Emotion

Tweeting Under Uncertainty: The Relationship Between Uncertain Language and Negative Emotions in the wild

Marc-Lluís Vives, Daantje de Bruin, Jeroen M. van Baar, and Oriel FeldmanHall

Online First Publication, July 18, 2024. https://dx.doi.org/10.1037/emo0001376

AMERICAN PSYCHOLOGICAL ASSOCIATION

Emotion

ISSN: 1528-3542

https://doi.org/10.1037/emo0001376

Tweeting Under Uncertainty: The Relationship Between Uncertain Language and Negative Emotions in the wild

Marc-Lluís Vives[1], Daantje de Bruin[2], Jeroen M. van Baar[2], and Oriel FeldmanHall[2, 3]
[1] Institute of Psychology, Leiden University
[2] Department of Cognitive, Linguistic, and Psychological Sciences, Brown University
[3] Carney Institute for Brain Science, Brown University

Despite decades of research characterizing the relationship between uncertainty and emotion, little is known about how these constructs interact in the wild. Using naturalistic, large-scale language produced on Twitter, we ask whether increases in environmental uncertainty and associated aversive emotional reactions can be captured by the millions of digital traces of people sharing their thoughts online. Analyzing more than 20 million tweets from more than 7.5 million unique users, we find that uncertainty expressions peak when environmental uncertainty is high. This effect, however, is modulated by the type of trigger that increases uncertainty. Pandemics (COVID-19 in 2020) and national U.S. elections (2021) exhibit an increase in uncertainty language and negative sentiment in the real world, illustrating the well-documented relationship between uncertainty and aversive emotional reactions acting in lockstep. In contrast, when uncertain events involve a moral violation (i.e., the 2021 U.S. Capitol attack), specific negative emotions (i.e., anger, fear, and moral outrage) sharply increase, while uncertainty language abruptly decreases. This reveals that in the real world, uncertainty and emotion have a more complex relationship than originally assumed.

Keywords: Uncertainty, emotion, Twitter, social media, moral outrage

Supplemental materials: https://doi.org/10.1037/emo0001376.supp

A global pandemic, a highly disputed national election, and a violent attack on the seat of the government share a common feature: They are fraught with uncertainty. When will the pandemic end? Who will win the election? How far will the violence spread? These questions embody how global and national events create widespread uncertainty that can leave the public searching for answers.

Research aimed at understanding how people experience and respond to uncertain events typically operates on a smaller scale. Researchers recruit subjects and measure the amount of uncertainty and negative affect that an event elicits within an individual (De Martino et al., 2006; FeldmanHall et al., 2016; Sokol-Hessner et al., 2009). In these tightly controlled laboratory settings, it is well documented that uncertainty, measured in various ways (i.e., skin conductance, self-report, pupil dilation), is associated with negative affective experiences (Bar-Anan et al., 2009; Carleton, 2016; Critchley et al., 2001). Highly uncertain situations typically cause aversive responses, and for some, this can include pathological rumination and anxiety (Carleton, 2016; Hirsh et al., 2012; Maner et al., 2007). The intimate coupling between negative affect and uncertainty is now a well-known and well-researched topic (Anderson et al., 2019; Bar-Anan et al., 2009; Carleton, 2016; Morriss et al., 2023; van den Bos, 2009).

With the advent and widespread popularization of social media platforms, the manner in which people communicate has changed, making it possible to observe people's responses to uncertain events on an unprecedented scale. On platforms like Twitter, users spend hours organically sharing their thoughts online every hour of every day, providing researchers with unfettered access to millions of people's real-time reactions through the language they post online. These online platforms can also be mined to measure expressions of uncertainty language (Simchon et al., 2021), enabling scientists to

C. Daryl Cameron served as action editor.

Oriel FeldmanHall https://orcid.org/0000-0002-0726-3861

Oriel FeldmanHall received funding from Grant 2123469 from the National Science Foundation. Marc-Lluís Vives and Daantje de Bruin contributed equally to this article and share first authorship.

The experiment data are publicly available at TweetsKB or directly scraped off Twitter. The code used to analyze the data is available at https://osf.io/pbzyc/?view_only=6a9b45baf09f4d67a2f5a4694bab9a8f. The analyses pipeline was not preregistered.

Marc-Lluís Vives played a lead role in conceptualization, data curation, investigation, and visualization, a supporting role in project administration, and an equal role in formal analysis, methodology, writing–original draft, and writing–review and editing. Daantje de Bruin played a lead role in data curation, formal analysis, and methodology, a supporting role in investigation, writing–original draft, and writing–review and editing, and an equal role in conceptualization and visualization. Jeroen M. van Baar played a supporting role in data curation and formal analysis. Oriel FeldmanHall played a lead role in conceptualization, investigation, methodology, project administration, resources, supervision, validation, and writing–review and editing, a supporting role in formal analysis, and an equal role in writing–original draft.

Correspondence concerning this article should be addressed to Oriel FeldmanHall, Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, Box 1821, Metcalf Research Building, 190 Thayer Street, Providence, RI 02912, United States. Email: Oriel.Feldmanhall@Brown.edu

take a more naturalistic approach to understanding how people experience and respond to uncertain events.

We capitalize on this real-world, large-scale naturalistic experiment to assay human responses to changes in environmental uncertainty brought on by sudden, disruptive global events. We first ask whether the language used in brief online social media posts reflects and tracks these changes in environmental uncertainty. If we can demonstrate that social media platforms such as Twitter serve as a barometer for environmental uncertainty, we can interrogate important open psychological questions regarding the emotional experience of uncertainty, especially given that Twitter posts leave permanent psychological fingerprints that can be leveraged to measure how people feel in real-time—a method that has already been used to track people's mental health (Bathina et al., 2021; Reece et al., 2017).

