximum temperature and temperature range bin, for each 1cm precipitation bin, and for each 20 percentage point bin of cloud cover and relative humidity (represented here by $f()$, $g()$, and $h()$ respectively). This procedure allows for flexible estimation of the association between our meteorological variables and expressed sentiment [26–28].

Unobserved geographic or temporal factors may influence sentiment in a way that correlates with weather. For example, people may be happier on average in cities that have better infrastructure or on days when they are likely to have more leisure time. Further, there may be unobserved, city-specific trends, such as changes in amount of daylight throughout the year or evolution in city-level economic conditions over time, that influence the expressed sentiment of a city. To ensure that these factors do not bias our estimates of the association between weather and expressed sentiment, we include in [Eq 1](#) ν_{jm} and γ_t to represent city-by-calendar-month and calendar date indicator variables, respectively. These variables control for all

constant unobserved characteristics for each city across its seasons and for each unique day in the data [29,30].

We adjust for within-city and within-day correlation in ϵ_{jmt} by employing heteroskedasticity-robust standard errors clustered on both city-year-month and day [31]. We exclude non-climatic control variables from Eq 1 because of their potential to generate bias in our parameters of interest [30,32]. Finally, we weight the city-level regression by the number of underlying social media posts for each city-day.

We omit the 20°C–25°C maximum temperature, the 0°C–5°C temperature range, 0cm precipitation, 0–20% cloud cover, and 40–60% humidity indicator variables when estimating Eq 1. We interpret our estimates as the percentage point change in positive or negative expressed sentiment associated with a particular meteorological observation range relative to these baseline categories.

Exclusion of weather terms. Our first analyses examine the sentiment of all expressions contained within our data, inclusive of terms that may refer directly to the weather. As weather discussion may not necessarily reflect changes in individuals' underlying emotional states, however, we use a large dictionary of weather terms to filter out posts in our Twitter data that contain a plausible reference to the weather, and again run the models in Eq 1 and Equation A in S1 File on the messages that do not contain these weather related terms (because we no longer have access to the raw Facebook posts, we were unable to re-calculate our non-weather expressed sentiment metrics for our Facebook corpus). Approximately 4% of tweets contained one or more of our weather terms. To investigate the effectiveness of this filter, we manually classified a random sample of 1,000 posts that included a weather term and determined that approximately 28% of that sample were about the weather. We also manually classified a second random sample of 1,000 non-weather term posts and determined that only 0.2% were about weather-related constructs.

Effect sizes in context. To contextualize our estimates, we compare the effect size of the association between below freezing temperatures and non-weather related positive expressed sentiment to the effect size of a number of plausibly negative events that occurred over the time span of our Twitter data. To do so, we again estimate both Eq 1 and Equation A in S1 File. But, in addition to the terms in those models, we include indicator terms for each of the comparison events to estimate these parameters simultaneously alongside our meteorological variables. These indicator terms isolate the specific location and the specific dates of the event so that they are not collinear with the fixed effects in our models. For example, for the effect size of the San Francisco Bay Area earthquake, we create an indicator variable for the San Francisco/Oakland metropolitan area on August 24th 2014, the date of the earthquake.

Results

Descriptive statistics

We present the descriptive statistics associated with our main variables in Table 1.

All expressed sentiment

The results of estimating Eq 1 for the association between meteorological conditions and positive and negative sentiment on Facebook indicate that temperature, precipitation, humidity, and cloud cover each significantly relate to expressions of sentiment (see Fig 1 panels (a) and (b) for the maximum temperature and precipitation results). Positive expressions increase up to maximum temperatures of 20°C and decline past 30°C. The impact of temperature on negative expressions is qualitatively the opposite of its impact on positive expressions, though smaller in magnitude. Precipitation worsens expressed sentiment. Daily temperature ranges

Table 1. Summary statistics of dependent and independent variables.

