amse
école d'économie d'aix-marseille
aix-marseille school of economics

Working Papers / Documents de travail

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

Sara Colella
Frédéric Dufourt
Vincent A. Hildebrand
Rémi Vivès

WP 2022- Nr 18 v2

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

First version: July 2022; Revised version: August 2023

Abstract

We use a distinctive methodology that leverages a fixed population of Twitter users located in France to gauge the mental health effects of repeated lockdown orders. To do so, we derive from our population a mental health indicator that measures the frequency of words expressing anger, anxiety and sadness. Our indicator did not reveal a statistically significant mental health response during the first lockdown, while the second lockdown triggered a sharp and persistent deterioration in all three emotions. Our estimates also show a more severe deterioration in mental health among women and younger users during the second lockdown. These results suggest that successive stay-at-home orders significantly worsen mental health across a large segment of the population. We also show that individuals who are closer to their social network were partially protected by this network during the first lockdown, but were no longer protected during the second, demonstrating the gravity of successive lockdowns for mental health.

JEL classification: C81, I12, I18, I31

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

---

*Forthcoming in *Economics and Human Biology*. 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 Roland Pongou and three anonymous referees for their valuable recommendations, and to Boriana Miloucheva, Philippe Van Kerm and Bruno Ventelou for 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. The benefit of these policies is clear: their success in minimizing the virus' transmission was 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 both the individual and population level.

It is well-documented that the implementation of the most restrictive containment measures coincided with deteriorating mental health conditions at the onset of the pandemic. For example, a number of studies using 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 disproportionally 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, France and the United States (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 at the onset of the pandemic revealed by Google Trends data provide further evidence of deteriorating mental health during lockdowns (Brodeur et al., 2021; Silverio-Murillo et al., 2021). Interestingly, the 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.

It is highly likely that one day in the future, we will face another pandemic. In order to prepare, it would be prudent to take the lessons we have learned from the COVID-19 pandemic to determine how to contain the spread of future pathogens while minimizing the negative effects of containment measures on the population. This paper aims to contribute to this discussion by examining the costs of the first two lockdown episodes in France to individuals' psychological well-being. To the best of our knowledge, this study is among the few to explicitly document the deleterious mental health effects of *successive* lockdowns. To assess this cost, we use a unique methodology that leverages a fixed population of Twitter users located in France.

France provides an interesting case study for two main reasons. First, it belongs to a set of Western European 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 deteriorated during these two lockdown episodes in France, it is much 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

[1] See section 2 for further details.

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 an analysis of the first two lockdowns, that the cost to mental health of successive waves of containment orders sharply increased 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 and women.

The first contribution of this paper is methodological: we create a unique longitudinal dataset of a fixed Twitter population and their tweets covering the year of the first stages of the pandemic (2020) and the year prior (2019). We measure mental health with indicators built using the Linguistic Inquiry and Word Count (LIWC) lexicon (Pennebaker et al., 2015) to identify words in a given tweet related to negative emotions such as anger, anxiety or sadness.[2] We then use a machine learning algorithm developed by Wang et al. (2019) to infer the age and gender of the Twitter users.[3] We further refine our population, distinguishing between high- and low-frequency users – an indicator of social connectedness that may be relevant to assessing the effect of social distancing. Inferring these individual characteristics enables us to assess whether mental health variations are heterogeneous across women and men, different age groups, and more or less socially connected (through social media) segments of the population. Although we build this dataset to study the mental health effects of pandemic lockdowns, it offers additional future possibilities for exploring the effects of various types of events on a fixed population of Twitter users.

The second contribution lies in the differences between our study and other COVID-19 related studies that use “digital traces" to measure the effects of lockdown on emotional well-being. First, unlike digital data such as Google Trends, our methodology allows us to follow a large fixed group of individuals across time while inferring their individual characteristics. Second, contrary to the vast majority of COVID-19 related Twitter studies, 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

[2]LIWC is a dictionary-based text analysis to measure the psychological meaning of word use. It allows us to measure the emotional tone of a text by calculating the frequency of words belonging to the lexical field of specific emotions.

[3]This model, known as the M3 inference model, is a multilingual, multimodal and multi-attribute deep learning system that uses both images and textual data of the users' profiles to predict their demographics.

well-being more accurately than does a subset of their tweets that only includes pandemic-specific keywords. These factors that distinguish our study from other COVID-19-related studies using digital traces shed additional light on the value of our methodological contribution more generally, and highlight the benefits of studying the effects of particular events by using the digital traces of a fixed population as opposed to keyword searches or Google Trends. In this vein, 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.

Finally, we contribute to the literature on COVID-19 more generally by providing new evidence that repeated lockdown policies are associated with increasing mental health costs of varying magnitude across gender and age groups and their interaction (see, e.g., Banks and Xu, 2020; Haesebaert et al., 2020; Pierce et al., 2020; Davillas and Jones, 2021; Pongou et al., 2022, among others). 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 exacerbates mental health inequalities between socio-economic groups. This result has significant potential implications for how we choose to manage future epidemics/pandemics.

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, illustrating the potential relevance of such a tool for monitoring mental health conditions in real time. We base our econometric analysis on three empirical strategies: a Difference-in-Difference estimation, an event study analysis, and a Regression Discontinuity Design with a Difference-in-Difference model, where the pandemic year (2020) is the treatment and the year 2019 serves as the control period to account for the confounding effect of a secular trend and seasonal patterns in mental health. 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 deterioration in their mental health 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).

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, and the resulting indicator can be monitored almost in real time.[4] 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 such an 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 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 11, 2020, following the rapid spread of novel coronavirus (COVID-19). This declaration prompted numerous governments to enact 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 14. On March 16, 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 11. Following that date, France entered a period of progressive deconfinement organized in two stages, proceeding differently by sub-region ("département")

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

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 11 to June 2), 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 2 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 30), there were 541 new daily positive cases compared to 5,233 at the peak of the first wave (April 3). 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 11, 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 5, Paris was put on "maximum alert" due to the spike in COVID-19 cases in the city and its suburbs.[5] On October 14, President Macron ordered the enforcement of local curfews on October 17 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). As COVID-19 cases continued to rapidly spread across the country, overnight curfews were extended to 38 more departments on October 22, affecting a total number of 54 departments and 46 million people (around 70% of the population) nationwide. Finally, on October 27, President Macron announced a second nationwide lockdown effective on October 30 and expected to last until December 1. 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 able to work from home. Unlike the first lockdown,

[5]In 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.

visits to public outdoor spaces such as parks, beaches, and hiking trails were permitted.

On November 24, the government announced that the originally set date of December 1 to end the second stay-at-home order was to be extended to December 15, and would then be replaced by a nighttime curfew, which lasted until June 20, 2021.[6]

For the purpose of this study, we date the first lockdown period as beginning on March 14, 2020 when schools, universities and all non-essential businesses were closed, and ending on June 2, with their reopening and the lift of the travel ban, which effectively ended the period of social isolation. We date the beginning of the second lockdown to October 22 when nighttime curfews were extended to about 70% of the metropolitan population because this event reinforced anticipation of new nationwide mobility restrictions. We date this lockdown as ending on December 15, when travel restrictions were lifted, as was 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 for disentangling the mental health cost of repeated containment policies because France adopted stringent nationwide containment policies that were rapidly supported by generous economic measures to support all economic actors.

On March 26, 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 workweek caused by a partial or a full temporary closure of operations.[7] 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, which amounted to roughly 84% of their net income. In April 2020, a peak of 8.6 millions of workers (roughly 30% of the working-age population) benefited from this scheme. 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.

Due to these support measures, the total cost of which was estimated at €120B, the vast majority of households in France did not suffer significant losses in income during the COVID-19

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

[7] On February 1, 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.

pandemic. According to the French National Institute of Statistics (INSEE, 2021), the disposable income of households actually slightly increased by 0.4% in real terms between 2019 and 2020. Since many work-related and leisure-related expenses (such as transportation, meals outside of the home, cultural and sports activities, etc.) were also strongly reduced during this period, the saving rate of households also increased from 15.1% to 21.4% of disposable income between 2019 and 2020. All this substantiates our claim that the fear of income losses was not a likely driver of the mental health deterioration observed during lockdown episodes.

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.

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 28, 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 (#LockdownDayXx).

These papers contributed to the rapidly expanding field of research on 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) and 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 Internet Archive's Twitter Stream Grab database. We then retrieve all tweets posted by each user in our sample with the Twitter Application Programming Interface (API) over a period spanning from January 2019 to March 1, 2021.[8] 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, 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 Internet Archive's Twitter Stream Grab database. This archive stores a 1 percent random sample of all tweets posted since September 2011.[9] 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.[10]

[8] March 1, 2021 is end point for the dataset used in this paper given our focus on the mental health impact of the first two lockdowns. The data collection for our population is, however, still ongoing and available at `https://twitterspehereobservatory.github.io/`.

[9] The Internet Archive's Twitter Stream Grab database stopped collecting data in June 2021.

[10] Twitter Archive database made data available with a 6-month lag.

To ensure that the reported location is in France, we cross-reference it against the French GeoNames database after normalizing both the twitter data and GeoNames.[11] 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 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 their 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.[12] 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 API. To complete this step, for each user, we scrape data backwards in time up to January 1, 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 1, 2019. Including data covering the year 2019 provides an

[11] `http://geonames.org/`. Retr. June, 2020.

[12] See, `https://ec.europa.eu/eurostat/web/nuts/background`, for further details on the NUTS classification system.

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 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 multilingual, multimodal and multi-attribute deep learning system exploits reported usernames, profile names, profile descriptions, 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.

Finally, we use Botometer to distinguish the accounts held by individuals from bots; see Sayyadiharikandeh et al. (2020) for further details. Botometer is recognized as being one of the most reliable procedures for bot detection. The procedure is based on supervised machine learning trained on very large bot datasets in multiple languages. We follow Keller and Klinger (2019) and Jemielniak and Krempovych (2021) and set a Complete Automation Probability score at 0.76. The procedure classified 939 users as being bots. The total number of remaining users in the dataset is 39,031. We provide descriptive statistics in Section 4.1.

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 (Coppersmith et al., 2014; Park et al., 2012). 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, stopwords and websites.[13] We also convert the text to lower case. Second, we run an emotion analysis based on the LIWC dictionary in French (Garcia and Rimé, 2019; Pennebaker et al., 2015; Piolat et al., 2011) 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 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 incorrectly generate an indicator of sadness of 1/5.[14] The systematic treatment of commonly used words to negate ideas has the benefit of significantly minimizing potential misclassification of emotions.[15] For each individual user i posting tweets on 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

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

[14]In the formal analysis, we also remove stopwords. 'je', 'ne' and 'suis' are removed.

[15]Our approach does not control for sarcasm, which is a pitfall of most natural language processing (NLP) applications. Detecting sarcasm is a challenging task, not only for NLP applications but also for human beings: see for instance Camp (2012) and Ilavarasan et al. (2020) for further details. We believe that the bias induced by not controlling for sarcasm does not affect our results negatively. In fact, since sarcasm is primarily defined as a positive text with an underlying negative sentiment, not controlling for sarcasm must result in underestimating levels of anger, anxiety and sadness. Hence, controlling for sarcasm would rather reinforce our results instead of mitigating them.

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.[16]

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 tweets capture all responses to another tweet, whereas a retweet is the re-posting of a tweet that can be your own or someone else's. 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 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 10,322,655 daily observations from 52,341,001 tweets posted by 39,031 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 1, 2019 and March 1, 2021.[17][18]

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.

[16] `https://twittersphereobservatory.github.io/`. The measures are updated weekly with a one-week lag.

[17] Remember that our mental health indicators include the words in all the tweets of user i on day t. We have 10,322,655 daily observation following this aggregation. In other words, one daily observation can include many tweets.

[18] The list of users with their corresponding predicted demographic information is available upon request.

About 76% 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 under 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 10.1% more closely mirrors 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 points 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 more closely mirrors 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, just 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 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 0.014%, and is 0.539% for the keywords used in Banda et al. (2021). Our dataset thus tweets other than those 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, to investigate econometrically the effects of each lockdown on our users' mental health, we extract two subsamples from our dataset 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 1 to June 2 in both 2019 and 2020. June 2, 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 3,482,864 daily observations capturing the change in emotions of 32,063 active unique users.

Sample 2 covers the period spanning from July 1 to December 15 in both 2019 and 2020. On July 1, 2020, France reopened its borders with non-European Union countries, while the second lockdown officially ended on December 15, 2020. Likewise, the matching period in 2019 is used as the comparison period in the regression analysis. Sample 2 includes 4,711,643 daily observations from 38,040 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 1 2019 and December 31 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 behavioral response to lockdown orders.

Figure 2 depicts the 7-day moving average of our aggregate mental health indicator in 2020. 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 slightly higher than were 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 28.[19]

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 implementation of a 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 14. Interestingly, extending these nighttime curfews on October 22 to a total of 54 departments affecting approximately 70% of the French population also coincides with a period of sharp increase in the mental health indicator that persisted over the entire lockdown 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 15. It remained, however, at a higher level than before the first lockdown, 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

[19] As 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 2.

health indicator. They reveal an increase in the level of anxiety during the pandemic year compared to 2019 that gradually intensifies during the lockdown period. 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 lockdown 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 lockdown 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 explore the impact of enforced lockdowns on mental health as well as the existence of heterogeneous responses across gender, age groups and social connectedness.

5 Statistical Analysis

5.1 Empirical Strategy

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. In the spirit of Brodeur et al. (2021), we implement this strategy using three complementary estimation methods: a Difference-in-Difference (DID) estimation, an event study, and a combined Regression Discontinuity Design (RDD) in time with a Difference-in-Difference (DID) model.

5.1.1 Difference-in-Difference Estimator

Our DID model uses the calendar year prior to the onset of the COVID-19 pandemic (2019) as the comparison year to account for the confounding effect of a secular trend and seasonal patterns in mental health. This model enables us to estimate the average effect of mobility restrictions on mental health over the entire lockdown period.