To understand the relationship between uncertainty and emotional reactions in the wild, we measure emotion and uncertainty language contained in Twitter posts during sudden, large-scale uncertain events. Although our goal is to map the relationship between uncertainty and emotion in naturalistic contexts, the range of situations in which uncertain events arise suggests that this relationship between uncertainty and negative emotion might be more nuanced than originally assumed (Anderson et al., 2019; Bar-Anan et al., 2009; Carleton, 2016; DeSteno et al., 2000). On one hand, decades of laboratory studies find that people exhibit negatively valenced and arousing responses to decisions with uncertain outcomes (Brosschot et al., 2016; FeldmanHall et al., 2016; Grupe & Nitschke, 2013; Lovallo & Kahneman, 2000; Sokol-Hessner et al., 2009). On Twitter, this mapping should be reflected by increases in environmental uncertainty coinciding with increases in negative emotional language. For example, based on the well-known links between uncertainty and fear (Buhr & Dugas, 2009), increases of uncertainty expressions should also be associated with increases in fear language.

On the other hand, since the dawn of Twitter, there have been numerous uncertain historical events that engender far more acute and extreme negative emotions, such as anger and outrage (Brady et al., 2017). For instance, in the wake of a moral violation where there is uncertainty about what will happen next (e.g., insurrections, hate crimes), Twitter often explodes with moral outrage (Crockett, 2017). Unlike fear or surprise, moral outrage is largely driven by anger (Goodenough, 1997; Hoffman, 2000; Montada & Schneider, 1989), which is defined by a subjective appraisal of increased confidence and certainty that can potently motivate action or approach-like tendencies (Lerner & Tiedens, 2006). That feelings of anger can reduce uncertainty (Lerner & Keltner, 2001) by facilitating psychological closure (Nabi, 1999, 2002; Skitka et al., 2004) suggests that, in some cases, the relationship between uncertainty and some specific negative emotions—namely, anger and moral outrage—might be inverted. In other words, it is possible that the increased confidence and certainty stemming from feelings of anger and moral outrage (Lerner & Keltner, 2001) actually *attenuate* the experience of uncertainty. This would be reflected by increasing expressions of acute emotions correlating with *less* uncertainty language online.

To map the relationship between uncertainty and emotional reactions, we investigated three events that spanned just over a year: the beginning of COVID-19, the 2020 U.S. national elections, and the attack on the Capitol on January 6, 2021. We first investigated one of the most uncertain periods in recent history, the COVID-19 pandemic. Especially in its early stages in March 2020, COVID-19 was associated with a spike in anxiety (Heffner et al., 2021), which waned over time (Bavolar et al., 2023; Robinson et al., 2022). However, since the COVID-19 pandemic was a once-in-a-generation event, it is possible that observing a spike in uncertainty language on social media may not generalize to more conventional and repeated fluctuations of uncertainty. For example, national events such as the 2020 U.S. presidential election occur on a regular basis, yet the unpredictability of their outcomes qualifies them as uncertain (Bloom, 2014). The 2020 election in particular took an unprecedentedly long time to be called. Therefore, as a test of the boundary conditions of our finding, we analyzed tweets posted on the days leading up to the 2020 U.S. election until the election was ultimately called a week later.

We also investigated the window of time when the U.S. Capitol was under attack. Although violent uprisings are considered events that increase uncertainty about the immediate future (Bloom, 2014), they also carry a different psychological weight compared to elections and pandemics. Specifically, brief violent events (e.g., shootings, terrorist attacks, or the attack on the U.S. Capitol on January 6, 2021) are often construed as moral violations perpetrated by a particular individual or group. Moral violations typically elicit feelings of blame, anger, and the belief that the perpetrators deserve punishment (Salerno & Peter-Hagene, 2013). This suite of emotional responses is the defining feature of moral outrage (Batson et al., 2007, 2009; Haidt, 2003; Montada & Schneider, 1989). Given that feelings of moral outrage and anger are believed to increase confidence and certainty (Lerner & Tiedens, 2006), it is possible that acute emotional responses stemming from moral violations help to resolve the subjective feelings of uncertainty imposed by the environment.

Leveraging these naturally occurring events, in which large amounts of uncertainty were injected into the human experience on a national and global scale, we test (a) whether the language in posts on social media reflects drastic changes in environmental uncertainty; (b) whether increases in uncertainty language correlate with increases in general negative affect; and (c) whether acute and specific negative emotional experiences (i.e., anger/moral outrage) are associated with less uncertainty language used in online posts.

Method

Data Collection

The tweets used were either taken from the TweetsKB corpus (Fafalios et al., 2018) or scraped from Twitter using the Python package Twarc2. For COVID-19, we analyzed a random sample of the TweetsKB corpus (~15%) between February and April 2020 (data from 2019 were used as a contrast window), totaling N = 12,616,321 tweets. For the 2020 presidential election, we analyzed a random sample of the TweetsKB corpus (again, ~15%) in the week before and after election day (total N = 1,666,975). For the attack on the U.S. Capitol, we scraped 500,000 random tweets on January 6, 2021, on an hour-×-hour basis from 9 a.m. until 9 p.m. Eastern Standard Time (total N = 6,484,109), since TweetsKB did not have enough hourly temporal resolution. When scraping Twitter, we searched for tweets containing one of the 10 most common English words ("the," "of," "and," "a," "to," "in," "is," "you," "that," "it") while filtering out retweets. Using the Twarc2 package in Python, these tweets were archived as JavaScript Object Notation data. The jsonlines and Natural Language Toolkit packages in Python were used to convert the tweets into words and characters that can then be analyzed.

Windows of Analysis

Our goal was to measure naturalistic language on Twitter immediately before, during, and after major events believed to be perceived as uncertain. We selected these three events given that these types of events are linked to uncertainty surges (Bloom, 2014). However, given that the selected events unfolded at different timescales, we used the following rationales to define analysis windows, with the knowledge that we would include some time right before and after the uncertain event in order to establish a baseline, capture a trend change, and test whether uncertainty language returned to baseline after the uncertain event.