	Facebook City Mean	Facebook City Std. Dev.	Twitter City Mean	Twitter City Std. Dev.
Pos. Rate	41	4.39	34.93	2.72
Neg. Rate	21.4	3.25	18.54	2.43
Max. Temperature	20.38	10.79	21.24	10.83
Diurnal Temp. Range	11.45	4.31	11.34	4.16
Precipitation	0.25	0.78	0.26	0.82
Cloud Cover	36.5	27.41	40.09	27.09
Humidity	68.23	18.52	67.64	18

<https://doi.org/10.1371/journal.pone.0195750.t001>

exceeding 15°C are associated with significant increases in positive expressions (coefficient: 0.104, p: 0.002, n: 85,801) and reductions in negative expressions (coefficient: -0.177, p: < 0.001, n: 85,801). Levels of relative humidity exceeding 80% decrease positive expressions (coefficient: -0.084, p: 0.001, n: 85,801) and increase negative expressions (coefficient: 0.039, p: 0.038, n: 85,801), as do days with high cloud cover, reducing positive expressions (coefficient: -0.2, p: < 0.001, n: 85,801) and increasing negative ones (coefficient: 0.168, p: < 0.001, n: 85,801) (see Table C in S1 File for further details on these results).

Panels (c) and (d) of Fig 1 display the results of estimating Eq 1 on the Twitter city-level data. The nature of the impact of temperature and precipitation on sentiment expression is quite similar to the associations in the Facebook data, though attenuated in magnitude. For example, the association between below freezing temperatures and city-level positive sentiment expressions on Twitter is approximately 45% the size of this parameter in the Facebook data. The effect sizes of precipitation, temperature range, cloud cover, and humidity on expressed sentiment on city-level Twitter data retain statistical significance but are similarly attenuated in magnitude as compared to Facebook (see Table C in S1 File for more details).

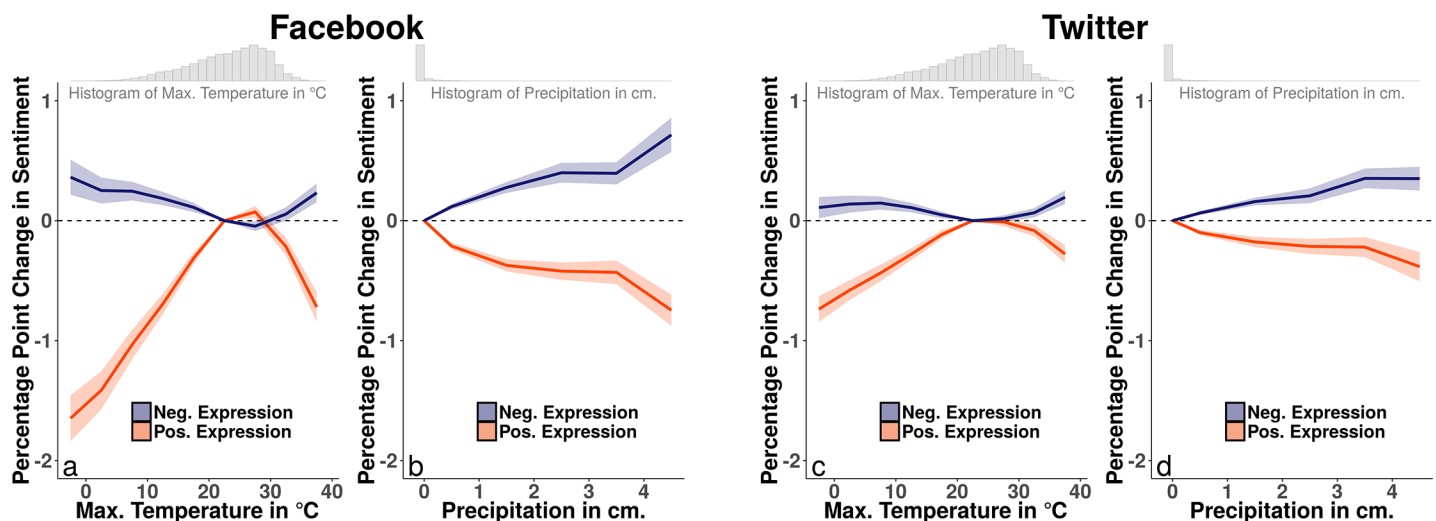


Fig 1. Facebook and Twitter analyses for all message types. Panel (a) depicts the relationship between daily maximum temperatures and the rates of expressed sentiment of approximately 2.4 billion Facebook status updates from 2009–2012, aggregated to the city-level. It draws from the estimation of Eq 1 and plots the predicted change in expressed sentiment associated with each maximum temperature bin. Panel (b) depicts the relationship between daily precipitation and the rates of sentiment expression of Facebook status updates, also drawing on estimation of Eq 1. Panels (c) and (d) replicate these analyses for nearly 1.1 billion Twitter posts between 2013 and 2016 aggregated to the same city. Shaded error bounds represent 95% confidence intervals calculated using heteroskedasticity-robust standard errors clustered on both city-year-month and day.

