



HAL
open science

Mental health effects of COVID-19 lockdowns: a Twitter-based analysis

Sara Colella, Frédéric Dufourt, Vincent A Hildebrand, Rémi Vivès

► **To cite this version:**

Sara Colella, Frédéric Dufourt, Vincent A Hildebrand, Rémi Vivès. Mental health effects of COVID-19 lockdowns: a Twitter-based analysis. 2022. hal-03740701v1

HAL Id: hal-03740701

<https://amu.hal.science/hal-03740701v1>

Preprint submitted on 29 Jul 2022 (v1), last revised 2 Oct 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mental health effects of COVID-19 lockdowns: a Twitter-based analysis

Sara Colella
Frédéric Dufourt
Vincent A. Hildebrand
Rémi Vivès

WP 2022- Nr 18

Mental health effects of COVID-19 lockdowns: a Twitter-based analysis*

Sara Colella[†]

Aix Marseille Univ, CNRS, AMSE,
Marseille, France

Frédéric Dufourt[‡]

Aix Marseille Univ, CNRS, AMSE,
Marseille, France

Vincent A. Hildebrand[§]

Glendon College, York University
Dalla Lana School of Public Health, University of Toronto

Rémi Vivès[¶]

Glendon College, York University

July 2022

Abstract

We derive a mental health indicator measuring the frequency of words expressing anger, anxiety and sadness from a fixed population of Twitter users located in France. During the first COVID-19 lockdown, our indicator did not reveal a statistically significant mental health response, while the second lockdown triggered a sharp and persistent deterioration in all three emotions. In addition, DID and event study estimates show a more severe mental health deterioration among women and younger users during the second lockdown. Our results suggest that successive stay-at-home orders significantly worsen mental health across a large segment of the population.

JEL classification: C81, I12, I18, I31

Keywords: COVID-19, lockdown, mental health, Twitter data, well-being

*The project leading to this publication has received funding from the French government under the “France 2030” investment plan managed by the French National Research Agency (reference: ANR-17-EURE-0020) and from Excellence Initiative of Aix-Marseille University - A*MIDEX. The authors are grateful to Boriana Miloucheva and Philippe Van Kerm for their helpful comments and suggestions. All errors remain our own.

[†]Address: 5-9 Bd Maurice Bourdet, 13001 Marseille, France., e-mail address: sara5colella@gmail.com

[‡]Address: 5-9 Bd Maurice Bourdet, 13001 Marseille, France., e-mail address: frederic.dufourt@univ-amu.fr

[§]Address: York Hall 363, 2275 Bayview Ave., Toronto, ON, M4N 3M6 Canada, e-mail address: vincenth@yorku.ca

[¶]Address: York Hall 341A, 2275 Bayview Ave., Toronto, ON, M4N 3M6 Canada, e-mail address: rvives@yorku.ca

1 Introduction

The COVID-19 pandemic spurred many countries around the world to adopt drastic lockdown policies to contain the spread of the virus and reduce the number of new daily cases. The benefit of these policies is clear: their success in minimizing the spread of the virus has been undeniable and widely observed in many countries around the world (Alfano and Ercolano, 2020). Yet containment measures also come at various costs, both from an economic standpoint and in terms of mental health at the individual and population level.

It is well-documented that deteriorating mental health conditions at the onset of the pandemic coincided with the implementation of the most restrictive containment measures. For example, a number of studies exploiting survey data have reported a significant mental health deterioration in the British population (Davillas and Jones, 2021; Oreffice and Quintana-Domeque, 2021; Pierce et al., 2020) that disproportionately affected women and young people (Banks and Xu, 2020). Revealed measures of psychological distress, evidenced by sizeable increase in calls to national helplines during the most restrictive phases of lockdowns, have been documented in several countries including Austria, Germany and France and the US (Arendt et al., 2020; Armbruster and Klotzbücher, 2020; Brülhart et al., 2021). Increased search volumes for terms related to boredom, loneliness, and sadness in Europe and the United States from Google trend data at the onset of the pandemic provide further evidence of that deteriorating mental health during lockdowns (Brodeur et al., 2021; Silverio-Murillo et al., 2021). Interestingly, aforementioned studies hint that worsening mental health appears to be mostly driven by feelings of anxiety, loneliness and fear of social isolation, rather than financial concerns or fear of contracting the virus.

Properly assessing costs is often a difficult but crucial step in any cost-benefit analysis, and the stakes are particularly important if one wishes to evaluate the social desirability of imposing new stay-at-home orders as COVID-19 strains with diverse degrees of virulence continue to emerge. This paper addresses this question by examining the costs in terms of psychological well-being during the first two lockdown episodes in France. To the best of our knowledge, this is the first study that explicitly documents the costs of successive confinements on psychological distress in France.

France provides an interesting case study for two main reasons. First, France is representative of "European-style" countries with a large welfare state, in which the individual economic costs and risks associated with the pandemic were largely covered by the government.¹ As a result, if mental health deteriorates during these two lockdown episodes in France, it is much

¹See section 2 for further details.

more likely that this deterioration can be attributed to the effects of social isolation *per se* than to the other potentially significant individual economic risks and costs resulting from stay-at-home orders in countries with smaller government safety nets. Second, except for the first two weeks of the second lockdown, during which the lockdown policies were gradually implemented over the entire territory, the first two lockdown episodes in France were for the most part *nationwide*.

Our main findings suggest, based on the analysis of the first two lockdowns, that the mental health costs of successive waves of containment orders sharply increase over time. In contrast to aforementioned studies, we find no significant mental health response during the first lockdown order in France. However, we find a sharp and statistically significant deterioration during the second lockdown that is of especially large magnitude among the younger population.

To reach this conclusion, we create a unique longitudinal dataset of active Twitter users located in France. We measure mental health with indicators built using the Linguistic Inquiry and Word Count (LIWC) lexicon to identify words in a given tweet related to negative emotions such as anger, anxiety or sadness. We use a machine learning algorithm to infer the age and gender of the Twitter users. Inferring this demographic data enables us to assess whether mental health variations are heterogeneous across women and men and different age groups of the population.

The key factor distinguishing our dataset from other COVID-19 related Twitter studies is that we do not filter tweets based on COVID-19-related keywords (see among others [Balech et al. 2020](#) and [Chen et al. 2020](#)). By working instead with a fixed population of users and their tweet archives, we minimize sampling bias. For example, selecting tweets according to pandemic-related keywords may in fact overrepresent negative emotions, since tweets mentioning COVID-19 may be more likely to convey negative emotions. A set of tweets filtered by keywords may also not provide a representative sample of the Twittersphere. In addition, the entire Twitter history of a user likely reflects their psychological well-being more accurately than does a subset of their tweets that only includes pandemic-specific keywords.

Our analysis is based on both a descriptive and an econometric approach. In the descriptive approach, we draw informal conclusions by simply inspecting the evolution of our mental health indicators before and during lockdown episodes, emphasizing the potential relevance of such a tool for monitoring mental health conditions in real time. We base our econometric approach on two models: a difference-in-difference estimation and an event-study analysis, considering the year 2019 as the control period for both. This econometric analysis enables us to exploit the unique longitudinal aspect of our dataset to test more formally whether the changes in mental health during the two lockdown episodes were significantly different from the changes over the same period a year before. In addition, since our dataset includes the predicted gender

of Twitter users and their predicted age group, we also test whether women experienced a greater mental health deterioration during lockdown episodes than men (the asymmetric "mental load" hypothesis) and whether young people were disproportionately more affected by these lockdown episodes than older people (the age difference hypothesis).

This work contributes to the literature related to COVID-19 in three distinctive ways. First, we provide new evidence from France that repeated lockdown policies are associated with increasing mental health costs. Along the same lines, exploiting a natural experiment arising from differences in the length of lockdown restrictions between England and Scotland, [Serrano-Alarcón et al. \(2021\)](#) suggest that stretching out lockdown restrictions worsens mental health inequalities between socio-economic groups. This result has significant potential implications that should be taken into consideration in the management of future variants or the management of future epidemics/pandemics in general. Our second contribution is methodological: we create a unique longitudinal dataset of a fixed Twitter population, along with their demographic attributes and their tweets covering the year of the first stages of the pandemic (2020) and the year prior (2019). Our dataset differentiates this contribution from other COVID-19-related Twitter studies because it exploits a more representative sample of the Twittersphere and the strength of its longitudinal nature allows us to exploit within-individual variations in the users' psychological well-being before and after each lockdown. Finally, we argue that a dataset of Twitter users such as ours could form the foundation for a relevant tool for monitoring changes in mental health in a particular geographic population, both during pandemic episodes and potentially during other large-scale events affecting the population.

Our mental health indicator is, of course, not free of limitations. First, Twitter users are not a representative sample of the population. Beyond this, mental health is very broadly assessed through the use of words related to negative emotions, rather than finer descriptions that can be derived from questionnaires and surveys usually deployed in medical or psychological studies. A twitter-based indicator such as ours, however, does confer several advantages. Data collection from Twitter is fast, easy and essentially costless. Once the procedure for constructing the indicator is implemented, the procedure can handle a large number of observations drawn from a potentially very large number of users,² and the resulting indicator can be monitored almost in real time. For example, the sharp degradation of mental health during the second lockdown episode could have been detected in a matter of days, or at most one or two weeks, had our indicator been available in real time.

The remainder of this paper is organized as follows. Section 2 provides the chronology of the COVID-19 pandemic in France for the period under study. Section 3 documents the

²Our main limitation is Twitter restrictions on the number of tweets that can be retrieved from the past.

construction of our Twitter dataset and mental health indicators. Section 4 describes our sample and presents descriptive statistics. Section 5 presents the empirical strategy together with the results. Section 6 discusses our main findings and Section 7 offers some concluding remarks.

2 Background

Timeline of COVID-19 Restrictions

The World Health Organization declared a global pandemic on March 11th, 2020, following the rapid spread of novel coronavirus (COVID-19). This declaration spurred numerous governments into enacting orders that restricted individuals' movement to help curb the spread of the disease. In France, schools, universities and all non-essential businesses (restaurants, cafes, movie theatres, etc.) were shut down on March 14th. On March 16th, the French government announced the implementation of a lockdown with strict mobility restrictions effective at noon the next day for "at least 15 days". Individuals were only allowed to leave their homes for specific reasons disclosed on an official form ("attestation de déplacement dérogatoire") that had to be carried at all times while outside of the home. To ensure compliance with the stay-at-home order, the government deployed 100,000 police officers who were empowered to randomly stop pedestrians and ask to see their attestations. Offenders caught without the form were fined €135 and up to €3,750 in the event of a repeat offense. Permitted reasons for leaving the home were limited to essential activities, including assisting vulnerable individuals, going to work (if it was not possible to work from home), and purchasing needed groceries or medication. Leaving home for physical exercise was permitted, but only for an hour per day, and within a maximum radius of one kilometer around the home.

The strict lockdown measures of the first stay-at-home order were extended twice, ending approximately two months later on May 11th. Following that date, France entered a period of progressive deconfinement organized in two stages, proceeding differently by sub-region ("departement") depending on the number of COVID-19 cases, their testing capacity and the saturation level of local emergency departments.

This deconfinement strategy proceeded in tandem with a massive testing campaign. The first stage of deconfinement lasted just under a month (May 11th to June 2nd), comprising a gradual reopening of schools and stores. Individuals could leave their homes without having to carry an attestation and were permitted to travel up to 100 kilometers from their residence. The second stage of deconfinement started on June 2nd and lasted until the end of the month. During this stage, restaurants, cafes and bars reopened, as did cultural, sport and tourist venues

such as museums and hotels.

The first stay-at-home order was successful in reducing the spread of COVID-19 cases. At the end of the second stage of deconfinement (June 30th), there were 541 new daily positive cases compared to 5,233 at the peak of the first wave (April 3rd). The number of new daily positive cases remained low until mid-August but then began to steadily rise again, possibly due to individuals returning to work and the mandatory in-person return to school after summer break. On September 11th, the government announced that the virus was actively circulating in 42 of France's 101 departments, which reinforced the belief that a second wave of the coronavirus epidemic was already well under way.

On October 5th, Paris was put on "maximum alert" due to the spike in COVID-19 cases in the city and its suburbs.³ On October 14th, faced with the inability to curb the spread of the disease, President Macron ordered the enforcement of local curfews on October 17th for at least 4 weeks in the Paris region (Île-de-France) and 8 other cities (Grenoble, Lille, Lyon, Aix-Marseille, Saint-Étienne, Rouen, Montpellier and Toulouse). In the following days, as COVID-19 cases continued to spread across the country, curfews were extended to 54 "départements". Finally, on October 27th, President Macron announced a second nationwide lockdown effective on October 30th and expected to last until December 1st. Newly imposed mobility restrictions were to be reassessed every two weeks.

