Breaking Down the Lockdown: The Impact of Stay-At-Home Mandates on Uncertainty and Sentiments

Carolina Biliotti, Falco J. Bargagli-Stoffi, Nicolò Fraccaroli, Michelangelo Puliga, Massimo Riccaboni

Submitted to: Journal of Medical Internet Research
on: July 23, 2024

Table of Contents

Original Manuscript 5
Supplementary Files 29
Figures 30
Figure 1 31
Figure 2 32
Figure 3 33
Multimedia Appendixes 34
Multimedia Appendix 1 35

Breaking Down the Lockdown: The Impact of Stay-At-Home Mandates on Uncertainty and Sentiments

Carolina Biliotti[1]; Falco J. Bargagli-Stoffi[2]; Nicolò Fraccaroli[3]; Michelangelo Puliga[4]; Massimo Riccaboni[1]

[1]IMT School for Advanced Studies Lucca IT
[2]Harvard University Cambridge US
[3]Brown University, W.R. Rhodes Center for International Economics and Finance Providence US
[4]Bloomag, Inc. San Francisco US

Corresponding Author:
Nicolò Fraccaroli
Brown University, W.R. Rhodes Center for International Economics and Finance
111 Thayer Street, Box 1970
Providence
US

Abstract

Background: Since the spread of the SARS-CoV-2 virus and the lockdown measures went hand-in-hand, it is difficult to distinguish how public opinion reacted to the lockdown measures from the reactions to COVID-19.

Objective: We analyze the causal effect of COVID-19 lockdown policies on sentiment and uncertainty using the Italian lockdown in February 2020 as a quasi-experiment. Communities inside and just outside the lockdown area were equally confronted with COVID-19 at the time of the implementation of the policy, offering a form of random allocation of the lockdown. The two areas had also balanced socioeconomic and demographic characteristics before the lockdown, indicating that the definition of the boundaries of the area under strict lockdown approximates a randomized experiment. This allows to identify the causal impact of lockdowns on public emotions, disentangling the changed due to the policy itself, from the changes induced by the spread of the novel virus.

Methods: We employ Twitter data, natural language models (N = 24,261), and a difference-in-differences approach to compare sentiment changes within (n=1,567) and outside (n=22,694) the lockdown areas before and after the beginning of the lockdown. Tweets are classified into four categories—economics, health, politics, and lockdown policy—to analyze the corresponding emotional responses.

Results: We find that the lockdown had no significant effect on economic uncertainty (b=0.005, SE=0.007, t(125)=0.70, P =.48) or economic negative sentiment (b=-0.011, SE=0.0089, t(125)=-1.32, P =.19), but increased uncertainty about health (b=0.036, SE=0.0065, t(125)=5.55, P<.001) and the lockdown policy (b=0.026, SE=0.006, t(125)=4.47, P<.001) and negative sentiment towards politics (b=0.025, SE=0.011, t(125)=2.33, P =.02), suggesting that lockdowns have wide externalities beyond health.

Conclusions: Our results emphasize the need for authorities to use these findings to improve future policies and communication efforts to mitigate uncertainty and social panic.

(JMIR Preprints 23/07/2024:64667)
DOI: https://doi.org/10.2196/preprints.64667

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?
✓ **Please make my preprint PDF available to anyone at any time (recommended).**
Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.
Only make the preprint title and abstract visible.
No, I do not wish to publish my submitted manuscript as a preprint.
2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?
✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain v
Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in <a href="http

Original Manuscript

Original Paper

Breaking Down the Lockdown: The Impact of Stay-At-Home Mandates on Uncertainty and Sentiments

Carolina Biliotti (IMT School for Advanced Studies, Lucca, Italy),
Falco J. Bargagli-Stoffi (Harvard University, Cambridge MA, USA),
Corresponding author: Nicolò Fraccaroli (Brown University, W.R. Rhodes Center for International Economics and Finance 111 Thayer Street, Box 1970, Providence, RI, USA 02912-1970 E-mail: nicolo_fraccaroli@brown.edu),
Michelangelo Puliga (Lead Data Scientist at Bloomag, Inc., San Francisco, California, US),
Massimo Riccaboni (IMT School for Advanced Studies, Lucca, Italy).

Abstract

Background: Since the spread of the SARS-CoV-2 virus and the lockdown measures went hand-in-hand, it is difficult to distinguish how public opinion reacted to the lockdown measures from the reactions to COVID-19.
Objective: We analyze the causal effect of COVID-19 lockdown policies on sentiment and uncertainty using the Italian lockdown in February 2020 as a quasi-experiment. Communities inside and just outside the lockdown area were equally confronted with COVID-19 at the time of the implementation of the policy, offering a form of random allocation of the lockdown. The two areas had also balanced socioeconomic and demographic characteristics before the lockdown, indicating that the definition of the boundaries of the area under strict lockdown approximates a randomized experiment. This allows to identify the causal impact of lockdowns on public emotions, disentangling the changed due to the policy itself, from the changes induced by the spread of the novel virus.
Methods: We employ Twitter data, natural language models (N = 24,261), and a difference-in-differences approach to compare sentiment changes within (n=1,567) and outside (n=22,694) the lockdown areas before and after the beginning of the lockdown. Tweets are classified into four categories—economics, health, politics, and lockdown policy—to analyze the corresponding emotional responses.
Results: We find that the lockdown had no significant effect on economic uncertainty (b=0.005, SE=0.007, t(125)=0.70, P =.48) or economic negative sentiment (b=-0.011, SE=0.0089, t(125)=-1.32, P =.19), but increased uncertainty about health (b=0.036, SE=0.0065, t(125)=5.55, P<.001) and the lockdown policy (b=0.026, SE=0.006, t(125)=4.47, P<.001) and negative sentiment towards politics (b=0.025, SE=0.011, t(125)=2.33, P =.02), suggesting that lockdowns have wide externalities beyond health.
Conclusions: Our results emphasize the need for authorities to use these findings to improve future policies and communication efforts to mitigate uncertainty and social panic.

Keywords: Lockdown Policy; Sentiment Analysis; Uncertainty; Social Media; Quasi-experiment

Introduction

Background

The spread of the SARS-CoV-2 virus and the associated COVID-19 disease was accompanied by an extraordinary rise in uncertainty and negative emotions, resulting in significant economic and social costs [1–4]. Against this backdrop, most governments had implemented lockdown measures by April 2020 to contain the spread of the virus among their populations. Lockdown policies, or stay-at-home mandates, are temporary restrictions that prevent residents from leaving their homes, except for indispensable tasks or work in essential businesses. While these policies are considered effective in slowing down the spread of COVID-19, as backed by recent evidence [5–7], their implications on public sentiments remain unclear. Lockdown policies might worsen public uncertainty and negative emotions by intensifying health concerns, influencing the expectations of negative consequences for economic and social activities, and triggering political backlash. On the other hand, they could signal the commitment to control the virus, reduce information asymmetry, and improve public sentiment.