COVID-19

At the beginning of the pandemic, when information was just emerging about the virus, little was known about what the virus was, how it was transmitted, or in what geographic regions it was present. The first confirmed case in the United States was on January 20, 2020 (Holshue et al., 2020), although news of this only broke later. On February 11, the World Health Organization announced the official name of the disease, and on March 11, it declared a global pandemic (Taylor, 2020). This declaration ushered in a wave of "stay-at-home" orders, travel bans, sealed borders, and school closures spanning the globe. More than a month later, on April 24, certain states in the United States began to partially reopen. We focused on this early 2-month time frame because it was during this window that uncertainty was the greatest, and very little was known about (a) the disease, (b) what the symptoms were, (c) how to prevent infection, (d) what life would be like under the stay-at-home orders, and (e) how long life would continue under such duress. We stopped our analysis once the United States began to return to more normal operations. Our main analysis was therefore conducted between January 1, 2020, and April 30, 2020 (see Supplemental Material Results for a trend analysis conducted under other specifications).

2020 U.S. Election

We focused our analyses during the days leading up to the election (end of October) until the election was finally called a week later, on November 7 (Wikipedia Contributors, 2023a). We focused on this time frame given that November 1–2 was the deadline for early polling or mail-in ballots (respectively) in a number of key swing states (e.g., Florida, Vermont). While November 3 was election day, once the polls closed just after midnight, there was a lack of clarity over whether there was a clear victor, due to some media outlets trumpeting "election fraud and missing ballots." On November 4, Donald Trump announced he was filing lawsuits to halt the vote count in Michigan and Georgia, and requested a recount in Wisconsin. At 6 p.m. that day, the Associated Press projected that Joe Biden was six electoral votes shy of the necessary votes to win the election. On November 5, Biden urged people to "remain calm," while Trump continued to claim voter fraud, further adding to the uncertainty of the election outcome. Finally, on November 7, Pennsylvania called for Biden, placing Biden above the required 270 electoral votes to win the election. The window of analysis stopped once the election was called, which was before the accusations of a stolen election by Trump gained groundswell. Thus, our main analysis covered the time from October 29, 2020, to November 8, 2020 (see Supplemental Material Results for a trend analysis conducted under other specifications).

2021 U.S. Capitol Attack

Given that the events unfolded over just one afternoon on January 6, 2021, rather than using weeks or days, here we focused on an hour-×-hour time frame. At 11 a.m., a contingent of Proud Boys left the "March for Trump" rally and marched toward the Capitol building (Wikipedia Contributors, 2023b). At noon, Trump began to publicly allege that the election was stolen and claimed he would walk with the crowd to the Capitol. Close to 1 p.m., the first police barricade was breached, and at 2:10 p.m., the mob breached the final barricade and began to enter the building by smashing windows and opening doors with hammers. At 2:20 p.m., the House began to evacuate, while rioters continued to breach the Capitol. The National Guard was called in, and by 6:15 p.m., the police were able to establish a perimeter outside the U.S. Capitol. By 8 p.m., the police declared the Capitol building secure. The analysis therefore covered a window ranging from 9 a.m. to 9 p.m. on January 6, 2021 (see Supplemental Material Results for a trend analysis conducted under other specifications).

Because researchers face several forking paths in their analytical decisions that can lead to bias, we used other time windows to assess the degree to which the results varied across different time specifications (see Supplemental Material Results). Even with multiple other windows of analyses (i.e., seven other windows of time were analyzed), we observe similar results: Uncertainty language use spiked at the same moment for COVID-19 and the attack on the Capitol. The trend change was less robust for the 2020 national election.

Transparency and Openness

All applied methods and analyses pipelines have been detailed in the article. Data used for COVID-19 and 2020 national elections are readily available at https://data.gesis.org/tweetskb/. Due to Twitter/now X's policies around sharing the content of tweets, we cannot directly share tweets through a repository. See https://developer.twitter.com/en/developer-terms/agreement-and-policy for more legal information. Specific codes for analysis are available at https://osf.io/pbzyc/?view_only=6a9b45baf09f4d67a2f5a4694bab9a8f.

Data Analysis

General Sentiment

Sentiment analysis was conducted using the sentimentr package (Rinker, 2017) in RStudio, which was designed specifically to process polarity in large language data sets. This package considers valence shifters, that is, words that alter the semantic orientation of other words or sentences, including (de-)amplifiers and negators, which results in a more nuanced measurement of sentiment. Even though sentiment dictionaries are simple and straightforward methods for analyzing text, they can show similar performance when compared to more complex methods (Macanovic & Przepiorka, 2022), although see Rathje et al. (2023) for an additional review of how large language models can often outperform dictionary methods. For a more sustained discussion about the pros and cons of various unsupervised machine learning approaches, please see Supplemental Material. The sentiment analysis dictionary has 11,710 words that range between negatively valenced (−1; e.g., "abhor") and positively valenced (1; e.g., "admire"), where 0 reflects neutral valence (e.g., "it's like," "we're like"). For each tweet, general sentiment and emotion scores were computed by combining the sentiment and emotion scores of all words used in the tweet. For these analyses, the default dictionaries based on the lexicon package were implemented (sentiment: hash_sentiment_jockers_rinker, hash_valence_shifters; emotion: hash_nrc_emotions). To avoid an overlap between the sentiment dictionary and the uncertainty dictionary, we removed the words that were present in the uncertainty dictionary (see below), such that the sentiment dictionary totaled 11,637 words. General sentiment was assessed by looking at the unbounded polarity scores for each tweet returned by the sentiment_by() function, in which scores below 0 represent negative sentiment and scores above 0 denote positive sentiment. We then averaged these sentiment scores per day (COVID-19, 2020 U.S. national election) or per hour (2021 attack to the U.S. Capitol), resulting in an average measure of sentiment for each timepoint.

