

# HOW EPIDEMIC PSYCHOLOGY WORKS ON TWITTER: EVOLUTION OF RESPONSES TO THE COVID-19 PANDEMIC IN THE U.S.

LUCA MARIA AIELLO<sup>1,2</sup>, DANIELE QUERCIA<sup>1,3,\*</sup>, KE ZHOU<sup>1</sup>, MARIOS CONSTANTINIDES<sup>1</sup>, SANJA SCEPANOVIC<sup>1</sup>,  
AND SAGAR JOGLEKAR<sup>1</sup>

ABSTRACT. Disruptions resulting from an epidemic might often appear to amount to chaos but, in reality, can be understood in a systematic way through the lens of “epidemic psychology”. According to Philip Strong, the founder of the sociological study of epidemic infectious diseases, not only is an epidemic biological; there is also the potential for three psycho-social epidemics: of fear, moralization, and action. This work empirically tests Strong’s model at scale by studying the use of language of 122M tweets related to the COVID-19 pandemic posted in the U.S. during the whole year of 2020. On Twitter, we identified three distinct phases. Each of them is characterized by different regimes of the three psycho-social epidemics. In the refusal phase, users refused to accept reality despite the increasing number of deaths in other countries. In the anger phase (started after the announcement of the first death in the country), users’ fear translated into anger about the looming feeling that things were about to change. Finally, in the acceptance phase, which began after the authorities imposed physical-distancing measures, users settled into a “new normal” for their daily activities. Overall, refusal of accepting reality gradually died off as the year went on, while acceptance increasingly took hold. During 2020, as cases surged in waves, so did anger, re-emerging cyclically at each wave. Our real-time operationalization of Strong’s model is designed in a way that makes it possible to embed epidemic psychology into real-time models (e.g., epidemiological and mobility models).

## INTRODUCTION

In our daily lives, our dominant perception is of order. But every now and then chaos threatens that order: epidemics dramatically break out, revolutions erupt, empires suddenly fall, and stock markets crash. Epidemics, in particular, present not only collective health hazards but also special challenges to mental health and public order that need to be addressed by social and behavioral sciences [Van Bavel et al., 2020]. Almost 30 years ago, in the wake of the AIDS epidemic, Philip Strong, the founder of the sociological study of epidemic infectious diseases, reflected: “*the human origin of epidemic psychology lies not so much in our unruly passions as in the threat of epidemic disease to our everyday assumptions.*” [Strong, 1990] In the recent COVID-19 pandemic [Brooks et al., 2020] (an ongoing pandemic of a coronavirus disease), it has been shown that the main source of uncertainty and anxiety has indeed come from the disruption of what Alfred Shutz called the “routines and recipes” of daily life [Schutz et al., 1973] (e.g., every simple act, from eating at work to visiting our parents, takes on new meanings).

Yet, the chaos resulting from an epidemic turns out to be more predictable than what one would initially expect. Philip Strong observed that any new *health* epidemic resulted into three *psycho-social* epidemics: of fear, moralization, and action. The epidemic of fear represents the fear of catching the disease, which comes with the suspicion against alleged disease carriers, which, in turn, may spark panic and irrational behavior. The epidemic of moralization is characterized by moral responses both to the viral epidemic itself and to the epidemic of fear, which may result in either positive reactions (e.g., cooperation) or negative ones (e.g., stigmatization). The epidemic of action accounts for the rational or irrational changes of daily habits that people make in response to the disease or as a result of the two other psycho-social epidemics. Strong was writing in the wake of the AIDS/HIV crisis, but he based his model on studies that went back to Europe’s Black Death in the 14<sup>th</sup> century. Importantly, he showed that these three psycho-social epidemics are created by language and incrementally fed through it: language transmits the *fear* that the infection is an existential threat to humanity and that we are all going to die; language depicts the epidemic as a verdict on human failings and as a *moral* judgment on minorities; and language shapes the means through which people collectively intend to, however pointless, *act* against the threat.

There have been numerous studies of how information propagated on social media during epidemic outbreaks occurred in the past decade, such as Zika [Fu et al., 2016, Wood, 2018, Sommariva et al., 2018], Ebola [Oyeyemi

---

<sup>1</sup>NOKIA BELL LABS CAMBRIDGE, UNITED KINGDOM

<sup>2</sup>IT UNIVERSITY COPENHAGEN, DENMARK

<sup>3</sup>CENTRE FOR URBAN SCIENCE AND PROGRESS, KING’S COLLEGE LONDON, UNITED KINGDOM

\*CORRESPONDING AUTHOR: QUERCIA@CANTAB.NET

et al., 2014], and the H1N1 influenza [Chew and Eysenbach, 2010]. Similarly, with the outbreak of COVID-19, people around the world have been collectively expressing their thoughts and concerns about the pandemic on social media. As such, researchers studied this epidemic from multiple angles: social media posts have been analyzed in terms of content and behavioral markers [Hou et al., 2020, Li et al., 2020], and of tracking the diffusion of COVID-related information [Cinelli et al., 2020] and misinformation [Pulido et al., 2020, Ferrara, 2020, Kouzy et al., 2020, Yang et al., 2020]. Search queries have suggested specific information-seeking responses to the pandemic [Bento et al., 2020]. Psychological responses to COVID-19 have been studied mostly through surveys [Wang et al., 2020, Qiu et al., 2020].

Hitherto, there has never been any large-scale empirical study of whether the use of language during an epidemic reflects Strong’s model. With the opportunity of having a sufficiently large-scale data at hand, we set out to test whether Strong’s model did hold on Twitter during COVID-19, and did so at the scale of an entire country, that of the United States. By running our study on Twitter, we expose the interpretation of our results to a number of limitations, notably to issues of representativeness [Li et al., 2013], self-presentation biases [Waterloo et al., 2018], and data noise [Saif et al., 2012, Ferrara et al., 2016]. Indeed, recent surveys estimated that only 22% of U.S. adults use Twitter [Wojcik and Hughes, 2019], and that the characteristics of these users deviate from the general population’s: compared to the average adult in the U.S., Twitter users are much younger, are more likely to have college degrees, and are slightly more likely to identify with the Democratic Party. Despite such limitations, social media represents the response of a relevant part of the general population to global events, and it does so at a scale and a granularity that have been so far unattainable by publicly-available data sources.

After operationalizing Strong’s model by using lexicons for psycholinguistic text analysis, upon the collection of 122M tweets about the epidemic from February 1<sup>st</sup> to December 31<sup>st</sup>, we conducted a quantitative analysis on the differences in language style and a thematic analysis of the actual social media posts. The temporal scope of our study does not capture the full pandemic as it is still ongoing at the time of writing but it includes the three major contagion waves in the U.S., characterizing the entire year of 2020. The first wave captured the initial diffusion of the virus in the world, and then its arrival and first peak in the U.S. The following two waves captured subsequent periods of alarming diffusion.

The three psycho-social epidemics, as theorized by Strong, evolve concurrently over time. In the particular case of our period of study, we found that this concurrent evolution resulted into three regimes or phases, which are not part of Strong’s theoretical framework and experimentally emerged. In the first phase (the *refusal* phase), the psycho-social epidemic of fear began. Twitter users in the U.S. refused to accept reality: they feared the uncertainty created by the disruption of what was considered to be “normal”; focused their moral concerns on others in an act of distancing oneself from others; yet, despite all this, they refused to change the normal course of action. After the announcement of the first death in the country, the second phase (the *anger* phase) began: the psycho-social epidemic of fear intensified while the epidemics of morality and action kicked-off abruptly. Twitter users expressed more anger than fear about the looming feeling that things were about to change; focused their moral concerns on oneself in an act of reckoning with what was happening; and suspended their daily activities. After the authorities imposed physical-distancing measures, the third phase (the *acceptance* phase) took over: the epidemic of fear started to fade away while the epidemics of morality and action turned into more constructive and forward-looking social processes. Twitter users expressed more sadness than anger or fear; focused their moral concerns on the collective and, in so doing, promoted pro-social behavior; and found a “new normal” in their daily activities, which consisted of their past daily activities being physically restricted to their homes and neighborhoods. The phase of *acceptance* dominated Twitter conversations for the rest of the year, although the *anger* phase re-emerged cyclically with the rise of new waves of contagion. In particular, we observed two peaks of anger: when the death toll in the U.S. reached 100,000 people (second contagion wave), and when President Trump tested positive to COVID-19 (third contagion wave).

## DATASET

From an existing collection of COVID-related tweets [Chen et al., 2020], we gathered 554,941,519 tweets posted between February 1<sup>st</sup> up to December 31<sup>st</sup>. We focused our analysis on the United States, the country where Twitter penetration is highest. To identify Twitter users living in it, we parsed the free-text location description of their user profile (e.g., “San Francisco, CA”). We did so by using a set of custom regular expressions that match variations for the expression “United States of America”, as well as the names of 333 cities and 51 states in the U.S. (and their combinations). Albeit not always accurate, matching location strings against known location names is a tested approach that yields good results for a coarse-grained localization at state or country-level [Dredze et al., 2013]. Overall, we were left with 6,271,835 unique users in the U.S. who posted 122,320,155 tweets in English.

Before analyzing how our language categories unfolded over time, we experimentally tested whether the number of data points at hand was sufficient to compute our metrics. It was indeed the case as the number of active users per day varied from a minimum of 72k on February 2<sup>nd</sup> to a maximum of 1.84M on March 18<sup>th</sup>, with an average of 437k. A small number of accounts tweeted a disproportionately high number of times, reaching a maximum of 15,823 tweets; those were clearly automated accounts, which were discarded by our methodology. As we shall discuss in Methods, we normalized all our aggregate temporal measures so that they were not affected by the fluctuating volume of tweets over time.

## METHODS

**Coding Strong’s model.** Back in the 1990s, Philip Strong was able not only to describe the psychological impact of epidemics on social order but also to model it. He observed that the early reaction to major fatal epidemics is a distinctive psycho-social form and can be modeled along three main dimensions: fear, morality, and action. During a large-scale epidemic, basic assumptions about social interaction and, more generally, about social order are disrupted, and, more specifically, they are so by: the fear of others, competing moralities, and the responses to the epidemic. Crucially, all these three elements are created, transmitted, and mediated by language: language transmits fears, elaborates on the stigmatization of minorities, and shapes the means through which people collectively respond to the epidemic [Strong, 1990, Goffman, 2009, Cap, 2016].

As opposed to existing attempts to model psychological and social aspects of epidemic crises [McConnell, 2005, Khan and Huremović, 2019], Strong’s model meets our three main choice criteria:

- (1) *Well-grounded.* It proposes a comprehensive, highly-cited, and still relevant theoretical model, which is based on an extensive review of studies of past large-scale epidemics that did span centuries and, as such, were of different nature, speaking to the generalizability of the framework.
- (2) *Focused on psycho-social aspects.* Its main goal is to characterize people’s psychological and social responses to epidemics rather than describing how an epidemic unfolds over time.
- (3) *Directly operationalizable from language use.* The description of the psycho-social responses provided by Strong lends itself to operationalization, as its defining concepts have been mapped to language markers by previous literature (as Table 1 shows).