The second stay-at-home order was not as restrictive as the first one, but it nonetheless remained very strict. University classes were moved entirely online, even though some university libraries remained open. Elementary and secondary schools remained open but classes were shut down as soon as one pupil came into contact with an infected person. The need to carry an attestation to leave home was reinstated with a wider range of valid reasons for leaving home. Visits to nursing homes were permitted with strict hygiene requirements. Remote work was still strongly encouraged, but not required for those who were not able to work from home. Unlike the first lockdown, visits to public outdoor spaces such as parks, beaches, and hiking trails were permitted.

On November 24th, the government announced that the originally set date of December 1st to end the second stay-at-home order was to be extended to December 15th, and would then be replaced by a nighttime curfew,⁴ which lasted until June 20th, 2021.

Based on this description of events, for the purpose of this study, we date the first lockdown

³In France, a locality is designated "maximum alert" area when i) the infection rate in a locality exceeds 250 cases per 100,000 people; ii) the incidence rate among people over age 65 surpasses 100 cases per 100,000 people; iii) and at least 30% of intensive care beds are reserved for Covid-19 patients.

⁴There was an exemption to the nighttime curfew on Christmas eve, but it was strictly enforced for the New Year.

period as originating on March 14th, 2020 with the closure of schools, universities and all non-essential businesses, and ending on June 2nd, with their reopening and the lift of the travel ban, which effectively ended the period of social isolation. We date the second lockdown as beginning on October 14th, with the announcement of the enforcement of progressive curfews, and ending on December 15th, with the removal of travel restrictions and the requirement to carry an “attestation” when leaving home outside of imposed curfew hours.

Economic Support

We argue that France provides a pertinent case study to disentangle the mental health cost of repeated containment policies based on the argument that France adopted stringent nationwide containment policies rapidly supported by generous economic measures to support all economic actors.

On March 26th, to mitigate the deterioration of the business environment caused by the first lockdown, and to avoid economic layoffs, the French government eased the eligibility criteria for employers to qualify for a Partial Reduction of Activity scheme (PA). Employers benefiting from the COVID-19 amended PA scheme were to retain and compensate their employees for the number of working hours falling below the standard legally mandated workweek⁵ caused by a partial or a full temporary closure of operations. Under that scheme, employees were guaranteed to receive a grant from their employer of at least 70% of their gross earnings (or €8.03 net per hour) for each work hour lost. The state contribution of the grant was capped to 4.5 times the legal minimum wage, which indicates that employees were fully covered by the state for salaries up to €6924 a month (Hubbard and Strain, 2020; Foki, 2021). In addition, to further support businesses, France also offered tax deferrals and loans and provided grants to industries hit hardest by the pandemic; see Blanchard et al. (2020); Cahuc (2022) for further details.

3 Twitter Dataset and Twitter-Based Indicator of Mental Health

Since the emergence of the pandemic, a plethora of Twitter datasets related to COVID-19 have been created and openly shared; see, e.g., Balech et al. (2020), Chen et al. (2020), Gruzd and Mai (2020), Banda et al. (2021) and Gupta et al. (2021), among others.

⁵On February 1st, 2000, the statutory workweek in France was set a 35 hours for all companies with more than 20 employees, and extended on January 2002 for the rest.

These datasets are typically constructed by first filtering tweets using a set of predefined COVID-19 related keywords. For example, [Banda et al. \(2021\)](#) collected 1.12 billion tweets posted in English, French, German, Russian and Spanish that match keywords such as “2019CoV”, “WuhanVirus” and “pneumonia”. Likewise, [Chen et al. \(2020\)](#) collected 123 million real-time tweets starting on January 28th, 2020 that track tweets containing a list of COVID-19 related keywords that was gradually extended over time. [Balech et al. \(2020\)](#) created a dataset of primarily French tweets selected by using the hashtag #ConfinementJourXx (#Confinement-DayXx).

These papers contributed to the rapidly expanding area of COVID-19 research by creating datasets that were readily available to the research community from the onset of the pandemic. These early contributions generally adopted a “test-and-learn” approach by adapting their methodology along the way by, for example, extending the list of keywords used to retrieve tweets; see, e.g., [Chen et al. 2020](#), [Banda et al. 2021](#).

Our study further contributes to this literature by constructing a singular dataset that allows us to track the emotional well-being – hereafter referred to as “mental health” – of a sample of French users, which is as representative of the Twittersphere as possible. One key assumption of such approach is that the change in mental health of the population of France can at least be partially captured by the temporal variation of the textual content of our Twitter sample. To achieve this objective, the construction of our dataset differs from aforementioned contributions in several important ways.

3.1 Data Collection and Variable Refinement

Rather than initiating the data collection by filtering tweets using targeted keywords, we draw a fixed random population of active users from the 1 percent random sample of all tweets made publicly available by Twitter to the research community.⁶ We then retrieve all tweets posted by each user in our sample with the Twitter API over a period spanning from January 2019 to March 1st, 2021.⁷ This approach allows us to build a unique longitudinal sample that tracks the textual content of all tweets from a representative sample of the Twitter population of France.

Drawing a population of users as opposed to first selecting tweets by keywords offers several advantages. First, a user’s psychological health may be reflected in the textual content of their tweets without necessarily using pandemic-related keywords in their posts. By the same token,

⁶See, Twitter academic research portal, <https://developer.twitter.com/en/use-cases/do-research/academic-research>., for further details

⁷March 1st, 2021 is end point for the dataset used in this paper given our focus on the mental health impact of the first two confinements. The data collection for our population is, however, still ongoing and available at <https://twittersphereobservatory.github.io/>.

a user might tweet negatively about the pandemic, but those tweets might not reflect their overall well-being. In addition, a sample drawn by filtering tweets on a list of keywords may simply not be representative of the Twitter population; see for instance [Bruns et al. \(2017\)](#) and [King et al. \(2017\)](#). Drawing a random sample of Twitter users circumvents these selection issues. In this regard, our data collection is comparable to [Su et al. \(2020\)](#), who first sampled users located in Wuhan and Lombardy from Weibo and Twitter respectively, and then collected their tweets two weeks before and after the first lockdown in each region. However, the Twitter dataset constructed by the authors includes 14,269 tweets posted by 188 users, which because of its small size is limited in the statistical inference that can be drawn from it. Each step leading to the construction of our dataset is further detailed below.

Account selection

The first step of data collection consists in selecting active user accounts from the Twitter Archive database. This archive stores a 1 percent random sample of all tweets posted since September 2011. From this database, we extract a 5 percent random sample of users who tweeted between August and December 2019 and self-reported the location of their account as being in France. In selecting users based on self-reported geographic location, we follow [Mislove et al. \(2011\)](#) and [Durazzi et al. \(2021\)](#). We choose to extract users who tweeted between August and December 2019 for two reasons. First, we wanted to ensure that our users were active in 2019. Second, the latest available data of the Twitter Archive database at the beginning of this project were from that time period.⁸

To ensure that the reported location is in France, we cross-reference it against the French GeoNames database⁹ after normalizing both the twitter data and GeoNames. This database consists of a list of names of French regions, sub-regions, cities and towns, together with their geographical information such as population, GPS coordinates, alternative names, etc. We classify an account as being in France if the self-reported location matches a French location from the GeoNames database.

Misclassification errors may arise due to the lack of accuracy of self-reported (location) data. For example, while “Us” (“us” when normalized) is the name of a French town, a user who enters “us” in the location field might be referring to the United States. As another example, “Rue” (“rue” when normalized) is the name of a French town, but it is also a French word that means “street”, so a user might simply use that word in the sense of “street” in the location field. Furthermore, several French cities share their names with cities abroad, e.g. Montreal, France

⁸Twitter Archive database made data available with a 6-month lag.

⁹<http://geonames.org/>. Retr. June, 2020.

and Montreal, Canada.

We address the sources of potential misclassification errors in a number of ways. We exclude accounts whose reported city names have fewer than four characters. We eliminate similarly named cities that have larger populations outside of France. As an example, this procedure would remove St. Louis, France from our data, which has significantly fewer inhabitants than Saint Louis, United States, without excluding Paris, France, which has a significantly larger population than Paris, Texas (United States) or Paris, Ontario (Canada).

It is common in this literature to remove cities that are sparsely populated to minimize noise and to improve computing power. For instance, [Durazzi et al. \(2021\)](#) eliminate cities having less than 30,000 inhabitants. We take a more conservative approach and drop cities that have 500 inhabitants or less.

Finally, we restrict our data to users who reported as location the name of a region, a sub-region, or a city, rather than simply “France”. This user-generated geographic information is then used to assign users to their corresponding NUTS-1 region.¹⁰ For users who entered multiple locations, we assume the first reported location to be the primary location of residence.

Tweet Collection

In a second step, we collect all available tweets associated with each randomly drawn account using the Twitter Application Programming Interface (API). To complete this step, for each user, we scrape data backwards in time up to January 1st, 2019. The Twitter API only allows us to scrape up to 3,200 tweets from a user within a selected time period, which means that for some very active users, our data may not extend all the way back to January 1st, 2019. Including data covering the year 2019 provides an essential comparison group that was not exposed to the pandemic and could be used to control for the potential presence of seasonal patterns of mental health.

Data Refinement

We finalize our dataset with the following data refinements. We only keep tweets written in French, because our study focuses on analysing the textual content of tweets written in French. This filtering also brings the additional benefit of further minimizing potential user misclassification and/or geolocation error. For each remaining account, we infer basic demographic information (i.e. age group, gender, and whether the account is held by an individual or an organization) using the M3-Inference tool ([Wang et al., 2019](#)). This state-of-the-art deep learning

¹⁰See, <https://ec.europa.eu/eurostat/web/nuts/background>, for further details on the NUTS classification system.

model exploits reported usernames, profile names, profile description, and profile images to infer demographic characteristics of user accounts. The M3 model was trained to operate with major European languages, including French. Because we are primarily interested in measuring individual well-being, we remove 3,278 accounts identified as being held by organizations.

3.2 Textual Analysis and Construction of a Mental Health Indicator

An important contribution of our data work is the construction of daily indicators capturing each user’s emotional state based on the textual analysis of their tweets using the Linguistic Inquiry and Word Count (LIWC) corpus. These indicators form the basis of our analysis.

LIWC dictionaries are widely used in the field of language psychology and well-validated to infer behavioural outcomes (Boyd and Schwartz, 2020). A number of psycholinguistic studies have exploited LIWC dictionaries to classify Twitter users along psychological conditions such as depression, bipolar disorder, and post-traumatic stress disorder (Park et al., 2012; Coppersmith et al., 2014). Since the advent of the COVID-19 pandemic, there has been a surge of studies leveraging the LIWC corpus to assess the impact of the pandemic on mental health outcomes using the textual content of tweets. For example, Zhang et al. (2020) use the LIWC corpus to examine the change in mental health of depressed users during the pandemic. Aiello et al. (2021) use the LIWC corpus to identify whether psycho-social responses to the pandemic occur in phases (refusal, anger, and acceptance). Dyer and Kolic (2020) track the relationship between the progression of the pandemic, and the public’s perception of its risk.

In this paper, we derived our emotion indicators in two steps. First, we preprocess the textual content of each tweet following a commonly adopted procedure: we remove numbers, punctuation, hashtag signs, mentions, URL, emojis, stopwords¹¹ and websites. We also convert the text to lower case. Second, we run an emotion analysis based on the LIWC dictionary in French (Tausczik and Pennebaker, 2010; Piolat et al., 2011; Garcia and Rimé, 2019) to classify the textual content of each tweet standardized in step 1 into three negative emotions that make up our indicator of mental health, namely anger, anxiety and sadness.

An additional innovation of our approach over previous contributions using LIWC dictionaries is the systematic treatment of negation. In our study, if the word ‘pas’ (‘not’ in French) is placed right before/after a word with a match in LIWC, then the emotion associated with this word is reclassified as neutral. In the sentence, ‘je ne suis pas triste’ (I am not sad), ‘pas’ is

¹¹Stopwords are words that are commonly used and do not convey useful information in the case of this study, such as articles, pronouns and prepositions.

right before ‘triste’, which belongs to the lexical field of sadness. Because the negation (‘pas’) appears right before the word indicating an emotion (‘triste’), the emotional tone of the tweet is negated, and the resulting sadness index would be 0. In this example, failing to control for negation would wrongly generate an indicator of sadness of 1/5.¹² The systematic treatment of commonly used words to negate ideas has the benefit of significantly minimizing potential misclassification of emotions. For each individual user i posting tweets in day t , we use the outcome of the LIWC classification to derive daily indicators of anger, anxiety and sadness. Examples of French-language tweets and their associated LIWC matches can be found in Table 3 in the Appendix.