<https://doi.org/10.1371/journal.pone.0195750.g001>

Expressed sentiment of non-weather messages

Analyzing all posts provides useful descriptive characteristics of the ways in which aggregate sentiment associates with the weather. However, in order to employ expressed sentiment as a better proxy for underlying emotional states, it is useful to exclude direct references to the weather itself, as more extreme weather conditions significantly alter both the rate of weather discussion and the sentiment of this discussion (see Figs D and E in [S1 File](#)). We present the results of estimating [Eq 1](#) on our corpus of Twitter posts excluding weather terms in [Fig 2](#).

Panels (a) and (b) of [Fig 2](#), illustrate that the effect sizes associated with maximum temperature and precipitation on non-weather sentiment at the city-level are slightly smaller than the effect sizes seen in the all-posts sample ([Fig 1](#) panels (c) and (d)). As an example, the association between below freezing temperatures and positive sentiment, non-weather-term expressions is approximately 84% the size of this association in the all-posts model. The effect sizes of temperature range, high humidity, and cloud cover retain significance in this model, though are also attenuated in size compared to the all-posts model (see [Table E](#) in [S1 File](#) for details).

Effect sizes in context

To understand the relative magnitude of these predicted changes, it can be helpful to look at the effect sizes associated with other types of events on expressed sentiment. We chose five salient events for their theoretically negative association with expressed positive, non-weather sentiment coupled with their diversity of type: 1) the end of daylight saving time in the cities