We operationalized Strong’s epidemic psychology theoretical framework in two steps. First, three authors hand-coded Strong’s seminal paper [Strong, 1990] using *line-by-line* coding [Gibbs, 2007] to identify keywords that characterize the three psycho-social epidemics. For each of the three psycho-social epidemics, the three authors generated independent lists of keywords that were conservatively combined by intersecting them. The words that were left out by the intersection were mostly synonyms (e.g., “catching disease” as a synonym for “contagion”), so we did not discard any important concept. According to Strong, the three psycho-social epidemics are intertwined and, as such, the concepts that define one specific psycho-social epidemic might be relevant to the remaining two as well. For example, *suspicion* is an element of the epidemic of fear but is tightly related to *stigmatization* as well, a phenomenon that Strong describes as typical of the epidemic of moralization. In our coding exercise, we adhered as much as possible to the description in Strong’s paper and obtained a strict partition of keywords across psycho-social epidemics. In the second step, the same three authors mapped each of these keywords to *language categories*, namely sets of words that reflect how these concepts are expressed in natural language (e.g., words expressing *anger* or *trust*). We took these categories from four existing language lexicons widely used in psychometric studies:

- *Linguistic Inquiry Word Count (LIWC)* [Tausczik and Pennebaker, 2010]. A lexicon of words and word stems grouped into over 125 categories reflecting emotions, social processes, and basic functions, among others. The LIWC lexicon is based on the premise that the words people use to communicate can provide clues to their psychological states [Tausczik and Pennebaker, 2010]. It allows written passages to be analyzed syntactically (how the words are used together to form phrases or sentences) and semantically (an analysis of the meaning of the words or phrases).
- *Emolex* [Mohammad and Turney, 2013]. A lexicon that classifies 6k+ words and stems into the eight primary emotions of Plutchik’s psychoevolutionary theory [Plutchik, 1991].
- *Moral Foundation Lexicon* [Graham et al., 2009]. A lexicon of 318 words and stems, which are grouped into 5 categories of moral foundations [Graham et al., 2013]: harm, fairness, in-group, authority, and purity. Each of which is further split into expressions of virtue or vice.

- *Pro-social behavior* [Frimer et al., 2014]. A lexicon of 146 pro-social words and stems, which have been found to be frequently used when people describe pro-social goals [Frimer et al., 2014].

The three authors grouped similar keywords together and mapped groups of keywords to one or more language categories. This grouping and mapping procedure was informed by previous studies that investigated how these keywords are expressed through language. These studies are listed in Table 1.

**Language categories over time.** We considered that a tweet contained a language category  $c$  if at least one of the tweet’s words or stems belonged to that category. The tweet-category association is binary and disregards the number of matching words within the same tweet. That is mainly because, in short snippets of text (tweets are limited to 280 characters), multiple occurrences are rare and do not necessarily reflect the intensity of a category [Russell, 2013]. For each language category  $c$ , we counted the number of users  $U_c(t)$  who posted at least one tweet at time  $t$  containing that category. We then obtained the fraction of users who mentioned category  $c$  by dividing  $U_c(t)$  by the total number of users  $U(t)$  who tweeted at time  $t$ :

$$(1) \quad f_c(t) = \frac{U_c(t)}{U(t)}.$$

Computing the fraction of users rather than the fraction of tweets prevents biases introduced by exceptionally active users, thus capturing more faithfully the prevalence of different language categories in our Twitter population. This also helps discounting the impact of social bots, which tend to have anomalous levels of activity (especially retweeting [Bessi and Ferrara, 2016]).

Different categories might be verbalized with considerably different frequencies. For example, the language category “I” (first-person pronoun) from the LIWC lexicon naturally occurred much more frequently than the category “death” from the same lexicon. To enable a comparison across categories, we standardized all the fractions:

$$(2) \quad z_c(t) = \frac{f_c(t) - \mu_{[0,T]}(f_c)}{\sigma_{[0,T]}(f_c)},$$

where  $\mu(f_c)$  and  $\sigma(f_c)$  represent the mean and standard deviation of the  $f_c(t)$  scores over the whole time period, from  $t = 0$  (February 1<sup>st</sup>) to  $t = T$  (April 16<sup>th</sup>). These  $z$ -scores ease also the interpretation of the results as they represent the relative variation of a category’s prevalence compared to its average: they take on values higher (lower) than zero when the original value is higher (lower) than the average.

**Other behavioral markers.** To assess the validity of our operationalization of Strong’s model, we compared its results with the output of alternative state-of-the-art text-mining techniques, and with real-world mobility patterns.

*Interaction types.* We compared the results obtained via word-matching with a state-of-the-art deep learning tool for Natural Language Processing designed to capture fundamental types of social interactions from conversational language [Choi et al., 2020]. This tool uses Long Short-Term Memory neural networks (LSTMs) [Hochreiter and Schmidhuber, 1997] that take in input a 300-dimensional GloVe representation of words [Pennington et al., 2014] and output a series of confidence scores in the range  $[0, 1]$  that estimate the likelihood that the text expresses certain types of social interactions. The classifiers exhibited a very high classification performance, up to an Area Under the ROC Curve (AUC) of 0.98. AUC is a performance metric that measures the ability of the model to assign higher confidence scores to positive examples (i.e., text characterized by the type of interaction of interest) than to negative examples, independent of any fixed decision threshold; the expected value for random classification is 0.5, whereas an AUC of 1 indicates a perfect classification.

Out of the ten interaction types that the tool can classify [Deri et al., 2018], three were detected frequently with likelihood  $> 0.5$  in our Twitter data: *conflict* (expressions of contrast or diverging views [Tajfel et al., 1979]), *social support* (giving emotional or practical aid and companionship [Fiske et al., 2007]), and *power* (expressions that mark a person’s power over the behavior and outcomes of another [Blau, 1964]).

Given a tweet’s textual message  $m$  and an interaction type  $i$ , we used the classifier to compute the likelihood score  $l_i(m)$  that the message contains that interaction type. We then binarized the confidence scores using a threshold-based indicator function:

$$(3) \quad l_i^\theta(m) = \begin{cases} 1, & \text{if } l_i(m) \geq \theta_i \\ 0, & \text{otherwise} \end{cases}$$

Following the original approach [Choi et al., 2020], we used a different threshold for each interaction type, as the distributions of their likelihood scores tend to vary considerably. We thus picked conservatively  $\theta_i$  as the value of the 85<sup>th</sup> percentile of the distribution of the confidence scores  $l_i$ , thus favoring precision over recall. Last, similar to

	Keywords	Supporting literature	Lang. categories	Max. peak		
				1 <sup>st</sup>	2 <sup>nd</sup>	3 <sup>rd</sup>
Fear	emotional maelstrom	These LIWC categories have been used to analyze complex emotional responses to traumatic events (PTSD) and to characterize the language of people suffering from mental health [Coppersmith et al., 2014]	swear (liwc) anger (liwc) negemo (liwc) sadness (liwc)	.03 .07 .09 -.09	<b>.54</b> .14 <b>.17</b> .04	.14 <b>.16</b> .02 <b>.19</b>
	fear	Fear-related words like the ones included in Emolex have been often used to measure fear of both tangible and intangible threats [Kahn et al., 2007, Gill et al., 2008]	fear (emolex) death (liwc)	<b>.07</b> <b>.45</b>	.05 .16	.03 .03
	anxiety, panic	The <i>anxiety</i> category of LIWC has been used to study different forms of anxiety in social media [Shen and Rudzicz, 2017]	anxiety (liwc)	.31	<b>.46</b>	-.15
	disorientation	By definition, the <i>tentative</i> category of LIWC expresses uncertainty [Tausczik and Pennebaker, 2010]	tentative (liwc)	.02	<b>.10</b>	.03
	suspicion	Suspicion is often formalized as lack of trust [Deutsch, 1958]	trust (liwc)	-.04	.03	<b>.12</b>
	irrationality	The <i>negate</i> category of LIWC has been used to measure cognitive distortions and irrational interpretations of reality [Simms et al., 2017]	negate (liwc)	.08	<b>.15</b>	.07
	religion	Religious expressions from LIWC have been used to study how people appeal to religious entities during moments of hardship [Shaw et al., 2007]	religion (liwc)	.12	.17	<b>.22</b>
	contagion	These LIWC categories were used to study the perception of diseases in cancer support groups, people affected by eating disorder, and alcoholics [Alpers et al., 2005, Wolf et al., 2013, Kornfield et al., 2018]	body (liwc) feel (liwc)	.01 -.04	<b>.27</b> <b>.34</b>	.13 .03
Moralization	warn, risk avoidance, risk perception	A LIWC category used to model risk perception connected to epidemics [Hou et al., 2020]	risk (liwc)	.04	<b>.15</b>	.08
	polarization, segregation	Different personal pronouns have been used to study in- and out-group dynamics and to characterize language markers of racism [Arguello et al., 2006]; personal pronouns and markers of differentiation have been considered in studies on racist language [Figea et al., 2016]	I (liwc) we (liwc) they (liwc) differ (liwc)	-.10 -.09 .03 .01	<b>.49</b> <b>.24</b> .02 <b>.08</b>	.26 .22 <b>.10</b> .05
	stigmatization, blame, abuse	Pronouns <i>I</i> and <i>they</i> were used to quantify blame in personal [Borelli and Sbarra, 2011] and political context [Windsor et al., 2014]. Hate speech is associated with they-oriented statements [ElSherief et al., 2018]	(same categories as previous line)			
	cooperation, coordination, collective consciousness	The moral value of <i>care</i> expresses the will of protecting others [Graham et al., 2009]. Cooperation is often verbalized by referencing the in-group and by expressing <i>affiliation</i> [Rezapour et al., 2019]	affiliation (liwc) care (moral virtue) prosocial (prosocial)	-.16 -.09 -.10	<b>.25</b> .12 .07	.22 <b>.19</b> <b>.25</b>
	faith in authority	The moral value of <i>authority</i> expresses the will of playing by the rules of a hierarchy versus challenging it [Rezapour et al., 2019]	authority (moral virtue)	-.04	.11	<b>.13</b>
	authority enforcement	The <i>power</i> category of LIWC expresses exertion of dominance [Tausczik and Pennebaker, 2010]	power (liwc)	-.01	.02	<b>.14</b>
	Action	restrictions, travel, privacy	Daily habits concern mainly people’s experience of <i>home</i> , <i>work</i> , <i>leisure</i> , and <i>movement</i> between them [Gonzalez et al., 2008]	motion (liwc) home (liwc) work (liwc) social (liwc) leisure (liwc)	.02 -.15 -.05 -.06 .05	.04 <b>.38</b> .04 <b>.09</b> <b>.16</b>

TABLE 1. Operationalization of the Strong’s epidemic psychology theoretical framework. From Strong’s paper, three annotators extracted *keywords* that characterize the three psycho-social epidemics and mapped them to relevant *language categories* from existing language lexicons used in psychometric studies. Category names are followed by the name of their corresponding lexicon in parenthesis. We support the association between keywords and language categories with examples of *supporting literature*. To summarize how the use of the language categories varies across the three temporal states, we computed the *peak* values of the different language categories (days when their standardized fractions reached the maximum), and reported the percentage increase at peak compared to the average over the whole time period; in each row, the maximum value is highlighted in bold.

how we constructed temporal signals for the language categories, we counted the number of users  $U_i(t)$  who posted at least one tweet at time  $t$  that contains interaction type  $i$ . We then obtained the fraction of users who mentioned interaction type  $i$  by dividing  $U_i(t)$  by the total number of users  $U(t)$  who tweeted at time  $t$ :

$$(4) \quad f_i(t) = \frac{U_i(t)}{U(t)}.$$

Last, we min-max normalized these fractions, considering the minimum and maximum values during the whole time period  $[0, T]$ :

$$(5) \quad f_i(t) = \frac{f_i(t) - \min_{[0, T]}(f_i)}{\max_{[0, T]}(f_i) - \min_{[0, T]}(f_i)}.$$

*Mentions of medical entities.* We used a state-of-the-art deep learning method for medical entity extraction to identify medical symptoms on Twitter in relation to COVID-19 [Scepanovic et al., 2020]. When applied to tweets, the method extracts  $n$ -grams representing medical symptoms (e.g., “feeling sick”). This method is based on the Bi-LSTM sequence-tagging architecture [Huang et al., 2015] in combination with GloVe word embeddings [Pennington et al., 2014] and RoBERTa contextual embeddings [Liu et al., 2019]. To optimize the entity extraction performance on noisy textual data from social media, we trained its sequence-tagging architecture on the Micromed database [Jimeno-Yepes et al., 2015], a collection of tweets manually labeled with medical entities. The hyperparameters we used are: 256 hidden units, a batch size of 4, and a learning rate of 0.1 which we gradually halved whenever there was no performance improvement after 3 epochs. We trained for a maximum of 200 epochs or before the learning rate became too small ( $\leq .0001$ ). The final model achieved an F1-score of .72 on Micromed. The F1-score is a performance measure that combines precision (the fraction of extracted entities that are actually medical entities) and recall (the fraction of medical entities present in the text that the method is able to retrieve). We based our implementation on Flair [Akbik et al., 2019] and Pytorch [Paszke et al., 2017], two popular deep learning libraries in Python.