Each daily indicator measures the frequency of words in the *standardized* text of all tweets posted by user i on day t that can be matched to the lexical field of anger, anxiety and sadness. For instance, the daily indicator of individual anger is defined as:

$$Anger_{i,t} = \frac{\text{Number of words matched to the lexical field of anger in tweets posted by } i \text{ on day } t}{\text{Number of words in all tweets sent by } i \text{ on day } t}$$

Finally, we derive a daily indicator of mental health for each user, $MH_{i,t}$, defined as the ratio of the sum of all words belonging to the lexical field of anger, anxiety and sadness over the sum of all words (in all tweets) posted by user i on day t . In other words, our indicator of mental health simply measures the frequency of all words that are matched to the lexical field of anger, anxiety and sadness in all tweets posted by each user daily. These daily indicators provide the measures of users’ psychological well-being exploited in our econometric analysis.

To provide readily available measures of psychological well-being, we derive aggregate indicators calculated as the daily average of individual indicators over all sampled users – $Anger_t$, $Anxiety_t$, $Sadness_t$, MH_t . We provide these measures to the public at the Twittersphere Observatory website.¹³

Lastly, we classify tweets into three categories: i) original tweets, ii) replies to tweets and iii) retweets. Original tweets include all original tweets posted by users. Replies to tweet capture all responses to another tweet, whereas a retweet is the re-posting of a tweet that can be your own or from someone else. In this paper, we only use original and replies to tweets to measure mental health, as their text is most likely authored by the account owner and therefore most closely reflects the user’s mental state. Garcia and Rimé (2019) adopted a similar approach to measure the emotional response of French Twitter users to the Paris terrorist attack of November

¹²In the formal analysis, we also remove stopwords. ‘je’, ‘ne’ and ‘suis’ are removed.

¹³<https://twittersphereobservatory.github.io/>. The measures are updated weekly with a one-week lag.

2015. Key descriptive statistics of our sample are reported in Table 1 and discussed in the next section.

4 Descriptive Statistics and Descriptive Analysis

4.1 Descriptive Statistics

Our refined data includes 10,438,153 daily observations from 52,885,834 tweets¹⁴ posted by 39,970 active Twitter accounts held by individuals located in France with information on predicted gender (female or male) and age group (<18, 18-28, 29-39, >39), regional location, and daily indicators of anger, anxiety, sadness and mental health covering the period between January 1st, 2019 and March 1st, 2021.¹⁵

It is well documented that Twitter users are not representative of the overall population. Mislove et al. 2011 show that Twitter users tend to be younger and more educated than the average American and more likely to live in a dense area. Sloan et al. 2015 report comparable results for the UK. Demographic variables of our sample reported in Table 1 largely corroborate these findings.

About 75% of users in our sample are less than 29 years old, which is about twice as large as in the general population in 2020. Adults over 39 years old represent just over 14% of our sample while accounting for about 53% of the total population of France in 2020 (Insee, 2021). At the same time, the share of users between 29 and 39 at 10.2% mirrors more closely the estimated 13.5% of the general population in 2020 (Insee, 2021).

The share of our sample located in “Île-de-France” (Paris Region) is about 14 percentage point higher than the general population, which corroborates findings for the US and the UK that Twitter users tend to be disproportionately concentrated in urban areas (Mislove et al., 2011; Sloan et al., 2015; Mellon and Prosser, 2017). Interestingly, the share of our sample residing in other regions mirrors more closely patterns in the general population, differing by between 0.5 and 3.5 percentage points. Note that the ordering of regions in our sample according to population size largely reflects the 2020 French census. Finally, under 40% of users in our Twitter sample are female, which is about 12 percentage points lower than in the general population (Insee, 2021).

We then use the list of keywords used in Balech et al. (2020) (which uses French data) and

¹⁴Remember that our mental health indicators include the words in all the tweets of user i on day t . We have 10,438,152 daily observation following this aggregation. In other words, one daily observation can include many tweets.

¹⁵The list of users with their corresponding predicted demographic information is available upon request.

[Banda et al. \(2021\)](#) (which is a broader dataset) and count the number of tweets in our dataset that mention those keywords. We find that the proportion of tweets in our dataset that include the keywords used in [Balech et al. \(2020\)](#) is 0.036%, and is 1.552% for the keywords used in [Banda et al. \(2021\)](#). Our dataset thus does not only include tweets that directly mention the pandemic, and so may capture a fuller representation of our users' mental health, as users may indicate their well-being in tweets unrelated to the pandemic.

Finally, we extract two subsamples from our dataset to investigate econometrically the effects of each lockdown on our users' mental health that include each lockdown period, several weeks leading up to the lockdown, and the same span of time one year prior (i.e., in 2019).

Sample 1 covers the period spanning from January 1st to June 2nd in both 2019 and 2020. June 2nd, 2020 marks the progressive reopening of restaurants, cinemas and in-person shopping. The matching period in 2019 data is used as the comparison period in the regression analysis. Sample 1 includes 3,525,131 daily observations capturing the change in emotions of 32,774 active unique users.

Sample 2 covers the period spanning from July 1st to December 15th in both 2019 and 2020. On July 1st, 2020, France reopened its borders with non-European Union countries, while the second lockdown officially ended on December 15th, 2020. Likewise, the matching period in 2019 is used as the comparison period in the regression analysis. Sample 2 includes 4,763,401 daily observations from 38,904 active unique users (see [Table 2](#) in the appendix for further details).

4.2 Descriptive Analysis

Our analysis focuses on the indicators capturing the emotions expressed in original tweets and replies to tweets.

[Figure 1](#) depicts the 7-day moving average of the total number of original tweets and replies to tweets posted between January 1st 2019 and December 31st 2020. It shows an upward trend in tweeting activity with two distinctly large temporary increase arising over the periods covered by each lockdown. These two surges reveal a clear behavioural response to lockdown orders.

[Figure 2](#) depicts the 7-day moving average of our aggregate mental health indicator in 2020, where the shaded areas cover the lockdown periods defined in [section 2](#). This figure allows convenient visual comparisons of the mental health response before and after the implementation of each stay-at-home order, as well as their relative magnitude.

A close inspection of the data reveals unremarkable differences in the levels of our indicator before and after the first stay-at-home order. In particular, our mental health indicator follows a V-shaped trajectory at the onset of the pandemic, exhibiting improving mental health at first,

evidenced by a decrease in our indicator, followed by a consistent deterioration during the second half of March, evidenced by a decrease in the indicator. The level of our indicator, however, reached a plateau in early April at levels not substantially higher than observed during the pre-lockdown period, and then rapidly fell back to its lowest level of the year following the announcement of the lifting of most restrictions on May 28th.¹⁶

In contrast to the first lockdown order, Figure 2 shows a significantly larger and more persistent deterioration of mental health, as we measure it, following the enforcement of second lockdown order. Our mental health indicator reached one of its lowest levels a few days before French President Macron announced the imposition of local curfews on October 14th. Interestingly, the progressive implementation of these curfews between October 17th and October 27th also coincides with a sharp increase in the mental health indicator that persisted over the entire confinement period. This observation may indicate that the second lockdown was more deleterious to mental health than the first one. Note that our indicator sharply decreased with the implementation of the progressive removal of restrictions on December 15th. It remained, however, at a higher level than before the first confinement, underscoring a possible lasting adverse impact on the average emotional well-being of the population.

Figure 3 displays the emotional responses measured by all our indicators over the periods covered by Samples 1 and 2. The top-left panels of Figure 3 confirm the absence of a significant mental health response overall during the first lockdown compared to the same period in 2019. The bottom three panels on the left depict the changes in each item entering our aggregate mental health indicator. They reveal an increase in the level of anxiety during the pandemic year compared to 2019, that gradually intensifies during the period of confinement. They also reveal that the V-shaped trajectory of our aggregate mental health indicator at the onset of the pandemic is largely driven by expressed anger and sadness. Taken together, Figure 3 shows little evidence supporting the idea that the first confinement in France lead to a major deterioration in mental health from our Twitter population, in comparison to the baseline period in 2019.

The insignificant emotional response captured by all our indicators during the first confinement is in stark contrast with the sharp deterioration in mental health following the announcement of the second. The bottom three panels of Figure 3 show that the mental health response to the second lockdown is driven by all emotions underlying the construction of our aggregate mental health indicator. Overall, these observations reinforce the preliminary conclusion that the second lockdown generated a significant deterioration in the population's mental health.

In the next section, we fully exploit the longitudinal dimension of our sample to further

¹⁶As discussed in section 2, the removal of restrictions includes the reopening of bars, restaurants, sport and cultural activities as well as the full removal of travel ban on June 2nd.

explore the impact of enforced lockdowns on mental health as well as the existence of heterogeneous responses across gender and age-groups.

5 Statistical Analysis

5.1 Empirical Strategy

5.1.1 Difference-in-Difference Estimator

To evaluate the impact of COVID-19 lockdowns on individual mental health, we adopt an identification strategy that explicitly exploits the unique longitudinal dimension of our sample to account for unobserved individual fixed effects. To do so, we compare within individual variations in sadness, anxiety and anger before and after each lockdown order with the corresponding period a year prior. This approach can be viewed as a difference-in-differences (DID) model in which the calendar year prior to the onset of the COVID-19 pandemic serves as a counterfactual outcome, and which allows us to also account for the confounding effect of a secular trend and seasonal patterns.

More specifically, to measure the impact of lockdown orders on mental health, we estimate separately for each lockdown the following equation model:

$$MH_{it} = \beta L_t + \delta (L_t \times Y_{20}) + \gamma Z_{t-1} + \alpha_i + \mu_t + \epsilon_{it} \quad (1)$$

whereby the dependent variable MH_{it} is the mental health indicator of user i on day t or one of the three emotions underlying our mental health indicator (anger, anxiety or anger). L_t is a binary variable taking the value 1 either on the days following the announcement of a lockdown order in 2020 or on the days matching the corresponding period a year prior (2019). That is, $L_t = 1$ either for the period spanning March 14th to June 2nd and zero otherwise (Sample 1: first lockdown) or for the period spanning October 14th to December 15th (Sample 2: second lockdown). Y_{20} is a binary variable taking the value one the year of the pandemic (2020) and zero the year prior (2019). Z_{t-1} controls for the lagged number of new daily deaths from COVID-19 per million. The parameter α_i absorbs all confounding unobserved individual fixed effects and μ_t absorbs time fixed effects such as secular trends and seasonal patterns through the inclusion of year, month and day of the week indicators. ϵ_{it} is an error term absorbing all other determinants of users' dimensions of mental health not captured by our model.

δ is our parameter of interest. It measures the average impact of the stay-at-home orders on the emotional state of users. The validity of our estimation results rests on the standard parallel

trend assumption that no significant shocks other than those related to the lockdown impacted individual users' average emotional state during the pandemic year and the control period a year prior.

5.1.2 Event-Study Estimator

We then further explore the adaptation and persistence of individual emotional states throughout each lockdown while still controlling for fixed differences across individual users and time by estimating the following event study model:

$$E_{iwt} = Y_{20} \sum_{\substack{\tau=-q \\ \tau \neq -1}}^m \delta_{\tau} \mathbb{1}(w - w^* = \tau) + \sum_{\substack{\tau=-q \\ \tau \neq -1}}^m \kappa_{\tau} \mathbb{1}(w - w^* = \tau) + \gamma Z_{t-1} + \alpha_i + \mu_t + \epsilon_{it} \quad (2)$$

whereby $\mathbb{1}(t - t^* = \tau)$ are indicator variables capturing the number of weeks τ relative to week w^* , which marks the onset of a lockdown order. The period covered by our data includes 10 (15) weeks leading up to the first (second) lockdown and 11 (9) weeks after. The week prior to the implementation of each lockdown order is the omitted category. As a result, each estimate of δ_{τ} measures the weekly change in mental health¹⁷ relative to the comparable week in 2019 as measured from the week prior to the implementation of each lockdown order.

We first estimate equations (1) on all users and then separately by gender and age groups to further explore the heterogeneous impact of lockdown orders. We then estimate equation (2) to shed further light on the strength and persistence of estimated emotional responses. All standard errors are heteroskedastic-robust and clustered at the individual level.

5.2 Main Results

5.2.1 Difference-in-Difference Estimation Results

Coefficient estimates measuring the effects between lockdown orders and emotional well-being are summarized graphically in Figure 4. The vertical dotted line of each plot is the line of null effect. Actual estimates and standard errors are fully reported in Table 4. These estimates reinforce previously discussed descriptive findings of diverging emotional responses between the first and the second lockdown.

¹⁷As captured by expressed anger, anxiety and sadness in tweets

We do not find statistically significant evidence that the first lockdown triggered worsening mental health of our Twitter population. While the announcement of the first lockdown order coincides with a sharp increase in tweet volume,¹⁸ which may indicate a stronger need to communicate and share emotions during lockdown episodes, we do not find statistically significant changes in anxiety and sadness compared to the baseline period in 2019. Interestingly, we find a statistically significant *decline* in anger of just under 4%.¹⁹ Combined with the moderate decline in anxiety reported by the point estimates, this drop in anger leads to a statistically significant decline in our aggregate indicator, suggesting slightly improved overall mental health conditions. As discussed below, while arguably surprising, this result is in line with survey data from different sources conducted in France over the same period (see *infra*).