Goal of the paper

This paper investigates whether lockdown measures exacerbate or alleviate uncertainty and negative sentiment. We consider multiple important dimensions of the public concern by proposing a nuanced measure of uncertainty and sentiments to assess the heterogeneous effect of lockdowns on different topics of the public's interest, such as economics, health and politics. It is crucial to understand whether the economic and social costs of lockdowns outweigh their health benefits.

Prior Work

However, it is challenging to determine the causal effect of lockdowns on public emotions due to issues of simultaneity and endogeneity. Because authorities in most countries introduced lockdown measures immediately after the virus was detected, the COVID-19- induced deterioration of public sentiment and the implementation of lockdown measures are concomitant. Endogeneity is also an issue, as lockdown measures were put in place when COVID-19 cases were discovered. These areas are likely the same areas where people perceive a higher risk of pandemic-related costs and where sentiment is therefore lowest. Given these limitations, most studies on changes in sentiments during COVID- 19 lockdowns provide associational evidence [8–12] that may over- or underestimate the impact of lockdowns on public emotions, leading to inaccurate conclusions.

Previous studies reported on the rising uncertainty and negativity during COVID-19 but did not break down how much of this could be attributed to restrictive measures [13,14].
Moreover, many focused on shifts in aggregate uncertainty and sentiments during COVID-19 [15-17], or on single dimensions, such as economic uncertainty [8,13], political polarization [18,19], or health-related negative emotions [17, 20], without jointly accounting for changes across multiple relevant dimensions of the public discourse at the time.
The distinction of the effect of the measures across multiple topics, such as politics, economics or health, is valuable, as the effects of the lockdown may be heterogeneous across different dimensions.

Our Study

To address this, we focus on the Italian lockdown of February 2020, which was the first in a high-income economy. This case allows for quasi-experimental analysis since the lockdown was initially imposed on communities where the first COVID-19 cases were detected, offering a form of random assignment.

The transmission of the virus was similar inside and outside the lockdown area at the time the policy was implemented [21], and health and socioeconomic conditions were balanced, allowing us to apply causal inference methods to assess the impact of lockdowns on public opinion. Specifically, the timing of these lockdowns, before the availability of vaccines and with little knowledge about the SARS-CoV-2 virus, offers a distinct opportunity to isolate and measure the direct influence of lockdown measures on public emotions, such as sentiments and uncertainty.

We use Twitter data to measure uncertainty and sentiment by applying deep learning and natural language processing techniques. For uncertainty, we want to distinguish tweets asking questions, expressing hesitation, irresolution, confusion, anxiety, tweets searching for clarity on topics, from tweets that instead express assurance and confidence. By measuring sentiment, we want to capture other emotions such as anger, disillusionment, and disapproval, which we broadly define as *negative sentiments*. We distinguish between tweets from users inside the lockdown area (the so called 'red zone') and those outside, including neighboring cities (referred to as the 'orange zone').

This data collection is done before and after the lockdown. To measure uncertainty and negative sentiments in each tweet, we fine-tune the AlBERTo model [22], a natural language processing model based on the Bidirectional Encoder Representations from Transformers (BERT) model [23] customized for Italian text data. BERT is able to efficiently contextualize words, offering significant advantages relative to existing strategies that measure emotions based on the simple bag-of-words approaches [24,25]. BERT has a strong track record of achieving top results in various natural language processing tasks, including sentiment analysis, fake news detection, and analyzing public opinions on Twitter during the COVID-19 pandemic [15, 26–29]. Given all the existing evidence on the accuracy of text analysis with a BERT-based model, we believe that our measures of sentiment and uncertainty constructed via AlBERTo are representative of the emotions expressed in the tweets.

To better understand the impact of the lockdown on public emotions, we break down uncertainty and sentiments into four main dimensions: economy, health, politics, and lockdown policy. Identifying health and economic-related tweets is crucial in order to understand how lockdown policies differently affected the perception of health and economic risks and investigate the health-economic trade-off [30]. The political dimension is also considered, as lockdowns could be affecting sentiments toward politics, influencing attitudes towards incumbent politicians [31-36], or leading to increased political polarization [18,19]. The fourth dimension (which we refer to as *policy* for simplicity) allows us to capture the noise generated by uncertainty and negative feelings associated with the behavioral guidelines of the restrictions. This helps distinguish whether uncertainty and negative sentiments stem from the authorities implementing the policy (political uncertainty) or from uncertainties about the policy's details (uncertainty around lockdown policies). Overall, this categorization helps to capture changes in public sentiment in response to the COVID-19 pandemic and the government's actions.
We use manually created dictionaries for each dimension to categorize the tweets and create a matrix of document features to identify words associated with each dimension in the tweets.

Using a Difference-in-Differences (DiD) specification, we estimate the causal impact of the lockdown on economic, health, political, and policy uncertainty and negative sentiment. Causal identification is achieved as the *treatment* assignment—the lockdown policy—came as an exogenous shock to the public's uncertainty and negative sentiment, independent of potential outcomes. Our results are robust to a placebo test and a battery of robustness tests.

Our findings reveal that, in the Italian context, the lockdown had no significant effect on public concerns about the economy, implying that lockdown-induced negative emotions did not impact economic worries. Economic concerns did not increase for those subjected to the lockdown, which is relevant in the ongoing debate about the economic effects of lockdowns. Stay-at-home mandates indeed have economic effects [37], but they are derived from other, more direct channels such as activity closures.

By contrast, the stay-at-home mandate did affect the health- and politics-related emotions of individuals within the lockdown area, where users expressed higher uncertainty when discussing health and the policy, and more negative sentiments about politics than those outside the lockdown area. The lockdown exacerbated concerns about health- related risks, rather than reassuring about the authorities' commitment to fight the virus. This increased uncertainty may have led to greater awareness and compliance with measures to contain the virus, which could have a positive impact [38].