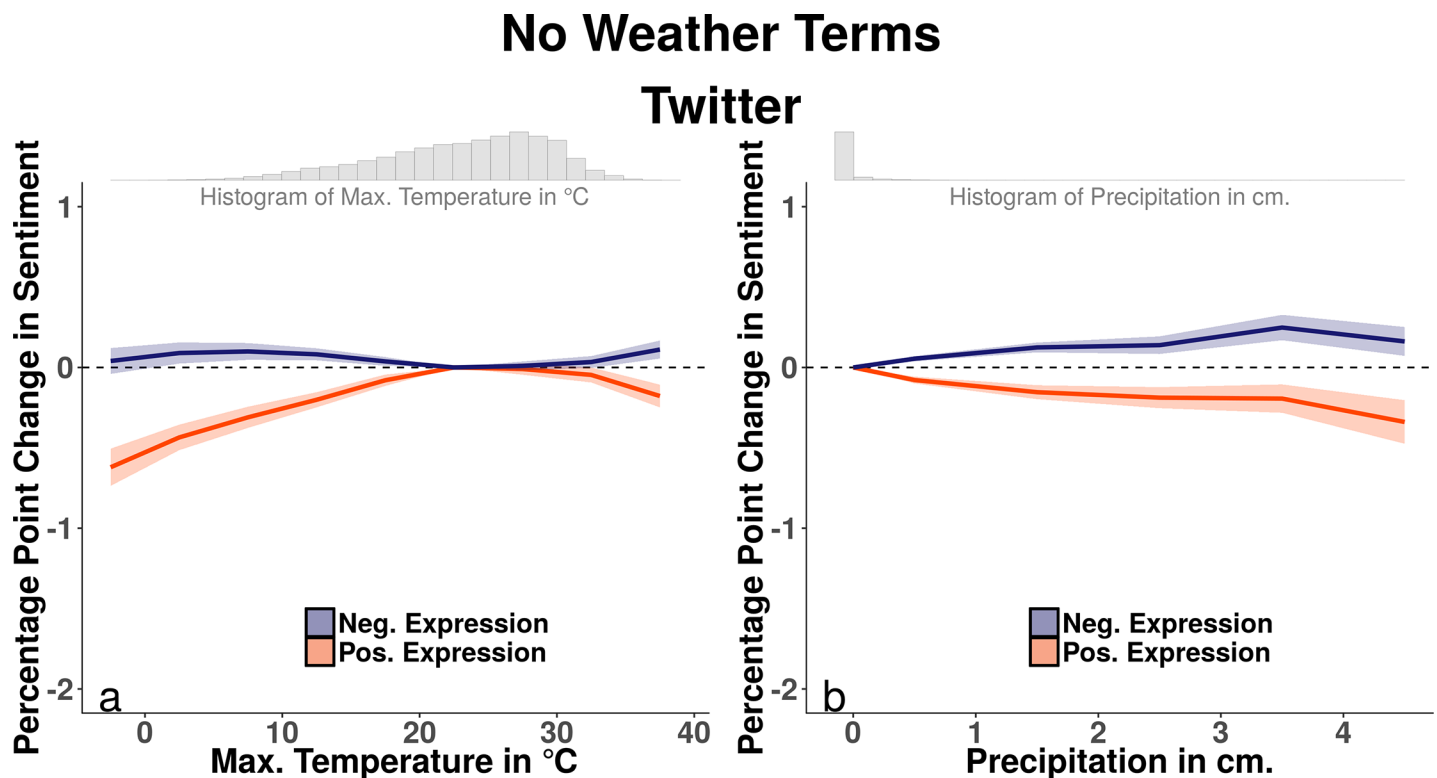


Fig 2. Twitter analyses for posts without weather terms. Panel (a) depicts the relationship between daily maximum temperatures and the rates of expressed sentiment for non-weather posts, aggregated to the city-level. It draws from the estimation of [Eq 1](#) and plots the predicted change in expressed sentiment associated with each maximum temperature bin. Panel (b) depicts the relationship between daily precipitation and the rates of sentiment expression of non-weather posts, also drawing on estimation of [Eq 1](#). Shaded error bounds represent 95% confidence intervals calculated using heteroskedasticity-robust standard errors clustered on both city-year-month and day.

<https://doi.org/10.1371/journal.pone.0195750.g002>




Fig 3. Effect sizes in context. Comparisons between the effect size of below freezing temperatures on positive, non-weather, expressed sentiment with the effect sizes of other locale-specific events over the course of our data on the same sentiment metric at the Twitter city-level. The effect size of freezing temperatures compares in magnitude to other significant events. Error bars represent 95% confidence intervals calculated using heteroskedasticity-robust standard errors clustered on both city-year-month and day.

<https://doi.org/10.1371/journal.pone.0195750.g003>

above the median latitude of our sample of cities 2) the annual anniversary of the September 11th terrorist attacks in the New York City metropolitan area 3) the December 2015 terrorist attacks in the San Bernardino metropolitan area 4) the October 2015 flooding in the Carolinas in the Charlotte metropolitan area (the term ‘flood’ and its derivatives are part of our excluded weather terms) and 5) the San Francisco Bay Area earthquake in August 2014 in the San Francisco/Oakland metropolitan area.

Fig 3 indicates that each of these events is significantly associated with reductions in non-weather, expressed positive sentiments in the areas local to the events. Further, a day of below freezing temperature in our sample is substantively meaningful. For example, the effect size of a day of below freezing temperature on positive expressed sentiment is 62% the effect size of the 2015 Carolina floods on aggregate expressed sentiment in Charlotte. We present the effect of these events on negative sentiment in Fig F in S1 File. Of important note, by conducting this comparison, we do not intend to comment on the overall impact of these events on society or on individual well-being. There may be many reasons why, for example, the San Bernardino shootings produced a less substantial reduction in positive expressed sentiment on Twitter than did the SF Bay Area earthquake. Speculation on the underlying factors that drive the association between these specific events and expressed sentiment is beyond the scope of this study.

Discussion

There are several considerations important to the interpretation of our results. First, while we have data on millions of individuals’ expressed sentiment as reflected by their social media

posts, optimal data would also include these individuals' daily self-reported emotional states. As mentioned above, while sentiment expressions on social media can be reflective of underlying emotions [17], the linguistic measures we employ here represent an imperfect and noisy proxy of emotional states. Future studies are needed to improve the psychometric validity of sentiment metrics.