A key identifying assumption of this approach is that, in the absence of lockdowns, the emotional expression in Tweets during the pandemic year would have been comparable to those in 2019 (common trend assumption). Hence, to measure the impact of lockdown orders on mental health, we estimate separately for each lockdown the following model:

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

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 on the days covering the period of each lockdown in 2020 and on the days covering the same time period in the comparison year (2019). Y_{20} is a binary variable taking the value 1 the year of the pandemic (2020) and 0 for the comparison year (2019). Z_{t-1} controls for the lagged number of new daily deaths from COVID-19 per million to account for the potential confounding effect of the severity of the pandemic that may also affect individual emotional health. 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.

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:

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

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 (16) weeks leading up to the first (second) lockdown and 11 (8) 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 health relative to the comparable week in 2019 as measured from the week prior to the implementation of each lockdown order.[20]

[20]Our measure of mental heatlh captures expressed anger, anxiety and sadness in the tweets.

5.1.3 RDD-DID

Finally, to identify potential structural breaks in mental well-being following each lockdown announcement, we follow Brodeur et al. (2021) and combine a Regression Discontinuity Design (RDD) in time with a Difference-in-Difference (DID) model and estimate the following RDD-DID model:

$$
\begin{aligned}
MH_{it} = \; & \delta(L_t \times Y20) + \psi\, f\,(D_{it}) * L_t \times \, Y20 + \theta\, f(D_{it})(1 - L_t) \times \, Y20 \\
& + \phi\, f(D_{it}) \, \times \, L_t + \lambda\, f(D_{it})(1 - L_t) + \beta\, L_t \\
& + \gamma Z_{t-1} + \alpha_i + \mu_t + \epsilon_{it} \qquad (3)
\end{aligned}
$$

whereby D_{it} is the distance in days from the announcement of the lockdown or cut-off dates. This distance is calculated for each individual i tweeting on day t. D_{it} is positive for the days following the lockdown announcement ($L_t = 1$) and negative for the days prior to the announcement ($L_t = 0$). $f(.)$ is a polynomial function of the distance in days from the lockdown announcement that captures the trends in the outcomes of interest on either side of the cut-off dates. Our regression models use a linear and a quadratic function for $f(.)$.

As discussed *supra*, the cut-off dates retained for the first and second lockdown are March 14 and October 22, respectively. Z_{t-1}, μ_t and ϵ_{it} are defined as in equation (1). δ is our parameter of interest and can be interpreted as a measure of the immediate effect of the lockdown announcement on emotional well-being. A key identifying assumption is the absence of other policy changes at the cut-off point, in addition to the aforementioned common trend assumption. All reported standard errors are heteroskedastic-robust and clustered at the individual level.[21]

5.2 Main Results

5.2.1 Difference-in-Difference Estimation Results

5.2.1.1 Overall Impact of Lockdowns on Emotional Well-being

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

[21] As a robustness check, we also estimate a Donut-RDD-DID, in which we remove observations between the first local curfews on October 14, 2020 and their widespread extension on October 22 to capture the potential anticipatory effects of the progressive imposition of local nighttime curfews over this period, Results are provided in the supplementary online Appendix.

the second lockdown.

We do not find statistically significant evidence that the first lockdown is associated with worsening mental health in our Twitter population. While the announcement of the first lockdown order coincides with a sharp increase in tweet volume, 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.[22] By contrast, we find a statistically significant *decline* in anger leading 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 e.g. Perrona and Senik, 2021).

In contrast, we find that the second lockdown episode is associated with a significant deterioration in mental health, evidenced by a statically significant increase of about 15% in all expressed emotions compared to the baseline period in 2019, 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 may increasingly be more harmful to individual well-being, a result that is of significant importance for the cost-benefit analysis of these policies on population health.

5.2.1.2 Impact of Lockdowns by Age, Gender and Social Connectedness

Environmental factors can explain unequal health outcomes across different segments of the population, which are unexplained by biological differences. For example, Pongou (2013) shows that the preconception environment explains a significant part of the difference in mortality rates between boys and girls. In the context of lockdowns, different mental health effects between men and women can result from a variety of biological reasons and environmental factors associated with a disproportionate burden of household chores and childcare (the “mental load" hypothesis) and a higher risk of intimate partner violence (see, e.g.,Almeida et al. 2020 for a review). Likewise, younger individuals may experience stronger increases in psychological distress during lockdowns (despite being less vulnerable to COVID-19 infection) due to a larger exposure to social isolation and other socioconomic considerations like a higher probability of living alone and in smaller spaces (see, e.g.,Haesebaert et al. (2020); Pierce et al. (2020); Vacchiano (2022))

One advantage of our longitudinal study is that we observe additional individual characteristics (beyond gender and age) that standard keyword-based studies cannot observe. For example, we observe the number of tweets posted each day by a single user and/or the frequency at which this user uses Twitter as a communication device in a given period of time. As emphasized in earlier studies

[22] As 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 11, 2020 (see Figure 1, section 4.2).

(see e.g. Riedl et al. 2013), these measures adequately proxy the degree of "social connectedness" of an individual, at least on social media, which may provide a relevant indicator for assessing the differentiated mental health effects of lockdowns. We thus explore these heterogeneous mental heath responses to enforced lockdowns across age groups, gender and social connectedness within our Twitter sample.

Age Groups

We consider two broad age categories: users aged 28 and under (younger users) and users older than 28 (older users). We find evidence of a differentiated mental health response among age groups unfavorable to younger users (see Figure 5 and point estimates reported in Table 4). The difference is moderate during the first lockdown but markedly more pronounced during the second lockdown. During the latter, both age groups experienced statistically significant increases in anger, anxiety and sadness compared to the year before, but these increases are larger for younger users. 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 (see Figure 6 and coefficient estimates in Table 5). 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. Overall, our results of a moderately better emotional response for women than for men to the first lockdown episode in France contrasts with the findings of most earlier studies conducted in other countries that found a larger deterioration in mental health for women than for men during the first months of the pandemic.

However, our results for the second lockdown are markedly different. Figure 6 reveals that while mental health significantly deteriorated for both men and women during this episode in comparison to the year before, this deterioration was greater for women than for men (mostly driven by a disproportionately larger increase in expressed sadness). Overall, these findings 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.

Social connectedness

According to a recent study (Annan and Archibong, 2022), communication, or more generally the ability to stay connected with relatives, friends, etc., meaningfully improves mental well-being. To test this hypothesis within our framework, we split our sample into two broad categories of users based on the number of tweets they posted during the period prior to the first lockdown. Our main hypothesis is that users who tweet frequently have, on average, a larger social network and potentially larger possibilities to communicate with others (following Riedl et al. 2013), thereby mitigating the effects of physical social distancing compared to less active users. We thus label the group of users who posted more tweets than the median user in the comparison year "high-frequency users", and we label the complementary group "low-frequency users". Results are presented in Figure 7.

During the first lockdown, our results largely confirm those obtained by Annan and Archibong (2022), though in a very different context. High-frequency users appear to experience a significant improvement in the state of their mental health compared to the year before, while in contrast, low-frequency users experience a statistically significant deterioration. Yet, this difference completely vanishes during the second lockdown, which tends to confirm that the second lockdown affected the mental health of a much broader segment of the population.[23]

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 8, 9 and 10 for the pooled sample of all users across gender and age groups, respectively.

Event study results reinforce our DID findings. No sizeable variation in mental health over the course of the first stay-at-home order is found, except for a small and transitory peak in sadness observed roughly four weeks after the beginning of the lockdown. 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. Event study plots across age groups and gender also confirm that the mental health of younger users and of women deteriorated more during the second lockdown than that of older users and of men.

Compared to our DID findings, an interesting aspect of our event study estimates is that they provide an assessment of the degree of persistence of the effects of lockdowns on emotional well-

[23]One potential pitfall of this analysis is that the differentiation between high and low-frequency users may partially overlap with the former distinctions by age or by users' gender. For example, younger or male users could post proportionately more tweets than older or female users. To control for this bias, we repeated our estimations by combining in various ways the different individual characteristics of users. Estimation results are presented in Table A.1 of the supplementary Appendix. Our conclusion that high-frequency users experienced a less severe mental health deterioration than low-frequency users during the first lockdown is overwhelmingly confirmed across each sub-segment of the population.

being. Figure 8 shows that the deterioration in mental health during the second lockdown persisted over the entire period covering this lockdown, with the aggregate mental health indicator reaching a plateau roughly three weeks after the initial implementation.

In this regard, our results usefully complement those obtained by studies relying on other sources of digital traces such as Google Trends to infer the mental health consequences of stay-at-home orders. For example, Brodeur et al. (2021) find long-lasting effects of lockdowns on searches for boredom in the US and several Western European countries, but only transitory effects on searches for loneliness and contrasted effects on searches for sadness. By contrast, searches for stress, anxiety and sadness remained high throughout the entire lockdown period for 11 Latin American countries studied by Silverio-Murillo et al. (2021). Our results suggest that, to the extent that mental health starts to deteriorate following a new stay-at-home order, this deterioration is likely to be a long-lasting one since all three emotions underlying the construction of our aggregate mental health indicator (anger, anxiety and sadness) increase and remain high throughout the entire lockdown period.

5.2.3 RDD-DID Results

A visual inspection of the immediate effects of the two lockdown announcements on emotional health are presented descriptively in exploratory RDD plots in Appendix Figures B.1 and B.2. Consistent with earlier findings, these plots provide suggestive evidence that a structural break associated with a deterioration in mental health only occurred after the announcement of the second lockdown, while we do not observe any discontinuity during the same period in the control year 2019. By contrast, no noticeable break is observed after the announcement of the first lockdown.

RDD-DID estimation results are reported in Appendix Tables A.2 and A.3 for different choices of polynomial functions. They show that the announcement of the first lockdown is associated with a statistically significant improvement in mental health largely driven by male users whereas the announcement of the second lockdown is associated with a statistically significant deterioration in all emotions underlying our mental health indicator. Furthermore, also in line with our DID and event study results, RDD-DID estimates suggest that the structural break at the announcement of the second lockdown is more pronounced among females and younger users. These statistically significant local treatment effect estimates provide additional evidence that the association between lockdown and the deterioration of mental health may have a causal interpretation.

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 to 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 entire 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, European and Latin American countries using either digital traces such as Google Trends (Brodeur et al., 2021; Knipe et al., 2020; Silverio-Murillo et al., 2021) or 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.[24] 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) report 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

[24] See CoviPrev survey on `https://www.santepubliquefrance.fr`

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 COVID-19 lockdowns had a deteriorating effect on 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, which reports 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 8). 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 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 such a degree that many respondents were characterized as presenting severe anxiety and depression symptoms according to standard clinical measures. At the same time, it is interesting to observe that in spite of these differences, our mental health indicator and the sub-indices associated with anxiety, sadness and anger all correlate quite well with the anxiety index of the Coviprev survey. For example, the correlation between our aggregate mental health indicator and the Coviprev anxiety index is 0.67. See Table A.6 in the supplementary Appendix for details.

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 among younger people and/or students compared to the rest of the population during lockdowns; see Husky et al. (2020), Haesebaert et al. (2020) for France, Banks and Xu (2020); Davillas and Jones (2021); Pierce et al. (2020) for the UK, Pongou et al. (2022) for Canada, and 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 in comparison to the previous year. By contrast, during the second lockdown, the age gap is explained by a greater deterioration 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 disproportionally more affected by lockdown orders than non-students (see Arsandaux et al., 2021; Husky et al., 2020; Macalli et al., 2021). A finer analysis of our Twitter population that distinguished between young users who are students and those who are not 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, Pongou et al. (2022) for Canada, Vloo et al. (2021) for the Netherlands and the references therein.[25] 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 CoviPrev 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 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 in the CoviPrev depression index over this period, albeit with a milder amplitude than the anxiety index. 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, and 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 limitations that merit 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

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

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 a population's emotional response to the kind of large scale major events a society can experience. To validate the use of these types of indicators as public health monitoring tools that can robustly inform policy, more empirical evidence that they are accurate must be gathered by comparing the mental health of individuals as derived through Twitter-based indicators to validated ground truth (see, e.g., Di Cara et al. 2023).