In contrast, we find a strong deterioration in mental health during the second lockdown episode, evidenced by a statically significant increase of about 15% in all expressed emotions compared to the baseline period in 2019, and translating into a statistically significant increase of 15.4% of our aggregate indicator (see again Figure 4 and Table 4 for estimates and standard errors). Taken together, these findings provide compelling evidence that repeated containment policies are increasingly more harmful to individual well-being, a result that is of significant importance for the cost-benefit analysis of these policies on population health.

5.2.1.1 Impact of Lockdown by Age and Gender

Recent studies, mostly exploiting mental health measures derived from survey data, found unequal distribution of the mental health burden caused by the pandemic across various segments of the population. In particular, studies underscore that women and younger respondents generally experienced higher rates of mental distress than other population groups; see, e.g., [Adams-Prassl et al. \(2022\)](#) for the US, [Macalli et al. \(2021\)](#) for France, [Lucchini et al. \(2021\)](#) for Italy, [Banks and Xu \(2020\)](#); [Davillas and Jones \(2021\)](#); [Oreffice and Quintana-Domeque \(2021\)](#); [Pierce et al. \(2020\)](#) for the UK, [Pedersen et al. \(2022\)](#) in Denmark, among others. We further explore these heterogeneous mental health responses to enforced lockdowns across age groups and gender within our Twitter sample.

Age Groups

To examine the heterogeneous impact of lockdown orders on different age groups, we consider two broad age categories. The first group includes all users age 28 and under (younger

¹⁸As discussed earlier, we observe a sharp increase in tweet volume following the first lockdown order for about 2 months before quickly returning to its secular trend following the announcement of the progressive removal of restrictions on May 11th, 2020 (see Figure 1, section 4.2).

¹⁹Calculated as the ratio of the DID point estimate over the mean indicator of anger during the baseline lockdown period in 2019.

users) and the second group, all users older than 28 (older users). Coefficient estimates and associated 90% confidence intervals are displayed in Figure 5, and actual point estimates are reported in Table 4.

We find moderate evidence of a differentiated mental health response among age groups during the first lockdown unfavorable to younger users. Consistent with our results above, Figure 5 reveals that older users experienced a statistically significant improvement in overall mental health conditions compared to the year prior (associated with a decline in our mental health indicator), whereas no statistically significant change is found for younger users. In direct age group comparisons, Table 4 shows that this heterogeneous mental health response is mainly driven by a statistically significant different variation in expressed sadness for older users than for younger users.

In comparison, we find a similar but markedly more pronounced divergent mental health response across age groups during the second lockdown. Coefficient estimates show that both age groups experienced statistically significant increases in anger, anxiety and sadness compared to the year before, and also that these increases are larger for younger users than for older ones. Direct age group comparisons reported in Table 4 show that these differences translate into a larger statistically significant increase in the aggregate mental health indicator for younger users (reflecting aggravated mental health conditions) and that this difference is mostly driven by statistically larger increases in anxiety and sadness. Overall, as discussed below, our finding of a more severe mental health response for young people is largely consistent with survey-based results reported in several studies for various countries that use different methodologies to measure mental health and its variations over time.

Gender

Emotional responses across gender during lockdowns are more subtle and contrasted. Figure 6 reveals that, according to our indicator, the general mental health of women *improved* during the first lockdown compared to the year before, while no statistically significant change is observed for men. This improved psychological well-being for women is mostly driven by statistically significant decreases in expressed anxiety and anger. Coefficient estimates reported in Table 5 confirm that in direct gender comparisons, women experienced a statistically significant larger decline in expressed anxiety than men during the first lockdown, while the variations in expressed anger and sadness are not statistically significantly different across gender. Overall, our results point toward a moderately better emotional response for women than for men to the first lockdown episode in France. As discussed below, this finding contrasts with the results of most earlier studies conducted in other countries, which found a larger deterioration in mental

health for women than for men during the first months of the epidemic.

Our assessment of a gender difference in the emotional responses to the second lockdown are markedly different. Figure 6 reveals that mental health significantly deteriorated for both men and women during this episode in comparison to the year before, and that this deterioration was greater for women than for men. Point estimates show that women experienced a slightly larger increase in sadness and anxiety than men compared to the earlier period, while the increase in anger is roughly similar across genders. Results from direct gender comparisons reported in Table 5 show that the larger increase in expressed sadness for women is statistically significant, while the differences in anger and anxiety have the expected sign but are not statistically significant. Interestingly, from a dynamic perspective, our results imply that women experienced a stronger deterioration in all negative emotions underlying our mental health indicator between the first and the second lockdown, since a slightly favorable emotional response during the first lockdown for women translated into a strongly unfavorable response during the second one. Taken together, these findings do provide some support, albeit with some differences, to aforementioned studies conducted in other countries, which reported a more severe deterioration in mental health for women than for men during lockdown episodes. See the discussion below.

5.2.2 Event Study Results

Estimated coefficients and associated 95% confidence interval from estimating equation (2) are reported on event study plots in Figures 7, 8 and 9 for the pooled sample of all users across gender and age groups, respectively.

Unsurprisingly, in line with DID estimates, event study plots show no sizeable variation in mental health over the course of the first stay-at-home order except for the first week following the lockdown implementation. In this first week, a statistically significant decrease in the aggregate mental health indicator (associated with improved psychological well-being) is observed, mostly driven by a decrease in anger and sadness (see left panels of Figure 7). In addition, a peak in sadness is observed roughly four weeks after the beginning of the lockdown. All other weekly point estimates are statistically insignificant.

In sharp contrast, estimates for the second lockdown show a rapid and statistically significant rise in anger, anxiety and sadness on the days following the imposition of local curfews which persisted over the entire period covered by our sample. Interestingly, all measures gradually increase over the first three weeks succeeding “day zero”, during which the second lockdown was progressively implemented over the national territory,²⁰ and reach a plateau around the third

²⁰Starting with the regions with the highest incidence rates.

week, which includes the day at which the lockdown became nationwide (October 30th, 2020). These results are consistent with DID findings of a much larger mental health deterioration over the course of the second lockdown than the first one.

Event study plots across age groups also reinforce DID findings that mental health worsened much more acutely among younger users during the second lockdown (see Figure 8). We note a growing mental health gap initiated by disproportionately larger increases in anger and sadness for younger users over the first four weeks following the implementation of this second lockdown, which is then accompanied by a similarly disproportionately larger increase in anxiety over the last weeks of this lockdown.

Event study plots across gender show that women experienced, at each point in time, worse mental health conditions than men during the first months of the coronavirus pandemic (see Figure 9). However, this gender gap appears to slightly tighten following the first lockdown implementation, which is again consistent with our DID findings. In contrast, plots indicate that women experienced larger increases in all emotions during the second lockdown. Overall, results from the event study largely corroborate those obtained from the DID model, showing that the additional mental health burden experienced by women was more acute during the second lockdown than during the first one.

6 Discussion

It is useful at this stage to summarize the main results obtained from our Twitter-based mental health indicators applied to France and compare them with findings obtained with other methodologies and/or for other countries. First, we find significant differences in the mental health response to the first and second lockdowns. While we find no significant mental health deterioration during the first lockdown (with even a slight improvement for female and for older users), we find a strong and statistically significant decline in psychological well-being for all users over the course of the second lockdown. Second, we find that the deterioration of mental health is of significantly greater magnitude for young users and for women. Our results thus suggest that the increasing mental health costs of successive waves of stay-at-home orders disproportionately affect different segments of the population.

Our results regarding the first lockdown differ somewhat from earlier studies on different countries using various methodologies. Many studies have documented worsening mental health at the onset of the pandemic in the US and in several European countries using validated instruments to detect psychological distress, such as the 12-Item General Health Questionnaire (GHQ-12) from survey data ([Armbruster and Klotzbücher, 2020](#); [Arendt et al., 2020](#); [Banks and](#)

Xu, 2020; Lucchini et al., 2021). Studies exploring the mental health response to the nationwide lockdowns enforced in France are scarce, but it is interesting to observe that our results are largely in line with those reported by the two main large-scale studies based on survey data conducted in France.

For example, Santé Publique France, the French public health agency, documents an increase in overall life satisfaction and a decrease in anxiety levels (starting from an initial high) over the course of the first lockdown based on of a representative sample of 2000 households aged 18+ living in France.²¹ Likewise, the CAMME survey conducted by INSEE/CEPREMAP reports an increasing trend in self-declared life satisfaction between March and June 2020, with a peak reached in the latter date when the removal of restrictions was implemented (see e.g. Perrona and Senik, 2021). In contrast, using a longitudinal cohort study initially set up to study home, school and leisure injuries, Ramiz et al. (2021) reports lower self-rated mental health during the first months of the pandemic, with increased risk among women, young and elderly respondents. Yet the comparison is undertaken with responses obtained on average 4.8 years earlier, which makes it difficult to infer at precisely which date the deterioration in mental health occurred. Conversely, in another longitudinal study conducted in Germany in which mental health status was recorded much more regularly (every three months before the pandemic), increasing mental health scores and decreasing numbers of daily hassles are reported for the vast majority of respondents over the 8 weeks covering the first German lockdown (Ahrens et al., 2021). This results is largely consistent with the two aforementioned large-scale studies conducted in France.

Our results for the second lockdown are largely in line with international evidence showing that coronavirus lockdowns deteriorated mental health. Interestingly, negative mental health effects of the second stay-at-home order in France are also consistent with findings from the CAMME survey documenting a strong decrease in the life satisfaction index over the course of this lockdown (see again Perrona and Senik, 2021). Yet these results differ somewhat from those of the CoviPrev survey reporting a roughly stable anxiety index between the first and the second lockdown.

Two points are worth mentioning here. First, our results indicate that among the three indices underlying our aggregate mental health indicator (anger, anxiety and sadness), anxiety is by far the variable that increased *the least* during the second lockdown (see Figure 7). Second, the anxiety index of the CoviPrev survey is based on the fraction of respondents presenting an anxiety score on the HAD scale greater than 10, i.e. respondents with mild to severe anxiety and depression levels, while our anxiety indicator is based on the fraction of words associated with

²¹See CoviPrev survey on <https://www.santepubliquefrance.fr>

anxiety in the tweets posted by our entire population of Twitter users. It is therefore possible that the anxiety level did increase in the general population during the second lockdown, but not to the point that many respondents were characterized as presenting severe anxiety and depression symptoms according to standard clinical measures.

The differences in the emotional responses to the two lockdowns for different age groups are also worth discussing. Several studies found that mental health deteriorated more for younger people and/or for students compared to the rest of the population during lockdowns; see [Husky et al. \(2020\)](#) for France, [Banks and Xu \(2020\)](#); [Davillas and Jones \(2021\)](#); [Pierce et al. \(2020\)](#) for the UK, [Lucchini et al. \(2021\)](#) for Italy, for example. Our results for France are largely in line with these results, even though we underscore differences between the first and the second lockdown.

During the first lockdown, the response differences between age groups are explained by a mental health *improvement* for older users, while the mental health of younger users did not significantly change compared to the year before. By contrast, during the second lockdown, the age gap is explained by larger deteriorations in mental health outcomes for younger users. It also merits mentioning that our sample of Twitter users does not allow us to distinguish between young people who are students and those who are not. It would be relevant to make this distinction, since recent evidence suggest that students were significantly disproportionately more affected by lockdown orders than non-students see (see [Arsandaux et al., 2021](#); [Husky et al., 2020](#); [Macalli et al., 2021](#)). A finer analysis of our Twitter population that distinguished between student and non-student young users would enable us to shed some light on this recent piece of evidence.

Finally, our results that show differences across gender over the course of the first and the second lockdown are also in line with the bulk of international evidence; see, e.g., [Adams-Prassl et al. \(2022\)](#) for the US, [Banks and Xu \(2020\)](#); [Davillas and Jones \(2021\)](#); [Oreffice and Quintana-Domeque \(2021\)](#); [Pierce et al. \(2020\)](#) for the UK,²² [Vloo et al. \(2021\)](#) for the Netherlands and the references therein. But the differences we observe between the first and the second lockdown are worth stressing. Our indicator suggests that the mental health of women slightly improved during the first lockdown in France, while it strongly deteriorated (and deteriorated to a larger extent than for men) during the second one. Again, while perhaps surprising, our findings regarding the first lockdown in France are consistent with those reported in the CoviPre survey. For example, according to their aforementioned anxiety index, the fraction of women presenting an anxiety score greater than 10 on the HAD scale *decreased* by 12.3 percentage points (from 31.6% to 19.3%) between the first wave of the survey on March 23-25, 2020 and

²²Conflicting evidence also exists for the UK ([Serrano-Alarcón et al., 2021](#)).

the ninth wave of the survey on May 27-29, 2020, close to the end of the lockdown. It also decreased for men over the same period, but only by 7.9 percentage points (from 21.3% to 13.4%). A similar larger decrease is observed for women regarding the CoviPrev depression index over this period, albeit with a milder amplitude. Nonetheless, the following months showed that this favorable mental health evolution for women was transitory and specific to the first lockdown in France, since the mental health of women significantly deteriorated during the second lockdown to a larger extent than it did for men.