The increase of negative sentiments towards politics represents a clear cost that policymakers must take into account. These political costs could deter elected officials from implementing lockdowns in future pandemics, regardless of their effectiveness in controlling the spread of viruses. The increased discontent with politics could reflect the increase in political polarization, which was on the rise during COVID-19 [18,19], but it could also be a sign of a change in attitude towards
the elected politician in office, for which there is mixed evidence [31-33,39]. It is not clear to what extent the changes in attitudes to politics reported in the literature can be directly attributed to lockdown measures. We show that once the effect of the lockdown is causally identified, the policy worsens negative sentiments toward politics, possibly leading to a delay, a restriction to narrower regions or a weakening of policy responses when lockdown measures are necessary.

Methods

Identification Strategy

Following the discovery of the first transmission case of COVID-19 in the small town of Codogno, the Italian government announced on February 22 a decree imposing strict quarantine measures in the town of Codogno, as well as in nine neighboring municipalities in the province of Lodi (Castiglione d'Adda, Casalpusterlengo, Fombio, Maleo, Somaglia, Bertonico, Terranova dei Passerini, Castelgerundo e San Fiorano), starting the next day (February 23, 2020) [40]. The lockdown was enforced with police on the streets, without the possibility of entering or leaving the lockdown area.

On March 8, an "orange zone" (*zona arancione*) was established to cover the municipalities near the red zone. Lombardy and 14 other cities outside that region (Modena, Parma, Piacenza, Reggio nell'Emilia, Rimini, Pesaro e Urbino, Alessandria, Asti, No- vara, Verbano-Cusio-Ossola, Vercelli, Padua, Treviso, Venezia), accounting for 16 million people, were placed under *partial* lockdown restrictions [41]. In the orange zone, people were *invited* to avoid transit in and out of the municipality of residence, but economic activities were allowed to continue. The measure of March 8 will be extended nationwide on March 9, ending the strict *red zone* in the province of Lodi [42]. Orders implementing the nationwide full lockdown were announced on March 22.

We exploit the exogenous shock of the unexpected lockdown measures enacted in Italy on February 23, 2020, to assess the causal impact of lockdown restrictions on public emotions. We argue that the detection of the first European COVID-19 case in the municipality of Codogno, and not in other nearby municipalities -- see Figure 1 -- was purely coincidental. When physicians learned that the patient, who showed symptoms attributable to COVID-19, had contact with a friend who had recently been on a business trip to China, they decided to force the existing protocol of testing only Italians and foreigners who had been in China [43]. As a consequence of the first positive case, hospitals began testing people for COVID- 19 in the area of Codogno. This decision was made simply by identifying the first COVID-19 patient in Codogno.

Recent retrospective epidemiological studies on the transmission risk of coronavirus in Lombardy in February 2020 revealed a homogeneous transmission potential of the virus in the different provinces of Lombardy at the time of the discovery of the first case in Codogno [21]. This means that COVID-19 incidence between the red zone and the neighboring municipalities in the orange zone was balanced at the time of the first lockdown. The red zone was locked down, while nearby orange zone locations was not. Therefore, people in the orange zone serve as a suitable control group to assess the impact of the lockdown measures on the emotions of people in the red zone.

Causal identification is achieved through the exogenous shock represented by the policy, so that we can rule out selection bias of units under lockdown, which are balanced and comparable to the control group. To establish a quasi-experimental approach, we consider municipalities subjected to the lockdown as the treated group, specifically the ``red zone", which was at the epicenter of the COVID-19 outbreak. The control group consisted of municipalities surrounding Codogno, designated as the orange zone. To ensure comparability, only the closest non-urban areas to the red zone were selected.

In Figure 2 we show that complete randomization of the lockdown assignment holds for municipality level pre-lockdown covariates. The quantiles represent the acceptance region of our randomization test with $\alpha = 0.15$, using the standardized covariate mean differences as the test statistic. For each iteration, the lockdown is randomly assigned to units and the standardized mean difference is calculated. Complete randomization holds for all the considered covariates, as the observed standardized difference is well within the acceptance bounds [44]. The two areas have comparable pre-treatment demographic and socio-economic characteristics, providing support for random allocation.

Moreover, no other relevant event took place during the sample period in the area, which does not represent a major hub of the national or regional economy, so that we can exclude any bias induced by unaccounted major events influencing public reaction.

Additionally, the influence of news outlets and media coverage on reactions in the red zone was limited, as there was no local media in the red zone, and the national media had no access to it. Those inside and just outside the red zone can be considered equally exposed to the same partial and limited media coverage of the lockdown policies.

Overall, our study design makes a strong case for causal identification of the (lockdown) treatment effect, as the orange zone serves as a suitable counterfactual for the red zone. Because the observations we identify as controls have the same demographic, geographic, social, and economic composition, as well as the same exposure to the virus, the trends in aggregated and topic-related uncertainty and negative sentiment should not systematically differ between treated and control.

Figure 1. *Red* zone (lockdown of February 23, 2020 in Lombardy, Italy) and the surrounding selected control municipalities included in the analysis (in orange). We found tweets from 8 out of 10 cities in the red zone. The control area includes the municipalities surrounding Codogno in a circle with a radius of 42 km. The area we considered is homogeneous in terms of demographics, socioeconomics and virus exposure.

Figure 2. Standardized covariate mean difference along with 7.5% and 92.5% complete randomization quantiles using 2000 permutations. We check for covariate balance among the red zone and control municipalities of the orange zone matching the user location of tweets featured in the main analysis (8 red zone and 118 orange zone municipalities). We considered Istat data [45-49] on social, economic, and demographic characteristics: number of residents at January 1rst, 2020, January, February and March average number of deaths over 2015-2019, share of residents in occupational categories, age cohorts, and education levels. We also collect information on industry and services from 2017, including total output and value added (in euros), number of employees, total staff, and local units. We impute with the group mean missing values on industry and services for two treated cities (Castelgerundo, 1,473 residents, and Bertonico, 1,059 residents) and missing entries on monthly average deaths for two controls (Calendasco, 2,409 residents, and Cerro all'Ambro, 5\,149 residents) and one treated city (Bertonico, 1,059 residents). Before testing, we standardize the covariates on mean mortality, industry, services, and total residents.

Figure 3. Word clouds of the first 100 most frequent terms in tweets, classified as aggre- gated (a) *uncertainty* and (b) *negative sentiment*. In terms of our thematic categories, the most prominent terms in Figure (a) relate to health (e.g. *coronavirus, virus, cases (casi), deaths (morti), patient (paziente), asintomatic (asintomatico)*) and politics (e.g. *Europe (europa), parties (partito)*), and the lockdown policy (*supplies/food (viveri necessari), closures (chiuderlo)*). In Figure (b), the most frequently occurring words belong to the category health (*coronavirus, deaths (morti), ambulances (ambulanze)*), and to the cat- egory lockdown policy (*closures (chiuderlo), provisions/food supplies (viveri necessari)*). The complete set of words belonging to each thematic category is listed and translated in Tables 1 and 2 of the Supplementary Material.