Second, and relatedly, our chosen LIWC sentiment metrics may imperfectly measure the sentiment of expressions on social media. We examine the robustness of our findings to the use of other sentiment classification tools with our Twitter data in Figs G, H, I, J in [S1 File](#). In those analyses we employ both the SentiStrength and Hedonometer algorithms and find that our results are quite robust across all three of our employed sentiment metrics [33,34]. However, because all three of our metrics likely have idiosyncratic errors associated with them, our measurement of the sentiment of expressions remains imperfect. Future studies should endeavor to investigate the effect of meteorological conditions on specific changes in the distribution of the corpora of human lexical expressions [34–38].

Third, measurement error may exist between observed weather and the weather users actually experience, possibly attenuating the magnitude of our estimates [39]. Issues of right-hand-side measurement error may be particularly salient with respect to our measures of cloud cover and humidity, as they are derived from reanalysis data rather than directly from station observations [30]. Further, automated accounts in our Twitter data that may somewhat bias or attenuate our effect estimates for that platform. Future studies are needed to investigate the role of automated Twitter accounts in shaping overall sentiment on the platform.

Fourth, our analysis is conducted on individuals who self-select into participation in social media. Our results may not apply to demographics that do not use either Facebook or Twitter, such as older generations. Because the elderly are less common users of social media and because they may be particularly vulnerable to adverse weather conditions [32], the results we present here may underestimate the true population associations between weather and expressed sentiment.

Fifth, our data are restricted to observations from one country with a predominately temperate climate and with one of the highest rates of air conditioning prevalence in the world [40]. It is critical to repeat this analysis where possible in poorer countries with more extreme climates as they may see even greater alterations in overall expressed sentiment due to meteorological conditions.

Ultimately, given the ubiquity of our exposure to varying weather conditions, understanding the influence they may have on our emotional states is of high importance. Here we provide a window into this relationship via the measurement of expressed sentiment on social media. We observe that the weather is associated with statistically significant and substantively meaningful changes in expressed sentiment for posts both inclusive of and exclusive of weather-related terms. We find substantial evidence that less ideal weather conditions relate to worsened sentiment. To the extent that the sentiment of expressions serves as a valid proxy for underlying emotions, we find some observational evidence that the weather may functionally alter human emotional states.

Supporting information

S1 File. Supporting information file. Equation A in S1 File: User-level regression equation.

Fig A in S1 File: Twitter user-level analyses for all message types.

Fig B in S1 File: Twitter user-level analyses for posts without weather terms.

Fig C in S1 File: User-level comparisons between the effect size of below freezing temperatures on positive, non-weather, expressed sentiment with the effect sizes of other locale-specific

events over the course of our data on the same sentiment metric.

Fig D in S1 File: Effects of the weather on city-level frequency and user-level probability of weather speech.

Fig E in S1 File: Effects of the weather on expressed sentiment of posts that contain weather terms.

Fig F in S1 File: Effect sizes in context for negative sentiment.

Fig G in S1 File: Replication using SentiStrength classification.

Fig H in S1 File: Replication using SentiStrength classification.

Fig I in S1 File: Replication using Hedonometer classification.

Fig J in S1 File: Replication using Hedonometer classification.

Fig K in S1 File: Correlation between alternative positive sentiment state metrics.

Fig L in S1 File: Correlation between alternative positivenegative sentiment state metrics.

Table A in S1 File: Summary statistics of main dependent and independent variables.

Table B in S1 File: Examples of LIWC sentiment encoding.

Table C in S1 File: City-level weather and expressed sentiment, all posts.

Table D in S1 File: User-level weather and expressed sentiment, Twitter all posts.

Table E in S1 File: City-level weather and expressed sentiment, Twitter no weather terms.

Table F in S1 File: User-level weather and expressed sentiment, Twitter no weather terms. (PDF)

Acknowledgments

We thank the San Diego Supercomputer Center, Zachary Steinert-Threlkeld, Jason J. Jones for their assistance and thank Uri Simonsohn, Tal Yarkoni, and our anonymous reviewers for their thorough and helpful feedback.

Author Contributions

Conceptualization: Patrick Baylis, Nick Obradovich.

Data curation: Nick Obradovich, Yury Kryvasheyev, Haohui Chen, Lorenzo Coviello.

Formal analysis: Patrick Baylis, Nick Obradovich.

Funding acquisition: Nick Obradovich, Esteban Moro.