Specific Negative Emotions

Using the same package, we also analyzed two specific emotions commonly linked with uncertainty: anger (low uncertainty appraisal) and fear (high uncertainty appraisal; Roseman, 1984). To conduct this language analysis, we used the emotions_by() function from the sentimentr library. For each tweet, this function returns a score for "anger" and "fear" ranging from 0 (= no words in the tweet reflect fear or anger sentiment) to 1 (= all words in the tweet reflect fear or anger sentiment), which represents the proportion of the words that are associated with each emotion, taking into account valence shifters (e.g., "*un*happy" being treated as "not happy"). According to a review of sentiment computation methods in R, sentimentr is the package that most successfully accounts for negators (Naldi, 2019), which made it the package of choice for our work. We then computed an average score for each timepoint (day or hour, depending on the study) to assess changes in anger and fear expressions over time and their relationship with uncertainty language. To quantify moral outrage, we used the methods developed by Brady et al. (2021): We computed the probability of each tweet containing moral outrage sentiment as indexed by a supervised machine learning classifier (i.e., a pretrained deep gated recurrent unit), which was trained with tweets during events that sparked moral outrage (e.g., Brett Kavanaugh confirmation hearing). As before, we averaged sentiment per timepoint (day/hr) to assess changes of moral outrage over time and its relationship with uncertainty language.

Uncertainty Language

To compute uncertainty language within a specific time frame (e.g., 1 day or 1 hr), we used the sentimentr package in conjunction with a dictionary developed to measure uncertainty (Loughran & McDonald, 2015, available at https://sraf.nd.edu/loughranmcdonald-master-dictionary). There were 297 words classified as relating to uncertainty (e.g., uncertain, unlikely, improbable, etc.; from Loughran-McDonald's dictionary). An uncertainty sentiment score was computed for each tweet by indexing the degree to which uncertainty language was used in an unbounded fashion (where negative values denote low uncertainty and positive values denote high uncertainty). We then computed an average uncertainty score per day (COVID-19, 2020 U.S. national election) or hour (2021 attack to the U.S. Capitol), resulting in an average measure of uncertainty at each timepoint. This method was inspired by previous research showing that fluctuations in societal uncertainty can be captured by computing the frequency of occurrence of the word "uncertainty" in newspapers (Bloom, 2014).

Trend Analysis

To test for significant changes in uncertainty language over time, we used the trend package in Rstudio (Pohlert et al., 2016). We tested for a shift in word usage tendency with Lanzante's nonparametric procedure for single change-point detection with a Wilcoxon–Mann–Whitney test. By testing for a significant change-point in the central tendency of the time series, this procedure effectively identifies whether a significant trend change has occurred in the selected time period.

Correlational Analysis

To investigate the relationship between uncertainty, general sentiment, and specific emotions (anger, fear, and moral outrage) over time, we computed Pearson's correlations on the normalized data. However, since time-series data are autocorrelated, we used a phase randomization analysis (Theiler et al., 1992), which creates surrogate data with the same autocorrelated structure as the original time-series dataset, and then performs 10,000 permutations to test the degree to which the observed correlation is a consequence of the autocorrelated structure of the data (nltools package; Chang et al., 2019). If the correlation survives phase randomization testing, then the effect is assumed not to be caused by the autocorrelated structure.

Results

We first scraped more than 12 million tweets posted during February–April 2020, the beginning of the COVID-19 pandemic, and we contrasted the presence of uncertainty language during this period with the language of tweets written during the same period 1 year earlier (i.e., February–April of 2019; N = 12,616,321).

Results reveal a steep increase in uncertainty language during the first stages of the pandemic, as denoted by a trend analysis ($W = 172$, $p < .001$; Figure 1A). In particular, the specific word "uncertainty" rose dramatically, resulting in a significant spike that increased by a factor of seven ($W = 484$, $p < .001$; Figure 1B). This jump in uncertainty language on Twitter precedes by a few days the moment when the president of the United States officially recognized the pandemic on March 13, 2020. We next probed the relationship between emotion and uncertainty language. We found that increases in uncertainty language were significantly correlated with increases in negative sentiment ($r = 0.59$, $p < .001$; Figure 1C) and fear expressions ($r = 0.41$, $p = .003$; Figure 1C) during the same time period. We did not observe the same relationship with anger or

Figure 1

Uncertainty Language During the Early Stages of the COVID-19 Pandemic

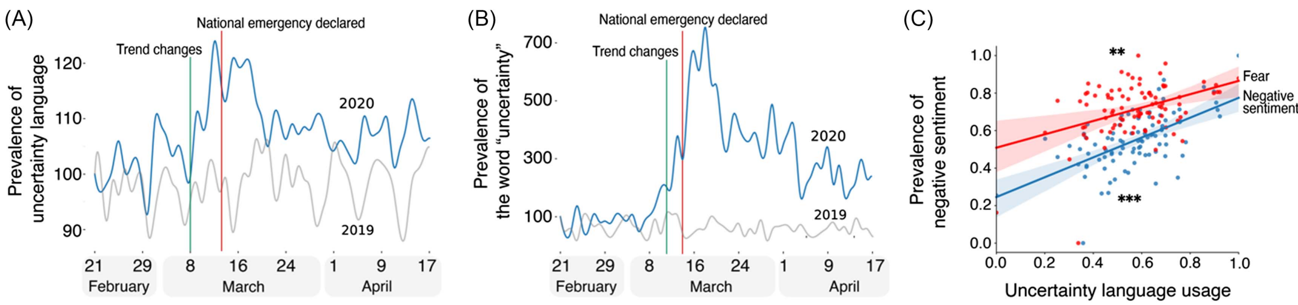

Note. (A) There was a sharp increase in uncertainty language on Twitter during the first few months of the pandemic. (B) This increase was especially significant for the word "uncertainty" and its variants ("uncertain"), which increased by a factor of seven. Uncertainty language and sentiment analysis data were normalized to 100 for the period within each year. (C) This uptick in uncertainty language was significantly associated with an increase in negative sentiment and fear expressed on Twitter. For ease of comparison, the plots in Panels A and B were indexed by taking all datapoints/first datapoint × 100, a common and effective means of normalizing data to a common starting point to see how variables change over time relative to one another. This is the case for all figures in the article. Data were normalized between 0 and 1 in Panel C. Only significant correlations are depicted. See the online article for the color version of this figure.
** $p < .01$. *** $p < .001$.

moral outrage (anger: $r = -0.13$, $p = .35$, moral outrage: $r = -0.44$, $p = .11$).