For each unique medical entity  $e$ , we counted the number of users  $U_e(t)$  who posted at least one tweet at time  $t$  that mentioned that entity. We then obtained the fraction of users who mentioned medical entity  $e$  by dividing  $U_e(t)$  by the total number of users  $U(t)$  who tweeted at time  $t$ :

$$(6) \quad f_e(t) = \frac{U_e(t)}{U(t)}.$$

Last, we min-max normalize these fractions, considering the minimum and maximum values during the whole time period  $[0, T]$ :

$$(7) \quad f_e(t) = \frac{f_e(t) - \min_{[0, T]}(f_e)}{\max_{[0, T]}(f_e) - \min_{[0, T]}(f_e)}.$$

*Mobility traces.* Foursquare is a local search and discovery mobile application that relies on the users’ past mobility records to recommend places user might like. The application uses GPS geo-localization to estimate the user position and to infer the places they visited. In response to the COVID-19 crisis, Foursquare made publicly available the data gathered from a pool of 13 million U.S. users. These users were “always-on” during the period of data collection, meaning that they allowed the application to gather geo-location data at all times, even when the application was not in use. The data (published through the [visitdata.org](http://visitdata.org) website) consists of the daily number of users  $v_{s,j}$  visiting any venue of type  $j$  in state  $s$ , starting from February 1<sup>st</sup> to the present day (e.g., 419,256 users visited schools in Indiana on February 1<sup>st</sup>). Overall, 35 distinct location categories were provided. To obtain country-wide temporal indicators, we first applied a min-max normalization to the  $v_{s,j}$  values:

$$(8) \quad v_{s,j}(t) = \frac{v_{s,j}(t) - \min_{[0, T]}(v_{s,j})}{\max_{[0, T]}(v_{s,j}) - \min_{[0, T]}(v_{s,j})}.$$

We then averaged the values across all states:

$$(9) \quad v_j(t) = \frac{1}{S} \sum_s v_{s,j}(t),$$

where  $S$  is the total number of states. By weighting each state equally, we obtained a measure that is more representative of the whole U.S. territory, rather than being biased towards high-density regions.

**Time series smoothing.** All our temporal indicators are affected by large day-to-day fluctuations. To extract more consistent trends out of our time series, we applied a smoothing function—a common practice when analyzing temporal data extracted from social media [O’Connor et al., 2010]. Given a time-varying signal  $x(t)$ , we apply a “boxcar” moving average over a window of the previous  $k$  days:

$$(10) \quad x^*(t) = \frac{\sum_{i=t-k}^t x(i)}{k};$$

We selected a window of one week ( $k = 7$ ). Weekly time windows are typically used to smooth out both day-to-day variations as well as weekly periodicities [O’Connor et al., 2010]. We applied the smoothing to all the time series: the language categories ( $z_c^*(t)$ ), the mentions of medical entities ( $f_e^*(t)$ ), the interaction types ( $f_i^*(t)$ ), and the foursquare visits ( $v_j^*(t)$ ).

**Change-point detection.** To identify phases characterized by different combinations of the language categories, we identified *change-points*—periods in which the values of all categories varied considerably at once. To quantify such variations, for each language category  $c$ , we computed  $\nabla(z_c^*(t))$ , namely the daily average *squared gradient* [Lütkepohl, 2005] of the smoothed standardized fractions of that category. To calculate the gradient, we used the Python function `numpy.gradient`. The gradient provides a measure of the rate of increase or decrease of the signal; we consider the absolute value of the gradient, to account for the magnitude of change rather than the direction of change. To identify periods of consistent change as opposed to quick instantaneous shifts, we apply temporal smoothing (Equation. 10) also to the time-series of gradients, and we denote the smoothed squared gradients with  $\nabla^*$ . Last, we average the gradients of all language categories to obtain the overall gradient over time:

$$(11) \quad \nabla(t) = \frac{1}{D} \sum_d \nabla^*(z_c^*(t)).$$

Peaks in the time series  $\nabla(t)$  represent the days of highest variation, and we marked them as change-points. Using the Python function `scipy.signal.find_peaks`, we identified peaks as the local maxima whose values is higher than the average plus two standard deviations, as it is common practice [Palshikar et al., 2009].

## RESULTS

**Language use until the first contagion wave.** During the first wave, U.S. residents experienced what the pandemic entailed for the first time, and did so by going through an entire life-cycle, making this contagion wave self-contained: arrival of an unknown virus, skepticism, isolation measures, full lock-down, and first reopening. Figure 1 (A-C) shows how the standardized fractions of all the language categories as per Formula (2) changed from February 1<sup>st</sup> to April 15<sup>th</sup>, the day in which restrictions in most states were lifted. The cell color encodes values higher than the average in red, and lower in blue. We partitioned the language categories according to the three psycho-social epidemics. Figure 1 D shows the value of the average squared gradient over time (as per Formula 11); peaks in the curve represent the days of high local variation. We marked the peaks above two standard deviations from the mean as change-points. We found two change-points that coincide with two key events: February 27<sup>th</sup>, the day of the announcement of the first infection in the country; and March 24<sup>th</sup>, the day of the announcement of the ‘stay at home’ orders. These change-points identify three phases, which are described next by dwelling on the *peaks* of the different language categories (days when their standardized fractions reached the maximum) and reporting the percentage increase at peak (the increase is compared to the average from February 1<sup>st</sup> to April 15<sup>th</sup>, and its peak is denoted by ‘max peak’ in Table 1). The first phase (*refusal phase*) was characterized by *anxiety* and *fear*. *Death* was frequently mentioned, with a peak on February 11 of +45% compared to its average during the whole time period. The pronoun *they* was used in this temporal state more than average; this suggests that the focus of discussion was on the implications of the viral epidemic on ‘others’, as this was when no infection had been discovered in the U.S. yet. All other language categories exhibited no significant variations, which reflected an overall situation of ‘business-as-usual.’

The second phase (*anger phase*) began on February 27<sup>th</sup> with an outburst of *negative emotions* (predominantly *anger*), right after the first COVID-19 contagion in the U.S. was announced. The abstract fear of death was replaced by expressions of concrete health concerns, such as words expressing *risk*, and mentions of how *body* parts did *feel*. On March 13<sup>th</sup>, the federal government announced the state of national emergency, followed by the enforcement of state-level ‘stay at home’ orders. During those days, we observed a sharp increase of the use of the pronoun *I* and of *swear* words (with a peak of +54% on March 18<sup>th</sup>), which hints at a climate of discussion characterized by conflict and polarization. At the same time, we observed an increase in the use of words related to the daily habits affected by the impending restriction policies, such as *motion*, *social* activities, and *leisure*. The mentions of words