Data Collection

For our main analysis, we collected Italian tweets both before and after the lockdown, from inside and just outside the red zone. Data collection occurred from December 1, 2019, to March 22, 2020, using the Twitter Stream API. This data included information about users' activity, user-defined locations, and the content of individual tweets.

Twitter-based indices have proven to be effective tools for understanding the emotional well-being of people during the COVID-19 pandemic [8, 50]. Social media-based measures are powerful tools to accurately identify uncertainty and negative sentiments following major events, such as COVID-19 [51]. As public's reactions to crisis is similar between those that use Twitter and those that do not, social media emotions become more representative of the emotions of the underlying population [52].

Tweets from the Codogno municipality and surrounding areas are collected using the Twitter Stream API with two filters: Italian language and geographical filtering. The latter allows to retrieve tweets with GPS coordinates within a maximal radius of 42km from a targeted center, in our case, Codogno. Therefore, we are able to extract tweets from the `red zone' locations and `orange zone' areas closest to Codogno. Then, we rely on the self-reported locations of the Twitter users to allocate tweets to municipalities either inside or outside the red zone. Self-reported locations are useful for fine-tuning user locations without relying solely on GPS coordinates to distinguish people inside and outside the red zone.[1] The text of the self-reported users locations is cleaned and matched with Istat data on Italian municipalities' geographical coordinates [54]. [2]

We clean tweets' text (removing URLs, hashtags and taggings) and remove re- peated tweets (tweets with the same cleaned text from the same users).
On March 8, all control units (areas that were not yet treated with lockdown measures) entered a *partial* lockdown, one day before the nationwide extension of the partial lockdown measures on March 9, 2020. As we focus only on the effect of the initial lockdown policy, we remove all observations from March 8, 2020. Additionally, to avoid any potential anticipation of the policy of February 23 around the time of the discovery of the first COVID-19 transmission case in Codogno (February 20, 2020), we drop observations spanning from February, 20 to February 22, 2020, and we remove all tweets from March 7, 2020—the day of the announcement of the first extension of the restrictions. At last, we drop all users that are information sources or business accounts.

We obtain a final sample of 24,261 unique tweets, 1,567 from eight out of the ten lockdown municipalities of the red zone and 22,694 unique tweets from 118 unique locations within the orange zone, with 1,124 unique users (60 and 1,064 from the red and orange zones respectively)[3]. Figure 1 highlights the cities with at least one active geo-referenced user account featured in our analysis.[4]

Uncertainty and Sentiment Classifier

To classify tweets into uncertainty and sentiment classes, we tune the pre-trained AlBERTo model [22], a Deep Learning natural language model trained on a large corpus of tweets in Italian (∼200 million, TWITA dataset)[5].

We fine-tune the pre-trained AlBERTo in a two-step process, using manually labeled tweets for training (n=6318). Manual labels indicate possible degrees of uncertainty (*neutral, uncertainty, certainty*) and sentiments (*neutral, negative* or *positive sentiment*).[6] The first step focuses on distinguishing neutral tweets from non-neutral ones, and the second step further classifies non-neutral tweets as uncertainty or certainty, and as positive or negative sentiment.

After fine-tuning the model and predicting the labels of all tweets, we obtain two binary variables named *uncertainty* and *negative sentiment* for each tweet in our dataset. The

[1]For a review of location-based methods, see [53].

[2] Manual checks on self-reported location of tweets also were carried out to check whether the user location was coherent with the contents of the tweets.

[3] We observe no tweets from two municipalities in the red zones, including San Fiorano (1,849 inhabitants) and Terranova dei Passerini (731 inhabitants).

[4] An additional, larger sample of tweets from all over Italy (774,407 tweets) was collected during the same period for placebo analyses, using language filtering and keyword-based queries related to COVID-19. We used these tweets for our placebo analyses detailed in Section H of the Supplementary Material. Using geographical coordinates of municipalities retrieved via Istat [54], we are able to identify additional unique tweets from the red zone located and from new orange zone municipalities in close proximity to Codogno. The new observations are added to the sample of tweets collected around Codogno.

[5] The model is available at https://github.com/marcopoli/AlBERTo-it.

[6] Authors 1, 2, and 3 manually assigned a subset of tweets to labels. The labels were then independently validated and verified by the authors.

variable *uncertainty* will equal 1 if the tweet expresses uncertainty – concern, questions, uneasiness – zero otherwise—i.e., the text is conveying a different emotional state, like certainty or indifference. The same reasoning applies to the *negative sentiment* categorical variable: it is set to 1 if the tweet expresses negative sentiments, zero if the text *does not* express negative sentiment—i.e., it is expressing positivity or neutrality. In Figure 3, we report the wordclouds with the first 100 most common terms features in tweets classified as uncertainty and negative sentiment. Detailed explanation of the classification procedure can be found in the Supplementary Material files, Section A.

Topics Classifier

To categorize tweets into topics, a dictionary-based classifier is used, which identifies topic-specific words in the tweets (Section B, Supplementary Material). We select four categories that are likely to be influenced by the policy: health, economics, politics and policy. A binary variable is created for each topic, indicating whether a tweet contains at least one term from the topic-specific dictionary. The overlap of topic labels in a single tweet is allowed, as tweets can cover multiple topics simultaneously. In Section C of the Supplementary Material files, we report a sample of tweets with the highest Shannon entropy by emotion-topic pair and discuss the performance of the machine learning classifier in identifying tweets related to each topic and emotion. Table 1 shows the number of tweets in our sample classified as uncertainty, or negative sentiment, aggregated and grouped by topics.

Table 1. Number of tweets classified as uncertainty, and negative sentiment, aggregated and grouped by topics (N = 24,261).[a]

	Aggregated, n (%)	Economics, n (%)	Health, n (%)	Politics, n (%)	Policy, n (%)
Uncertainty[b]	5,270 (0.217)	369 (0.015)	1,063 (0.043)	301 (0.012)	304 (0.012)
Negative Sentiment[b]	8,451 (0.348)	497 (0.020)	1,029 (0.042)	738 (0.03)	253 (0.01)

[a] We left out the category of Leisure from the Table: uncertainty Leisure, n (%) = 3,233 (0.61), Negative Sentiment Leisure, n (%) = 5,934 (0.70)

[b] The remaining categories to measure uncertainty (negative sentiment), neutral and certainty (positive sentiment), are not reported.