7 Concluding remarks

Our results suggest that a Twitter-based indicator can be a useful tool for monitoring the mental health of the population, its evolution over time and its variations in response to large scale events. Obviously, a mental health indicator such as ours also has certain obvious limitations worthy of discussion. In particular, a Twitter-based sample is known to be unrepresentative of the general population. This lack of representativeness may in turn yield biased results. In our view, this issue does not compromise our overall findings. First, even though men are slightly over-represented compared to women in the sample and young people are over-represented compared to older people, we exploit the sheer size of our sample to explore the association of lockdown orders with the mental health of each group separately. Second, even if our sample is biased with respect to other unobserved characteristics such as the average income of users and their education level, the selection bias is important only to the extent that it significantly biases the conclusions drawn from the analysis of the obtained aggregate and disaggregated indices. The answer to this question is likely to be very context dependent. In our particular context of analyzing the mental health effects of lockdowns, it is fair to acknowledge that wealthier and more educated people may have been less exposed to mental stress than poorer and/or less educated ones. If this is the case, the consequence is that our aggregate indicator likely *underestimates* the average mental health deterioration of the population. Ultimately, our aim is not to claim that our indicator is by any means an unbiased estimate of the mental state of the population. We do claim, however, that because our indicator can be constructed and monitored in real time, it provides a versatile tool to rapidly detect the a population's emotional response to the kind of large scale major events a society can experience.

References

- Adams-Prassl, A., T. Boneva, M. Golin, and C. Rauh (2022, January). The impact of the coronavirus lockdown on mental health: evidence from the united states. *Economic Policy*. 18, 23
- Ahrens, K. F., R. J. Neumann, B. Kollmann, M. M. Plichta, K. Lieb, O. Tüscher, and A. Reif (2021). Differential impact of covid-related lockdown on mental health in germany. *World Psychiatry* 20(1), 140. 22
- Aiello, L. M., D. Quercia, K. Zhou, M. Constantinides, S. Šćepanović, and S. Joglekar (2021). How epidemic psychology works on twitter: evolution of responses to the covid-19 pandemic in the us. *Humanities and Social Sciences Communications* 8(1), 1–15. 11
- Alfano, V. and S. Ercolano (2020). The efficacy of lockdown against covid-19: a cross-country panel analysis. *Applied health economics and health policy* 18, 509–517. 2
- Arendt, F., A. Markiewitz, M. Mestas, and S. Scherr (2020, November). COVID-19 pandemic, government responses, and public mental health: Investigating consequences through crisis hotline calls in two countries. *Social Science & Medicine* 265, 113532. 2, 21
- Armbruster, S. and V. Klotzbücher (2020). Lost in lockdown? COVID-19, social distancing, and mental health in Germany. Discussion Paper Series 2020-04, University of Freiburg, Wilfried Guth Endowed Chair for Constitutional Political Economy and Competition Policy. 2, 21
- Arsandaux, J., I. Montagni, M. Macalli, N. Texier, M. Pouriel, R. Germain, A. Mebarki, S. Kinouani, M. Tournier, S. Schuck, and C. Tzourio (2021, August). Mental health condition of college students compared to non-students during COVID-19 lockdown: the CONFINS study. *BMJ Open* 11(8), e053231. 23
- Balech, S., C. Benavent, and M. Calciu (2020). The first french covid19 lockdown twitter dataset. *arXiv preprint arXiv:2005.05075*. 3, 7, 8, 13, 14
- Banda, J. M., R. Tekumalla, G. Wang, J. Yu, T. Liu, Y. Ding, E. Artemova, E. Tutubalina, and G. Chowell (2021). A large-scale covid-19 twitter chatter dataset for open scientific research—an international collaboration. *Epidemiologia* 2(3), 315–324. 7, 8, 14

- Banks, J. and X. Xu (2020, September). The mental health effects of the first two months of lockdown during the COVID-19 pandemic in the UK. *Fiscal Studies* 41(3), 685–708. 2, 18, 21, 23
- Blanchard, O., T. Philippon, and J. Pisani-Ferry (2020). A new policy toolkit is needed as countries exit covid-19 lockdowns. *Policy Contribution* 12. 7
- Boyd, R. L. and H. A. Schwartz (2020, October). Natural language analysis and the psychology of verbal behavior: The past, present, and future states of the field. *Journal of Language and Social Psychology* 40(1), 21–41. 11
- Brodeur, A., A. E. Clark, S. Fleche, and N. Powdthavee (2021, January). COVID-19, lockdowns and well-being: Evidence from google trends. *Journal of Public Economics* 193, 104346. 2
- Bruns, A., B. Moon, F. Münch, and T. Sadkowsky (2017). The australian twittersphere in 2016: Mapping the follower/followee network. *Social Media+ Society* 3(4), 2056305117748162. 9
- Brühlhart, M., V. Klotzbücher, R. Lalive, and S. K. Reich (2021, January). Mental health concerns during the covid-19 pandemic as revealed by helpline calls. *Nature* 600, 121–126. 2
- Cahuc, P. (2022). Iza covid-19 crisis response monitoring: France. 7
- Chen, E., K. Lerman, E. Ferrara, et al. (2020). Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR Public Health and Surveillance* 6(2), e19273. 3, 7, 8
- Coppersmith, G., M. Dredze, and C. Harman (2014). Quantifying mental health signals in twitter. In *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*, pp. 51–60. 11
- Davillas, A. and A. M. Jones (2021, April). The first wave of the COVID-19 pandemic and its impact on socioeconomic inequality in psychological distress in the UK. *Health Economics* 30(7), 1668–1683. 2, 18, 23
- Durazzi, F., M. Müller, M. Salathé, and D. Remondini (2021). Clusters of science and health related twitter users become more isolated during the covid-19 pandemic. *Scientific reports* 11(1), 1–11. 9, 10
- Dyer, J. and B. Kolic (2020). Public risk perception and emotion on twitter during the covid-19 pandemic. *Applied Network Science* 5(1), 1–32. 11

- Foki, G. (2021, October). Covid-19 : les différents systèmes d'indemnisation du chômage partiel en europe. Technical report, Unédic. [7](#)
- Garcia, D. and B. Rimé (2019, March). Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological Science* 30(4), 617–628. [11](#), [12](#)
- Gruzd, A. and P. Mai (2020, July). Going viral: How a single tweet spawned a COVID-19 conspiracy theory on twitter. *Big Data & S10.1177/0261927x20967028ociety* 7(2), 205395172093840. [7](#)
- Gupta, R., A. Vishwanath, and Y. Yang (2021). Covid-19 twitter dataset with latent topics, sentiments and emotions attributes. [7](#)
- Hubbard, G. and M. R. Strain (2020, October). Has the paycheck protection program succeeded? resreport 13808, IZA. [7](#)
- Husky, M. M., V. Kovess-Masfety, and J. D. Swendsen (2020, October). Stress and anxiety among university students in france during covid-19 mandatory confinement. *Comprehensive Psychiatry* 102, 152191. [23](#)
- Insee (2021). Estimation de population, (données provisoires et révisées arrêtées à fin février 2021). [13](#)
- King, G., P. Lam, and M. E. Roberts (2017, April). Computer-assisted keyword and document set discovery from unstructured text. *American Journal of Political Science* 61(4), 971–988. [9](#)
- Lucchini, M., T. Gerosa, M. Pancheva, M. Pisati, C. Respi, and E. Riva (2021, November). Differential effects of COVID-19 and containment measures on mental health: Evidence from ITA.LI—italian lives, the italian household panel. *PLOS ONE* 16(11), e0259989. [18](#), [22](#), [23](#)
- Macalli, M., N. Texier, S. Schück, S. M. Côté, and C. Tzourio (2021). A repeated cross-sectional analysis assessing mental health conditions of adults as per student status during key periods of the covid-19 epidemic in france. *Scientific reports* 11(1), 1–6. [18](#), [23](#)
- Mellon, J. and C. Prosser (2017, July). Twitter and facebook are not representative of the general population: Political attitudes and demographics of british social media users. *Research &10.1016/j.socscimed.2021.114040 Politics* 4(3), 205316801772000. [13](#)

- Mislove, A., S. Lehmann, Y.-Y. Ahn, J.-P. Onnela, and J. Rosenquist (2011). Understanding the demographics of twitter users. In *Proceedings of the International AAAI Conference on Web and Social Media*, Volume 5. 9, 13
- Oreffice, S. and C. Quintana-Domeque (2021, March). Gender inequality in COVID-19 times: evidence from UK prolific participants. *Journal of Demographic Economics*, 1–27. 2, 18, 23
- Park, M., C. Cha, and M. Cha (2012). Depressive moods of users portrayed in twitter. 11
- Pedersen, M. T., T. O. Andersen, A. Clotworthy, A. K. Jensen, K. Strandberg-Larsen, N. H. Rod, and T. V. Varga (2022, January). Time trends in mental health indicators during the initial 16 months of the COVID-19 pandemic in denmark. *BMC Psychiatry* 22(1). 18
- Perrona, M. and C. Senik (2021, Avril). Le bien-être des français - mars 2021 un an après : l’usure. techreport 2021-02, Observatoire du Bien-être du CEPREMAP. 22
- Pierce, M., H. Hope, T. Ford, S. Hatch, M. Hotopf, A. John, E. Kontopantelis, R. Webb, S. Wessely, S. McManus, and K. M. Abel (2020, October). Mental health before and during the COVID-19 pandemic: a longitudinal probability sample survey of the UK population. *The Lancet Psychiatry* 7(10), 883–892. 2, 18, 23
- Piolat, A., R. J. Booth, C. K. Chung, M. Davids, and J. W. Pennebaker (2011). La version française du dictionnaire pour le liwc: modalités de construction et exemples d’utilisation. *Psychologie française* 56(3), 145–159. 11
- Ramiz, L., B. Contrand, M. Y. R. Castro, M. Dupuy, L. Lu, C. Sztal-Kutas, and E. Lagarde (2021, March). A longitudinal study of mental health before and during COVID-19 lockdown in the french population. *Globalization and Health* 17(1). 22
- Serrano-Alarcón, M., A. Kentikelenis, M. Mckee, and D. Stuckler (2021, November). Impact of COVID-19 lockdowns on mental health: Evidence from a quasi-natural experiment in england and scotland. *Health Economics* 31(2), 284–296. 4, 23
- Silverio-Murillo, A., L. Hoehn-Velasco, A. R. Tirado, and J. R. B. de la Miyar (2021, July). COVID-19 blues: Lockdowns and mental health-related google searches in latin america. *Social Science & Medicine* 281, 114040. 2
- Sloan, L., J. Morgan, P. Burnap, and M. Williams (2015). Who tweets? deriving the demographic characteristics of age, occupation and social class from twitter user meta-data. *PloS one* 10(3), e0115545. 13

- Su, Y., J. Xue, X. Liu, P. Wu, J. Chen, C. Chen, T. Liu, W. Gong, and T. Zhu (2020). Examining the impact of covid-19 lockdown in wuhan and lombardy: a psycholinguistic analysis on weibo and twitter. *International journal of environmental research and public health* 17(12), 4552. 9
- Tausczik, Y. R. and J. W. Pennebaker (2010). The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology* 29(1), 24–54. 11
- Vloo, A., R. Alessie, J. Mierau, M. H. Boezen, J. O. Mierau, L. Franke, J. Dekens, P. Deelen, P. Lanting, J. M. Vonk, et al. (2021). Gender differences in the mental health impact of the covid-19 lockdown: longitudinal evidence from the netherlands. *SSM-population health* 15, 100878. 23
- Wang, Z., S. Hale, D. I. Adelani, P. Grabowicz, T. Hartman, F. Flöck, and D. Jurgens (2019). Demographic inference and representative population estimates from multilingual social media data. In *The world wide web conference*, pp. 2056–2067. 10
- Zhang, Y., H. Lyu, Y. Liu, X. Zhang, Y. Wang, and J. Luo (2020). Monitoring depression trend on twitter during the covid-19 pandemic. *arXiv preprint arXiv:2007.00228*. 11

8 Tables & Figures

Table 1: Descriptive Statistics

Gender (%)	
Female	39.75
Male	60.25
Age Group (%)	
<18	40.84
18-28	34.81
29-39	10.2
>39	14.16
Region (%)	
Auvergne-Rhone-Alpes	9.46
Bourgogne-Franche-Comté	3.09
Bretagne	3.59
Centre-Val-de-Loire	3.27
Corse	.57
Grand Est	9.47
Haut-de-France	11.8
Île-de-France	33.81
Normandie	3.3
Nouvelle-Aquitaine	5.84
Occitanie	6.66
Pays de La Loire	2.99
Provence-Alpes-Côte d'Azur	6.14
Total number of	
Tweets	52, 885, 834
Original Tweets	10, 472, 645
Replies to Tweets	14, 031, 947
Retweets	28, 381, 242
Unique Twitter users	39, 970
Daily Observations	10, 438, 153