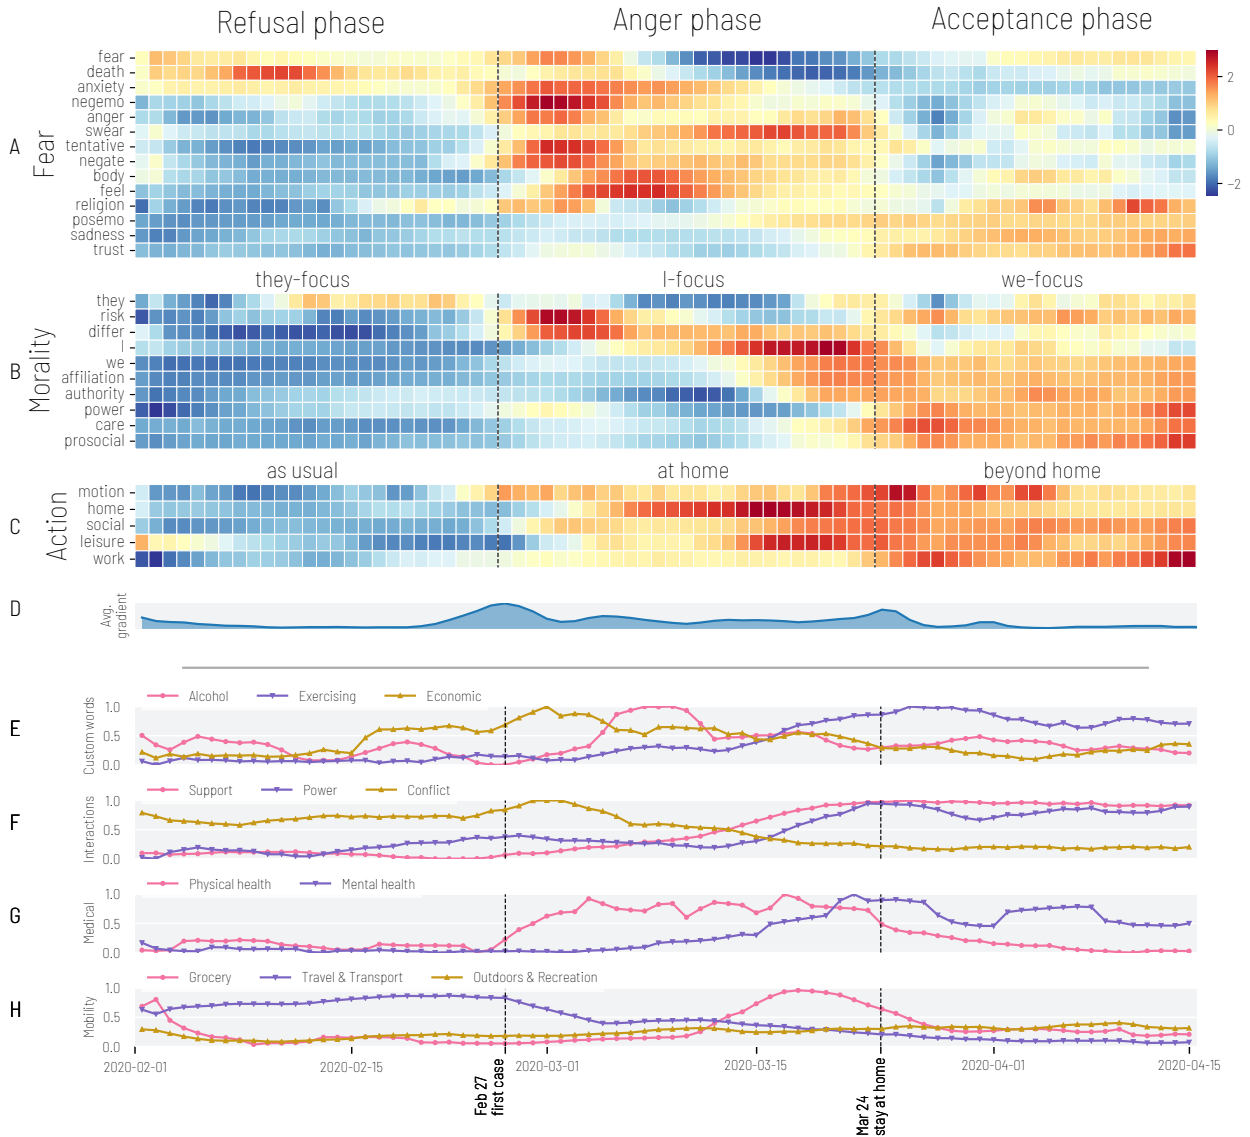


FIGURE 1. The Epidemic Psychology on Twitter during the first contagion wave. (A-C) evolution of the use of different language categories over time in tweets related to COVID-19. Each row in the heatmaps represents a language category (e.g., words expressing anxiety) that our manual coding associated with one of the three psycho-social epidemics. The cell color represents the daily standardized fraction of people who used words related to that category: values that are higher than the average are red and those that are lower are blue. Categories are partitioned in three groups according to the type of psycho-social epidemics they model: Fear, Morality, and Action. (D) average gradient (i.e., instantaneous variation) of all the language categories; the peaks of gradient identify *change-points*—dates around which a considerable change in the use of multiple language categories happened at once. The dashed vertical lines that cross all the plots represent these change-points. (E-H) temporal evolution of four families of indicators we used to corroborate the validity of the trends identified by the language categories. We checked internal validity by comparing the language categories with a custom keyword-search approach and two deep-learning NLP tools that extract types of social interactions and mentions of medical symptoms. We checked external validity by looking at mobility patterns in different venue categories as estimated by the GPS geo-localization service of the Foursquare mobile app. The timeline at the bottom of the figure marks some of the key events of the COVID-19 pandemic in the U.S. such as the announcements of the first infection of COVID-19 recorded.



related to *home* peaked on March 16<sup>th</sup> (+38%), the day when the federal government announced social distancing guidelines to be in place for at least two weeks.

The third phase (*acceptance phase*) started on March 24<sup>th</sup>, the day after the first physical-distancing measures were imposed by law. The increased use of words of *power* and *authority* likely reflected the emergence of discussion around the new policies enforced by government officials and public agencies. As the death toll raised steadily—hitting the mark of 1,000 deaths on March 26<sup>th</sup>—expressions of conflict faded away, and words of *sadness* became predominant. In those days of hardship, a sentiment of *care* for others and expressions of *prosocial* behavior became more frequent (+19% and +25%, respectively). Last, mentions of *work*-related activities peaked as many people either lost their job, or were compelled to work from home as result of the lockdown.

**Thematic analysis.** The language categories capture broad concepts related to Strong’s epidemic psychology theory, but they do not allow for an analysis of the fine-grained topics within each category. To study them, for each of the 87 combinations of language category and phase (29 language categories, for 3 phases), we listed the 100 most retweeted tweets (e.g., most popular tweets containing anxiety posted in the refusal phase). To identify overarching *themes*, we followed two steps that are commonly adopted in thematic analysis [Braun and Clarke, 2006, Smith and Shinebourne, 2012]. We first applied *open coding* to identify key concepts that emerged across multiple tweets; specifically, one of the authors read all the tweets and marked them with keywords that reflected the key concepts expressed in the text. We then used *axial coding* to identify relationships between the most frequent keywords to summarize them in semantically cohesive themes. Themes were reviewed in a recursive manner rather than linear, by re-evaluating and adjusting them as new tweets were parsed. Table 2 summarizes the most recurring themes, together with some of their representative tweets. In the refusal phase, statements of skepticism were re-tweeted widely (Table 2, row 1). The epidemic was frequently depicted as a “foreign” problem (r. 2) and all activities kept business as usual (r. 3).

In the anger phase, the discussion was characterized by outrage against three main categories: foreigners—especially Chinese individuals, as Supplementary Materials detail—(r. 4), political opponents (r. 5), and people who adopted different behavioral responses to the outbreak (r. 6). This level of conflict corroborates Strong’s postulate of the “*war against each other*”. Science and religion were two prominent topics of discussion. A lively debate raged around the validity of scientists’ recommendations (r. 7). Some social groups put their hopes on God rather than on science (r. 8). Mentions of people self-isolating at home became very frequent, and highlighted the contrast between judicious individuals and careless crowds (r. 9).

Finally, during the acceptance phase, the outburst of anger gave in to the sorrow caused by the mourning of thousands of people (r. 10). By accepting the real threat of the virus, our Twitter users were more open to find collective solutions to the problem and overcome fear with hope (r. 11). Although the positive attitude towards the authorities seemed prevalent, some users expressed disappointment against the restrictions imposed (r. 12). Those who were isolated at home started imagining a life beyond the isolation, especially in relation to reopening businesses (r. 13).

**Comparison with other behavioral markers.** To assess the validity of our approach, we compared the previous results with the output of alternative text-mining techniques applied to the same data (*internal validity*), and with real-world mobility traces (*external validity*).

*Comparison with other text mining techniques.* We processed the very same social media posts with three alternative text-mining techniques (Figure 1 E-G). In Table 4, we reported the three language categories with the strongest correlations with each behavioral marker.

First, to allow for interpretable and explainable results, we applied a simple word-matching method that relies on a custom lexicon containing three categories of words reflecting consumption of alcohol, physical exercising, and economic concerns, as those aspects have been found to characterize the COVID-19 pandemic [The Economist, 2020]. We measured the daily fraction of users mentioning words in each of those categories (Figure 1 E). In the refusal phase, the frequency of any of these words did not significantly increase. In the anger phase, the frequency of words related to economy peaked, and that related to alcohol consumption peaked shortly after that. Table 4 shows that economy-related words were highly correlated with the use of *anxiety* words ( $r = 0.73$ ), which is in line with studies indicating that the degree of apprehension for the declining economy was comparable to that of health-hazard concerns [Fetzer et al., 2020, Bareket-Bojmel et al., 2020]. Words of alcohol consumption were most correlated with the language dimensions of *body* ( $r = 0.70$ ), *feel* ( $r = 0.62$ ), *home* ( $r = 0.58$ ); in the period were health concerns were at their peak, home isolation caused a rising tide of alcohol use [Da et al., 2020, Finlay and Gilmore, 2020]. Finally, in the acceptance phase, the frequency of words related to physical exercise was significant;

Theme		Example tweets
<i>The refusal phase</i>		
1	denial	<i>“Less than 2% of all cases result in death. Approximately equivalent to seasonal flu. Relax people.”</i>
2	they-focus	<i>“We will continue to call it the #WuhanVirus, which is exactly what it is.”</i>
3	business as usual	<i>“Agriculture specialists at Dulles airport continue to protect our nation’s vital agricultural resources.”</i>
<i>The anger phase</i>		
4	anger vs. foreigners	<i>“Is there anything you won’t use to stir up hatred against the foreigner? #COVID19 is a global pandemic.”</i>
5	anger vs. political opponents	<i>“A new level of sickness has entered the body politic. The son of the monster mouthing off grotesque lies about Dems cheering #coronavirus and Wall Street crashing because we want an end to his father’s winning streak.”</i>
6	anger vs. each other	<i>“Coronavirus or not, if you are ill, stay the f**k home. You’re not a hero for going to work when you are unwell.”</i>
7	science debate	<i>“When it comes to how to fight #CoronavirusPandemic, I’m making my decisions based on healthcare professionals like Dr. Fauci and others, not political punditry”</i>
8	religion	<i>“no problem is too big for God to handle [...] with God’s help, we will overcome this threat.”</i>
9	I-focus, home	<i>“People get upset and annoyed at me when I tweet about the coronavirus, when I urge people to stay in and avoid crowds”, “I am in the high risk category for coronavirus so do me a favor [...] beg others to stay at home”</i>
<i>The acceptance phase</i>		
10	sadness	<i>“We deeply mourn the 758 New Yorkers we lost yesterday to COVID-19. New York is not numb. We know this is not just a number—it is real lives lost forever.”</i>
11	we-focus, hope	<i>“We are thankful for Japan’s friendship and cooperation as we stand together to defeat the #COVID19 pandemic.”, “During tough times, real friends stick together. The U.S. is thankful to #Taiwan for donating 2 million face masks to support our healthcare ”, “Now more than ever, we need to choose hope over fear. We will beat COVID-19. We will overcome this. Together.”</i>
12	authority	<i>“You can’t go to church, buy seeds or paint, operate your business, run on a beach, or take your kids to the park. You do have to obey all new ‘laws’, wear face masks in public, pay your taxes. Hopefully this is over by the 4th of July so we can celebrate our freedom.</i>
13	resuming work	<i>“We need to help as many working families and small businesses as possible. Workers who have lost their jobs or seen their hours slashed and families who are struggling to pay rent and put food on the table need help immediately. There’s no time to waste.”</i>

TABLE 2. Recurring themes in the three phases, found by the means of thematic analysis of tweets. Themes are paired with examples of popular tweets.

this happened at the same time when the use of positive words expressing togetherness was at its highest—*affiliation* ( $r = 0.95$ ), *posemo* ( $r = 0.93$ ), *we* ( $r = 0.92$ ). All these results match our previous interpretations of the peaks for our language categories.

Second, since it is unclear whether a simple word count approach is effective in studying how the three psychosocial epidemics unfolded over time, we additionally applied the deep-learning approach that extracts mentions of expressions of *conflict*, *social support*, and *power*. Figure 1 F shows the min-max normalized scores of the fraction of users posting tweets labeled with each of these three interaction types (as per Formula 5). In the refusal phase, *conflict* increased—this is when anxiety and blaming foreigners were recurring themes in Twitter. In the anger phase, conflict peaked (similar to *anxiety* words,  $r = 0.88$ ), yet, since the first lock-down measures were announced, initial expressions of power and of social support gradually increased as well. Finally, in the acceptance phase, social support peaked. Support was most correlated with the categories of *affiliation* ( $r = 0.98$ ), *positive emotions* ( $r = 0.96$ ), and *we* ( $r = 0.94$ ) (Table 4); power was most correlated with *prosocial* ( $r = 0.95$ ), *care* ( $r = 0.94$ ), and *authority* ( $r = 0.94$ ). Again, our previous interpretations concerning the existence of a phase of conflict followed by a phase of social support were further confirmed by the deep-learning tool, which, as opposed to our dictionary-based approaches, does not rely on word matching.