Investigation: Nick Obradovich.

Methodology: Patrick Baylis, Nick Obradovich.

Project administration: Nick Obradovich, Manuel Cebrian.

Writing – original draft: Nick Obradovich.

Writing – review & editing: Patrick Baylis, Nick Obradovich, Yury Kryvasheyev, Haohui Chen, Lorenzo Coviello, Esteban Moro, Manuel Cebrian, James H. Fowler.

References

1. Steptoe A, Wardle J, Marmot M. Positive affect and health-related neuroendocrine, cardiovascular, and inflammatory processes. *Proceedings of the National Academy of Sciences*. National Acad Sciences; 2005; 102: 6508–6512.
2. Nadler RT, Rabi R, Minda JP. Better mood and better performance learning rule-described categories is enhanced by positive mood. *Psychological Science*. Sage Publications; 2010; 21: 1770–1776. <https://doi.org/10.1177/0956797610387441> PMID: 20974709

3. Steptoe A, Dockray S, Wardle J. Positive affect and psychobiological processes relevant to health. *Journal of Personality*. Wiley Online Library; 2009; 77: 1747–1776. <https://doi.org/10.1111/j.1467-6494.2009.00599.x> PMID: 19796062
4. De Neve J-E, Oswald AJ. Estimating the influence of life satisfaction and positive affect on later income using sibling fixed effects. *Proceedings of the National Academy of Sciences*. National Acad Sciences; 2012; 109: 19953–19958.
5. Coviello L, Sohn Y, Kramer AD, Marlow C, Franceschetti M, Christakis NA, et al. Detecting emotional contagion in massive social networks. *PLOS ONE*. Public Library of Science; 2014; 9: e90315. <https://doi.org/10.1371/journal.pone.0090315> PMID: 24621792
6. Kramer AD, Guillory JE, Hancock JT. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*. National Acad Sciences; 2014; 201320040.
7. Persinger M. Lag responses in mood reports to changes in the weather matrix. *International Journal of Biometeorology*. Springer; 1975; 19: 108–114. PMID: 1165123
8. Schwarz N, Clore GL. Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*. American Psychological Association; 1983; 45: 513.
9. Denissen JJ, Butalid L, Penke L, Van Aken MA. The effects of weather on daily mood: A multilevel approach. *Emotion*. American Psychological Association; 2008; 8: 662. <https://doi.org/10.1037/a0013497> PMID: 18837616
10. Lucas RE, Lawless NM. Does life seem better on a sunny day? Examining the association between daily weather conditions and life satisfaction judgments. *Journal of Personality and Social Psychology*. American Psychological Association; 2013; 104: 872. <https://doi.org/10.1037/a0032124> PMID: 23607534
11. Li J, Wang X, Hovy E. What a nasty day: Exploring mood-weather relationship from Twitter. *ACM*; 2014; 1309–1318.
12. Klimstra TA, Frijns T, Keijsers L, Denissen JJ, Raaijmakers QA, Aken MA van, et al. Come rain or come shine: Individual differences in how weather affects mood. *Emotion*. American Psychological Association; 2011; 11: 1495. <https://doi.org/10.1037/a0024649> PMID: 21842988
13. Connolly M. Some like it mild and not too wet: The influence of weather on subjective well-being. *Journal of Happiness Studies*. Springer; 2013; 14: 457–473.
14. Hannak A, Anderson E, Barrett LF, Lehmann S, Mislove A, Riedewald M. Tweetin' in the rain: Exploring societal-scale effects of weather on mood. *ICWSM*. 2012.
15. Keller MC, Fredrickson BL, Ybarra O, Côté S, Johnson K, Mikels J, et al. A warm heart and a clear head the contingent effects of weather on mood and cognition. *Psychological Science*. SAGE Publications; 2005; 16: 724–731. <https://doi.org/10.1111/j.1467-9280.2005.01602.x> PMID: 16137259
16. Feddersen J, Metcalfe R, Wooden M. Subjective wellbeing: Why weather matters. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*. Wiley Online Library; 2016; 179: 203–228.
17. Settanni M, Marengo D. Sharing feelings online: Studying emotional well-being via automated text analysis of Facebook posts. *Frontiers in Psychology*. Frontiers Media SA; 2015; 6.
18. Baylis P. Temperature and temperament: Evidence from a billion tweets. *Energy Institute at Haas Working Paper*. 2015;
19. Greenwood S, Perrin A, Duggan M. *Social media update 2016*. Pew Research Center. 2016; 11.
20. Di Luzio M, Johnson GL, Daly C, Eischeid JK, Arnold JG. Constructing retrospective gridded daily precipitation and temperature datasets for the conterminous United States. *Journal of Applied Meteorology and Climatology*. 2008; 47: 475–497.
21. Kanamitsu M, Ebisuzaki W, Woollen J, Yang S-K, Hnilo J, Fiorino M, et al. NCEP–DOE AMIP-II reanalysis (R-2). *Bulletin of the American Meteorological Society*. 2002; 83.
22. Tausczik YR, Pennebaker JW. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*. Sage Publications; 2010; 29: 24–54.
23. Pennebaker JW, Chung CK, Ireland M, Gonzales A, Booth RJ. The development and psychometric properties of LIWC2007. 2007.
24. Golder SA, Macy MW. Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*. American Association for the Advancement of Science; 2011; 333: 1878–1881. <https://doi.org/10.1126/science.1202775> PMID: 21960633
25. Watson D, Clark LA, Tellegen A. Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*. American Psychological Association; 1988; 54: 1063. PMID: 3397865