Table 2: Descriptive Statistics by Samples

	Sample 1: Jan. 1 to Jun. 2		Sample 2: Jul. 1 to Dec. 15	
	2019	2020	2019	2020
Gender (%)				
Female	37.94	39.23	38.78	39.69
Male	62.06	60.77	61.22	60.31
Age Groups (%)				
<18	38.33	40.61	40.09	41.08
18-28	34.65	34.72	34.62	34.81
29-39	11.32	10.25	10.61	10.11
>39	15.7	14.42	14.67	14
Region (%)				
Auvergne-Rhone-Alpes	9.23	9.24	9.27	9.39
Bourgogne-Franche-Comté	3.15	3.05	3.09	3.01
Bretagne	4.38	3.85	4.04	3.73
Centre-Val-de-Loire	3.53	3.43	3.47	3.37
Corse	.68	.62	.61	.58
Grand Est	9.08	9.12	9.06	9.18
Haut-de-France	12.58	12.49	12.72	12.24
Île-de-France	31.04	32.63	31.75	33.58
Normandie	3.86	3.57	3.72	3.42
Nouvelle-Aquitaine	5.95	5.8	5.87	5.66
Occitanie	7.23	6.87	7.04	6.77
Pays de La Loire	3.37	3.19	3.26	3.01
Provence-Alpes-Côte d'Azur	5.93	6.15	6.12	6.07
Nb. of Unique Users				
	21, 316	30, 837	28, 281	36, 466
		32, 774		38, 904
Obs	1, 263, 323	2, 261, 808	2, 013, 413	2, 749, 988

Table 3: Examples of Tweets in Original Language

Emotion	Tweet	LIWC Dic.
Anger	<p>- "@— @— 0 patience, le moindre truc me tape sur les nerfs, purée quel enfer"</p> <p>- "Je suis à bout de nerfs, confinement de merde, maladie de merde, gens qui respectent pas les gestes barrières et le port du masque de merde"</p> <p>- "Test #Covid19 Révélations sur un nouveau scandale sanitaire https://t.co/HPuamXSHyT https://t.co/9a1zq1o8eG"</p> <p>- "@— Je ne suis pas spécialiste. Juste citoyen, en colère."</p> <p>- "@— C'est complètement dingue, un asile de fous ce pays."</p> <p>- "Ces 2 ados qui frappent sévèrement une infirmière dans un bus parce qu'elle leur demande de porter un masque me fout en colère. Cette bêtise brutale ! Mais que les dizaines de passagers adultes présents ne tentent absolument rien pour stopper ces 2 crétins me déprime beaucoup."</p> <p>- "Au bout d'un an à nous faire chier avec les règles de distanciation sociale le désastre psycho provoquer maintenant c'est masque FFP2 et 2 mètre de distance vous savez quoi aller vous faire foutre y en a marre on se restreint comme des cons et on peut crever en aller bosser https://t.co/xtEHTdhOOO"</p> <p>- "Ces politiciens qui sont ds leur voiture avec chauffeur, se baladent ds Paris à 22h et scandale, ya 2 autres pelés ds la rue; couvre feu ! Qu'ils sont cons ! https://t.co/DG51ToNI4J"</p> <p>- "Encore un #1er ministre qui n'a pas de couille. Qu'il le prenne leur foutue arrêter de #MasqueObligatoire dans l'espace public partout comme ça il arrête de nous faire chier avec leur #Covid_19 de merde"</p>	<p>tape, enfer merde, merde merde scandale</p> <p>colère fous</p> <p>fout, colère, brutale</p> <p>crétins chier</p> <p>foutre cons</p> <p>scandale, cons</p> <p>foutue</p> <p>chier, merde</p>
Anxiety	<p>- "@— Effectivement inquiétant , ce manque de maîtrise explique la situation de désordre dans lequel nous sommes ! n'était il pas auprès de Marisol Touraine ministre de la santé de Francois Hollande alors qu'il feint de découvrir les problèmes de l'hôpital soyons sérieux"</p> <p>- "\"J'ai bien peur que ce soit notre dernier marché de la saison. Je crains vraiment un nouveau confinement\" #Menerbes #Vaucluse https://t.co/YZG11kgSJK"</p> <p>- "Comment angoisser la planète? Écouter l' @OMS..."</p> <p>- "Je pense qu'il y a que moi qui suis inquiet de sortir à nouveau https://t.co/jljg5WkJT9"</p> <p>- "plus tard, Pujadas recevait la professeur épidémiologiste infectiologue à qui ils ont passé une vraie engueulade l'accusant de donner des chiffres inquiétants alors que eux avaient trouvé que le COVID était en baisse, c'était lunaire... Ils s'y sont mis à 3 comme des chiens.. https://t.co/m8hVSEDB2z"</p> <p>- "Crise d'angoisse bonjouuuuur"</p> <p>- "Je ne sais pas vraiment si c'est de l'anxiété que j'ai mais c'est horrible."</p> <p>- "C'était de la folie ce lundi midi devant le supermarché Leclerc de Guingamp, les gens se pressent de faire des provisions dans la crainte d'un possible #confinementtotal #CORONAVIRUSENFRANCE #COVID19france"</p> <p>- "@— @— La peur tue plus que le COVID ! La peur c'est la mort ! Il reste la litanie contre la peur des Bene Gesserit (Dune F. Herbert)"</p> <p>- "J'ai une connaissance qui a été deux fois cas contact ce mois-ci. Une fois à cause de sa mère "sceptique", une deuxième fois au boulot. Et sa femme est prof. De quoi se faire un bon ulcère."</p>	<p>inquiétant</p> <p>peur, crains</p> <p>angoisser inquiet</p> <p>inquiétants</p> <p>angoisse anxiété, horrible folie pressent crainte peur, peur peur</p> <p>sceptique</p>
Sadness	<p>- "@— Pathétique"</p> <p>- "@— Non mais oui abusé j'ai de ces moments de solitude parfois c'est violent"</p> <p>- "Et vous, comment ça va mal ?"</p> <p>- "Les pleurs des personnes âgées en EHPAD, enfermées et isolées jusqu'à l'absurde https://t.co/RdjEXvqGXn via @—"</p> <p>- "Enfermée dans ma chambre, 4 heures du mat jsuis perdue"</p> <p>- "J'espère que je pourrais aussi être visité comme pour les EHPAD parce que je suis un vieux monsieur seul et isolé"</p> <p>- "Ce n'est pas la peine de me faire sentir pitoyable parce que j'ai osé sortir alors que j'étais fatiguée. J'avais envie de les insulter. Trop fatiguée pour le faire ouvertement mais assez pour le faire intérieurement. Mon esprit est enfermé dans une boîte mal foutue."</p> <p>- "@— Je suis perdu"</p> <p>- "Vivement que je retrouve ma mère et mon frère parce que la solitude la c'est pesant"</p>	<p>pathétique violent mal pleurs, isolées</p> <p>perdue</p> <p>seul, isolé pitoyable fatiguée, fatiguée mal perdu pesant</p>

Table 4: DID Estimates for All Users and by Age Groups

Panel A: First Lockdown					
	Mental Health	Anger	Anxiety	Sadness	<i>N</i>
All Users	−0.357** (0.148)	−0.280*** (0.090)	−0.073 (0.062)	−0.003 (0.074)	3, 525, 131
Users <29	−0.265 (0.178)	−0.281*** (0.108)	−0.051 (0.074)	0.066 (0.088)	2, 691, 474
Users ≥ 29	−0.670** (0.261)	−0.288* (0.161)	−0.143 (0.113)	−0.238* (0.135)	833, 657
Δ	0.493 (0.305)	−0.001 (0.186)	0.088 (0.132)	0.406*** (0.156)	
Panel B: Second Lockdown					
All Users	2.377*** (0.132)	1.077*** (0.079)	0.496*** (0.057)	0.805*** (0.064)	4, 763, 401
Users <29	2.539*** (0.154)	1.124*** (0.091)	0.545*** (0.067)	0.869*** (0.073)	3, 674, 023
Users ≥ 29	1.841*** (0.248)	0.920*** (0.153)	0.330*** (0.105)	0.591*** (0.127)	1, 089, 378
Δ	0.697** (0.292)	0.204 (0.179)	0.215* (0.125)	0.278* (0.147)	

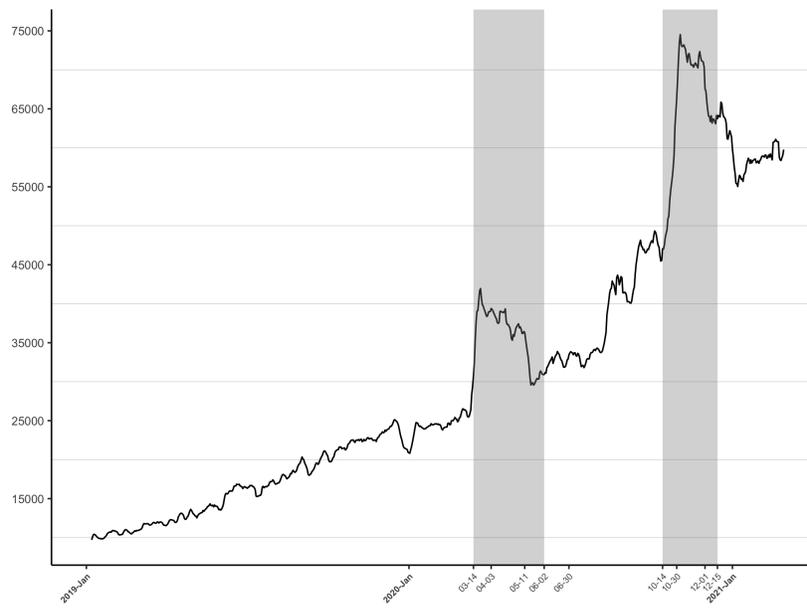
Note: This table reports conventional differences-in-differences estimates (DID). DID estimates measure the average impact of stay-at-home orders on anger, anxiety and sadness. *All Emotions* is a pooled indicator of these three emotions. All models control for individual and time fixed effects and the one-day lagged number of reported new deaths due to COVID-19. Robust standard errors are clustered at the individual level. *** $p < 0.01$, ** $p < 0.05$ and * $p < 0.1$. Standard errors in parentheses.

Table 5: DID Estimates across Gender and Age Groups (Within Group)

	Mental Health						Dependent variable:					
	Female		Male		Δ		Female		Male		Δ	
Panel A: First Lockdown												
All Users	-0.643*** (0.225)	-0.159 (0.197)	-0.485 (0.299)	-0.197 (0.124)	-0.207 (0.179)	-0.349*** (0.127)	-0.279*** (0.099)	0.070 (0.080)	0.039 (0.118)	-0.032 (0.095)	0.071 (0.151)	
Users <29	-0.540** (0.247)	-0.025 (0.252)	-0.516 (0.353)	-0.186 (0.159)	-0.206 (0.213)	-0.399*** (0.148)	-0.263** (0.107)	0.135 (0.102)	0.114 (0.129)	0.026 (0.120)	0.088 (0.176)	
Users \geq 29	-1.344** (0.524)	-0.479 (0.300)	-0.866 (0.604)	-0.226 (0.189)	-0.281 (0.351)	-0.292 (0.284)	-0.371 (0.255)	-0.080 (0.126)	-0.466* (0.276)	-0.173 (0.154)	-0.293 (0.316)	
Panel B: Second Lockdown												
All Users	2.582*** (0.203)	2.243*** (0.173)	0.338 (0.266)	1.085*** (0.107)	-0.018 (0.156)	0.139 (0.116)	0.579*** (0.090)	0.440*** (0.074)	0.936*** (0.103)	0.719*** (0.081)	0.217* (0.131)	
Users <29	2.689*** (0.221)	2.416*** (0.214)	0.273 (0.308)	1.141*** (0.132)	-0.034 (0.181)	0.091 (0.135)	0.595*** (0.098)	0.504*** (0.093)	0.988*** (0.112)	0.772*** (0.097)	0.216 (0.148)	
Users \geq 29	1.816*** (0.469)	1.848*** (0.289)	-0.032 (0.551)	0.955*** (0.180)	-0.173 (0.329)	0.166 (0.241)	0.460** (0.209)	0.294** (0.121)	0.574** (0.251)	0.598*** (0.146)	-0.025 (0.290)	

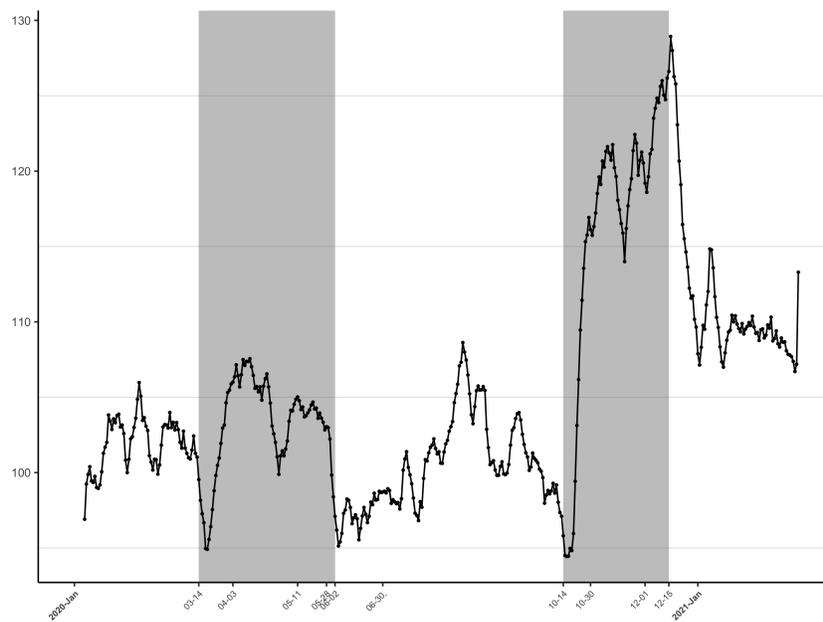
Note: This table reports conventional differences-in-differences estimates (DID) across gender and broad age groups. DID estimates measures the average impact of stay-at-home orders on anger, anxiety and sadness. *Mental Health* is a pooled indicator of these three emotions. Each regression controls for region fixed effects and time fixed effects (year, week and day of the week) and the one-day lagged number of reported new deaths due to COVID-19. Robust standard errors are clustered at the individual. *** $p < 0.01$, ** $p < 0.05$ and * $p < 0.1$. Standard errors in parentheses.