Phase	Top positive			Top negative		
Refusal	<b>death</b> (0.66)	they (0.06)	fear (0.04)	<b>I</b> (-1.51)	we (-1.27)	home (-1.22)
Anger	<b>swear</b> (2.17)	feel (1.51)	anxiety (1.46)	<b>death</b> (-0.70)	sadness (-0.51)	prosocial (-0.38)
Acceptance	<b>sad</b> (1.35)	affiliation (1.19)	prosocial (1.17)	<b>anxiety</b> (-1.62)	swear (-1.36)	I (-0.34)

TABLE 3. Top three positive and bottom negative beta coefficients of the logistic regression models for the three phases. The categories in bold are those included in our composite temporal score.

Third, we used the deep-learning tool that extracts mentions of medical entities from text [Scepanovic et al., 2020]. Out of all the entities extracted, we focused on the 100 most frequently mentioned and grouped them into two families of symptoms, respectively, those related to physical health (e.g., “fever”, “cough”, “sick”) and those related to mental health (e.g., “depression”, “stress”) [Brooks et al., 2020]. The min-max normalized fractions of users posting tweets containing mentions of these symptoms (as per Formula 7) are shown in Figure 1 G. In refusal phase, the frequency of symptom mentions did not change. In the anger phase, instead, physical symptoms started to be mentioned, and they were correlated with the language categories expressing panic and physical health concerns—*swear* ( $r = 0.83$ ), *feel* ( $r = 0.77$ ), and *negate* ( $r = 0.67$ ). In the acceptance phase, mentions of mental symptoms became most frequent. Interestingly, mental symptoms peaked when the Twitter discourse was characterized by positive feelings and prosocial interactions—*affiliation* ( $r = 0.91$ ), *we* ( $r = 0.88$ ), and *posemo* ( $r = 0.85$ ); this is in line with recent studies that found that the psychological toll of COVID-19 has similar traits to post-traumatic stress disorders and its symptoms might lag several weeks from the period of initial panic and forced isolation [Galea et al., 2020, Liang et al., 2020, Dutheil et al., 2020].

*Comparison with mobility traces.* To test for the external validity of our language categories, we compared their temporal trends with the mobility data from Foursquare. We picked three venue categories: *Grocery* shops, *Travel & Transport*, and *Outdoors & Recreation* to reflect three different types of fundamental human needs [Maslow, 1943]: the primary need of getting food supplies, the secondary need of moving around freely (or to limit mobility for safety), and the higher-level need of being entertained. In Figure 1 H, we show the min-max normalized number of visits over time (as per Formula 9). The periods of higher variations of the normalized number of visits match the transitions between the three phases. In the refusal phase, mobility patterns did not change. In the anger phase, instead, travel started to drop, and grocery shopping peaked, supporting the interpretation of a phase characterized by a wave of panic-induced stockpiling and a compulsion to save oneself—it co-occurred with the peak of use of the pronoun *I* ( $r = 0.80$ )—rather than helping others. Finally, in the acceptance phase, the panic around grocery shopping faded away, and the number of visits to parks and outdoor spaces increased.

**Embedding epidemic psychology in real-time models.** To embed our operationalization of epidemic psychology into real-time models (e.g., epidemiological models, urban mobility models), our measures need to work at any point in time during a new pandemic; yet, given their current definitions, they do not: that is because they are normalized values over the *whole* period of study (Figure 1 A-C). To fix that, we designed a new composite measure that does not rely on full temporal knowledge, and a corresponding detection method that determines which of the three phases one is in at any given point in time.

First, for each language category  $c$ , we computed the average value of  $f_c$  (as per Formula 1) during the first day of the epidemic, specifically  $\mu_{[0,1]}(f_c)$ . During the first day, 86k users tweeted. We experimented with longer periods (up to a week and 0.4M users), and obtained qualitatively similar results. We used the averages computed on this initial period as reference values for later measurements. The assumption behind this approach is that the modeler would know the set of relevant hashtags in the initial stages of the pandemic, which is reasonable considering that this was the case for all the major pandemics occurred in the last decade [Fu et al., 2016, Oyeyemi et al., 2014, Chew and Eysenbach, 2010]. Starting from the second day, we then calculated the percent change of the  $f_c$  values compared to the reference values:

$$(12) \quad \Delta\%_c(t) = \frac{f_c(t) - \mu_{[0,1]}(f_c)}{\mu_{[0,1]}(f_c)}.$$

For each phase, we defined a parsimonious measure composed only of two dimensions: the dimension most positively associated with the phase (expressed in percent change) minus that most negatively associated with it (e.g., (*death* - *I*) for the refusal phase). To identify such dimensions, we trained three logistic regression binary classifiers (one per phase). For each phase, we marked with label 1 all the days that were included in that phase and with 0 those

					Correlation with phases		
Marker		Most correlated language categories			Refusal	Anger	Acceptance
Custom words	Alcohol	body (0.70)	feel (0.62)	home (0.58)	-0.43	<b>0.46</b>	-0.12
	Economic	anxiety (0.73)	negemo (0.68)	negate (0.56)	-0.12	<b>0.37</b>	-0.53
	Exercising	affiliation (0.95)	posemo (0.93)	we (0.92)	-0.62	0.31	<b>0.89</b>
Interactions	Conflict	anxiety (0.88)	death (0.57)	negemo (0.54)	<b>0.58</b>	-0.24	-0.92
	Support	affiliation (0.98)	posemo (0.96)	we (0.94)	-0.68	0.37	<b>0.90</b>
	Power	prosocial (0.95)	care (0.94)	authority (0.94)	-0.48	0.18	<b>0.88</b>
Medical	Physical health	swear (0.83)	feel (0.77)	negate (0.67)	-0.66	<b>0.81</b>	-0.32
	Mental health	affiliation (0.91)	we (0.88)	posemo (0.85)	-0.65	0.36	<b>0.85</b>
Mobility	Travel	death (0.59)	anxiety (0.58)		<b>0.62</b>	-0.32	-0.82
	Grocery	I (0.80)	leisure (0.72)	home (0.64)	-0.77	<b>0.70</b>	0.29
	Outdoors	sad (0.68)	posemo (0.65)	affiliation (0.59)	-0.62	0.39	<b>0.72</b>

TABLE 4. (Left) Correlation of our language categories with behavioral markers computed with alternative techniques and datasets. For each marker, the three categories with strongest correlations are reported, together with their Pearson correlation values in parenthesis. (Right) Pearson correlation between values for our behavioral markers and “being” in a given phase or not. Values in bold indicate the highest values for each marker across the three phases. All reported correlations are statistically significant ( $p < 0.01$ ).

that were not. Then, we trained a classifier to estimate the probability  $P_{phase_i}(t)$  that day  $t$  belongs to phase  $i$ , out of the  $\Delta\%_c(t)$  values for all categories. During training, the logistic regression learned coefficients for each of the categories. On average, the classifiers were able to identify the correct phase for 98% of the days.

The regressions coefficients were then used to rank the language category by their predictive power. Table 3 shows the top three positive beta coefficients and bottom three negative ones for each of the three phases. The top and bottom categories of all phases belong to the LIWC lexicon. For each phase, we subtracted the top category from the bottom category without considering their beta coefficients, as these would require, again, full temporal knowledge:

$$\begin{aligned}
 \Delta\%_{Refusal} &= \Delta\%_{death} - \Delta\%_I \\
 \Delta\%_{Anger} &= \Delta\%_{swear} - \Delta\%_{death} \\
 \Delta\%_{Acceptance} &= \Delta\%_{sad} - \Delta\%_{anxiety}
 \end{aligned}
 \tag{13}$$

The resulting composite measure has the same change-points (Figure 3A) as the full-knowledge measure’s (Figure 1), suggesting that the real-time and parsimonious computation does not compromise the original trends. In a real-time scenario, transition between phases are captured by the changes of the dominant measure; for example, when the *refusal* curve is overtaken by the *anger* curve. In addition, we correlated the composite measures with each of the behavioral markers we used for validation (Figure 1 E-H) to find which are the markers most typically associated with each of the phases. We reported the correlations in Table 4. During the refusal phase, conflictual interactions were frequent ( $r = 0.58$ ) and long-range mobility was common ( $r = 0.62$ ); during the anger phase, as mobility reduced [Engle et al., 2020, Gao et al., 2020], some people hoarded groceries and alcohol [Da et al., 2020, Finlay and Gilmore, 2020] and expressed concerns for their physical health ( $r = 0.81$ ) and for the economy [Fetzer et al., 2020, Bareket-Bojmel et al., 2020]; last, during the acceptance phase, some people ventured outdoors, started exercising more, and expressed a stronger will to support each other ( $r = 0.90$ ), in the wake of a rising tide of deaths and mental health symptoms ( $r = 0.85$ ) [Galea et al., 2020, Liang et al., 2020, Dutheil et al., 2020].

**Language use after the first contagion wave.** After the first wave and the “stay-at-home” lock-down at the end of March, for the remaining part of the year, there were two other contagion waves (Figure 2A): one at the




FIGURE 2. The number of recorded infections in the U.S. (A), and the average gradient (i.e., instantaneous variation) of all the language categories (B). There were three contagion waves in the year (shaded areas in A) and, for each wave, there were peaks in the gradient (marked with circles in B), which identify *change-points*, that is, periods in which the use of language considerably changed.

beginning of June, and the other at the beginning of October. In a way similar to the first wave, these two others were both associated with significant changes in the use of language (Figure 2B): we had two change peaks for the second wave (i.e., May 27<sup>th</sup> and June 9<sup>th</sup>), and one for the third (October 2). This comes at no surprise as these two periods corresponded to two widely discussed events: the first 100K deaths in the U.S., and President Donald Trump testing positive for COVID-19. In particular, Figure 3B shows that these changes were both to do with the rumping up of the categories associated with the *anger phase*: discussions of changes in *mobility*, and posts characterized by *anger* and, more generally, by *negative emotions* were predominant. By contrast, the categories associated with *refusal* and *acceptance* diverged from each other: unsurprisingly, throughout the year, *refusal* gradually died off, while *acceptance* increasingly took hold. Overall, we observed two classes of pattern (Figure 3B). First, the three phases were not always orthogonal but blended together at times. During the second contagion wave, for example, *anger* and *acceptance* were both predominant, and had been so for several months. Second, the use of language had a cyclical nature. During each contagion wave, three consecutive local peaks (local maxima) were observed: *refusal* first, then *anger*, finally *acceptance*. This can be observed for all the three contagion waves. Such a cyclical nature was also reflected in our behavioral markers (Figure 4(E-H)). For each wave, mentions of *conflict* peaked to then be followed by mentions of *support* and *power* (Figure 4F). The same went for medical conditions: mentions of *physical health* peaked to then be followed by mentions of *mental health* (Figure 4G).

## DISCUSSION

**Findings beyond Strong’s model.** Strong’s theory offers a framework upon which to operationalize the three psycho-social epidemics from the use of language but does not specifically describe how these epidemics unfold. Based on our data-driven results, we confirmed that the three epidemics are indeed present in our social conversations spanning almost one year, and they unfolded in ways that allowed us to enrich what Strong initially hypothesized in relation to four main aspects.

First, Strong’s theory predicts the presence of the three psycho-social epidemics but does not describe how they would be related to each other over time. We found that the three epidemics, as expressed by the relative presence of their relevant language categories, simultaneously raised and fell over time, and did so in relation to one another, demarcating three specific temporal phases. Over time, we identified three specific combinations of the epidemics that generated three phases our Twitter users went through: an initial *refusal* phase, an *anger* phase, and a final *acceptance* phase. Since these temporal phases partly resemble the stages of grief by Kübler-Ross et al. [1972], a promising direction for future work is to explore the relationship between these phases and the stages of grief.