26. Graff Zivin JS, Neidell M. Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*. JSTOR; 2014; 32: 1–26.
27. Obradovich N. Climate change may speed democratic turnover. *Climatic Change*. Springer Netherlands; 2017; 140: 135–147.
28. Obradovich N, Fowler JH. Climate change may alter human physical activity patterns. *Nature Human Behaviour*. Nature Publishing Group; 2017; 1: 0097.
29. Wooldridge JM. *Econometric analysis of cross section and panel data*. MIT press; 2010.
30. Hsiang S. Climate econometrics. *Annual Review of Resource Economics*. 2016; 8: 43–75.
31. Cameron AC, Gelbach JB, Miller DL. Robust inference with multiway clustering. *Journal of Business & Economic Statistics*. 2011; 29.
32. Obradovich N, Migliorini R, Mednick SC, Fowler JH. Nighttime temperature and human sleep loss in a changing climate. *Science Advances*. American Association for the Advancement of Science; 2017; 3: e1601555. <https://doi.org/10.1126/sciadv.1601555> PMID: 28560320
33. Thelwall M, Buckley K, Paltoglou G, Cai D, Kappas A. Sentiment strength detection in short informal text. *Journal of the American Society for Information Science and Technology*. Wiley Online Library; 2010; 61: 2544–2558.
34. Dodds PS, Clark EM, Desu S, Frank MR, Reagan AJ, Williams JR, et al. Human language reveals a universal positivity bias. *Proceedings of the National Academy of Sciences*. National Acad Sciences; 2015; 112: 2389–2394.
35. Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM. Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter. *PloS one*. Public Library of Science; 2011; 6: e26752. <https://doi.org/10.1371/journal.pone.0026752> PMID: 22163266
36. Mitchell L, Frank MR, Harris KD, Dodds PS, Danforth CM. The geography of happiness: Connecting twitter sentiment and expression, demographics, and objective characteristics of place. *PloS one*. Public Library of Science; 2013; 8: e64417. <https://doi.org/10.1371/journal.pone.0064417> PMID: 23734200
37. Alajajian SE, Williams JR, Reagan AJ, Alajajian SC, Frank MR, Mitchell L, et al. The lexico-calorimeter: Gauging public health through caloric input and output on social media. *PloS one*. Public Library of Science; 2017; 12: e0168893. <https://doi.org/10.1371/journal.pone.0168893> PMID: 28187216
38. Reagan AJ, Danforth CM, Tivnan B, Williams JR, Dodds PS. Sentiment analysis methods for understanding large-scale texts: A case for using continuum-scored words and word shift graphs. *EPJ Data Science*. Springer; 2017; 6: 28.
39. Hausman J. Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *Journal of Economic Perspectives*. JSTOR; 2001; 57–67.
40. Auffhammer M, Baylis P, Hausman CH. Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the United States. *Proceedings of the National Academy of Sciences*. National Acad Sciences; 2017; 201613193.