As a test of the boundary conditions of our first finding, we next analyzed tweets ($N = 1{,}230{,}593$) posted on the days leading up to the 2020 U.S. election until the election was ultimately called a week later. As with the pandemic, we observed that the language used on Twitter reflects the uncertainty of the 2020 national election. There was an increase in uncertainty language in the days before the election result was called (with a significant uptick in uncertainty identified 2 days before the election on 11/01, $W = 2$, $p = .02$; Figure 2A). Furthermore, we replicated the correlation between negative emotional sentiment and the prevalence of uncertainty language expressed on Twitter: Negative emotional language increased as the amount of uncertainty language online increased ($r = 0.80$, $p = .04$; Figure 2B). However, unlike with the pandemic, we did not observe that this relationship between uncertainty and negative sentiment was driven by specific expressions of fear, anger, or moral outrage—none of which robustly scaled with uncertainty language (fear: $r = 0.49$, $p = .17$; anger: $r = 0.76$, $p = .06$; moral outrage: $r = -0.70$, $p = .13$). We note that even though there are large coefficients, the correlations are nonsignificant. This is likely the result of using phase randomization (see the Method section) to control for the autocorrelational structure

Figure 2

Uncertainty Language on Twitter During the 2020 U.S. Presidential Elections

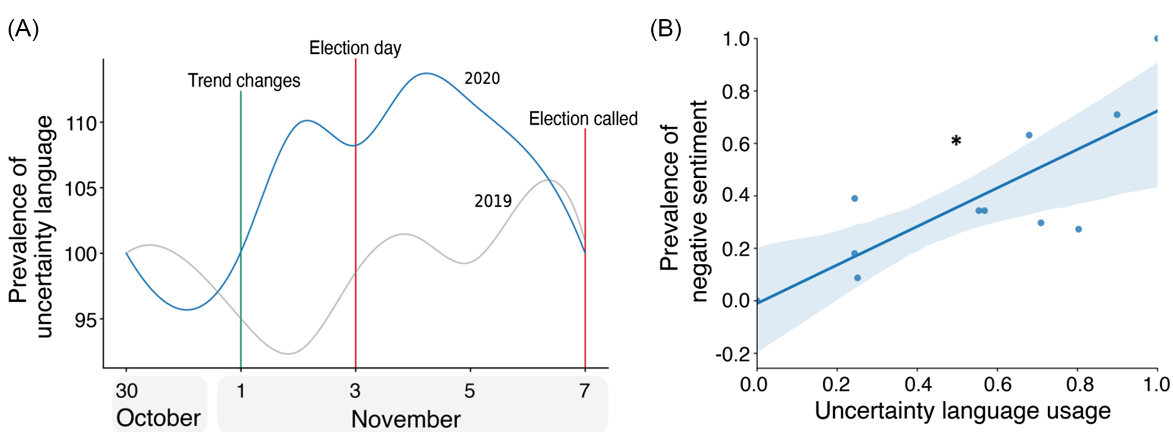

Note. (A) Uncertainty language on Twitter increased in the week leading up to the 2020 U.S. national election, peaking the day after the election but before the election was eventually called. Uncertainty language and sentiment analysis data were normalized to 100 for the period within each year. (B) This uptick in uncertainty language was associated with an increase in negative affective sentiment expressed on Twitter. Each datapoint represents 1 day during the time window. Data were indexed in Panel A using the first datapoint in the time window as a benchmark value. Data were normalized between 0 and 1 in Panel B. Only significant correlations are depicted. See the online article for the color version of this figure.
* $p < .05$.

of the data. In short, even during events that are perhaps less extraordinary and emotional than a worldwide pandemic, we still find that Twitter captures the experience of uncertainty stemming from a societally uncertain event, which increases together with negative affective sentiment.

Finally, we analyzed hour-×-hour tweets from 9 a.m. to 9 p.m. on the day of the Capitol attack ($N = 6{,}500{,}503$). We first confirmed that the attack on the Capitol was linked to greater moral outrage. A chi-square test comparing the prevalence of tweets containing moral outrage language during the attack on the Capitol, COVID-19, and 2020 national election confirmed our hypothesis (prevalence of moral outrage tweets during attack on the Capitol = 26.4%, COVID-19 = 17.9%, 2020 national election = 17.3%, x^2 = 152765.4, $p < .001$).

Although the fate of the Capitol was unknown, uncertainty language on Twitter did not increase when the attack started; rather, it substantially *decreased* by more than 20% (Figure 3A), a significant drop that started at 1 p.m. ($W = 40$, $p = .002$; Figure 3A). This drop in uncertainty language significantly correlated with an *increase* in negative sentiment: Even though people were expressing more negative emotions, these expressions were linked to a decrease in uncertainty language ($r = -0.88$, $p = .004$). This contrasts with the findings observed during the early stages of the pandemic and the 2020 U.S. national elections. More specifically, moral outrage language increased by almost 60% over the same time period (Figure 3B). Comparing uncertainty language and moral outrage sentiment revealed a negative relationship, such that in a given hour, the more moral outrage expressed on Twitter, the less uncertainty language expressed ($r = -0.88$, $p = .02$). The same relationship was found between uncertainty language and expressions of anger ($r = -0.83$, $p = .02$) and fear ($r = -0.91$, $p < .001$). This synchrony of escalating negative emotions was observed during the Capitol attack, a pattern we failed to find for COVID-19 and the 2020 U.S. national elections (see Supplemental Table 1). In short, the spike in moral outrage on Twitter surrounding the Capitol attack was associated with a steep reduction of uncertainty language and a sharp increase in negative emotional language—specifically fear, anger, and moral outrage.