Difference-in-Difference Model

To estimate the causal impact of the lockdown on sentiment and uncertainty we estimate a linear DiD model. Let *i* be a tweet from user *j* and *T* be the time period of analysis at which we observe tweet *i* of user *j*, $T_{ij}=t, t\in\{0,1,2\}$, before (t=0) and after (t=1) the implementation of the lockdown of February 23, 2020, and following the first nationwide lockdown of March 9, 2020 (t=2). We define D_{ij} as the treatment status indicator: $D_{ij}=1$ if tweet i is from a user j who is in the area subject to the lockdown of February 23, $D_{ij}=0$ otherwise. In period t=1 individuals are assigned to a treatment, while in t=2 all Italians are treated by the national lockdown, which was less restrictive than the lockdown of period t=1. All Italians remain untreated in t=0. Formally, we estimate the following DiD equation on repeated cross-sections of tweets:

$$Y_{ij,t}=\alpha+\beta_j+\sum_{t\in\{1,2\}}\lambda_t 1[T_{ij}=t]+\sum_{t\in\{1,2\}}\delta_t\left(1[T_{ij}=t]\times D_{ij}\right)+\varepsilon_{ij,t}$$

where $Y_{ij,t}$ is the binary outcome variable obtained from the NLP classification (i.e. uncertainty and negative sentiments, aggregated and grouped by topics). We include user-level fixed effects β_j and time-fixed effects λ_t, where $1[.]$ is the indicator function. Our coefficient of interest is given by δ_1, measuring the effect of the lockdown policy of February 23 on the treated group in the post-treatment period t=1.

In this DiD setting, the average treatment effect on the treated in period 1, denoted as ATT(1), is the causal estimand of interest. Under the assumptions of the DiD model, the OLS estimator $\hat{\delta}_1$ in (1) is a consistent estimator of ATT(1).[7]

As all suitable control units eventually become treated by partial lockdown restrictions on March 9, 2020, we cannot identify δ_2 as the average treatment effect on the treated units in period t=2. Instead, the OLS estimate of $\hat{\delta}_2$ can only inform us on the average difference in outcomes between the two groups in period t=2, once all are subjected to national partial restrictions.

Results

Table 2 and Table 3 report the DiD model estimates of the effect of the COVID-19 shock on uncertainty and negative sentiments, aggregated and broken down by topics—economics, health, politics, and lockdown policy. We cluster standard errors at the city level (126 clusters) using the Liang-Zeger formula [55]. Model estimates with White standard errors are provided in the Supplementary Material, Section J.

[7] The assumptions of the DiD model are discussed in more detail in Section D of the Supplementary Material.

Table 2: DiD Regression table for *Uncertainty*, aggregated and grouped by topics with user fixed effects. Standard errors (SE) are clustered at the municipality level (in parenthesis).

	(1) Aggregate		(2) Economics		(3) Health		(4) Politics		(5) Policy	
	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)
post=1	0.02	[-3e-6,0.04]	-0.0006	[-0.005,0.003]	0.037	[0.027,0.046]	-0.0077	[-0.027,0.001]	0.01	[0.007,0.014]
	(0.01)	(0.05)	(0.002)	(0.800)	(0.005)	(<.001)	(0.004)	(0.088)	(0.002)	(<.001)
post=2	0.022	[-0.01,0.056]	0.004	[-6e-4,0.01]	0.042	[0.03,0.052]	-0.006	[-0.016,0.003]	0.01	[0.005,0.015]
	(0.017)	(0.21)	(0.002)	(0.084)	(0.005)	(<.001)	(0.004)	(0.194)	(0.0025)	(<.001)
redzone=1 x post=1	0.15	[0.09,0.22]	0.005	[-0.01,0.02]	0.036	[0.023,0.05]	0.012	[4e-4,0.023]	0.026	[0.015,0.038]
	(0.03)	(<.001)	(0.007)	(0.485)	(0.0065)	(<.001)	(0.0058)	(0.042)	(0.006)	(<.001)
redzone=1 x post=2	0.03	[-0.07,0.13]	-0.004	[-0.02,0.016]	-0.023	[-0.04,-0.0037]	0.0138	[-0.012,0.015]	0.005	[-0.02,0.028]
	(0.05)	(0.550)	(0.01)	(0.703)	(0.01)	(0.02)	(0.007)	(0.845)	(0.011)	(0.653)
Constant	0.95	[0.85,1.04]	-0.0008	[-0.02,0.018]	-0.019	[-0.035,-0.002]	0.005	[-0.005,0.015]	-0.015	[-0.04,0.007]
	(0.047)	(<.001)	(0.01)	(0.933)	(0.008)	(0.026)	(0.005)	(0.338)	(0.011)	(0.183)
Observations	24261		24261		24261		24261		24261	

Standard errors in parentheses

Table 3: DiD Regression table for *Negative Sentiment*, aggregated and grouped by topics with user fixed effects. Standard errors (SE) are clustered at the municipality level (in parenthesis).

	(1) Aggregate		(2) Economics		(3) Health		(4) Politics		(5) Policy	
	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)	Est. (SE)	95 % CI (P value)
post=1	-0.044	[-0.09,-5e-4]	-0.004	[-0.015,0.006]	0.036	[0.024,0.05]	-0.015	[-0.03,0.001]	0.0058	[0.002,0.01]
	(0.022)	(0.047)	(0.005)	(0.419)	(0.0063)	(<.001)	(0.008)	(0.063)	(0.0017)	(<.001)
post=2	-0.056	[-0.10,-0.01]	0.003	[-0.007,0.015]	0.038	[0.026,0.05]	-0.02	[-0.037,-0.003]	0.0066	[0.003,0.01]
	(0.023)	(0.016)	(0.0057)	(0.517)	(0.0057)	(<.001)	(0.008)	(0.023)	(0.0018)	(<.001)
redzone=1 x post=1	-0.072	[-0.125,-0.018]	-0.011	[-0.03,0.006]	0.027	[-0.015,0.07]	0.025	[0.003,0.05]	0.024	[-0.015,0.064]
	(0.027)	(0.008)	(0.0089)	(0.190)	(0.021)	(0.208)	(0.011)	(0.021)	(0.02)	(0.223)
redzone=1 x post=2	-0.049	[-0.146,0.047]	-0.006	[-0.03,0.019]	0.0048	[-0.038,0.048]	0.028	[0.006,0.05]	0.018	[-0.012,0.048]
	(0.049)	(0.312)	(0.0126)	(0.632)	(0.022)	(0.826)	(0.011)	(0.014)	(0.015)	(0.234)
Constant	0.105	[0.02,0.19]	0.0024	[-0.02,0.024]	-0.042	[-0.084,-0.001]	-0.008	[-0.02,0.01]	-0.024	[-0.05,0.005]
	(0.043)	(0.015)	(0.011)	(0.834)	(0.021)	(0.046)	(0.007)	(0.253)	(0.015)	(0.101)
Observations	24261		24261		24261		24261		24261	