References

Adams-Prassl, A., T. Boneva, M. Golin, and C. Rauh (2022). The impact of the coronavirus lockdown on mental health: evidence from the united states. *Economic Policy*. 25

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. 24

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

Almeida, M., A. D. Shrestha, D. Stojanac, and L. J. Miller (2020, December). The impact of the COVID-19 pandemic on women's mental health. *Archives of Women's Mental Health 23*(6), 741–748. 19

Annan, F. and B. Archibong (2022). The value of communication for mental health. *Brookings Global Working Paper 177*. 21

Arendt, F., A. Markiewitz, M. Mestas, and S. Scherr (2020). COVID-19 pandemic, government responses, and public mental health: Investigating consequences through crisis hotline calls in two countries. *Social Science & Medicine 265*, 113532. 2, 23

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, 23

Arsandaux, J., I. Montagni, M. Macalli, N. Texier, M. Pouriel, R. Germain, A. Mebarki, S. Kinouani, M. Tournier, S. Schuck, and C. Tzourio (2021). Mental health condition of college students compared to non-students during COVID-19 lockdown: the CONFINS study. *BMJ Open 11*(8), e053231. 25

Balech, S., C. Benavent, and M. Calciu (2020). The first french covid19 lockdown twitter dataset. *arXiv preprint arXiv:2005.05075*. 3, 8, 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. 8, 14

Banks, J. and X. Xu (2020). 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, 4, 23, 24, 25

Blanchard, O., T. Philippon, and J. Pisani-Ferry (2020). A new policy toolkit is needed ascountries exit covid-19 lockdowns. *Policy Contribution 12*. 7

Boyd, R. L. and H. A. Schwartz (2020). 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). COVID-19, lockdowns and well-being: Evidence from google trends. *Journal of Public Economics 193*, 104346. 2, 16, 18, 22, 23

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

Brülhart, M., V. Klotzbücher, R. Lalive, and S. K. Reich (2021). 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

Camp, E. (2012). Sarcasm, pretense, and the semantics/pragmatics distinction. *Noûs 46*(4), 587–634. 12

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, 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). 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, 4, 24, 25

Di Cara, N. H., V. Maggio, O. S. Davis, and C. M. Haworth (2023). Methodologies for monitoring mental health on twitter: systematic review. *Journal of Medical Internet Research 25*, e42734. 26

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. 12

Foki, G. (2021). Covid-19 : les différents systèmes d'indemnisation du chômage partiel en europe. Technical report, Unédic. 7

Garcia, D. and B. Rimé (2019). Collective emotions and social resilience in the digital traces after a terrorist attack. *Psychological Science 30*(4), 617–628. 12, 13

Gruzd, A. and P. Mai (2020). Going viral: How a single tweet spawned a COVID-19 conspiracy theory on twitter. *Big Data & S10.1177/0261927x20967028ociety 7*(2), 205395172093840. 8

Gupta, R., A. Vishwanath, and Y. Yang (2021). Covid-19 twitter dataset with latent topics, sentiments and emotions attributes. 8

Haesebaert, F., J. Haesebaert, E. Zante, and N. Franck (2020). Who maintains good mental health in a locked-down country? a french nationwide online survey of 11,391 participants. *Health & Place 66*, 102440. 4, 19, 24

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). Stress and anxiety among university students in France during covid-19 mandatory confinement. *Comprehensive Psychiatry 102*, 152191. 24, 25

Ilavarasan, E. et al. (2020). A survey on sarcasm detection and challenges. In *2020 6th International Conference on Advanced Computing and Communication Systems (ICACCS)*, pp. 1234–1240. IEEE. 12

Insee (2021). Estimation de population, (données provisoires et révisées arrêtées à fin février 2021). 14

INSEE (2021). *"France - Portrait social"*. INSEE References - Edition 2021. "En 2020, malgré la crise sanitaire, le pouvoir d'achat des ménages résiste et leur épargne augmente", pp. 57–68. 8

Jemielniak, D. and Y. Krempovych (2021). An analysis of astrazeneca covid-19 vaccine misinformation and fear mongering on twitter. *Public Health 200*, 4–6. 11

Keller, T. R. and U. Klinger (2019). Social bots in election campaigns: Theoretical, empirical, and methodological implications. *Political Communication 36*(1), 171–189. 11

King, G., P. Lam, and M. E. Roberts (2017). Computer-assisted keyword and document set discovery from unstructured text. *American Journal of Political Science 61*(4), 971–988. 9

Knipe, D., H. Evans, A. Marchant, D. Gunnell, and A. John (2020, June). Mapping population mental health concerns related to covid-19 and the consequences of physical distancing: a google trends analysis. *Wellcome Open Research 5*, 82. 23

Lucchini, M., T. Gerosa, M. Pancheva, M. Pisati, C. Respi, and E. Riva (2021). 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. 23, 24

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. 25

Mellon, J. and C. Prosser (2017). 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. 14

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, 14

Oreffice, S. and C. Quintana-Domeque (2021). Gender inequality in COVID-19 times: evidence from UK prolific participants. *Journal of Demographic Economics*, 1–27. 2, 25

Park, M., C. Cha, and M. Cha (2012). Depressive moods of users portrayed in twitter. 11

Pennebaker, J. W., R. L. Boyd, K. Jordan, and K. Blackburn (2015). The development and psychometric properties of liwc2015. Technical report. 3, 12

Perrona, M. and C. Senik (2021). Le bien-être des français - mars 2021 un an après : l'usure. techreport 2021-02, Observatoire du Bien-être du CEPREMAP. 19, 23, 24

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). 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, 4, 19, 24, 25

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. 12

Pongou, R. (2013, November). Why is infant mortality higher in boys than in girls? a new hypothesis based on preconception environment and evidence from a large sample of twins. *Demography 50*(2), 421–444. 19

Pongou, R., B. O. Ahinkorah, S. Maltais, M. C. Mabeu, A. Agarwal, and S. Yaya (2022, November). Psychological distress during the COVID-19 pandemic in canada. *PLOS ONE 17*(11), e0277238. 4, 24, 25

Ramiz, L., B. Contrand, M. Y. R. Castro, M. Dupuy, L. Lu, C. Sztal-Kutas, and E. Lagarde (2021). A longitudinal study of mental health before and during COVID-19 lockdown in the french population. *Globalization and Health 17*(1). 23

Riedl, C., F. Köbler, S. Goswami, and H. Krcmar (2013, January). Tweeting to feel connected: A model for social connectedness in online social networks. *International Journal of Human-Computer Interaction 29*. 20, 21

Sayyadiharikandeh, M., O. Varol, K.-C. Yang, A. Flammini, and F. Menczer (2020). Detection of novel social bots by ensembles of specialized classifiers. In *Proceedings of the 29th ACM international conference on information & knowledge management*, pp. 2725–2732. 11

Serrano-Alarcón, M., A. Kentikelenis, M. Mckee, and D. Stuckler (2021). 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, 25

Silverio-Murillo, A., L. Hoehn-Velasco, A. R. Tirado, and J. R. B. de la Miyar (2021). COVID-19 blues: Lockdowns and mental health-related google searches in latin america. *Social Science & Medicine 281*, 114040. 2, 22, 23

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, 14

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

Vacchiano, M. (2022, May). How the first COVID-19 lockdown worsened younger generations' mental health: Insights from network theory. *Sociological Research Online 28*(3), 884–893. 19

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. 25

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. 3, 11

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

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

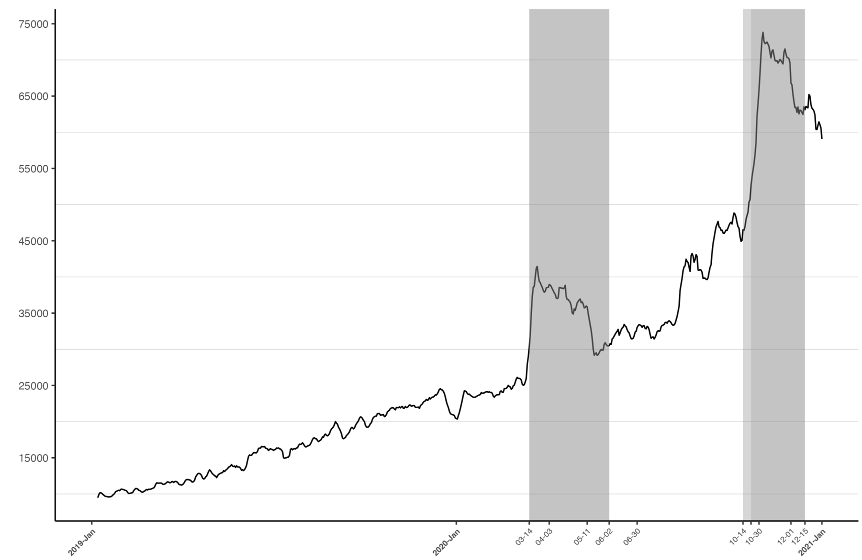

Note: Figure plots the 7-day moving average of the number of original tweets and replies to tweets from January 1 2019 to December 31 2020. The shaded areas indicate the first and second lockdown period, March 14 to June 2 2020 and October 22 to December 15 2020, respectively. The lighter shaded area covers the progressive enforcement period of local curfews.

Fig. 2. Daily Change in Mental Health

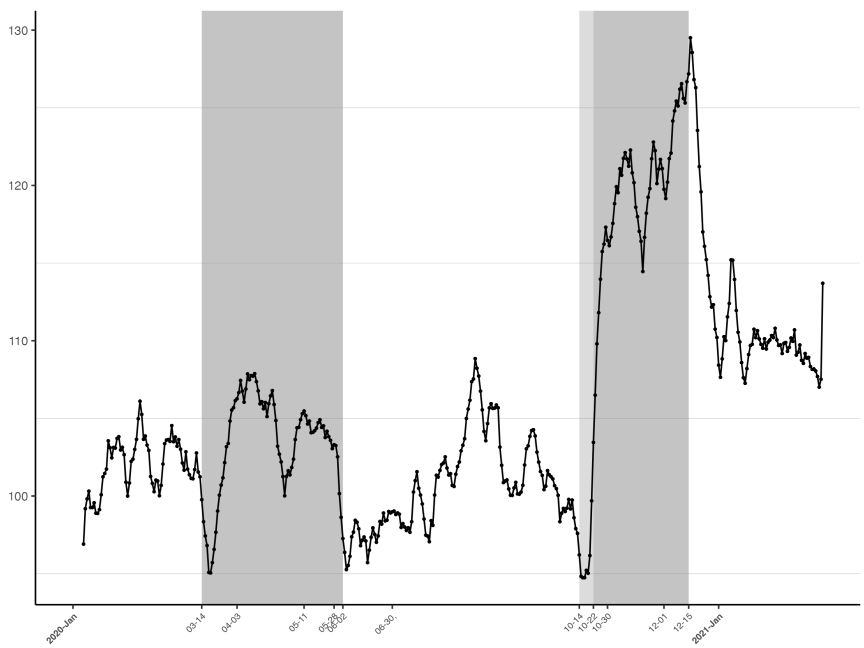

Note: Figure plots the 7-day moving average of the mental health indicator in 2020. Shaded areas cover the periods of the first and second lockdown order, March 14 to June 2 2020 and October 22 to December 15 2020, respectively. The lighter shaded area covers the progressive enforcement period of local curfews.

Fig. 3. Daily Change in Mental Health

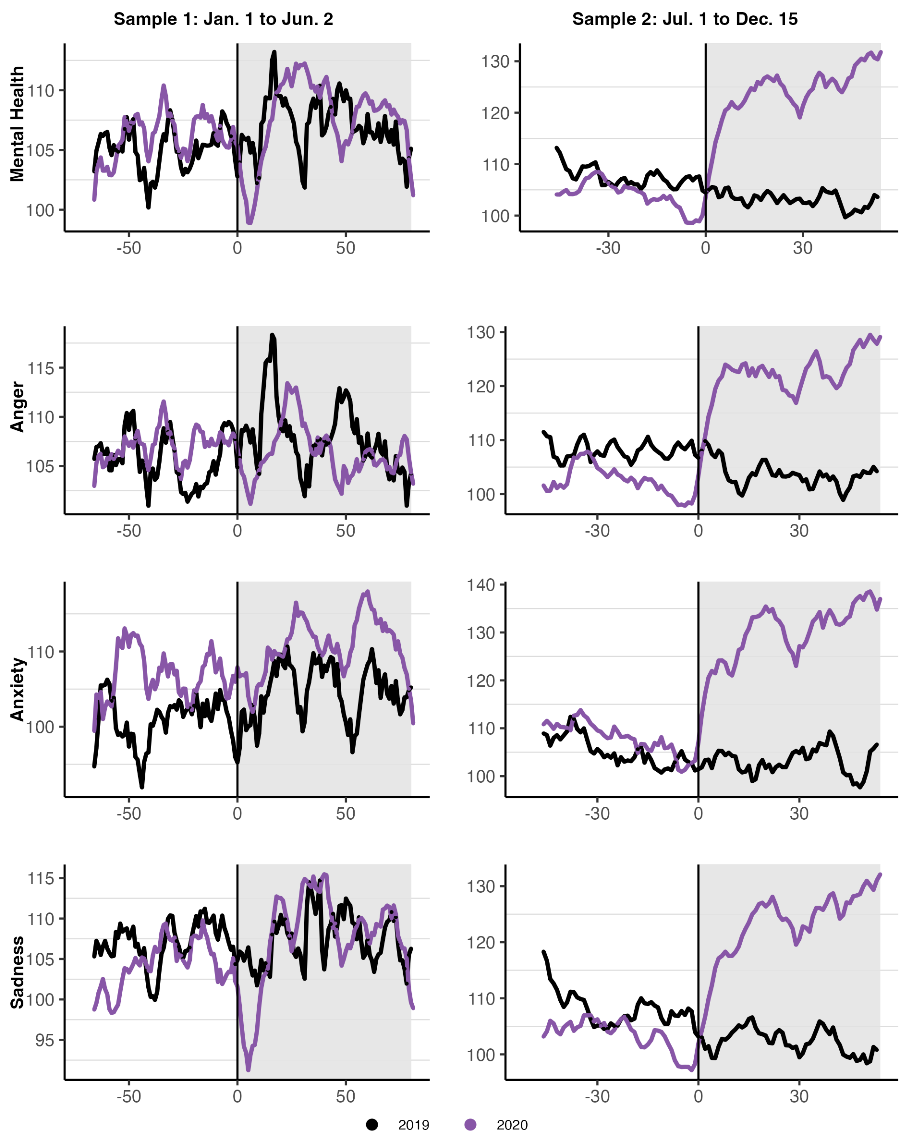

Note: Figure depicts the 7-day moving average of each indicator of emotions for sample 1 (left) and sample 2 (right). For Sample 1, the vertical line passing through the x-axis at "day-zero" of each plot corresponds to March 14. For Sample 2, "day-zero" corresponds to October 22, the day President Macron announced the extension of nighttime curfews.