Discussion

Although Twitter is often perceived as a forum for trivial daily musings and polarized content (Brady et al., 2017), it also reflects humanity's subjective experience of national and global events. Here, we demonstrate that people express uncertainty on Twitter in lockstep with major societal disruptions. Globally surprising events, such as the COVID-19 pandemic, increase the expression of uncertainty on Twitter by a large magnitude (700% when analyzing the use of the word "uncertainty"). Uncertainty expressions also increase, albeit to a lesser degree, for events that happen every few years but still have unpredictable outcomes, such as national elections. As expected, given a large body of theoretical work (Loewenstein et al., 2001), we found a strong positive relationship between uncertain events and expressions of negative affect. However, this relationship reverses for uncertain events that stoke more moral outrage, such as the 2021 attack on the U.S. Capitol. We observed nearly a 60% surge in language expressing feelings of moral outrage, anger, and fear during the 2021 insurrection. In parallel, the prevalence of uncertainty language decreased by 20%. What could explain this observed flip between negative emotional expression and uncertainty language? Feelings of moral outrage are typically expressed angrily and confidently, with little uncertainty (Hutcherson & Gross, 2011). Thus, it is possible that emotions arising from moral indignation are associated with a decrease in feelings of uncertainty.

Traditionally, the relationship between uncertainty and negative affect has been understood as linear: The more uncertainty increases, the more negative affect increases (Bar-Anan et al., 2009; FeldmanHall et al., 2016)—an effect illustrated many times over with the emotion of fear (Roseman, 1984; Rosen & Donley, 2006).

Figure 3

Uncertainty Language on Twitter During the January 6th Capitol Attack

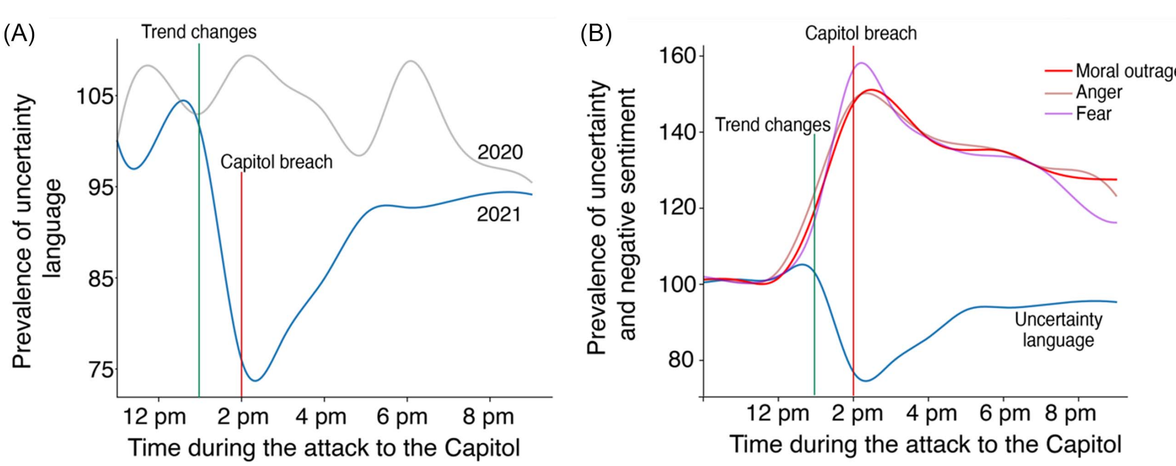

Note. (A) Uncertainty language on Twitter decreased during the 2021 attack on the U.S. Capitol. (B) In contrast, anger, fear, and moral outrage language increased by nearly 60%. The uncertainty and emotional language analyses were normalized to 100 for the period. Only the emotions that significantly correlate with uncertainty language are visualized. Data were indexed for ease of comparison, using the first datapoint in the time window as a benchmark value. See the online article for the color version of this figure.

While we replicate this effect in the wild with COVID-19, emotional appraisal theory argues that specific emotions like anger (and, by association, moral outrage) are associated with a sense of certainty and perceptions of control (Bar-Anan et al., 2009; Lerner & Keltner, 2001). Thus, feelings of anger should theoretically decrease the subjective experience of uncertainty instead of increasing it. While some data in the laboratory documents this relationship (Roseman, 1984; Rosen & Donley, 2006), we only observed this asymmetry between fear, anger, and uncertainty during a global pandemic. In contrast, when the situation is defined by a clear moral violation, such as the violent attack on the U.S. Capitol, anger, fear, and moral outrage all scaled together and in direct opposition to the amount of uncertainty language expressed on Twitter. In short, it appears that the relationship between uncertainty and negative affect might hinge on context more than on the emotion experienced.

It is possible that our findings are not explained by the specificity of the emotion (i.e., anger vs. fear) but instead are caused by emotional extremity, such that acute, extreme emotions are associated with perceptions of certainty and control. If this were the case, then any extreme emotion may dampen the amount of uncertainty experienced. This would mean that the relationship between emotion and uncertainty is less clear-cut than originally presumed, and may not be linked to specific negative emotions or even linear in nature. To put it succinctly, increasing uncertainty may scale with negative affect up to a certain point. Once the emotion becomes sufficiently extreme, it may correlate with decreasing amounts of uncertainty. Future work can explore whether there is a dynamic relationship between emotion extremity and the subjective experience of uncertainty.

Constraints on Generality

Analyzing language on Twitter is a powerful tool to circumvent generalizability limitations from laboratory experiments conducted solely with university students, as we can sample the thoughts and reactions of millions of people in real time as major societal events unfold. Our findings, however, are constrained to the specificity of the events investigated, as they are extraordinary, and are unlikely to repeat in the exact manner in the future. It is crucial, then, that future research determine the fundamental psychological components that an event must trigger to reproduce the effects observed here. Furthermore, users who expressed moral outrage and strong emotions during the attack to the Capitol on Twitter may have been more politically engaged than the average person, which, if true, might limit the generalizability of these findings. As the field is rapidly progressing in its development of sophisticated large language models to analyze text, future research can elucidate to what degree these tools iteratively improve the estimation of societal uncertainty and its relationship with emotion and affect.