Standard errors in parentheses

We first look at the effect of the lockdown on aggregate uncertainty. In column (1) of Table 2, the effect of treatment (red zone = 1 × post = 1) is positive and statistically significant at 0.001 (b=0.15, SE=0.01, t(125)= 4.82, P<.001). The estimated time trend does not show any significant increase in uncertainty following the implementation of the lockdown (post = 1, b=0.02, SE=0.01, t(125)= 1.98, P =.05) and the subsequent implementation of the partial nationwide measures (post = 2, b=0.022, SE=0.017, t(125)= 1.27,P =.20). The aggregate results indicate that the lockdown significantly increased aggregate uncertainty for those subjected to it. These results may not reflect the whole picture, as different forces may be at play in the aggregate expression of emotions.

The effects of the lockdown on economic uncertainty are shown in column (2) of Table 2. The lockdown does not have a significant impact on the economic uncertainty of those affected by the policy, with a small, positive estimated effect on economic uncertainty (b=0.005, SE=0.007, t(125)=0.70, P =.48). This indicates that the lockdown did not have a major impact on economic uncertainty.

Turning to health-related matters, column (3) of Table 2 shows that the lockdown significantly increased concern about health problems among people in the lockdown area (b=0.036, SE=0.0065, t(125)=5.55, P <.001). After the partial nationwide measures of March 9, people in the red zone expressed significantly less uncertainty in health discussions (b=- 0.023, SE=0.01, t(125)=-2.35, P =.02), as expected, as the measures became softer. On the other hand, the common trend of health-related uncertainty increased significantly after the start of the local lockdown (b=0.037, SE=0.005, t(125)=7.50, P <.001) and the partial nationwide measures (b=0.042, SE=0.0052, t(125)=8.00, P <.001).

The results related to the political uncertainty can be found in column (4) of Table 2. The lockdown significantly increased political uncertainty in the treated area (b=0.012, SE=0.0058, t(125)=2.05, P =.04). The effect disappears and is no longer statistically different from zero as soon as all units are subjected to the same nationwide measures (red zone = 1 × post = 2, b= 0.00138, SE=0.007, t(125)=0.20, P = 0.84).

Column (5) of Table 2 displays the results for uncertainty related to the lockdown policy itself. In Table 2, the lockdown significantly worsened the worries of those subjected to the measures (b=0.026, SE=0.006, t(125)=4.47, P <.001), although the coefficients of post = 1 (b=0.01, SE=0.002, t(125)=5.56, P <.001) and post = 2 (b= 0.01,SE=0.0025, t(125)=3.98, P <.001) show that the uncertainty of the common trend around the practical implications of the lockdown was significantly above baseline.

DiD estimates of the effect of the lockdown on aggregated negative sentiments are reported in Table 3, column (1). The lockdown significantly reduced negative sentiments for those in the red zone (b=-0.072, SE=0.027, t(125)=-2.67, P =.008). The probability of expressing negative sentiments decreases over time for both the treated and control groups, since the estimated coefficients of post = 1 (b=-0.044, SE=0.022, t(125)=-2.00, P =.047) and post = 2 (b=-0.056, SE=0.023, t(125)=-2.45, P =.016) are negative and statistically significant. The aggregate results indicate that the lockdown significantly decreased aggregated negative sentiments for those in the red zone.

In Table 3, column (2), we show the effect on negative sentiment in economics. Although not significantly different from zero, the estimated effect of a lockdown on negative sentiments in economic discussions is negative (b=-0.011, SE=0.0089, t(125)=-1.32, P =.19) and of relatively greater magnitude with respect to the effect estimated on economic uncertainty in column (2), Table 2.

Regression results in column (3) of Table 3 shows that the lockdown had no significant effect on the negative sentiment related to the health issues in the lockdown area (b= 0.027, SE=0.021, t(125)=1.27, P = .2). The general trend shows a significant worsening of sentiments about health issues both within and outside the lockdown area, as shown by the coefficients of post = 1 (b= 0.036, SE=0.0063, t(125)=5.70, P < 0.001) and post = 2 (b=0.038, SE=0.0057, t(125)=6.57, P <.001).

Column (4) of Table 3 displays the effect of the lockdown on negative political sentiment. We estimate a positive and significant coefficient for both interactions, indicating that sentiments toward politics turned more negative among the treated users when the first lockdown was implemented (b=0.025, SE=0.011, t(125)=2.33, P =.02), and after the nationwide measures (b=0.028, SE=0.011, t(125)=2.50, P =.01). In particular, the general trend in negative sentiment towards politics decreased significantly following the partial national lockdown (b=-0.02, SE=0.008, t(125)=-2.30, P =.02).

In Table 3, column (5), the red zone did not express negative sentiments significantly more often when discussing the policy (b=0.024, SE=0.02, t(125)=1.22, P =.22). However, negative sentiment towards the policy increases everywhere, both after the first lockdown policy (b=0.0058, SE=0.0017, t(125)=3.39, P <.001) and after the nationwide measures (b=0.0066, SE=0.0018, t(125)=3.60, P <.001).

All in all, our result leads to an important distinction regarding the impact of the lockdown on economic uncertainty and economic sentiments. While economic uncertainty increased due to the spread of COVID-19 cases (as shown in [24,56]), in our case study the lockdown itself did not worsen economic uncertainty and negative sentiments. As can be seen from Figure 3, the tweets at this time were less interested in purely economic discussions: the most frequently used terms in tweets expressing uncertainty or negative sentiment were not related to the economy. Although the estimated coefficient for the effect of lockdown on economic uncertainty is very close to zero and quite small, we cannot rule out that the lack of statistical significance of the relatively larger impact of lockdowns on negative sentiment in economics is not due to lack of power and the small sample size.[8]

Instead, the lockdown affected the emotions of treated individuals on other topics such as health and politics. Our findings show that the lockdown increased health- related uncertainty of people in the treated area. Higher uncertainty about health issues may reflect people's concern about their safety, but also increase public awareness of the health risks posed by the virus [38]. This shows that the health concerns arising from the implementation of a lockdown outweigh the reassurance that the authorities want to convey to the population by implementing the lockdown and displaying their commitment to contain the virus. However, our results suggest that the negative sentiment towards health issues was caused by COVID-19 and not by the lockdown.