Fig. 4. Estimated Impact of Lockdowns: All users

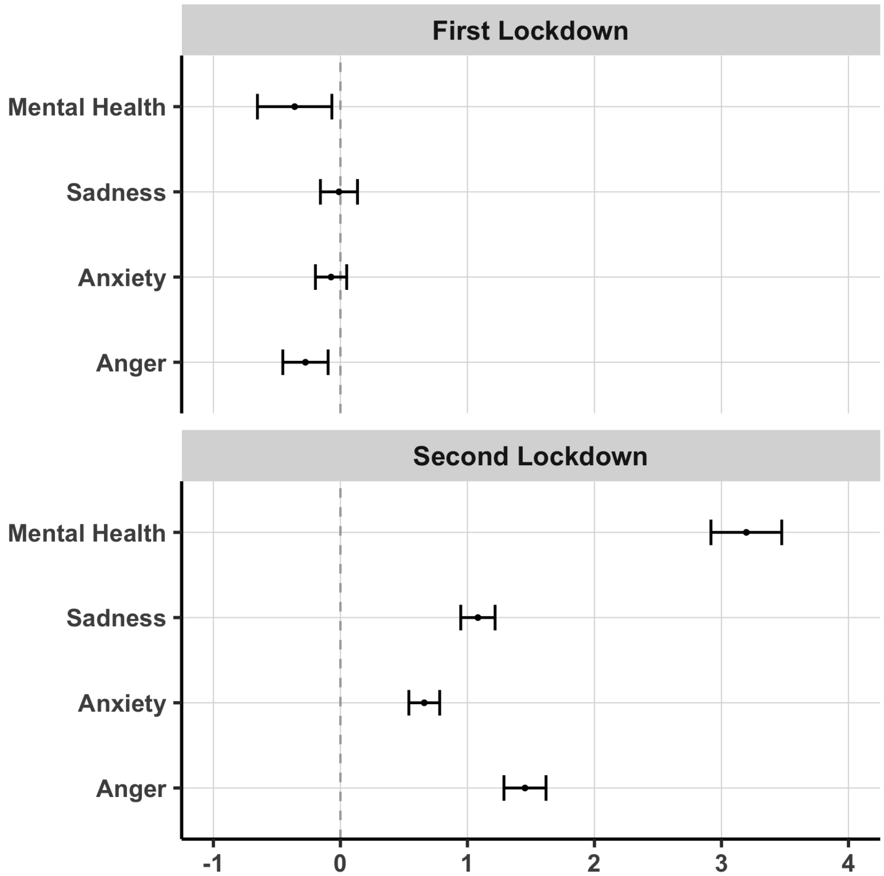

Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 95% confidence interval. The grey dotted line is the line of null effect. All estimates with clustered robust standard errors are reported in Table 4.

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

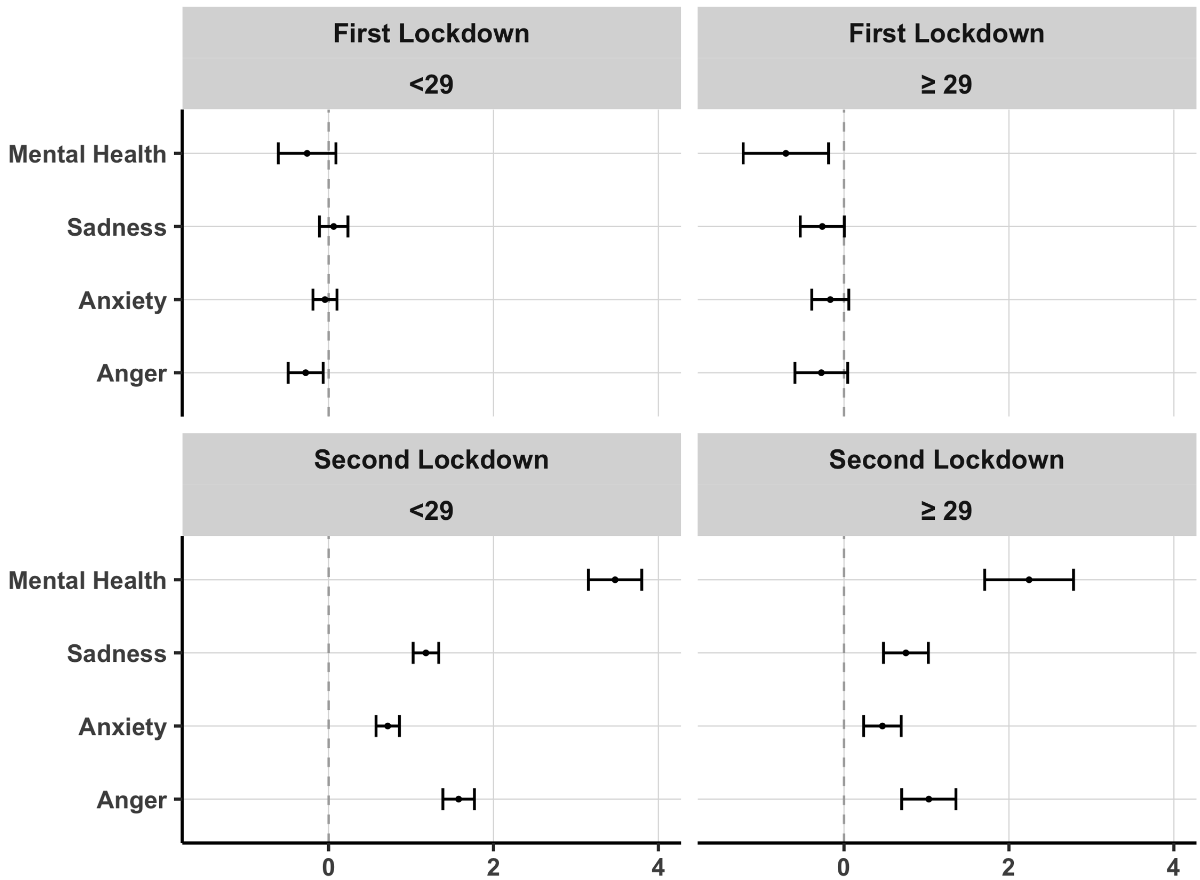

Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 95% confidence interval. The grey dotted line is the line of null effect. All estimates with associated clustered robust standard errors are reported in Table 4.

Fig. 6. Estimated Impact of Lockdowns by Gender

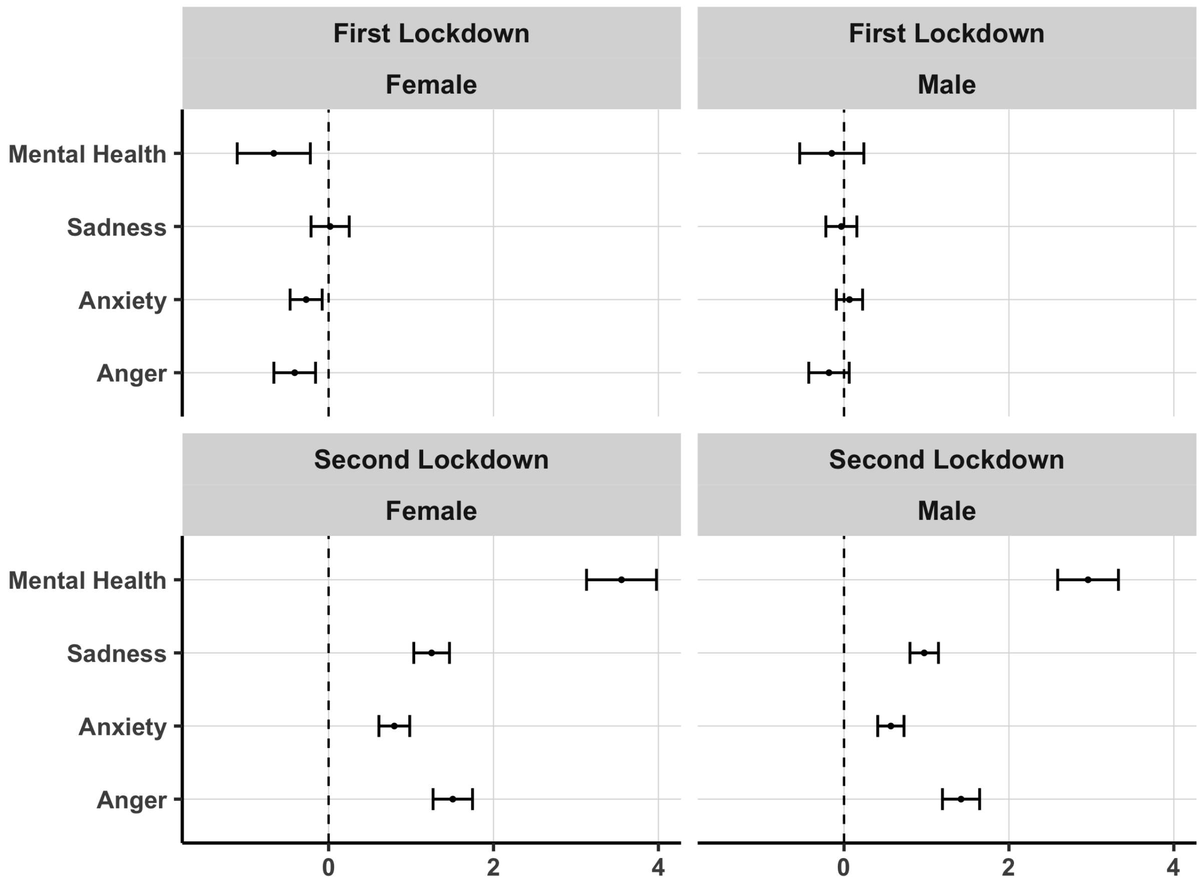

Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 95% confidence interval. The grey dotted line is the line of null effect. All estimates with clustered robust standard errors are reported in Table 5.

Fig. 7. Estimated Impact of Lockdowns by Users' Tweeting Intensity

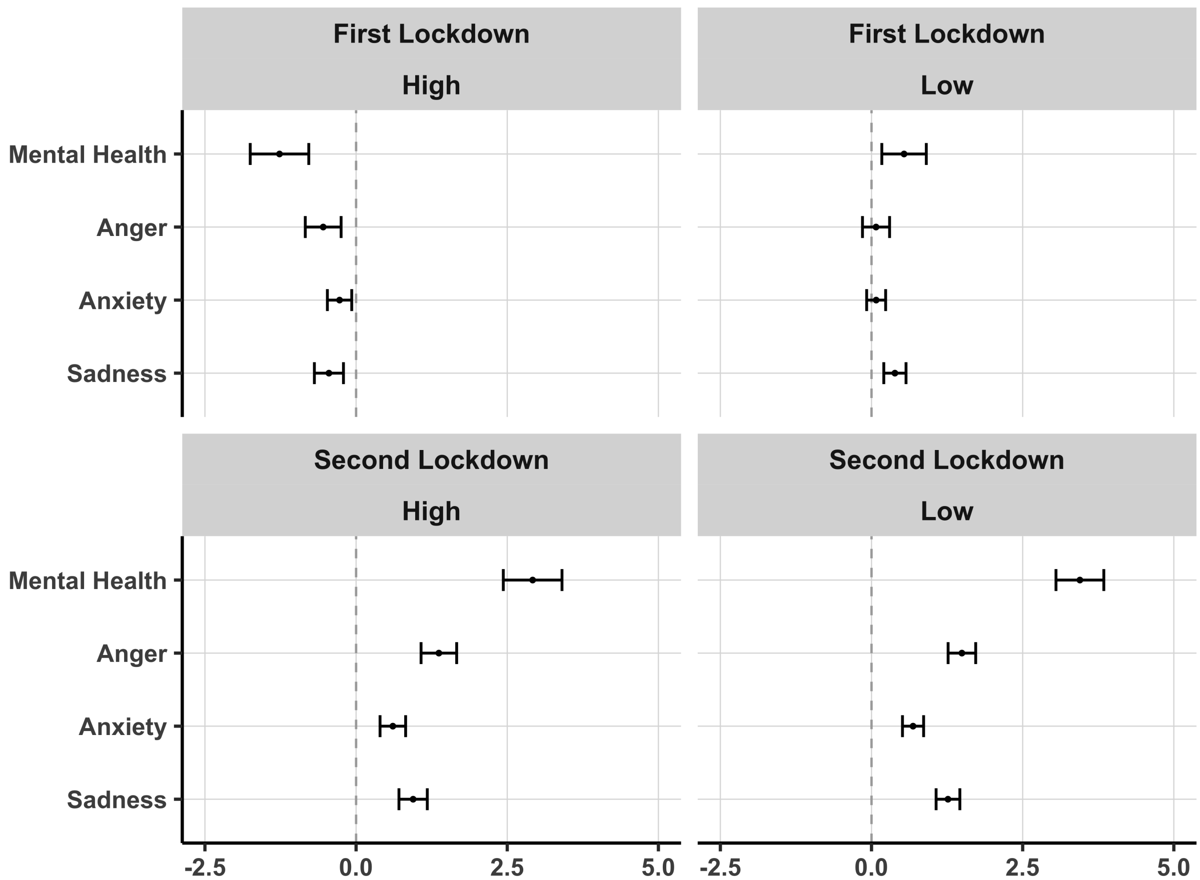

Note: DID estimates of the first (top panel) and second (bottom panel) lockdown on mental health, sadness, anxiety and anger with 95% confidence interval. The grey dotted line is the line of null effect. All estimates with clustered robust standard errors are reported in Table A.1.

Fig. 8. Change in Mental Health for All Users

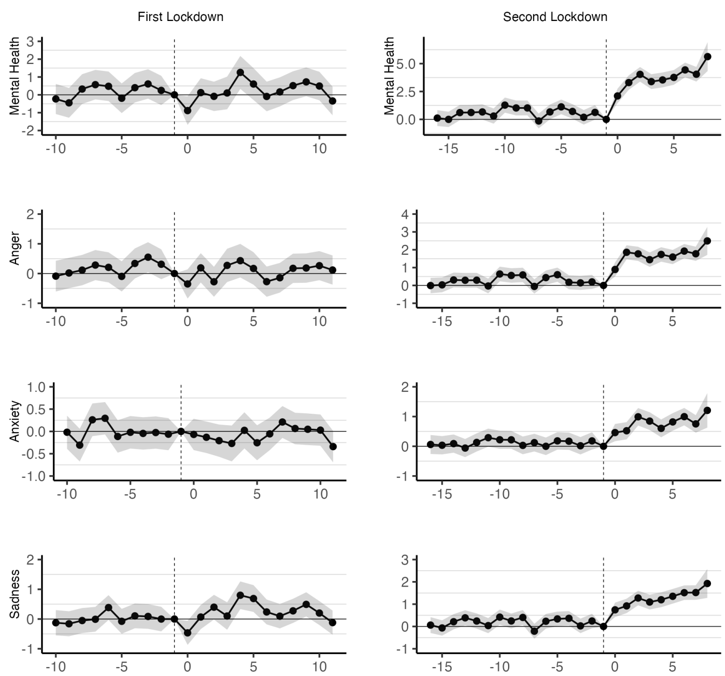

Note: Evolution of emotions weeks before and after the announcements of the first and the second lockdowns. This figure reports the estimated coefficients and associated 95% confidence interval from estimating equation (2).