Fig. 1. Number of Daily Original Tweets and Replies to Tweets



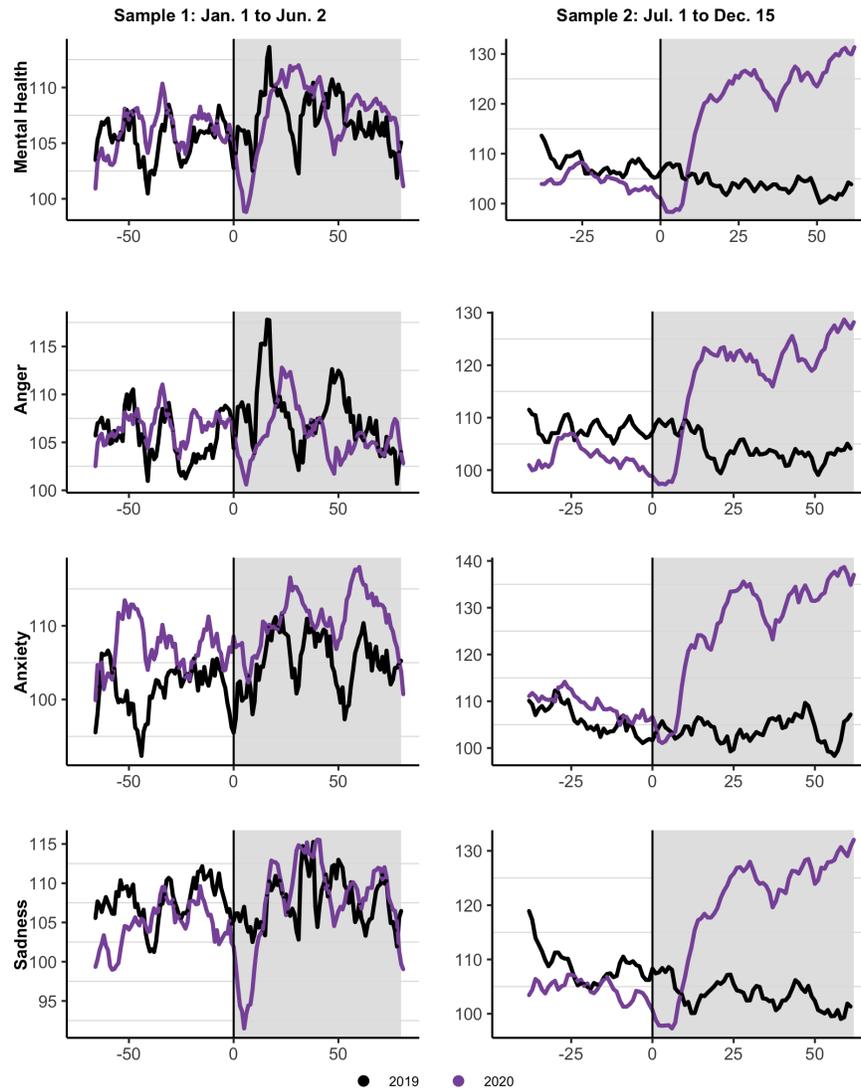
Note: Figure plots the 7-day moving average of the number of original tweets and replies to tweets from January 1st 2019 to December 31th 2020. The shaded areas indicate the first and second lockdown period, March 14th to June 2nd 2020 and October 14th to December 15th 2020, respectively.

Fig. 2. Daily Change in Mental Health



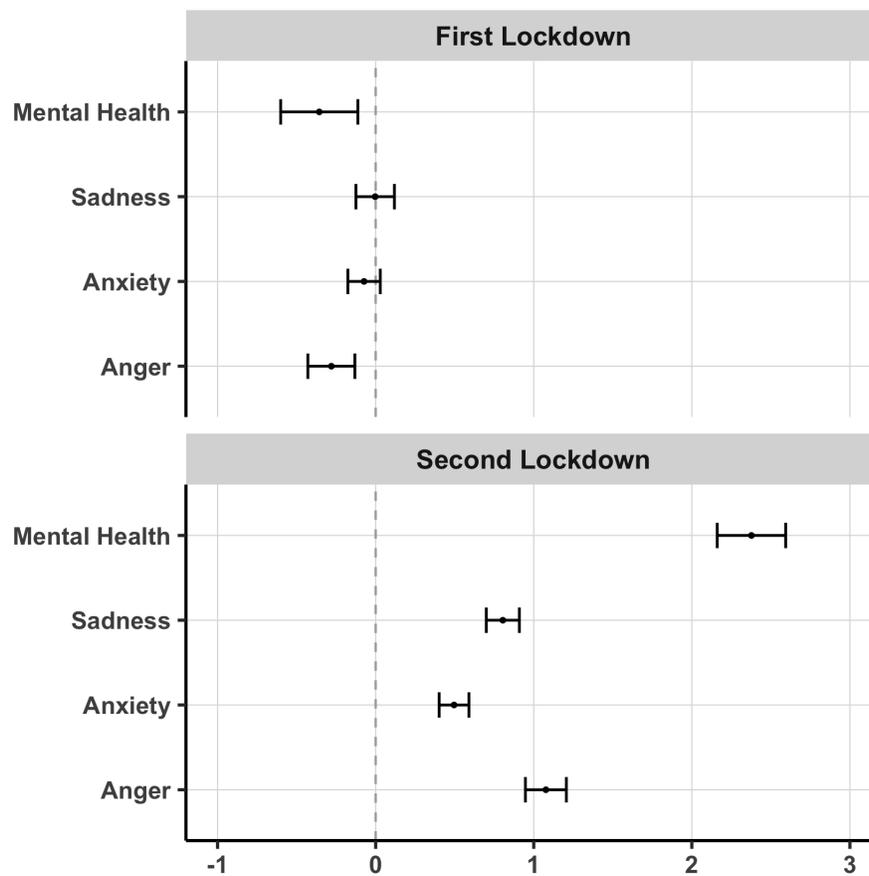
Note: Figure plots the 7-day moving average of the mental health indicator in 2020. Shaded areas cover the periods of the first and second lockdown order, March 14th to June 2nd 2020 and October 14th to December 15th 2020, respectively

Fig. 3. Daily Change in Mental Health



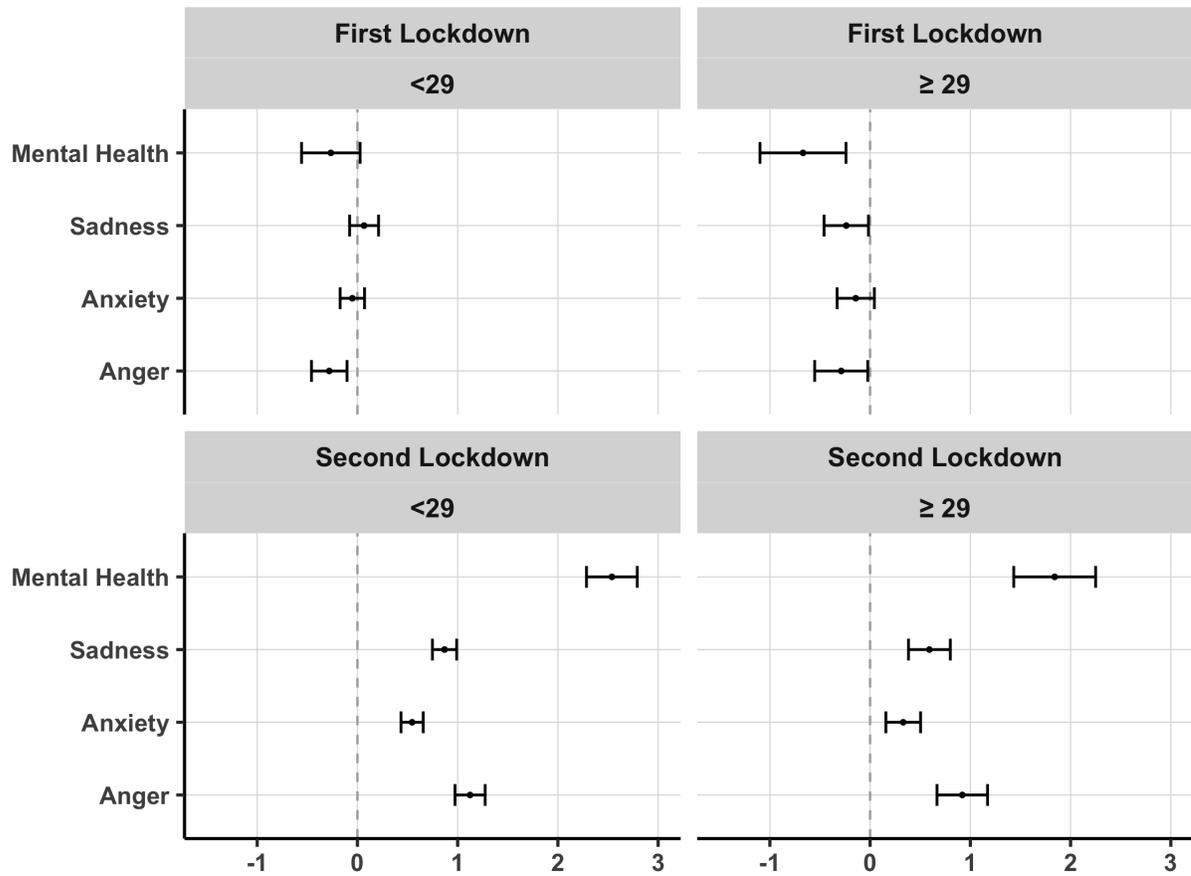
Note: Figure depicts the 7-day moving average of each indicator of emotions for sample 1 (left) and sample 2 (right). For Sample 1, the vertical line passing through the x-axis at “day-zero” of each plot corresponds to March 14th. For Sample 2, “day-zero” corresponds to October 14th, the day President Macron announced the progressive implementation of nighttime curfews. The shaded area captures the period covering the confinement period.