References

Anderson, E. C., Carleton, R. N., Diefenbach, M., & Han, P. K. J. (2019). The relationship between uncertainty and affect. *Frontiers in Psychology*, *10*, Article 2504. https://doi.org/10.3389/fpsyg.2019.02504

Bar-Anan, Y., Wilson, T. D., & Gilbert, D. T. (2009). The feeling of uncertainty intensifies affective reactions. *Emotion*, *9*(1), 123–127. https://doi.org/10.1037/a0014607

Bathina, K. C., Ten Thij, M., Lorenzo-Luaces, L., Rutter, L. A., & Bollen, J. (2021). Individuals with depression express more distorted thinking on social media. *Nature Human Behaviour*, *5*(4), 458–466. https://doi.org/10.1038/s41562-021-01050-7

Batson, C. D., Chao, M. C., & Givens, J. M. (2009). Pursuing moral outrage: Anger at torture. *Journal of Experimental Social Psychology*, *45*(1), 155–160. https://doi.org/10.1016/j.jesp.2008.07.017

Batson, C. D., Kennedy, C. L., Nord, L., Stocks, E. L., Fleming, D. A., Marzette, C. M., Lishner, D. A., Hayes, R. E., Kolchinsky, L. M., & Zerger, T. (2007). Anger at unfairness: Is it moral outrage? *European Journal of Social Psychology*, *37*(6), 1272–1285. https://doi.org/10.1002/ejsp.434

Bavolar, J., Kacmar, P., Hricova, M., Schrötter, J., Kovacova-Holevova, B., Köverova, M., & Raczova, B. (2023). Intolerance of uncertainty and reactions to the COVID-19 pandemic. *Journal of General Psychology*, *150*(2), 143–170. https://doi.org/10.1080/00221309.2021.1922346

Bloom, N. (2014). Fluctuations in uncertainty. *The Journal of Economic Perspectives*, *28*(2), 153–176. https://doi.org/10.1257/jep.28.2.153

Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. *Science Advances*, *7*(33), Article eabe5641. https://doi.org/10.1126/sciadv.abe5641

Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(28), 7313–7318. https://doi.org/10.1073/pnas.1618923114

Brosschot, J. F., Verkuil, B., & Thayer, J. F. (2016). The default response to uncertainty and the importance of perceived safety in anxiety and stress: An evolution-theoretical perspective. *Journal of Anxiety Disorders*, *41*, 22–34. https://doi.org/10.1016/j.janxdis.2016.04.012

Buhr, K., & Dugas, M. J. (2009). The role of fear of anxiety and intolerance of uncertainty in worry: An experimental manipulation. *Behaviour Research and Therapy*, *47*(3), 215–223. https://doi.org/10.1016/j.brat.2008.12.004

Carleton, R. N. (2016). Into the unknown: A review and synthesis of contemporary models involving uncertainty. *Journal of Anxiety Disorders*, *39*, 30–43. https://doi.org/10.1016/j.janxdis.2016.02.007

Chang, L., Sam, Jolly, E., Cheong, J. H., Burnashev, A., Chen, A., & Frey, S. (2019). *cosanlab/nltools: 0.3.14*. https://doi.org/10.5281/ZENODO.3251172

Critchley, H. D., Mathias, C. J., & Dolan, R. J. (2001). Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron*, *29*(2), 537–545. https://doi.org/10.1016/S0896-6273(01)00225-2

Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, *1*(11), 769–771. https://doi.org/10.1038/s41562-017-0213-3

De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, *313*(5787), 684–687. https://doi.org/10.1126/science.1128356

DeSteno, D., Petty, R. E., Wegener, D. T., & Rucker, D. D. (2000). Beyond valence in the perception of likelihood: The role of emotion specificity. *Journal of Personality and Social Psychology*, *78*(3), 397–416. https://doi.org/10.1037/0022-3514.78.3.397

Fafalios, P., Iosifidis, V., Ntoutsi, E., & Dietze, S. (2018). TweetsKB: A Public and Large-Scale RDF corpus of annotated tweets. In G. Goos (Ed.), *The semantic web. ESWC 2018. Lecture notes in computer science* (pp. 177–190). Springer. https://doi.org/10.1007/978-3-319-93417-4_12

FeldmanHall, O., Glimcher, P., Baker, A. L., & Phelps, E. A. (2016). Emotion and decision-making under uncertainty: Physiological arousal predicts increased gambling during ambiguity but not risk. *Journal of Experimental Psychology: General*, *145*(10), 1255–1262. https://doi.org/10.1037/xge0000205

Goodenough, W. H. (1997). Moral outrage: Territoriality in human guise. *Zygon*, *32*(1), 5–27. https://doi.org/10.1111/0591-2385.671997067

Grupe, D. W., & Nitschke, J. B. (2013). Uncertainty and anticipation in anxiety: An integrated neurobiological and psychological perspective.

Nature Reviews Neuroscience, *14*(7), 488–501. https://doi.org/10.1038/nrn3524

Haidt, J. (2003). The moral emotions. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 852–870). Oxford University Press.

Heffner, J., Vives, M.-L., & FeldmanHall, O. (2021). Anxiety, gender, and social media consumption predict COVID-19 emotional distress. *Humanities and Social Sciences Communications*, *8*(1), Article 140. https://doi.org/10.1057/s41599-021-00816-8

Hirsh, J. B., Mar, R. A., & Peterson, J. B. (2012). Psychological entropy: A framework for understanding uncertainty-related anxiety. *Psychological Review*, *119*(2), 304–320. https://doi.org/10.1037/a0026767

Hoffman, M. L. (2000). *Empathy and moral development: Implications for caring and justice*. Cambridge University Press. https://doi.org/10.1017/CBO9780511805851

Holshue, M. L., DeBolt, C., Lindquist, S., Lofy, K. H., Wiesman, J., Bruce, H., Spitters, C., Ericson, K., Wilkerson, S., Tural, A., Diaz, G., Cohn, A., Fox, L., Patel, A., Gerber, S. I., Kim, L., Tong, S., Lu, X., Lindstrom, S., … Pillai, S. K. (2020). First case of 2019 novel coronavirus in the United States. *The New England Journal of Medicine*, *382*(10), 929–936. https://doi.org/10.1056/NEJMoa2001191

Hutcherson, C. A., & Gross, J. J. (2011). The moral emotions: A social-functionalist account of anger, disgust, and contempt. *Journal of Personality and Social Psychology*, *100*(4), 719–737. https://doi.org/10.1037/a0022408

Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of Personality and Social Psychology*, *81*(1), 146–159. https://doi.org/10.1037/0022-3514.81.1.146