Fig. 9. Change in Mental Health by Age Groups

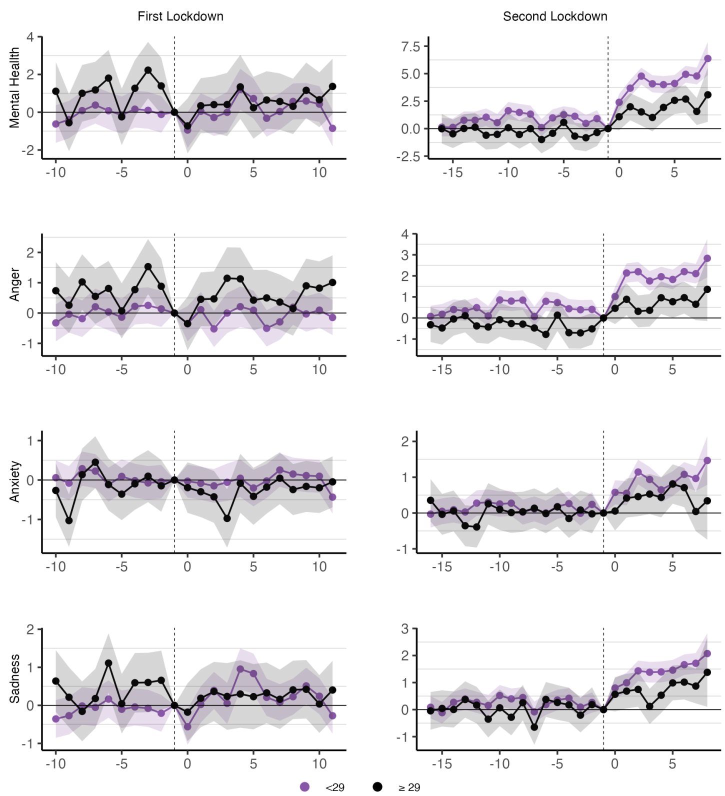

Note: Evolution of emotions weeks before and after the announcements of the first and the second lockdowns. This figure reports the estimated coefficients and associated 95% confidence interval from estimating equation (2).

Fig. 10. Change in Mental Health by Gender

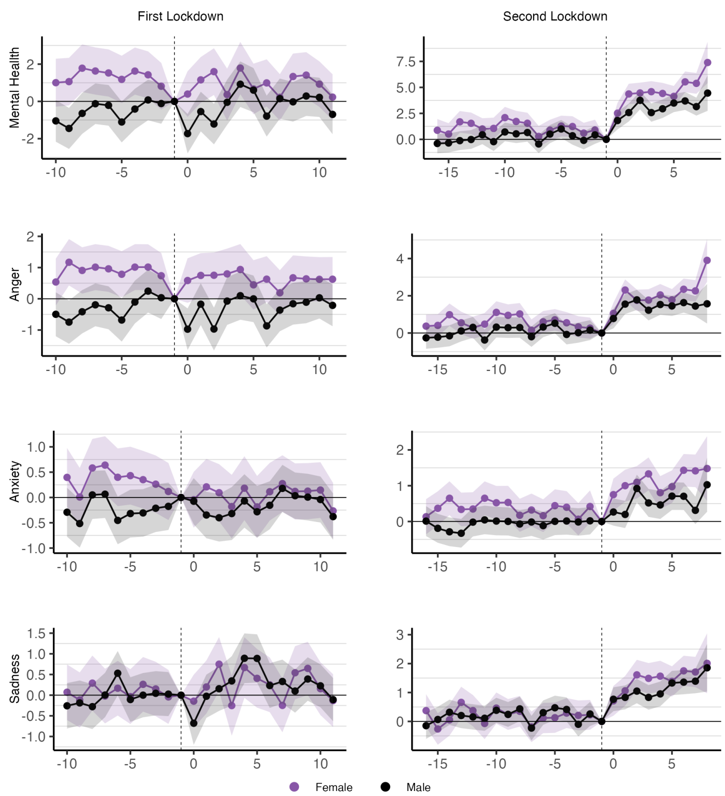

Note: Evolution of emotions weeks before and after the announcements of the first and the second lockdowns. This figure reports the estimated coefficients and associated 95% confidence interval from estimating equation (2).

Table 1: Descriptive Statistics of the Refined Dataset

Gender (%)	
Female	39.95
Male	60.05
Age Groups (%)	
<18	41.06
18-28	35.05
29-39	10.08
>39	13.8
Region (%)	
Auvergne-Rhone-Alpes	9.44
Bourgogne-Franche-Comté	3.09
Bretagne	3.63
Centre-Val-de-Loire	3.28
Corse	.58
Grand Est	9.47
Haut-de-France	11.92
Île-de-France	33.69
Normandie	3.33
Nouvelle-Aquitaine	5.81
Occitanie	6.64
Pays de La Loire	3.01
Provence-Alpes-Côte d'Azur	6.12
Total number of	
Tweets	52, 341, 001
Original Tweets	10, 245, 858
Replies to Tweets	13, 901, 582
Retweets	28, 193, 561
Unique Twitter users	39, 031
Daily Observations	10, 322, 655

Table 2: Descriptive Statistics of Sample 1 and Sample 2

	Sample 1		Sample 2	
	2019	**2020**	**2019**	**2020**
Gender (%)				
Female	38.09	39.43	38.94	39.82
Male	61.91	60.57	61.06	60.18
Age Groups (%)				
<18	38.63	40.87	40.35	41.24
18-28	34.92	34.93	34.9	34.95
29-39	11.18	10.14	10.49	10.04
>39	15.27	14.05	14.26	13.76
Region (%)				
Auvergne-Rhone-Alpes	9.2	9.22	9.27	9.36
Bourgogne-Franche-Comté	3.14	3.04	3.08	3
Bretagne	4.43	3.89	4.09	3.74
Centre-Val-de-Loire	3.54	3.45	3.49	3.38
Corse	.69	.63	.62	.59
Grand Est	9.11	9.11	9.06	9.17
Haut-de-France	12.7	12.59	12.86	12.32
Île-de-France	30.9	32.52	31.57	33.5
Normandie	3.89	3.6	3.74	3.45
Nouvelle-Aquitaine	5.89	5.75	5.82	5.65
Occitanie	7.21	6.87	7.02	6.76
Pays de La Loire	3.4	3.21	3.28	3.02
Provence-Alpes-Côte d'Azur	5.91	6.12	6.08	6.06
Nb. of Unique Users				
	20,874	30,219	27,561	35,944
	32,063		38,040	
Obs	1,244,807	2,238,057	1,986,141	2,725,502

Note: Sample 1 covers the period spanning from January 1 to June 2 in both 2019 and 2020. Sample 2 covers the period spanning from July 1 to December 15 in both 2019 and 2020.

Table 3: Examples of Tweets in Original Language

Emotion	Tweet	LIWC Dic.
Anger	- "@—— @—— 0 patience, le moindre truc me tape sur les nerfs, purée quel enfer"	tape, enfer
	- "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"	merde, merde merde
	- "Test #Covid19 Révélations sur un nouveau scandale sanitaire https://t.co/HPuamXSHyT https://t.co/9a1zq1o8eG"	scandale
	- "@—— Je ne suis pas spécialiste. Juste citoyen, en colère."	colère
	- "@—— C'est complètement dingue, un asile de fous ce pays."	fous
	- "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."	fout, colère, brutale crétins
	- "Au bout d'un an à nous faire chier avec les règles de distantiatikns social 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"	chier foutre cons
	- "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/DG51ToNl4J"	scandale, cons
	- "Encore un #1erministre qui n'a pas de couille. Qu'il le prenne leur foutue arrêter de #MasqueObligatoire dans l'espace publique partout comme ça il arrête de nous faire chier avec leur #Covid_19 de merde	foutue chier, merde
Anxiety	- "@—— 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"	inquiétant
	- "\"J'ai bien peur que ce soit notre dernier marché de la saison. Je crains vraiment un nouveau confinement\"#Menerbes #Vaucluse https://t.co/YZGl1kgSJk"	peur, crains
	- "Comment angoisser la planète? Écouter l'@OMS..."	angoisser
	- "Je pense qu'il y a que moi qui suis inquiet de sortir à nouveau https://t.co/jljg5WkJT9"	inquiet
	- "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"	inquiétants
	- "Crise d'angoisse bonjouuuuur"	angoisse
	- "Je ne sais pas vraiment si c'est de l'anxiété que j'ai mais c'est horrible."	anxiété, horrible
	- "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	folie pressent crainte
	- "@—— @—— 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)"	peur, peur peur
	- "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."	sceptique
Sadness	- "@—— Pathétique"	pathétique
	- "@—— Non mais oui abusé j'ai de ces moments de solitude parfois c'est violent"	violent
	- "Et vous, comment ça va mal ?"	mal
	- "Les pleurs des personnes âgées en EHPAD, enfermées et isolées jusqu'à l'absurde https://t.co/RdjEXvqGXn via @——"	pleurs, isolées
	- "Enfermée dans ma chambre, 4 heures du mat jsuis perdue"	perdue
	- "J'espère que je pourrais aussi être visité comme pour les EHPAD parce que je suis un vieux monsieur seul et isolé"	seul, isolé
	- "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."	pitoyable fatiguée, fatiguée mal
	- "@—— Je suis perdu"	perdu
	- "Vivement que je retrouve ma mère et mon frère parce que la solitude la c'est pesant"	pesant

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

	Mental Health	Anger	Anxiety	Sadness	*N*
Panel A: First Lockdown					
All Users	-0.360^{**}	-0.275^{***}	-0.074	-0.012	$3,482,864$
	(0.149)	(0.091)	(0.063)	(0.074)	
Users <29	-0.261	-0.278^{**}	-0.044	0.062	$2,673,535$
	(0.178)	(0.108)	(0.074)	(0.088)	
Users $\geq$ 29	-0.705^{***}	-0.275^{*}	-0.166	-0.264^{*}	$809,329$
	(0.264)	(0.163)	(0.115)	(0.136)	
Δ	0.525^{*}	-0.011	0.108	0.427^{***}	
	(0.309)	(0.188)	(0.133)	(0.158)	
Panel B: Second Lockdown					
All Users	3.196^{***}	1.453^{***}	0.660^{***}	1.083^{***}	$4,711,643$
	(0.142)	(0.085)	(0.062)	(0.069)	
Users <29	3.475^{***}	1.577^{***}	0.717^{***}	1.180^{***}	$3,651,338$
	(0.165)	(0.098)	(0.072)	(0.079)	
Users $\geq$ 29	2.245^{***}	1.029^{***}	0.466^{***}	0.750^{***}	$1,060,305$
	(0.275)	(0.168)	(0.116)	(0.139)	
Δ	1.230^{***}	0.549^{***}	0.251^{*}	0.430^{***}	
	(0.321)	(0.194)	(0.137)	(0.160)	

Note: This table reports Differences-in-Differences (DID) estimates for all users and by 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 the one-day lagged number of reported new deaths due to COVID-19, an individual and time fixed-effect (year, month and day of the week) 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 by Gender and Broad Age Groups

	Mental Health			Anger			Anxiety			Sadness		
	Female	Male	Δ	Female	Male	Δ	Female	Male	Δ	Female	Male	Δ
Panel A: First Lockdown												
All Users	-0.665***	-0.149	-0.516*	-0.411***	-0.182	-0.229	-0.272***	0.066	-0.339***	0.019	-0.033	0.051
	(0.226)	(0.198)	(0.301)	(0.129)	(0.125)	(0.180)	(0.099)	(0.081)	(0.128)	(0.118)	(0.096)	(0.152)
Users <29	-0.534**	-0.021	-0.513	-0.396***	-0.176	-0.220	-0.248**	0.135	-0.383***	0.110	0.021	0.090
	(0.248)	(0.254)	(0.355)	(0.142)	(0.160)	(0.214)	(0.107)	(0.103)	(0.149)	(0.130)	(0.120)	(0.177)
Users $\geq$ 29	-1.572***	-0.460	-1.112*	-0.537*	-0.201	-0.335	-0.427*	-0.094	-0.334	-0.608**	-0.165	-0.443
	(0.529)	(0.304)	(0.610)	(0.299)	(0.192)	(0.355)	(0.259)	(0.127)	(0.289)	(0.276)	(0.156)	(0.318)
Panel B: Second Lockdown												
All Users	3.553***	2.960***	0.593**	1.506***	1.419***	0.087	0.797***	0.568***	0.229*	1.249***	0.973***	0.277*
	(0.216)	(0.188)	(0.287)	(0.122)	(0.115)	(0.168)	(0.095)	(0.081)	(0.125)	(0.110)	(0.088)	(0.141)
Users <29	3.763***	3.237***	0.526	1.617***	1.547***	0.070	0.830***	0.622***	0.208	1.317***	1.068***	0.249
	(0.236)	(0.231)	(0.330)	(0.133)	(0.141)	(0.193)	(0.104)	(0.101)	(0.145)	(0.120)	(0.106)	(0.160)
Users $\geq$ 29	2.029***	2.306***	-0.277	0.706**	1.116***	-0.410	0.553**	0.441***	0.112	0.769***	0.749***	0.020
	(0.516)	(0.320)	(0.607)	(0.301)	(0.197)	(0.360)	(0.228)	(0.134)	(0.265)	(0.273)	(0.160)	(0.317)