The same logic applies to the impact of the lockdown on the emotions expressed when discussing the lockdown policy: concern about the guidelines and the consequences of the policy was high among those in the red zone, but this was not accompanied by a significant change in sentiments among the treated population. Sentiments in discussions related to the policy deteriorated in both the treated and the control area after the implementation of the measures, suggesting that the worsening of sentiments was probably due to COVID- 19 itself, rather than the lockdown.

Our findings reveal that political sentiments among Twitter users in the treated area deteriorated significantly. This indicates that lockdown policies come at the major cost of worsening political emotions. The rise in negative sentiment in the red zone may reflect increased political polarization as a direct result of the lockdown policy, consistent with earlier evidence on

[8] See Figures 2 of the Appendix for the sample size of tweets belonging to each topic category.

political polarization during COVID-19 [18,19]. A deterioration in negative sentiment towards politics could also be a sign of increasing dissatisfaction with politicians in office. Previous research has shown that attitudes towards politicians have improved during the pandemic [35, 39, 57]. This is consistent with the estimated significant decline in the overall trend following the March 9 nationwide measures, which followed the dramatic increase of deaths and COVID-19 cases over the national territory.

In addition, the lockdown has significantly worsened political uncertainty in the treated area. Nevertheless, as discussed in the next section, the estimated effect on political uncertainty is not robust to a series of checks on the validity of the assumptions of the DiD model.

The overall effect of the lockdown on uncertainty was positive, which is explained by the significant increase in uncertainty towards health, policy, and politics. On the other hand, aggregated negative sentiment decreased in the red zone. This effect was led by a decrease in negative sentiments when discussing topics different from those of interest -- mostly related to entertainment and sports (Section H of the Supplementary Material). However, the estimated effect on aggregate negative sentiment does not pass most of the robustness checks we conducted, as explained in the next section. This is to confirm the validity of the topic based analysis in order to better determine the impact of lockdowns on uncertainty and sentiment.

In general, our results shed light on the nuanced effects of lockdown policies on various emotional aspects and suggest that the impact on health, economics, and politics differ considerably. This distinction between the effects of the lockdown policy and broader emotional trends related to the pandemic is possible due to the quasi-experimental frame- work of the analysis. In the next section, we explain that the effects on aggregate negative sentiment and uncertainty towards politics are not robust when tested in a series of robustness checks.

Robustness Checks

We perform a battery of tests to assess the robustness of our results. Here we discuss the main takeaways. First, we test the differences in the pre-treatment trends between treated and control to check for violations of the “parallel trends” assumption in the DiD model (Section E.1 of the Supplementary Material). Pre-treatment trends in aggregated uncertainty, uncertainty related to health and the policy and negative sentiment towards politics are not statistically different between treated and control groups, supporting the parallel trends assumption. However, pre-treatment trends in aggregate negative sentiment and uncertainty related to politics do differ between the two groups, indicating a potential violation of parallel trends.

We also allow for different post-treatment violations of parallel trends (in line with [58]) and find robust significant effects for aggregated uncertainty, uncertainty related to health, the policy, and negative sentiments toward politics (Section E.2 of the Supplementary Material). The impact of the lockdown on uncertainty about politics and aggregated negative sentiment is not robust to violations of parallel trends.

Moreover, we estimate the DiD model by controlling for potential local spillovers of the treatment effect on the control units [59] (Section E.3 of the Supplementary Material). By assuming that spillovers were negligible past 30 kms from the treated area, the estimates of the total average effect of the treatment are in line with baseline estimates of the treatment effect.

Additionally, we jointly tests the p-values of the models using Benjamini-Hochberg adjusted p-values

[60] to account for multiple hypothesis testing (Section F of the Supplementary Material). This analysis confirms the original small p-values of the lockdown effect, even after adjusting for correlations between regression models, except for the p-value of the effect on uncertainty related to politics.

Moreover, we check whether our findings are robust against the exclusion of any of the administrative territorial units among the control group (province in Italian). We see that the significance of the effect of the lockdown on uncertainty related to politics and aggregate negative sentiment crucially depends on the administrative unit excluded from the control group (Section I of the Supplementary Material).

Finally, a placebo test is conducted to assess the effects of the lockdown when it is not unexpected, focusing on the partial national lockdown's effects (Section G of the Supplementary Material). The results of this test show that the national lockdown has no significant effect on public reactions in the placebo treatment group, reinforcing the robustness of the original findings to potential anticipatory effects.

Overall, the robustness tests support the key findings that the lockdown had significant effects on aggregated uncertainty, uncertainty related to health and the policy itself, as well as on negative sentiments surrounding politics. The effect on uncertainty related to politics and on aggregated negative sentiment fails to survive several of the tests.

Discussion

Principal Results

In this paper, we estimate the causal impact of lockdowns on different dimensions of uncertainty and sentiments. Due to its sudden and unexpected implementation, the first Italian lockdown of February 23, 2020 served as an ideal quasi-experimental setting to disentangle the effect of restrictions from the spread of the pandemic itself. We combined machine learning methods applied to textual data from Twitter with causal inference methodologies to estimate how emotions changed between people inside and outside the treated areas. The approach we propose in this paper provides a more nuanced definition of emotions, which so far in the literature has tended to focus on economic uncertainty and sentiments, that could be easily applied in other policy evaluation exercises. By doing so, we examine how emotions changed related to economics, health, politics, and the policy itself.

Overall, we find that lockdowns do indeed come at a cost, which is, however, somewhat unexpected. We did not find a causal impact on economic uncertainty and sentiments, suggesting that economic-related emotions were mainly driven by the increase in COVID- 19 cases rather than the policy itself. Instead, in the *red zone* areas under lockdown, users were more likely to express uncertainty about health and lockdown policy. The evidence highlights that lockdowns increase, rather than decrease, the uncertainty about health conditions among users, without causing a significant change in sentiments when discussing health and the handling of the sanitary emergency. Our results suggest that the worsening of sentiments in discussions related to health and the policy are caused by COVID-19, rather than the stay-at-home measures. Importantly, lockdowns have political costs, as treated areas experience a surge in negative sentiments related to politics. This could discourage future implementations of stay-at-home mandates to avoid political backlash.