Second, Strong’s narration slightly hints at a typical sequence of events according to which the epidemic of *fear* activates before that of *morality*, which is then followed by that of *action*. We indeed observed a similar set of events yet these events were not strictly sequential but rather cyclical. Interestingly, every new cycle started in conjunction with the same specific event: the diffusion rate of the virus reaching a local maximum. Shortly after every sharp increase in the diffusion rate, a cycle of refusal, anger, and acceptance unfolded among our Twitter users.

Third, Strong’s framework does not make any explicit distinction between the initial stages of an epidemic and its final stages. We found two regimes that considerably separated the initial stages from the later stages: in the initial cycle, the variation of the three epidemics was of a larger magnitude than those in the two subsequent cycles.




FIGURE 3. The evolution of the language categories associated with *refusal*, of those associated with *anger*, and of those associated with *acceptance*, zooming out from the first contagion wave (A) to the three waves during the entire 2020 (B).

Last, Strong’s theory should be used as a descriptive framework and, as such, cannot be used to explore short-term variations. By contrast, a data-driven experimental work such as ours is able to point out when such variations potentially took place, as Figure 4 shows. In future work, these points of change could be subject to qualitative inquiry, which might well enrich the original formulation of the theory.

**Implications.** New infectious diseases break out abruptly, and public health agencies try to rely on detailed planning yet often find themselves to improvise around their playbook. They are constantly confronting not only the health epidemic but also the three psycho-social epidemics. Measuring the effects of epidemics on societal dynamics and population mental health has been an open research problem for a long time, and multidisciplinary approaches have been called for [Holmes et al., 2020]. We contributed to this line of research by operationalizing Strong’s model and successfully testing it on Twitter in the U.S.. Since our methodology can be applied to any textual data, future work may well study alternative (even cross-cultural) population segments. Since our language categories are not tailored to a specific epidemic (e.g., they do not reflect any specific symptom an epidemic is associated with), our approach can be applied to a future epidemic, provided that the set of relevant hashtags associated with the epidemic is known; this is a reasonable assumption to make though, considering that the consensus on Twitter hashtags is reached quickly [Baronchelli, 2018], and that several epidemics that occurred in the last decade sparked discussions on Twitter since their early days [Fu et al., 2016, Oyeyemi et al., 2014, Chew and Eysenbach, 2010].

Our method complements the numerous cross-sectional studies on the psychological impact of health epidemics conducted on representative population samples [Shultz et al., 2015, Brooks et al., 2020], not least because it collects real-time statistics of implicit behavioral signals, which are orthogonal to survey responses.

For computer science researchers, our method could provide a starting point for developing more sophisticated tools for monitoring psycho-social epidemics. Furthermore, from the theoretical standpoint, our work provides the first operationalization of Strong’s model of the epidemic psychology, and widens its theoretical implications by observing cyclical phases of diffusion of the psycho-social epidemics.

Finally, the ability to systematically characterize the three psycho-social epidemics from the use of language on social media makes it possible to embed epidemic psychology into models currently used to tackle epidemics such as mobility models [Bansal et al., 2016]. To see how, consider that, in digital epidemiology [Salathe et al., 2012, Bauch and Galvani, 2013], some parameters of epidemic models are initialized or adjusted based on a variety of digital data to account for co-determinants of the spreading process that are hard to quantify with traditional data sources,

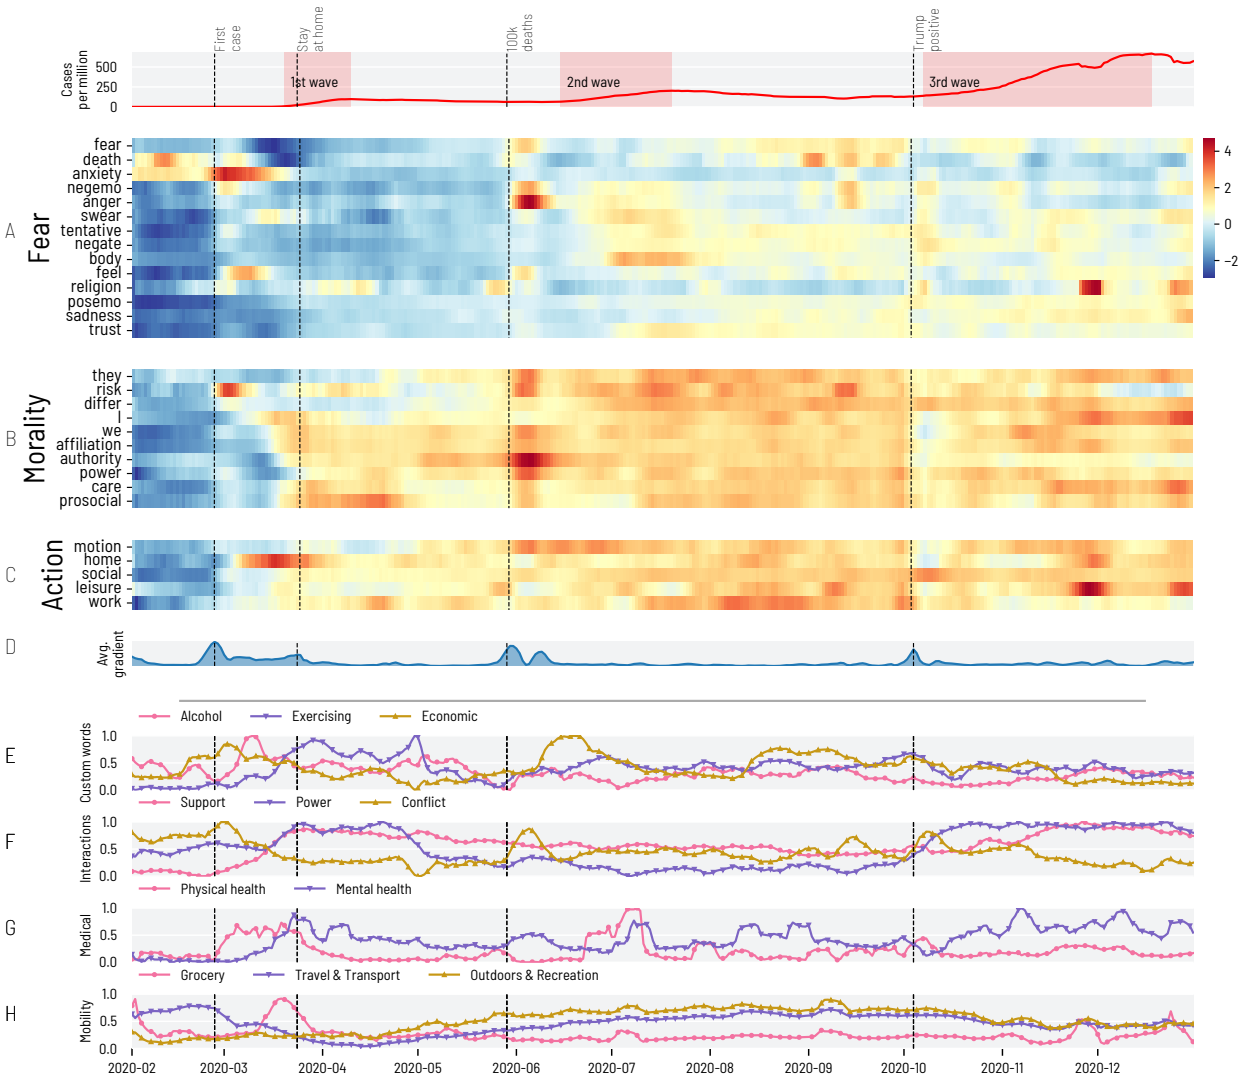


FIGURE 4. The Epidemic Psychology on Twitter during the three contagion waves (shown in the top most panel) during the entire year of 2020. (A-H) temporally expands (A-H) in Figure 1 with the only difference that the heatmaps in (A-C) here show values that were standardized using the mean and standard deviation calculated on the whole period of study, from February 1<sup>st</sup> to December 31<sup>st</sup>.

especially in the first stages of the outbreak. This is particularly useful when modeling social and psychological processes such as risk perception [Bagnoli et al., 2007, Moinet et al., 2018]. Interestingly, these approaches are designed to deal with partial data, therefore they can benefit even from digital data that is incomplete and—like in the case of our Twitter-based study—not necessarily representative of the whole population [Salathe et al., 2012].

**Limitations.** Future work could improve our work in five main aspects. First, we focused only on one viral epidemic, without being able to compare it to others. Yet, if one were to obtain past social media data during the outbreaks of diseases like Zika [Fu et al., 2016], Ebola [Oyeyemi et al., 2014], and the H1N1 influenza [Chew and Eysenbach, 2010], one could apply our methodology in those contexts as well, and identify similarities and differences. For example, one could study how mortality rates or speed of spreading influence the representation of Strong’s epidemic psychology on social media.

Second, our geographical focus was the entire United States and, as such, was coarse and limited in scope. In Supplementary Materials, we broke-down the analysis of temporal phases for individual states in the U.S., and observed no substantial differences across states. In the future, one could conduct a more systematic analysis on

a finer geographical granularity, relate differences between states to known events (e.g., a governor’s decisions, prevalence of cases, media landscape, and residents’ cultural traits). In particular, recent studies suggested that the public reaction to COVID-19 varied across the U.S. states depending on their political leaning [Painter and Qiu, 2020, Grossman et al., 2020]. One could also apply our methodology to other English-speaking countries, to investigate how cultural dimensions [Hofstede et al., 2005] and cross-cultural personality trait variations [Bleidorn et al., 2013] might influence the three psycho-social epidemics.

Third, the three psycho-social epidemics were not always orthogonal to each other but did blend together at times. Future work could focus on those particular periods of time to determine whether the use either of finer-grained categorizations of language or of event detection techniques other than our change-point detection [Aiello et al., 2013] could disentangle those periods in theoretically meaningful ways.

Fourth, our study is limited to Twitter, mainly because Twitter is the largest open stream of real-time social media data. The practice of using Twitter as a way of modeling the psychological state of a country carries its own limitations. Despite having a rather high penetration in the U.S. (around 20% of adults, according to the latest estimates [Perrin and Anderson]), its user base is not representative of the general population [Li et al., 2013]. Additionally, Twitter is notoriously populated by bots [Ferrara et al., 2016, Varol et al., 2017], automated accounts that are often used to amplify specific topics or view points. Bots played an important role to steer the discussion on several events of broad public interest [Bessi and Ferrara, 2016, Broniatowski et al., 2018], and it is reasonable to expect that they have a role in COVID-related discussions too, as some studies seem to suggest [Yang et al., 2020]. To partly discount their impact, since they tend to have anomalous levels of activity (especially retweeting [Bessi and Ferrara, 2016]), we performed two tests. First, we computed all our measures at user-level rather than tweet-level, which counter anomalous levels of activity. Second, we replicated our temporal analysis excluding retweets, and obtained very similar results. In the future, one could attempt to adapt our framework to different sources of online data, for example to web search queries—which have proven useful to identify different phases of the public reactions to the COVID-19 pandemic [Husnayain et al., 2020].

Last, as Strong himself acknowledged in his seminal paper: “*any sharp separation between different types of epidemic psychology is a dubious business.*” Our work has operationalized each psycho-social epidemic independently. In the future, modeling the relationships among the three epidemics might identify hitherto hidden emergent properties.