Dependent variable: Mental Health, Anger, Anxiety, Sadness

Note: This table reports Differences-in-Differences (DID) estimates by 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 the one-day lagged number of reported new deaths due to COVID-19, an individual and time fixed-effect (year, month and day of the week) Robust standard errors are clustered at the individual level.$^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

Online appendix

A Additional tables

A.1 Results by users' tweeting frequency

Table A.1: DID Estimates by Users' Tweeting Intensity

	Mental Health			Anger			Anxiety			Sadness		
	Low	High	Δ	Low	High	Δ	Low	High	Δ	Low	High	Δ
Panel A: First Lockdown												
All Users	0.538***	-1.267***	-1.805***	0.075	-0.544***	-0.619***	0.077	-0.273***	-0.350***	0.387***	-0.449***	-0.836***
	(0.188)	(0.247)	(0.310)	(0.114)	(0.151)	(0.190)	(0.080)	(0.103)	(0.130)	(0.094)	(0.123)	(0.154)
Users <29	0.579***	-1.203***	-1.781***	0.059	-0.579***	-0.637***	0.135	-0.287**	-0.422***	0.385***	-0.337**	-0.722***
	(0.217)	(0.306)	(0.375)	(0.132)	(0.187)	(0.229)	(0.092)	(0.127)	(0.157)	(0.108)	(0.151)	(0.186)
Users $\geq$ 29	0.344	-1.529***	-1.873***	0.127	-0.489**	-0.616*	-0.152	-0.244	-0.092	0.370*	-0.796***	-1.166***
	(0.367)	(0.395)	(0.539)	(0.226)	(0.244)	(0.332)	(0.165)	(0.167)	(0.234)	(0.189)	(0.202)	(0.276)
Female users <29	0.280	-1.770***	-2.049***	-0.174	-0.661**	-0.487	0.062	-0.749***	-0.810***	0.392**	-0.360	-0.752***
	(0.290)	(0.453)	(0.537)	(0.167)	(0.259)	(0.308)	(0.123)	(0.200)	(0.234)	(0.153)	(0.235)	(0.281)
Male users <29	0.879***	-0.775*	-1.654***	0.283	-0.519**	-0.803**	0.212	0.072	-0.140	0.384**	-0.327*	-0.711***
	(0.322)	(0.415)	(0.525)	(0.202)	(0.264)	(0.333)	(0.135)	(0.164)	(0.213)	(0.152)	(0.197)	(0.249)
Female users $\geq$ 29	-0.895	-2.555***	-1.660	-0.122	-1.015**	-0.893	-0.536	-0.632*	-0.096	-0.237	-0.908**	-0.671
	(0.741)	(0.781)	(1.076)	(0.401)	(0.458)	(0.608)	(0.364)	(0.368)	(0.518)	(0.364)	(0.442)	(0.572)
Male users $\geq$ 29	0.753*	-1.270***	-2.023***	0.213	-0.356	-0.569	-0.026	-0.147	-0.121	0.565**	-0.768***	-1.333***
	(0.423)	(0.454)	(0.620)	(0.269)	(0.282)	(0.390)	(0.183)	(0.187)	(0.262)	(0.221)	(0.227)	(0.317)
Panel B: Second Lockdown												
All Users	3.446***	2.920***	-0.526*	1.495***	1.369***	-0.126	0.687***	0.608***	-0.079	1.265***	0.943***	-0.321**
	(0.202)	(0.248)	(0.319)	(0.116)	(0.150)	(0.190)	(0.089)	(0.108)	(0.140)	(0.100)	(0.120)	(0.156)
Users <29	3.657***	3.209***	-0.448	1.549***	1.491***	-0.058	0.786***	0.636***	-0.151	1.322***	1.083***	-0.239
	(0.229)	(0.296)	(0.374)	(0.132)	(0.178)	(0.222)	(0.102)	(0.130)	(0.165)	(0.111)	(0.142)	(0.181)
Users $\geq$ 29	2.609***	2.077***	-0.532	1.278***	1.011***	-0.267	0.298	0.527***	0.229	1.032***	0.539**	-0.494
	(0.423)	(0.442)	(0.612)	(0.245)	(0.275)	(0.369)	(0.183)	(0.186)	(0.261)	(0.228)	(0.217)	(0.315)
Female users <29	3.871***	3.754***	-0.117	1.510***	1.615***	0.105	0.828***	0.919***	0.091	1.532***	1.219***	-0.313
	(0.311)	(0.441)	(0.540)	(0.170)	(0.253)	(0.305)	(0.141)	(0.191)	(0.238)	(0.160)	(0.228)	(0.278)
Male users <29	3.473***	2.801***	-0.673	1.590***	1.398***	-0.192	0.748***	0.420**	-0.328	1.135***	0.982***	-0.153
	(0.333)	(0.398)	(0.519)	(0.199)	(0.249)	(0.318)	(0.145)	(0.177)	(0.229)	(0.156)	(0.181)	(0.238)
Female users $\geq$ 29	3.060***	0.727	-2.333**	1.511***	-0.168	-1.679**	0.781**	0.129	-0.652	0.768*	0.766*	-0.003
	(0.769)	(0.853)	(1.148)	(0.457)	(0.482)	(0.664)	(0.311)	(0.403)	(0.509)	(0.448)	(0.438)	(0.626)
Male users $\geq$ 29	2.478***	2.406***	-0.072	1.207***	1.295***	0.087	0.146	0.626***	0.480	1.124***	0.485*	-0.639*
	(0.501)	(0.509)	(0.714)	(0.288)	(0.321)	(0.432)	(0.220)	(0.209)	(0.303)	(0.265)	(0.248)	(0.363)

Dependent variable: Mental Health, Anger, Anxiety, Sadness.

Notes: This table reports Differences-in-Differences (DID) estimates for high- and low-frequency users across gender and broad age groups. DID estimates measure 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 the one-day lagged number of reported new deaths due to COVID-19, an individual and a time fixed-effect (year, month and day of the week) Robust standard errors are clustered at the individual level.$^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

A.2 Regression Discontinuity Results

A.2.1 RDD-DID

Table A.2: RDD and RDD-DID Estimates (Linear Model)

	Mental Health			Anger			Anxiety			Sadness		
	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID
Panel A: First Lockdown												
All	-0.004	-0.841***	-0.752	0.275	-0.162	-0.525**	0.101	-0.201**	-0.252	-0.380**	-0.477***	0.025
	(0.367)	(0.235)	(0.386)	(0.227)	(0.147)	(0.244)	(0.159)	(0.102)	(0.166)	(0.186)	(0.117)	(0.192)
Female	-0.514	-0.690**	0.377	-0.065	-0.089	0.000	-0.190	-0.255	0.082	-0.260	-0.347	0.296
	(0.576)	(0.351)	(0.593)	(0.334)	(0.209)	(0.353)	(0.267)	(0.155)	(0.274)	(0.291)	(0.179)	(0.301)
Male	0.331	-0.949***	-1.512***	0.499	-0.218	-0.879***	0.294	-0.162	-0.471**	-0.462	-0.569***	-0.162
	(0.477)	(0.315)	(0.507)	(0.306)	(0.202)	(0.332)	(0.197)	(0.135)	(0.209)	(0.242)	(0.155)	(0.250)
Users <29	-0.147	-0.914***	-0.786	0.131	-0.189	-0.517	-0.030	-0.178	-0.164	-0.248	-0.547***	-0.105
	(0.445)	(0.273)	(0.463)	(0.274)	(0.171)	(0.292)	(0.191)	(0.119)	(0.198)	(0.224)	(0.133)	(0.229)
Users $\geq$ 29	0.413	-0.592	-0.640	0.696	-0.074	-0.531	0.492	-0.286	-0.540	-0.775**	-0.232	0.430
	(0.614)	(0.455)	(0.665)	(0.389)	(0.277)	(0.426)	(0.280)	(0.193)	(0.296)	(0.322)	(0.246)	(0.346)
Panel B: Second Lockdown												
All	-0.156	2.387***	2.922***	-0.127	1.304***	1.491***	0.100	0.460***	0.556***	-0.129	0.623***	0.875***
	(0.254)	(0.207)	(0.289)	(0.159)	(0.126)	(0.177)	(0.110)	(0.097)	(0.129)	(0.122)	(0.099)	(0.140)
Female	-0.926**	2.353***	3.565***	-0.468**	1.094***	1.669***	-0.226	0.604***	0.874***	-0.231	0.655***	1.022***
	(0.394)	(0.308)	(0.444)	(0.230)	(0.178)	(0.258)	(0.178)	(0.150)	(0.203)	(0.199)	(0.156)	(0.226)
Male	0.368	2.414***	2.491***	0.104	1.447***	1.373***	0.323**	0.363***	0.342**	-0.060	0.603***	0.776***
	(0.333)	(0.277)	(0.381)	(0.216)	(0.173)	(0.240)	(0.141)	(0.126)	(0.167)	(0.154)	(0.129)	(0.179)
Users <29	-0.189	2.670***	3.373***	-0.220	1.442***	1.766***	0.143	0.533***	0.624***	-0.113	0.696***	0.983***
	(0.300)	(0.240)	(0.340)	(0.185)	(0.146)	(0.207)	(0.131)	(0.113)	(0.152)	(0.141)	(0.114)	(0.162)
Users $\geq$ 29	-0.030	1.407***	1.368***	0.195	0.824***	0.546	-0.042	0.207	0.318	-0.183	0.376	0.504
	(0.457)	(0.401)	(0.526)	(0.300)	(0.245)	(0.334)	(0.195)	(0.183)	(0.235)	(0.240)	(0.200)	(0.276)

Dependent variable: Mental Health, Anger, Anxiety, Sadness.

Note: This table reports estimates of the structural break in in anger, anxiety, sadness and mental health following implementation of the lockdown order. RDD19 and RDD20 report estimates from a standard parametric RDD model for both 2019 and 2020. All regressions control for the one-day lagged number of reported new deaths due to COVID-19, individual and time fixed-effects (month and day of the week for RDD19 and RDD20 and year, month, day of the week for RDD-DID). Robust standard errors are clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

Table A.3: RDD and RDD-DID Estimates (Quadratic Model)

	Mental Health			Anger			Anxiety			Sadness		
	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID
Panel A: First Lockdown												
All	0.516	-1.136***	-1.595***	0.348	-0.307	-0.742**	0.224	-0.279	-0.434	-0.056	-0.551***	-0.419
	(0.543)	(0.348)	(0.592)	(0.338)	(0.216)	(0.368)	(0.231)	(0.152)	(0.254)	(0.277)	(0.172)	(0.301)
Female	0.390	-1.751***	-1.617	0.529	-0.643**	-1.036	-0.186	-0.578**	-0.294	0.047	-0.530**	-0.287
	(0.861)	(0.528)	(0.919)	(0.486)	(0.317)	(0.530)	(0.403)	(0.228)	(0.425)	(0.433)	(0.260)	(0.468)
Male	0.596	-0.709	-1.562**	0.226	-0.078	-0.541	0.496	-0.068	-0.514	-0.126	-0.563**	-0.507
	(0.699)	(0.462)	(0.775)	(0.461)	(0.293)	(0.502)	(0.276)	(0.204)	(0.315)	(0.360)	(0.229)	(0.393)
Users <29	-0.016	-1.237***	-1.320	0.070	-0.292	-0.552	-0.021	-0.253	-0.287	-0.066	-0.692***	-0.481
	(0.658)	(0.409)	(0.711)	(0.403)	(0.255)	(0.439)	(0.279)	(0.178)	(0.305)	(0.335)	(0.198)	(0.358)
Users $\geq$ 29	2.115**	-0.788	-2.387**	1.180**	-0.363	-1.313**	0.961**	-0.370	-0.903**	-0.026	-0.055	-0.171
	(0.898)	(0.638)	(1.018)	(0.601)	(0.385)	(0.652)	(0.387)	(0.287)	(0.437)	(0.466)	(0.346)	(0.541)
Panel B: Second Lockdown												
All	0.093	2.246***	2.185***	0.055	1.003***	1.150***	0.049	0.488***	0.241	-0.011	0.756***	0.794***
	(0.375)	(0.297)	(0.443)	(0.237)	(0.182)	(0.277)	(0.162)	(0.139)	(0.199)	(0.176)	(0.145)	(0.210)
Female	-0.890	1.760***	2.727***	-0.395	0.859***	1.468***	-0.190	0.364	0.350	-0.304	0.538**	0.910***
	(0.594)	(0.439)	(0.676)	(0.350)	(0.251)	(0.393)	(0.267)	(0.213)	(0.318)	(0.291)	(0.232)	(0.341)
Male	0.763	2.572***	1.817***	0.360	1.099***	0.936**	0.214	0.570***	0.164	0.190	0.902***	0.717***
	(0.483)	(0.400)	(0.584)	(0.320)	(0.252)	(0.380)	(0.203)	(0.183)	(0.254)	(0.219)	(0.185)	(0.266)
Users <29	-0.029	2.553***	2.615***	-0.141	1.274***	1.599***	0.090	0.523***	0.219	0.022	0.755***	0.797***
	(0.444)	(0.346)	(0.521)	(0.278)	(0.211)	(0.322)	(0.191)	(0.163)	(0.235)	(0.205)	(0.167)	(0.243)
Users $\geq$ 29	0.543	1.189**	0.688	0.737	0.069	-0.401	-0.082	0.364	0.312	-0.112	0.756***	0.776
	(0.663)	(0.574)	(0.808)	(0.442)	(0.353)	(0.534)	(0.297)	(0.257)	(0.353)	(0.333)	(0.289)	(0.414)

Dependent variable: (spanning all columns)

Note: This table reports estimates of the structural break in in anger, anxiety, sadness and mental health following implementation of the lockdown order. RDD19 and RDD20 report estimates from a standard parametric RDD model for both 2019 and 2020. All regressions control for the one-day lagged number of reported new deaths due to COVID-19, individual and time-fixed effects (month and day of the week for RDD19 and RDD20 and year, month, day of the week for RDD-DID). Robust standard errors are clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