Lerner, J. S., & Tiedens, L. Z. (2006). Portrait of the angry decision maker: How appraisal tendencies shape Anger's influence on cognition. *Journal of Behavioral Decision Making*, *19*(2), 115–137. https://doi.org/10.1002/bdm.515

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, *127*(2), 267–286. https://doi.org/10.1037/0033-2909.127.2.267

Loughran, T., & McDonald, B. (2015). The use of word lists in textual analysis. *Journal of Behavioral Finance*, *16*(1), 1–11. https://doi.org/10.1080/15427560.2015.1000335

Lovallo, D., & Kahneman, D. (2000). Living with uncertainty: Attractiveness and resolution timing. *Journal of Behavioral Decision Making*, *13*(2), 179–190. https://doi.org/10.1002/(SICI)1099-0771(200004/06)13:2<179::AID-BDM332>3.0.CO;2-J

Macanovic, A., & Przepiorka, W. (2022). *A systematic evaluation of text mining methods for short texts: Mapping individuals' internal states from online posts*. https://doi.org/10.31235/osf.io/cerzs

Maner, J. K., Richey, J. A., Cromer, K., Mallott, M., Lejuez, C. W., Joiner, T. E., & Schmidt, N. B. (2007). Dispositional anxiety and risk-avoidant decision-making. *Personality and Individual Differences*, *42*(4), 665–675. https://doi.org/10.1016/j.paid.2006.08.016

Montada, L., & Schneider, A. (1989). Justice and emotional reactions to the disadvantaged. *Social Justice Research*, *3*(4), 313–344. https://doi.org/10.1007/BF01048081

Morriss, J., Goh, K., Hirsch, C. R., & Dodd, H. F. (2023). Intolerance of uncertainty heightens negative emotional states and dampens positive emotional states. *Frontiers in Psychiatry*, *14*, Article 1147970. https://doi.org/10.3389/fpsyt.2023.1147970

Nabi, R. L. (2002). Anger, fear, uncertainty, and attitudes: A test of the cognitive-functional model. *Communication Monographs*, *69*(3), 204–216. https://doi.org/10.1080/03637750216541

Nabi, R. L. (1999). A cognitive-functional model for the effects of discrete negative emotions on information processing, attitude change, and recall. *Communication Theory*, *9*(3), 292–320. https://doi.org/10.1111/j.1468-2885.1999.tb00172.x

Naldi, M. (2019). *A review of sentiment computation methods with R packages*. ArXiv.

Pohlert, T., Pohlert, M. T., & Kendall, S. (2016). *Non-parametric trend tests and change-point detection*.

Rathje, S., Mirea, D.-M., Sucholutsky, I., Marjieh, R., Robertson, C., & Van Bavel, J. J. (2023). *GPT is an effective tool for multilingual psychological text analysis*. PsyArXiv. https://doi.org/10.31234/osf.io/sekf5

Reece, A. G., Reagan, A. J., Lix, K. L. M., Dodds, P. S., Danforth, C. M., & Langer, E. J. (2017). Forecasting the onset and course of mental illness with Twitter data. *Scientific Reports*, *7*(1), Article 13006. https://doi.org/10.1038/s41598-017-12961-9

Rinker, T. (2017). *Package 'sentimentr.'*. https://cran.fhcrc.org/web/packages/sentimentr/sentimentr.pdf

Robinson, E., Sutin, A. R., Daly, M., & Jones, A. (2022). A systematic review and meta-analysis of longitudinal cohort studies comparing mental health before versus during the COVID-19 pandemic in 2020. *Journal of Affective Disorders*, *296*, 567–576. https://doi.org/10.1016/j.jad.2021.09.098

Roseman, I. J. (1984). Cognitive determinants of emotions: A structural theory. In P. Shaver (Ed.), *Review of Personality and Social Psychology* (Vol. 5, pp. 11–36). Sage Publications.

Rosen, J. B., & Donley, M. P. (2006). Animal studies of amygdala function in fear and uncertainty: Relevance to human research. *Biological Psychology*, *73*(1), 49–60. https://doi.org/10.1016/j.biopsycho.2006.01.007

Salerno, J. M., & Peter-Hagene, L. C. (2013). The interactive effect of anger and disgust on moral outrage and judgments. *Psychological Science*, *24*(10), 2069–2078. https://doi.org/10.1177/0956797613486988

Simchon, A., Turkin, C., Svoray, T., Kloog, I., Dorman, M., & Gilead, M. (2021). Beyond doubt in a dangerous world: The effect of existential threats on the certitude of societal discourse. *Journal of Experimental Social Psychology*, *97*, Article 104221. https://doi.org/10.1016/j.jesp.2021.104221

Skitka, L. J., Bauman, C. W., & Mullen, E. (2004). Political tolerance and coming to psychological closure following the September 11, 2001, terrorist attacks: An integrative approach. *Personality and Social Psychology Bulletin*, *30*(6), 743–756. https://doi.org/10.1177/0146167204263968

Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(13), 5035–5040. https://doi.org/10.1073/pnas.0806761106

Taylor, D. B. (2020). A timeline of the coronavirus pandemic. *The New York Times*.

Theiler, J., Eubank, S., Longtin, A., Galdrikian, B., & Farmer, J. D. (1992). Testing for nonlinearity in time series: The method of surrogate data. *Physica D: Nonlinear Phenomena*, *58*(1–4), 77–94. https://doi.org/10.1016/0167-2789(92)90102-S

van den Bos, K. (2009). Making sense of life: The existential self trying to deal with personal uncertainty. *Psychological Inquiry*, *20*(4), 197–217. https://doi.org/10.1080/10478400903333411

Wikipedia Contributors. (2023a). 2020 United States presidential election. In *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/wiki/2020_United_States_presidential_election

Wikipedia Contributors. (2023b). January 6 United States Capitol attack. In *Wikipedia, The Free Encyclopedia*. https://en.wikipedia.org/wiki/January_6_United_States_Capitol_attack

Received July 27, 2023
Revision received February 12, 2024
Accepted February 21, 2024 ■