#### REFERENCES

- L. M. Aiello, G. Petkos, C. Martin, D. Corney, S. Papadopoulos, R. Skraba, A. Göker, I. Kompatsiaris and A. Jaimes. Sensing trending topics in Twitter. *IEEE Transactions on Multimedia*, 15(6):1268–1282, 2013.
- A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf. FLAIR: An easy-to-use framework for state-of-the-art NLP. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, pages 54–59, 2019.
- G. W. Alpers, A. J. Winzelberg, C. Classen, H. Roberts, P. Dev, C. Koopman, and C. B. Taylor. Evaluation of computerized text analysis in an internet breast cancer support group. *Computers in Human Behavior*, 21(2): 361–376, 2005.
- J. Arguello, B. S. Butler, E. Joyce, R. Kraut, K. S. Ling, C. Rosé, and X. Wang. Talk to me: foundations for successful individual-group interactions in online communities. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 959–968, 2006.
- F. Bagnoli, P. Lio, and L. Sguanci. Risk perception in epidemic modeling. *Physical Review E*, 76(6):061904, 2007.
- S. Bansal, G. Chowell, L. Simonsen, A. Vespignani, and C. Viboud. Big data for infectious disease surveillance and modeling. *The Journal of infectious diseases*, 214(suppl\_4):S375–S379, 2016.
- L. Bareket-Bojmel, G. Shahar, and M. Margalit. Covid-19-related economic anxiety is as high as health anxiety: Findings from the usa, the uk, and israel. *International Journal of Cognitive Therapy*, page 1, 2020.
- A. Baronchelli. The emergence of consensus: a primer. *Royal Society open science*, 5(2):172189, 2018.
- C. T. Bauch and A. P. Galvani. Social factors in epidemiology. *Science*, 342(6154):47–49, 2013.
- A. I. Bento, T. Nguyen, C. Wing, F. Lozano-Rojas, Y.-Y. Ahn, and K. Simon. Evidence from internet search data shows information-seeking responses to news of local covid-19 cases. *Proceedings of the National Academy of Sciences*, 2020.



- A. Bessi and E. Ferrara. Social bots distort the 2016 us presidential election online discussion. *First Monday*, 21(11-7), 2016.
- P. M. Blau. *Exchange and Power in Social Life*. Transaction Publishers, 1964.
- W. Bleidorn, T. A. Klimstra, J. J. Denissen, P. J. Rentfrow, J. Potter, and S. D. Gosling. Personality maturation around the world: A cross-cultural examination of social-investment theory. *Psychological science*, 24(12):2530–2540, 2013.
- J. L. Borelli and D. A. Sbarra. Trauma history and linguistic self-focus moderate the course of psychological adjustment to divorce. *Journal of Social and Clinical Psychology*, 30(7):667–698, 2011.
- V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006.
- D. A. Broniatowski, A. M. Jamison, S. Qi, L. AlKulaib, T. Chen, A. Benton, S. C. Quinn, and M. Dredze. Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American journal of public health*, 108(10):1378–1384, 2018.
- S. K. Brooks, R. K. Webster, L. E. Smith, L. Woodland, S. Wessely, N. Greenberg, and G. J. Rubin. The psychological impact of quarantine and how to reduce it: rapid review of the evidence. *The Lancet*, 2020.
- P. Cap. *The language of fear: Communicating threat in public discourse*. Springer, 2016.
- E. Chen, K. Lerman, and E. Ferrara. Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health Surveillance*, 6(2):e19273, May 2020. ISSN 2369-2960. doi: 10.2196/19273.
- C. Chew and G. Eysenbach. Pandemics in the age of twitter: content analysis of tweets during the 2009 h1n1 outbreak. *PloS one*, 5(11), 2010.
- M. Choi, L. M. Aiello, K. Z. Varga, and D. Quercia. Ten social dimensions of conversations and relationships. In *Proceedings of The Web Conference, WWW*. ACM, 2020.
- M. Cinelli, W. Quattrociocchi, A. Galeazzi, C. M. Valensise, E. Brugnoli, A. L. Schmidt, P. Zola, F. Zollo, and A. Scala. The covid-19 social media infodemic. *arXiv preprint arXiv:2003.05004*, 2020.
- G. Coppersmith, M. Dredze, and C. Harman. Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pages 51–60, 2014.
- B. L. Da, G. Y. Im, and T. D. Schiano. Covid-19 hangover: a rising tide of alcohol use disorder and alcohol-associated liver disease. *Hepatology*, 2020.
- S. Deri, J. Rappaz, L. M. Aiello, and D. Quercia. Coloring in the links: Capturing social ties as they are perceived. In *Proceedings of the ACM conference on Computer Supported Cooperative Work and Social Computing, CSCW*, pages 1–18. ACM, 2018.
- M. Deutsch. Trust and suspicion. *Journal of conflict resolution*, 2(4):265–279, 1958.
- M. Dredze, M. J. Paul, S. Bergsma, and H. Tran. Carmen: A twitter geolocation system with applications to public health. In *Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- F. Dutheil, L. Mondillon, and V. Navel. Ptsd as the second tsunamis of the sars-cov-2 pandemic. *Psychological Medicine*, pages 1–2, 2020.
- M. ElSherief, V. Kulkarni, D. Nguyen, W. Y. Wang, and E. Belding. Hate lingo: A target-based linguistic analysis of hate speech in social media. In *Twelfth International AAAI Conference on Web and Social Media*, 2018.
- S. Engle, J. Stromme, and A. Zhou. Staying at home: mobility effects of covid-19. *Available at SSRN*, 2020.
- E. Ferrara. #covid-19 on twitter: Bots, conspiracies, and social media activism. *arXiv preprint arXiv:2004.09531*, 2020.
- E. Ferrara, O. Varol, C. Davis, F. Menczer, and A. Flammini. The rise of social bots. *Communications of the ACM*, 59(7):96–104, 2016.

- T. Fetzter, L. Hensel, J. Hermle, and C. Roth. Coronavirus perceptions and economic anxiety. *Review of Economics and Statistics*, pages 1–36, 2020.
- L. Figea, L. Kaati, and R. Scrivens. Measuring online affects in a white supremacy forum. In *2016 IEEE conference on intelligence and security informatics (ISI)*, pages 85–90. IEEE, 2016.
- I. Finlay and I. Gilmore. Covid-19 and alcohol—a dangerous cocktail, 2020.
- S. T. Fiske, A. J. Cuddy, and P. Glick. Universal Dimensions of Social Cognition: Warmth and Competence. *Trends in cognitive sciences*, 11(2):77–83, 2007.
- J. A. Frimer, N. K. Schaefer, and H. Oakes. Moral actor, selfish agent. *Journal of personality and social psychology*, 106(5):790, 2014.
- K.-W. Fu, H. Liang, N. Saroha, Z. T. H. Tse, P. Ip, and I. C.-H. Fung. How people react to zika virus outbreaks on twitter? a computational content analysis. *American journal of infection control*, 44(12):1700–1702, 2016.
- S. Galea, R. M. Merchant, and N. Lurie. The mental health consequences of covid-19 and physical distancing: The need for prevention and early intervention. *JAMA internal medicine*, 180(6):817–818, 2020.
- S. Gao, J. Rao, Y. Kang, Y. Liang, and J. Kruse. Mapping county-level mobility pattern changes in the united states in response to covid-19. *SIGSPATIAL Special*, 12(1):16–26, 2020.
- G. R. Gibbs. Thematic coding and categorizing. *Analyzing qualitative data*. London: Sage, pages 38–56, 2007.
- A. J. Gill, R. M. French, D. Gergle, and J. Oberlander. The language of emotion in short blog texts. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, pages 299–302, 2008.
- E. Goffman. *Stigma: Notes on the management of spoiled identity*. Simon and Schuster, 2009.
- M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi. Understanding individual human mobility patterns. *nature*, 453(7196):779–782, 2008.
- J. Graham, J. Haidt, and B. A. Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029, 2009.
- J. Graham, J. Haidt, S. Koleva, M. Motyl, R. Iyer, S. P. Wojcik, and P. H. Ditto. Moral foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental social psychology*, volume 47, pages 55–130. Elsevier, 2013.
- G. Grossman, S. Kim, J. Rexer, and H. Thirumurthy. Political partisanship influences behavioral responses to governors’ recommendations for covid-19 prevention in the united states. *Available at SSRN 3578695*, 2020.
- J. R. Hebert, L. Clemow, L. Pbert, I. S. Ockene, J. K. Ockene. Social desirability bias in dietary self-report may compromise the validity of dietary intake measures. *International journal of epidemiology*, 24(2):389–398, 1995.
- S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- G. H. Hofstede, G. J. Hofstede, and M. Minkov. *Cultures and organizations: Software of the mind*, volume 2. McGraw-hill New York, 2005.
- E. A. Holmes, R. C. O’Connor, V. H. Perry, I. Tracey, S. Wessely, L. Arseneault, C. Ballard, H. Christensen, R. C. Silver, I. Everall, et al. Multidisciplinary research priorities for the covid-19 pandemic: a call for action for mental health science. *The Lancet Psychiatry*, 2020.
- Z. Hou, F. Du, H. Jiang, X. Zhou, and L. Lin. Assessment of public attention, risk perception, emotional and behavioural responses to the covid-19 outbreak: social media surveillance in china. *Risk Perception, Emotional and Behavioural Responses to the COVID-19 Outbreak: Social Media Surveillance in China (3/6/2020)*, 2020.
- Z. Huang, W. Xu, and K. Yu. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*, 2015.
- M. Hunt, J. Auriemma, A. Cashaw. Self-report bias and underreporting of depression on the BDI-II *Journal of personality assessment*, 80(1):26–30, 2003.