A.2.2 Donut-RDD-DID

A.2.3 RDD-DID

Table A.4: Donut-RDD and Donut-RDD-DID Estimates (Linear Model)

	Mental Health			Anger			Anxiety			Sadness		
	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID
Panel A: First Lockdown												
All	-0.546	-0.391	-0.031	-0.078	-0.069	-0.213	-0.019	-0.051	-0.061	-0.449**	-0.271**	0.244
	(0.372)	(0.247)	(0.417)	(0.231)	(0.153)	(0.263)	(0.160)	(0.106)	(0.179)	(0.186)	(0.125)	(0.206)
Female	-1.082	0.075	1.380**	-0.565	0.162	0.558	-0.287	-0.050	0.271	-0.230	-0.038	0.550
	(0.570)	(0.365)	(0.631)	(0.338)	(0.212)	(0.380)	(0.260)	(0.162)	(0.285)	(0.294)	(0.194)	(0.323)
Male	-0.194	-0.718**	-0.985	0.243	-0.231	-0.733**	0.160	-0.051	-0.282	-0.597**	-0.436***	0.029
	(0.491)	(0.333)	(0.555)	(0.311)	(0.214)	(0.359)	(0.204)	(0.140)	(0.230)	(0.240)	(0.163)	(0.267)
Users <29	-0.373	-0.398	-0.259	-0.090	-0.129	-0.335	-0.029	-0.005	-0.048	-0.254	-0.263	0.123
	(0.447)	(0.288)	(0.498)	(0.277)	(0.179)	(0.315)	(0.189)	(0.125)	(0.211)	(0.220)	(0.142)	(0.243)
Users $\geq$ 29	-1.071	-0.375	0.638	-0.055	0.138	0.177	0.011	-0.214	-0.128	-1.026***	-0.299	0.589
	(0.651)	(0.471)	(0.738)	(0.399)	(0.289)	(0.461)	(0.300)	(0.193)	(0.332)	(0.341)	(0.260)	(0.382)
Panel B: Second Lockdown												
All	-0.237	2.544***	2.978***	-0.107	1.425***	1.534***	0.086	0.486***	0.568***	-0.216	0.632***	0.876***
	(0.254)	(0.216)	(0.312)	(0.156)	(0.131)	(0.188)	(0.112)	(0.098)	(0.139)	(0.127)	(0.105)	(0.153)
Female	-1.247***	2.221***	3.626***	-0.614***	1.008***	1.678***	-0.249	0.660***	0.926***	-0.384	0.554***	1.021***
	(0.392)	(0.329)	(0.486)	(0.230)	(0.192)	(0.283)	(0.172)	(0.151)	(0.214)	(0.211)	(0.164)	(0.249)
Male	0.448	2.764***	2.543***	0.236	1.708***	1.438***	0.313**	0.370***	0.326	-0.102	0.686***	0.778***
	(0.333)	(0.285)	(0.407)	(0.210)	(0.176)	(0.250)	(0.148)	(0.129)	(0.183)	(0.158)	(0.137)	(0.194)
Users <29	-0.260	2.711***	3.300***	-0.127	1.455***	1.656***	0.092	0.574***	0.672***	-0.224	0.682***	0.972***
	(0.299)	(0.250)	(0.366)	(0.182)	(0.152)	(0.220)	(0.133)	(0.114)	(0.163)	(0.147)	(0.121)	(0.179)
Users $\geq$ 29	-0.152	1.957***	1.865***	-0.033	1.317***	1.111***	0.072	0.178	0.205	-0.191	0.462**	0.549
	(0.466)	(0.422)	(0.576)	(0.295)	(0.252)	(0.352)	(0.202)	(0.194)	(0.259)	(0.247)	(0.209)	(0.295)

Dependent variable: Mental Health, Anger, Anxiety, Sadness.

Notes: This table reports estimates of the structural break in in anger, anxiety, sadness and mental health following implementation of the lockdown order. of the lockdown order. RDD19 and RDD20 report estimates from a standard parametric RDD model for both 2019 and 2020. For the first lockdown, observations between covering the period between March 11 and March 17 was removed from the sample. For the second lockdown, observations covering the period between October 14 and October 22 was removed from the sample. All regressions control for the one-day lagged number of reported new deaths due to COVID-19, individual and time fixed-effects (month and day of the week for RDD19 and RDD20 and year, month, day of the week for RDD-DID). Robust standard errors are clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

Table A.5: Donut-RDD and Donut-RDD-DID Estimates (Linear Model)

	Mental Health			Anger			Anxiety			Sadness		
	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID	RD19	RD20	RDDID
Panel A: First Lockdown												
All	-0.546	-0.391	-0.031	-0.078	-0.069	-0.213	-0.019	-0.051	-0.061	-0.449**	-0.271**	0.244
	(0.372)	(0.247)	(0.417)	(0.231)	(0.153)	(0.263)	(0.160)	(0.106)	(0.179)	(0.186)	(0.125)	(0.206)
Female	-1.082	0.075	1.380**	-0.565	0.162	0.558	-0.287	-0.050	0.271	-0.230	-0.038	0.550
	(0.570)	(0.365)	(0.631)	(0.338)	(0.212)	(0.380)	(0.260)	(0.162)	(0.285)	(0.294)	(0.194)	(0.323)
Male	-0.194	-0.718**	-0.985	0.243	-0.231	-0.733**	0.160	-0.051	-0.282	-0.597**	-0.436***	0.029
	(0.491)	(0.333)	(0.555)	(0.311)	(0.214)	(0.359)	(0.204)	(0.140)	(0.230)	(0.240)	(0.163)	(0.267)
Users <29	-0.373	-0.398	-0.259	-0.090	-0.129	-0.335	-0.029	-0.005	-0.048	-0.254	-0.263	0.123
	(0.447)	(0.288)	(0.498)	(0.277)	(0.179)	(0.315)	(0.189)	(0.125)	(0.211)	(0.220)	(0.142)	(0.243)
Users $\geq$ 29	-1.071	-0.375	0.638	-0.055	0.138	0.177	0.011	-0.214	-0.128	-1.026***	-0.299	0.589
	(0.651)	(0.471)	(0.738)	(0.399)	(0.289)	(0.461)	(0.300)	(0.193)	(0.332)	(0.341)	(0.260)	(0.382)
Panel B: Second Lockdown												
All	0.148	2.895***	3.349***	-0.240	1.411***	1.528***	0.646***	0.697***	0.820***	-0.258	0.787***	1.002***
	(0.535)	(0.445)	(0.313)	(0.342)	(0.281)	(0.191)	(0.248)	(0.192)	(0.139)	(0.252)	(0.236)	(0.157)
Female	-0.560	2.968***	3.969***	-0.196	1.039***	1.571***	0.259	1.088***	1.103***	-0.622	0.840**	1.295***
	(0.797)	(0.651)	(0.477)	(0.526)	(0.393)	(0.280)	(0.364)	(0.311)	(0.214)	(0.390)	(0.334)	(0.250)
Male	0.622	2.855***	2.937***	-0.270	1.667***	1.503***	0.909***	0.432	0.628***	-0.017	0.755**	0.806***
	(0.715)	(0.602)	(0.415)	(0.450)	(0.389)	(0.258)	(0.334)	(0.244)	(0.182)	(0.330)	(0.325)	(0.201)
Users <29	0.387	3.207***	3.907***	-0.075	1.480***	1.773***	0.662**	0.758***	0.937***	-0.200	0.969***	1.197***
	(0.632)	(0.512)	(0.366)	(0.412)	(0.311)	(0.223)	(0.290)	(0.222)	(0.162)	(0.297)	(0.281)	(0.181)
Users $\geq$ 29	-0.642	1.823**	1.434**	-0.784	1.169	0.688	0.593	0.487	0.410	-0.451	0.167	0.336
	(0.964)	(0.894)	(0.589)	(0.561)	(0.640)	(0.360)	(0.465)	(0.383)	(0.262)	(0.460)	(0.403)	(0.307)

Dependent variable: Mental Health, Anger, Anxiety, Sadness.

Note: This table reports estimates of the structural break in in anger, anxiety, sadness and mental health following implementation of the lockdown order. of the lockdown order. RDD19 and RDD20 report estimates from a standard parametric RDD model for both 2019 and 2020. For the first lockdown, observations between covering the period between March 11 and March 17 was removed from the sample. For the second lockdown, observations covering the period between October 22 and October 30 was removed from the sample. All regressions control for the one-day lagged number of reported new deaths due to COVID-19, individual and time fixed-effects (month and day of the week for RDD19 and RDD20 and year, month, day of the week for RDD-DID). Robust standard errors are clustered at the individual level. $^{***}p < 0.01$, $^{**}p < 0.05$ and $^{*}p < 0.1$. Standard errors in parentheses.

Table A.6: Correlations Between our Indicators and CoviPrev - Anxiety

	Mental Health	Anxiety	Sadness	Anger
Mental Health				
Anxiety	0.97***			
Sadness	0.97***	0.91***		
Anger	0.99***	0.94***	0.94***	
CoviPrev - Anxiety	0.67***	0.61***	0.66***	0.69***

Note: This table reports the Pearson correlation coefficients. The CovidPrev data span from March 30 to December 16 2020. $^{***}\ p < 0.01,\ ^{**}\ p < 0.05,\ ^{*}\ p < 0.10$

B Additional figures

Fig. B.1. Regression Discontinuity Plots: First Lockdown Period for 2019 and 2020

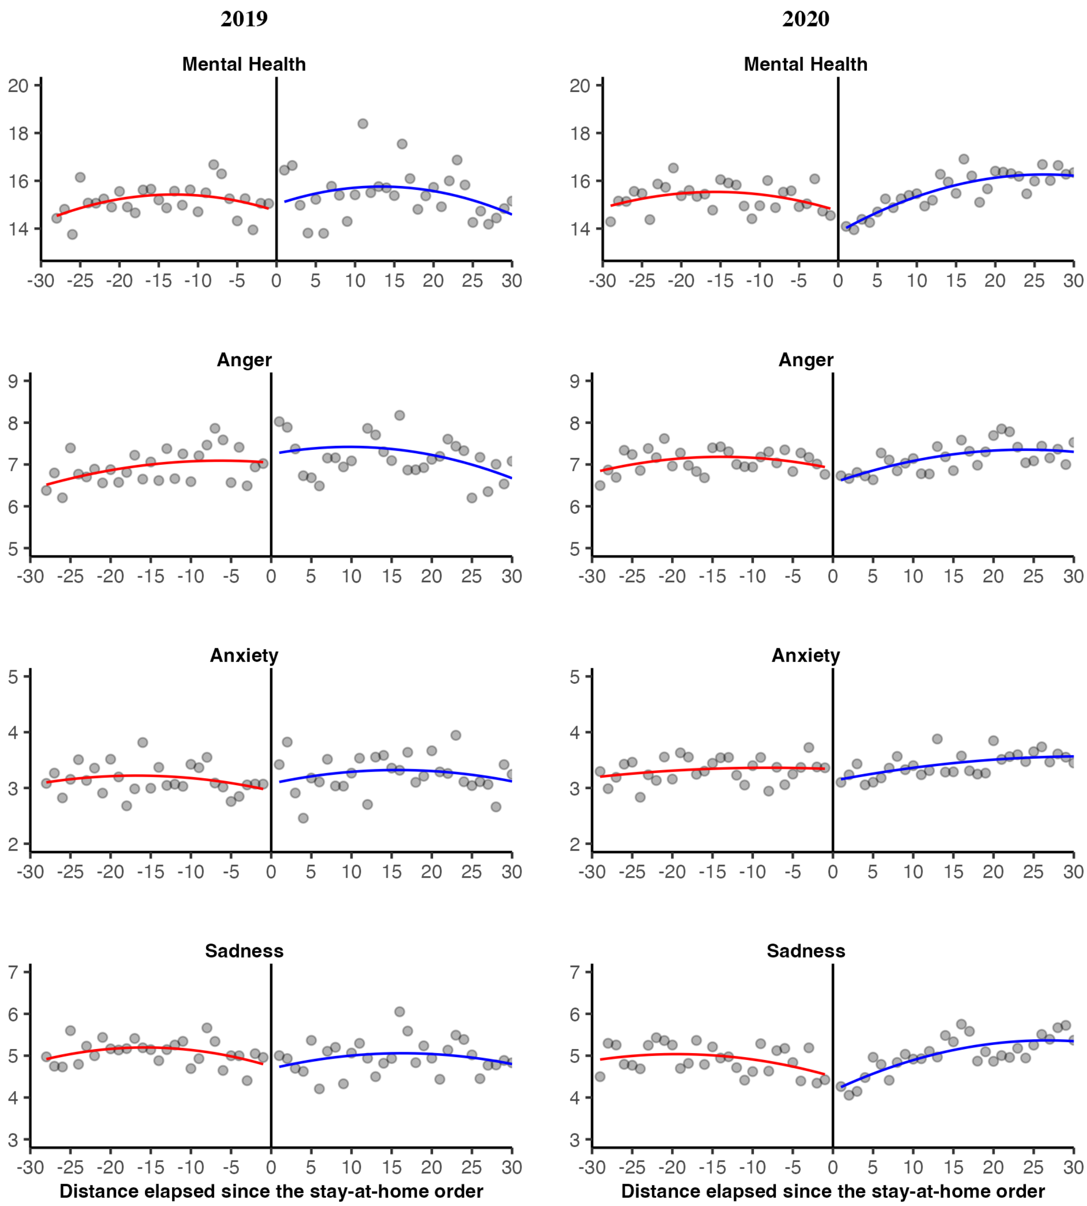

Note: RDD plots of measure emotional health before and after lockdown annoucements during the comparison year in 2019 (left) and the pandemic year (right).The dots provides the average daily emotion across all users by bins of one day. The lines show the fitted values of the indices (mental health, anger, anxiety, sadness) fitted with a second-order polynomial, before the annoucement in red and after the annoucement in blue.