Fig. 4. Estimated Impact of Lockdowns: All users



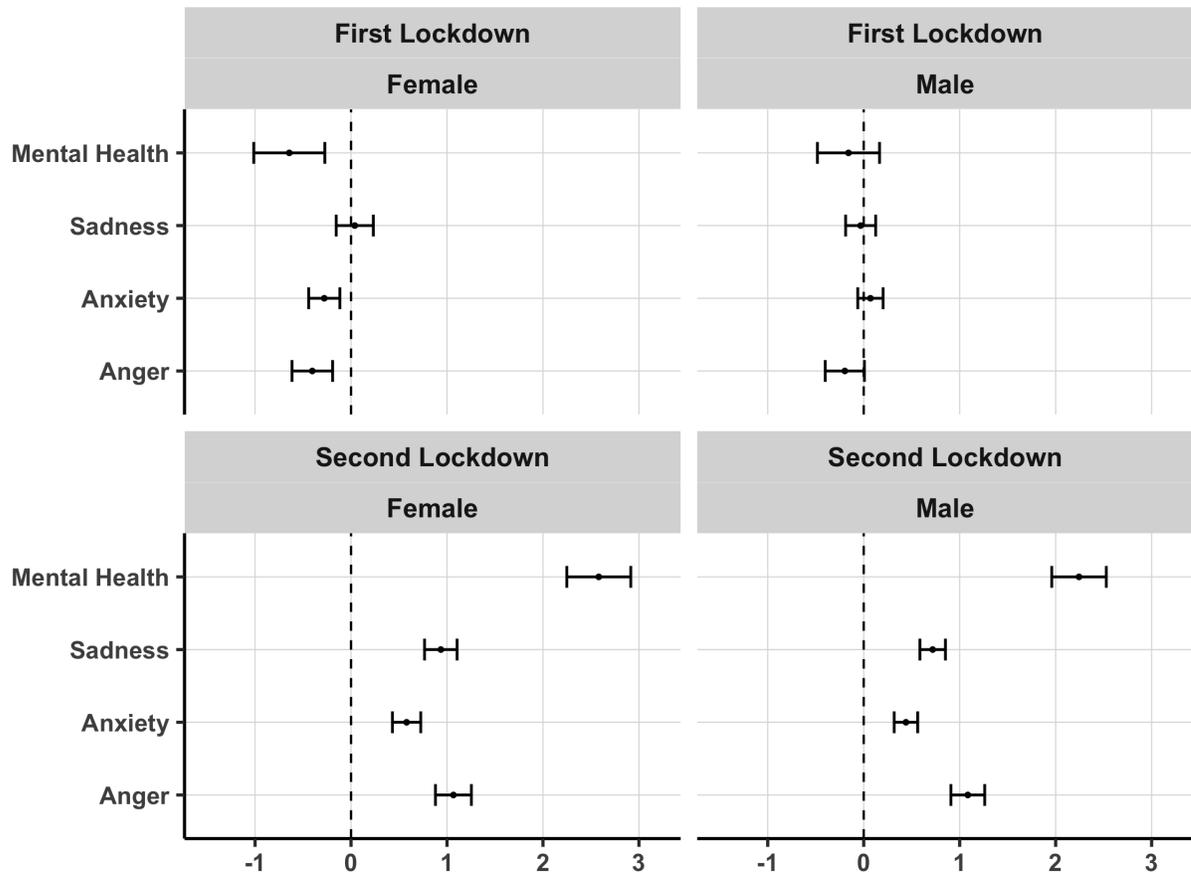
Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 90% confidence interval. The grey dotted line is the line of null effect. All estimates with clustered robust standard errors are reported in Table 4.

Fig. 5. Estimated Impact of Lockdowns by Broad Age Groups



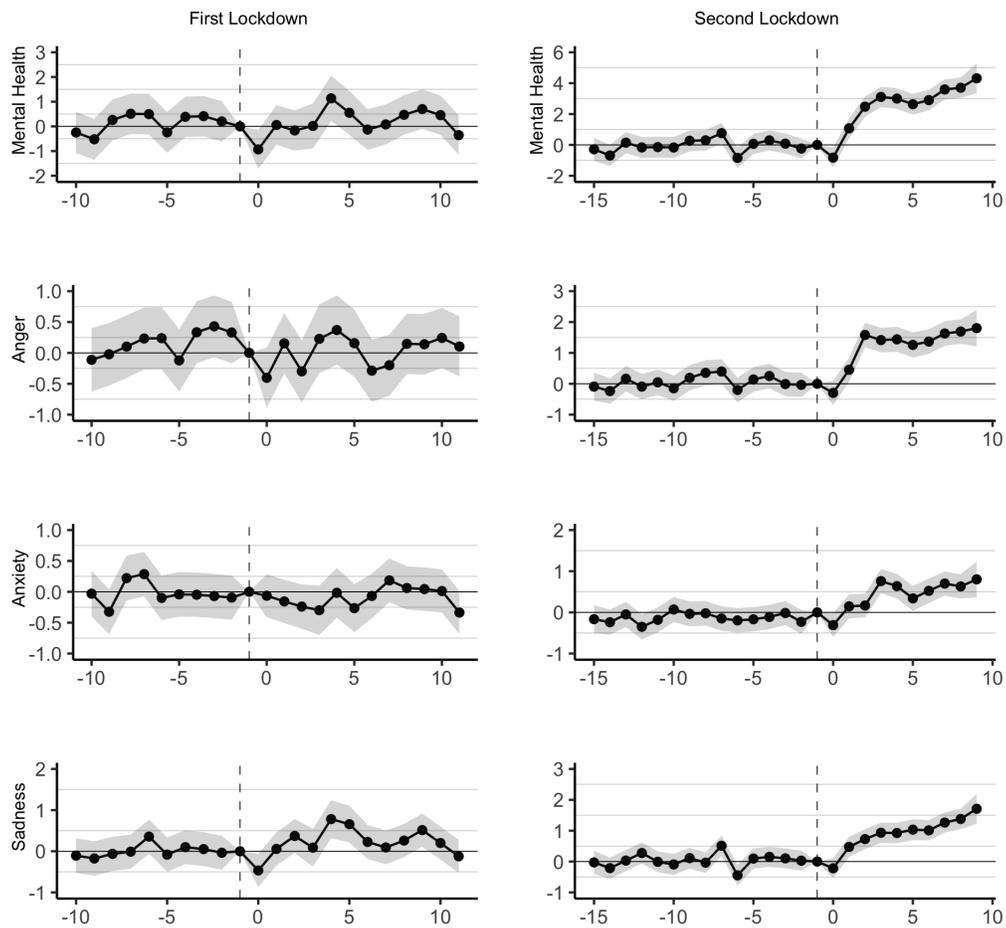
Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 90% confidence interval. The grey dotted line is the line of null effect. All estimates with associated clustered robust standard errors are reported in Table 4.

Fig. 6. Estimated Impact of Lockdowns by Gender



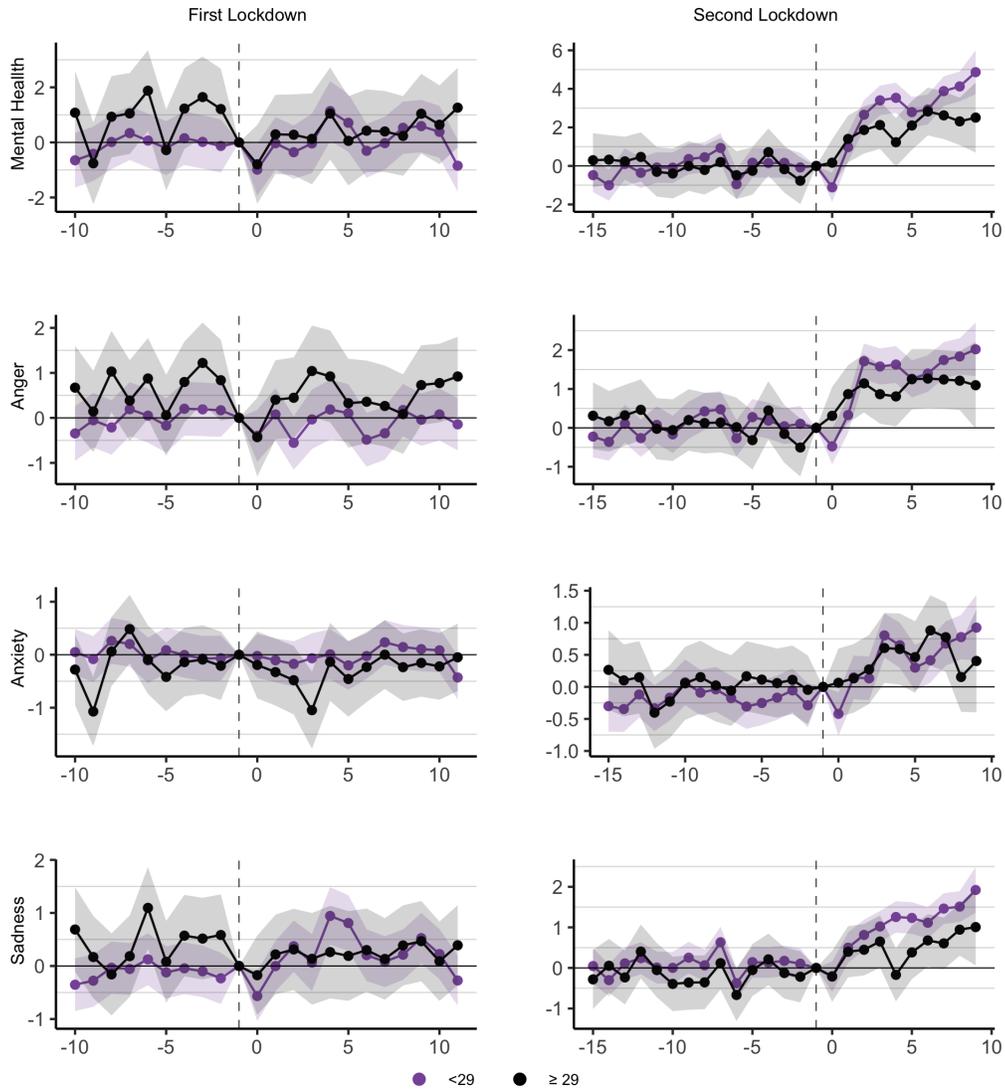
Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 90% confidence interval. The grey dotted line is the line of null effect. All estimates with clustered robust standard errors are reported in Table 5.

Fig. 7. Change in Mental Health for All Users



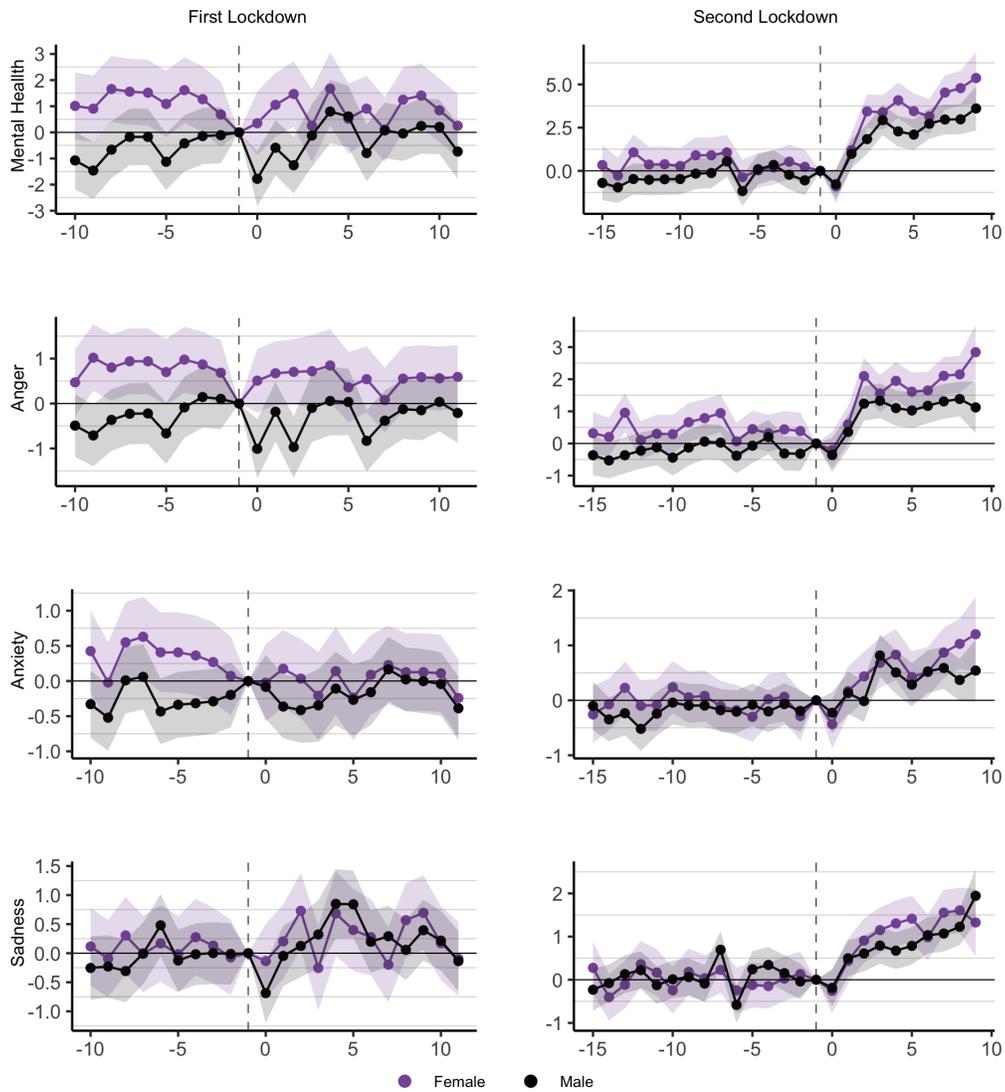
Note: Evolution of emotions weeks before and after the announcements of the first and the second stay-at-home orders.

Fig. 8. Change in Mental Health by Age Groups



Note: Evolution of emotions weeks before and after the announcements of the first and the second stay-at-home orders.

Fig. 9. Change in Mental Health by Gender



Note: Evolution of emotions weeks before and after the announcements of the first and the second stay-at-home orders.