- A. Husnayain, A. Fuad, and E. C.-Y. Su. Applications of google search trends for risk communication in infectious disease management: A case study of covid-19 outbreak in taiwan. *International Journal of Infectious Diseases*, 2020.
- A. Jimeno-Yepes, A. MacKinlay, B. Han, and Q. Chen. Identifying diseases, drugs, and symptoms in twitter. *Studies in Health Technology and Informatics*, 216:643, 2015.
- T. Johnson, M. Fendrich. Modeling sources of self-report bias in a survey of drug use epidemiology *Annals of epidemiology*, 15(5):381–389, 2005.
- J. H. Kahn, R. M. Tobin, A. E. Massey, and J. A. Anderson. Measuring emotional expression with the linguistic inquiry and word count. *The American journal of psychology*, pages 263–286, 2007.
- S. Khan and D. Huremović. Psychology of the pandemic. In *Psychiatry of Pandemics*, pages 37–44. Springer, 2019.
- R. Kornfield, C. L. Toma, D. V. Shah, T. J. Moon, and D. H. Gustafson. What do you say before you relapse? how language use in a peer-to-peer online discussion forum predicts risky drinking among those in recovery. *Health communication*, 33(9):1184–1193, 2018.
- R. Kouzy, J. Abi Jaoude, A. Kraitem, M. B. El Alam, B. Karam, E. Adib, J. Zarka, C. Traboulsi, E. W. Akl, and K. Baddour. Coronavirus goes viral: quantifying the covid-19 misinformation epidemic on twitter. *Cureus*, 12(3), 2020.
- E. Kübler-Ross, S. Wessler, and L. V. Avioli. On death and dying. *Jama*, 221(2):174–179, 1972.
- L. Li, M. F. Goodchild, and B. Xu. Spatial, temporal, and socioeconomic patterns in the use of twitter and flickr. *Cartography and geographic information science*, 40(2):61–77, 2013.
- S. Li, Y. Wang, J. Xue, N. Zhao, and T. Zhu. The impact of covid-19 epidemic declaration on psychological consequences: a study on active weibo users. *International journal of environmental research and public health*, 17(6):2032, 2020.
- L. Liang, H. Ren, R. Cao, Y. Hu, Z. Qin, C. Li, and S. Mei. The effect of covid-19 on youth mental health. *Psychiatric Quarterly*, pages 1–12, 2020.
- Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- H. Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005.
- A. H. Maslow. A theory of human motivation. *Psychological review*, 50(4):370, 1943.
- P. McConnell. Banks and avian flu: Planning for a possible pandemic. *Risk Trading Technology*, 2005.
- S. M. Mohammad and P. D. Turney. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013.
- A. Moinet, R. Pastor-Satorras, and A. Barrat. Effect of risk perception on epidemic spreading in temporal networks. *Physical Review E*, 97(1):012313, 2018.
- B. O’Connor, R. Balasubramanyan, B. R. Routledge, and N. A. Smith. From tweets to polls: Linking text sentiment to public opinion time series. In *Fourth international AAAI conference on weblogs and social media*, 2010.
- S. O. Oyeyemi, E. Gabarron, and R. Wynn. Ebola, twitter, and misinformation: a dangerous combination? *Bmj*, 349:g6178, 2014.
- M. Painter and T. Qiu. Political beliefs affect compliance with covid-19 social distancing orders. *Available at SSRN 3569098*, 2020.
- G. Palshikar et al. Simple algorithms for peak detection in time-series. In *Proc. 1st Int. Conf. Advanced Data Analysis, Business Analytics and Intelligence*, volume 122, 2009.
- A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *Proceedings of the Advances in Neural Information Processing Systems Autodiff Workshop*, 2017.

- J. Pennington, R. Socher, and C. Manning. GloVe: Global vectors for word representation. In *Proceedings of the conference on empirical methods in natural language processing*, pages 1532–1543. Association for Computational Linguistics, 2014.
- A. Perrin and M. Anderson. Share of u.s. adults using social media, including facebook, is mostly unchanged since 2018. URL <https://www.pewresearch.org/fact-tank/2019/04/10/share-of-u-s-adults-using-social-media-including-facebook-is-mostly-unchanged-since-2018/>.
- R. Plutchik. *The emotions*. University Press of America, 1991.
- C. M. Pulido, B. Villarejo-Carballido, G. Redondo-Sama, and A. Gómez. Covid-19 infodemic: More retweets for science-based information on coronavirus than for false information. *International Sociology*, page 0268580920914755, 2020.
- J. Qiu, B. Shen, M. Zhao, Z. Wang, B. Xie, and Y. Xu. A nationwide survey of psychological distress among chinese people in the covid-19 epidemic: implications and policy recommendations. *General psychiatry*, 33(2), 2020.
- R. Rezapour, S. H. Shah, and J. Diesner. Enhancing the measurement of social effects by capturing morality. In *Proceedings of the Tenth Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 35–45, 2019.
- M. A. Russell. *Mining the social web: data mining Facebook, Twitter, LinkedIn, Google+, GitHub, and more.* ” O’Reilly Media, Inc.”, 2013.
- H. Saif, Y. He, and H. Alani. Alleviating data sparsity for twitter sentiment analysis. CEUR Workshop Proceedings (CEUR-WS. org), 2012.
- M. Salathe, L. Bengtsson, T. J. Bodnar, D. D. Brewer, J. S. Brownstein, C. Buckee, E. M. Campbell, C. Cattuto, S. Khandelwal, P. L. Mabry, et al. Digital epidemiology. *PLoS Comput Biol*, 8(7):e1002616, 2012.
- S. Scepianovic, E. Martin-Lopez, D. Quercia, and K. Baykaner. Extracting medical entities from social media. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pages 170–181, 2020.
- A. Schutz, T. Luckmann, R. Zaner, and J. Engelhardt. *The Structures of the Life-world*. Number v. 1 in Northwestern University studies in phenomenology & existential philosophy. Northwestern University Press, 1973. ISBN 9780810106222. URL <https://books.google.co.uk/books?id=LGXBxIOXsh8C>.
- B. Shaw, J. Y. Han, E. Kim, D. Gustafson, R. Hawkins, J. Cleary, F. McTavish, S. Pingree, P. Eliason, and C. Lumpkins. Effects of prayer and religious expression within computer support groups on women with breast cancer. *Psycho-Oncology: Journal of the Psychological, Social and Behavioral Dimensions of Cancer*, 16(7): 676–687, 2007.
- J. H. Shen and F. Rudzicz. Detecting anxiety through reddit. In *Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology—From Linguistic Signal to Clinical Reality*, pages 58–65, 2017.
- J. M. Shultz, F. Baingana, and Y. Neria. The 2014 ebola outbreak and mental health: current status and recommended response. *Jama*, 313(6):567–568, 2015.
- T. Simms, C. Ramstedt, M. Rich, M. Richards, T. Martinez, and C. Giraud-Carrier. Detecting cognitive distortions through machine learning text analytics. In *2017 IEEE international conference on healthcare informatics (ICHI)*, pages 508–512. IEEE, 2017.
- J. A. Smith and P. Shinebourne. *Interpretative phenomenological analysis*. American Psychological Association, 2012.
- S. Sommariva, C. Vamos, A. Mantzarlis, L. U.-L. Đào, and D. Martinez Tyson. Spreading the (fake) news: exploring health messages on social media and the implications for health professionals using a case study. *American journal of health education*, 49(4):246–255, 2018.
- P. Strong. Epidemic psychology: a model. *Sociology of Health & Illness*, 12(3):249–259, 1990.
- H. Tajfel, J. C. Turner, W. G. Austin, and S. Worchel. An integrative theory of intergroup conflict. *Organizational Identity*, 1979.
- Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54, 2010.

- The Economist. With millions stuck at home, the online wellness industry is booming. <https://www.economist.com/international/2020/04/04/with-millions-stuck-at-home-the-online-wellness-industry-is-booming>.
- J. J. Van Bavel, K. Baicker, P. S. Boggio, V. Capraro, A. Cichocka, M. Cikara, M. J. Crockett, A. J. Crum, K. M. Douglas, J. N. Druckman, et al. Using social and behavioural science to support covid-19 pandemic response. *Nature Human Behaviour*, pages 1–12, 2020.
- O. Varol, E. Ferrara, C. A. Davis, F. Menczer, and A. Flammini. Online human-bot interactions: Detection, estimation, and characterization. In *Eleventh international AAAI conference on web and social media*, 2017.
- C. Wang, R. Pan, X. Wan, Y. Tan, L. Xu, C. S. Ho, and R. C. Ho. Immediate psychological responses and associated factors during the initial stage of the 2019 coronavirus disease (covid-19) epidemic among the general population in china. *International journal of environmental research and public health*, 17(5):1729, 2020.
- S. F. Waterloo, S. E. Baumgartner, J. Peter, and P. M. Valkenburg. Norms of online expressions of emotion: Comparing facebook, twitter, instagram, and whatsapp. *new media & society*, 20(5):1813–1831, 2018.
- L. C. Windsor, N. Dowell, and A. Graesser. The language of autocrats: Leaders’ language in natural disaster crises. *Risk, Hazards & Crisis in Public Policy*, 5(4):446–467, 2014.
- S. Wojcik and A. Hughes. Sizing up twitter users. *Washington, DC: Pew Research Center*, 2019.
- M. Wolf, F. Theis, and H. Kordy. Language use in eating disorder blogs: Psychological implications of social online activity. *Journal of Language and Social Psychology*, 32(2):212–226, 2013.
- M. J. Wood. Propagating and debunking conspiracy theories on twitter during the 2015–2016 zika virus outbreak. *Cyberpsychology, behavior, and social networking*, 21(8):485–490, 2018.
- K.-C. Yang, C. Torres-Lugo, and F. Menczer. Prevalence of low-credibility information on twitter during the covid-19 outbreak. *arXiv preprint arXiv:2004.14484*, 2020.

#### ACKNOWLEDGMENTS

We thank Sarah Konrath, Rosta Farzan, and Licia Capra for their useful feedback on the manuscript. This research was partly supported by the EU Grant “GO GREEN ROUTES” no. 869764.

#### DATA AVAILABILITY

The daily aggregates of the measurements are available at: <https://doi.org/10.6084/m9.figshare.14892642.v1>. The tweet IDs we used are available at <https://github.com/echen102/COVID-19-TweetIDs>. Other datasets and visualizations are available on the project’s site <http://social-dynamics.net/EpidemicPsychology>.

#### COMPETING INTERESTS

The author(s) declare no competing interests.

## SUPPLEMENTARY MATERIALS

**Spatial representativeness of Twitter data.** To make sure that the geographical penetration of our Twitter data reflects the spatial distribution of the U.S. population, we compared the user activity in each state with the census population. Figure S1 shows that both the number of Twitter users and the volume of their tweets correlate very strongly with the census population population estimates in year 2019 at the level of US states (<https://www.census.gov>).




FIGURE S1. Correlation of the census population of US states estimated in year 2019 with: the number of tweets posted in those states (left), and the number of unique Twitter users who posted at least one tweet (right) during the period of analysis. A linear fit is shown together with the coefficient of determination for a linear regression ( $R^2$ ), the Pearson correlation coefficient ( $r$ ), and the coefficient of the fit ( $\beta$ ). The number of tweets and that of Twitter users scale almost linearly with population estimates.

**Signals of racism in Twitter discussions.** To corroborate the qualitative intuition that the first phase was characterized by discussions that depicted the pandemic as a foreign problem, often using racist undertones, we measured the volume of tweets containing hashtags that were clear indicators of content that is either racist or aimed at antagonizing China. We manually selected those hashtags from the list of hashtags that appeared at least 100 times in the dataset. Figure S2 shows the min-max normalized volume of those hashtags over time; they peak during the first phase, when fear-related mentions and the focus of Twitter discussions on China and Asian foreigners were at their peaks.

**Breakdown of phases by state.** During the period of analysis, the country was fragmented into regions that were dealing with different stages of the outbreak, and the policies to contrast it varied across states. To assess whether those differences impacted the unfolding of the psycho-social epidemics across states, we computed the measures that represent the relative presence of the three phases over time ( $\Delta\%_{Refusal}$ ,  $\Delta\%_{Anger}$ ,  $\Delta\%_{Acceptance}$ ) for each state individually. The results are shown in Figure S3. Despite minor differences, the state-level curves indicate that: *i*) the three phases took place in all states; *ii*) the sequence of the phases is consistent across states (refusal is followed by anger, which is then followed by acceptance); and *iii*) there is very little variation of the times of transition between phases.




FIGURE S2. Min-max normalized frequency of hashtags connected to content that is xenophobic, racist, or generally aimed at antagonizing China during the COVID-19 pandemic. #WuhanVirus characterizes tweets that claimed either that the virus was voluntarily caused by China or that it was an issue limited to China. #QAnon denotes content supporting the so-called QAnon conspiracy that, during the first stages of the COVID pandemic, was amplifying narratives about Asian individuals being more vulnerable to the virus than white people, often using racist undertones.




FIGURE S3. Breakdown by US states of the temporal evolution of the relative frequency of a selected subset of the language categories associated with each of the three phases of refusal, anger, and acceptance.