Fig. B.2. Regression Discontinuity Plots: Second Lockdown Period for 2019 and 2020

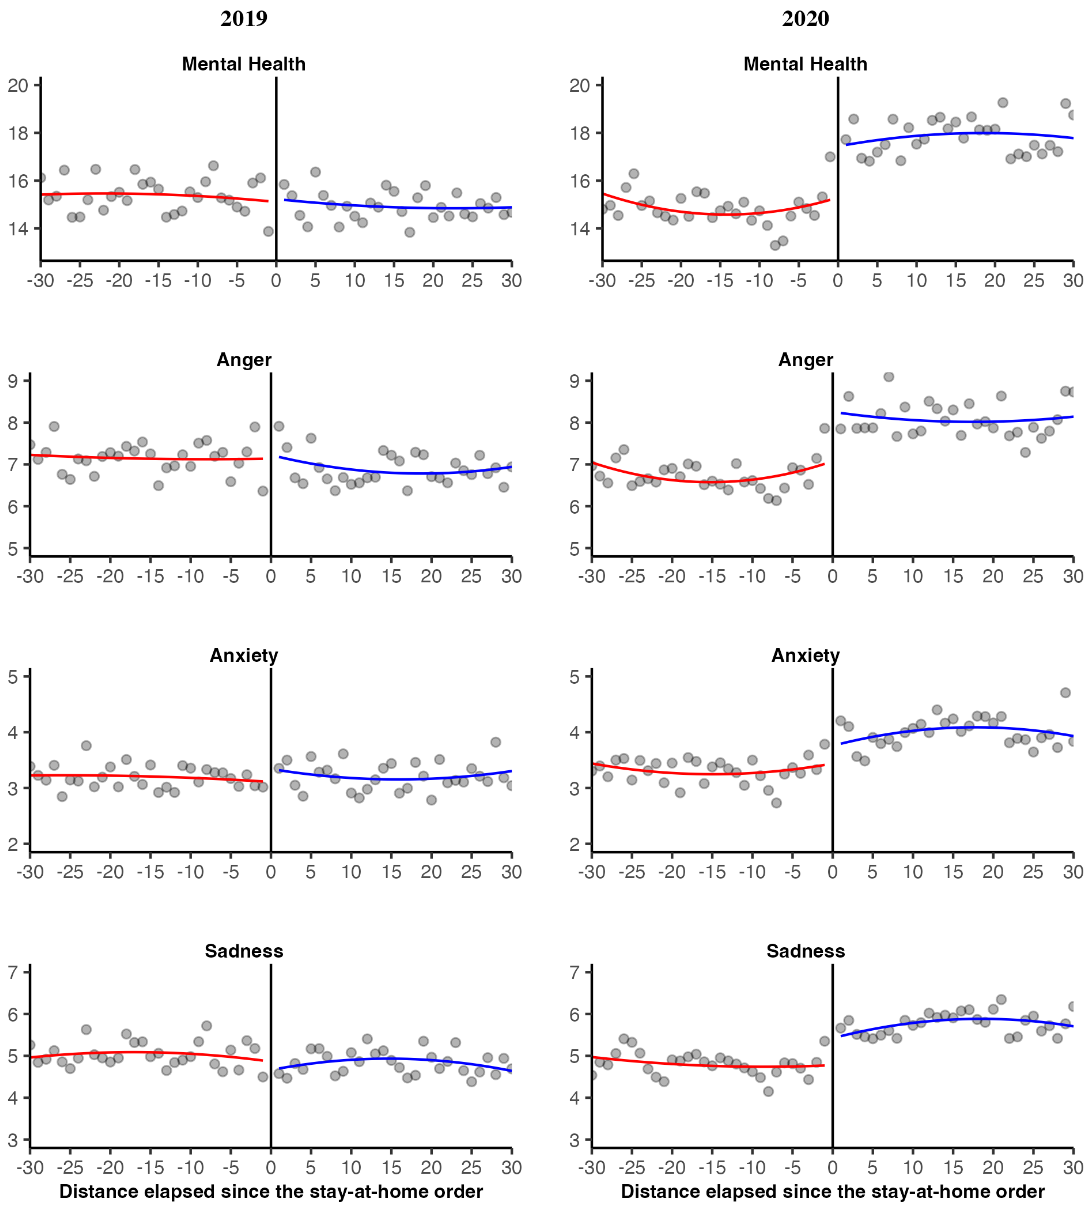

Note: RDD plots of measure emotional health before and after lockdown annoucements during the comparison year in 2019 (left) and the pandemic year (right).The dots provides the average daily emotion across all users by bins of one day. The lines show the fitted values of the indices (mental health, anger, anxiety, sadness) fitted with a second-order polynomial, before the annoucement in red and after the annoucement in blue.

Fig. B.3. Donut Regression Discontinuity Plots: Second Lockdown Period for 2019 and 2020

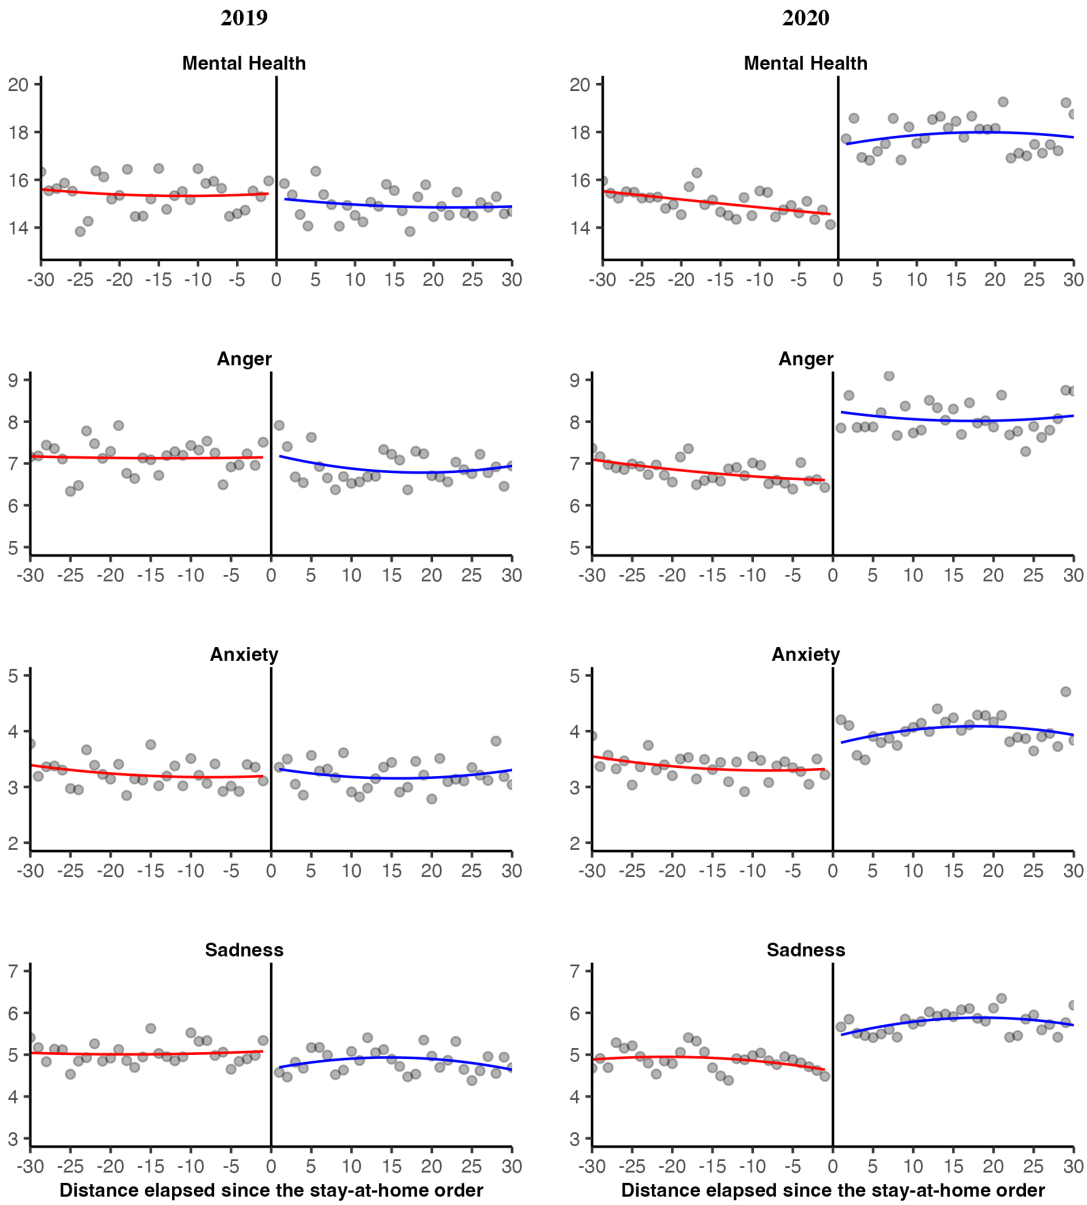

Note: Donut RDD plots of measure emotional health before and after lockdown annoucements during the comparison year in 2019 (left) and the pandemic year (right). Removed observations between October 14 and October 22. The dots provides the average daily emotion across all users by bins of one day. The lines show the fitted values of the indices (mental health, anger, anxiety, sadness) fitted with a second-order polynomial, before the annoucement in red and after the annoucement in blue.

Fig. B.4. Donut Regression Discontinuity Plots: Second Lockdown Period for 2019 and 2020

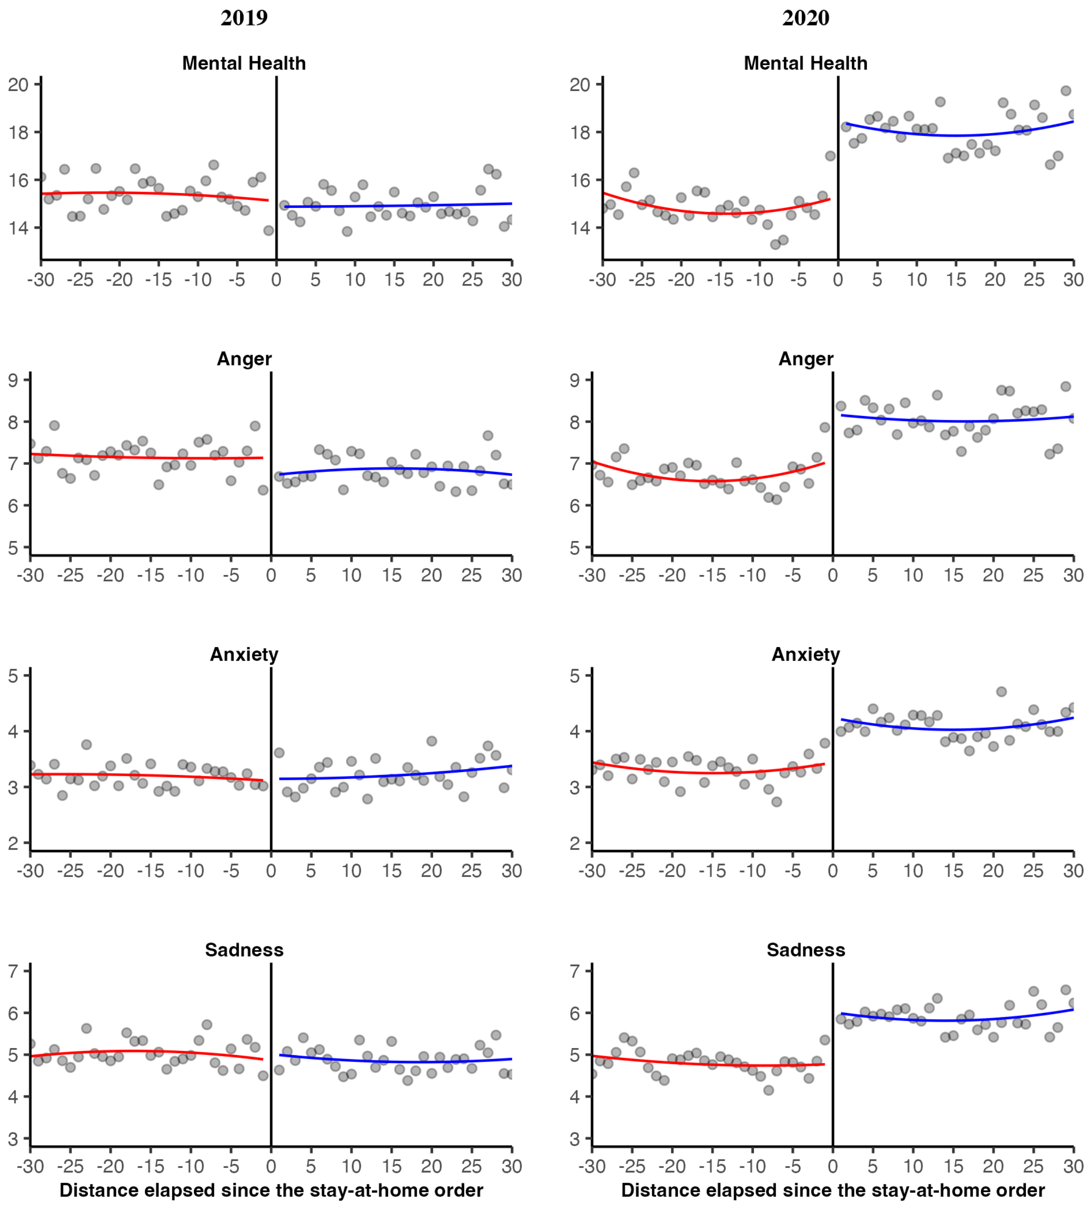

Note: Donut-RDD plots of measure emotional health before and after lockdown annoucements during the comparison year in 2019 (left) and the pandemic year (right). Removed observations between October 22 and October 30. The dots provides the average daily emotion across all users by bins of one day. The lines show the fitted values of the indices (mental health, anger, anxiety, sadness) fitted with a second-order polynomial, before the annoucement in red and after the annoucement in blue.