Limitations

It is important to note that all field experiments can suffer from lack of external validity. Our study is no exception, as a new lockdown today could be perceived differently from the unexpected lockdown we analyzed in this paper. One could argue that the public is much more informed today on the implications of a lockdown than it was in February 2020. People would adjust their strategy given what they have experienced, possibly changing the impact of the policy on uncertainty and sentiments. Moreover, given the negative political sentiments they generated, policymakers could be more reluctant to rely on such measures or at least to apply such a degree of strictness. Nevertheless, we show that lockdowns bear political and social costs in terms of worsening uncertainties and sentiments. While it is impossible to predict how the public and policymakers will react to a possible future lockdown, our results provide helpful guidance on the costs that policymakers will take into account before implementing new lockdown measures.

Comparison with Prior Work

The contribution of our work is twofold. First, it develops a new methodological framework to asses the impact of lockdown measures [31-33, 39, 61]. We focus on the causal effects on public emotions, for which, to our knowledge, there is only correlational evidence [10, 17, 39]. Understanding how stay-at-home mandates affect public emotions is critical for policymakers to assess the benefits and costs of such policies.

Second, we add to the literature on uncertainty and sentiment, which generally relies on non- traditional data such as textual data. Recent studies have documented particularly high levels of economic uncertainty [62] and health-related negative emotions [17,20] following the enactment of COVID-19 lockdown policies, as well as increased political polarization [18,19]. However, there is less clarity about what particular aspects have contributed to bringing about such changes. We contribute to filling this knowledge gap and show that uncertainty and negative sentiment did not increase uniformly across subjects after lock- down, but rather followed different dynamics. Thus, this is the first study to analyze the effects of stay-at-home mandates on multiple dimensions of uncertainty and sentiments.

Conclusions

Our findings suggest that communication around lockdown measures may have been confusing, leading to resentment among those affected toward the political class and failing to explicitly generate support for governmental plans. Policy makers must improve communication to reduce uncertainty and negative sentiments [63]. Our results emphasize the need for authorities to use these findings to improve future policies and communication efforts to mitigate uncertainty and social panic.

Acknowledgements

F.B.S. and M.R. conceived the research idea and designed the identification strategy; M.P. gathered data; C.B. conducted the experiment(s) and analysed the results.; C.B., F.B.S., N.F., and M.R. supervised investigation; C.B., F.B.S., N.F., M.P. and M.R. wrote and reviewed the manuscript. This work is supported in part by the Tuscany Health Ecosystem - THE, PNRR, NextGeneration EU. We benefited from helpful comments and suggestions from Danielle Braun, Abel Brodeur, Francesca Dominici, Kosuke Imai, Kevin Josey, Fabrizia Mealli, Giovanni Mellace, and Davide Viviano. Results previously presented at Harvard Data Science seminar (2022), Institute of Mathematical Statistics (IMS) International Conference on Statistics and Data Science (ICSDS) (2022), Joint Political Economy and Applied Microeconomics Workshop at the University of Bolzano (2022), at the Advanced Language Processing Winter School (2023), and at EEA-ESEM Congress (2023).

Conflicts of Interest

None declared.

Abbreviations

DiD: Difference-in-Difference
SE: Standard Error
JMIR: Journal of Medical Internet Research

Data Availability

The data underlying this article are available in Zenodo and can be accessed with DOI: 10.5281/zenodo.10927516.

Multimedia Supplementary Material 1

See Online Supplementary Material file.

References

1. Salah Abosedra, Nikiforos T. Laopodis, and Ali Fakih. "Dynamics and asymmetries between consumer sentiment and consumption in pre- and during-COVID-19 time". In: *The Journal of Economic Asymmetries* 24 (2021), e00227. issn: 1703-4949. doi: doi.org/10.1016/j.jeca.2021.e00227.

2. Giovanni Caggiano, Efrem Castelnuovo, and Richard Kima. "The global effects of Covid-19-induced uncertainty". In: *Economics Letters* 194 (2020), p. 109392. doi: 10.1016/j.econlet.2020.109392. url: https://doi.org/10.1016/j.econlet.2020.109392.

3. Micha l Brzoza-Brzezina and Grzegorz Weso lowski. *The Great Lockdown*. Working Papers 2021-060. Warsaw School of Economics, Collegium of Economic Analysis, 2021. url: https://ideas.repec.org/p/sgh/kaewps/2021060.html.

4. Christian Moser and Pierre Yared. *Pandemic Lockdown*. Staff Report 627. Federal Reserve Bank of Minneapolis, 2021. doi: 10.21034/sr.627.

5. Edoardo Di Porto, Paolo Naticchioni, and Vincenzo Scrutinio. "Lockdown, essential sectors, and Covid-19". In: *Journal of Health Economics* 81 (2022), p. 102572. issn: 0167-6296. doi: https:// doi. org/ 10 . 1016 / j.

jhealeco. 2021 . 102572. url: https://www.sciencedirect.com/science/article/pii/S0167629621001570.

6. James H Fowler et al. “Stay-at-home orders associate with subsequent decreases in COVID-19 cases and fatalities in the United States”. In: *PLoS One* 16.6 (2021), e0248849.

7. Renquan Zhang et al. “Evaluating the impact of stay-at-home and quarantine mea- sures on COVID-19 spread”. In: *BMC Infectious Diseases* 22.1 (2022), pp. 1–13.

8. Dave Altig et al. “Economic uncertainty before and during the COVID-19 pan- demic”. In: *Journal of Public Economics* 191 (2020), p. 104274.

9. Austan Goolsbee and Chad Syverson. “Fear, lockdown, and diversion”. In: *Journal of Public Economics* 193 (2021), p. 104311. doi: 10 . 1016 / j . jpubeco . 2020 . 104311.

10. Ishaani Priyadarshini et al. “A study on the sentiments and psychology of twitter users during COVID-19 lockdown period”. In: *Multimedia Tools and Applications* 81.19 (2021), pp. 27009–27031. doi: 10.1007/s11042-021-11004-w.

11. Jerome Tze-Hou Hsu and Richard Tzong-Han Tsai. “Increased Online Aggression During COVID-19 Lockdowns”. In: *J Med Internet Res* 24.8 (Aug. 2022), e38776. issn: 1438-8871. doi: 10.2196/38776. url: http://www.ncbi.nlm.nih.gov/ pubmed/35943771.

12. Athina Zervoyianni, Sophia Dimelis, and Alexandra Livada. “Economic Sentiment and the Covid-19 Crisis: Evidence from European Countries”. In: *Applied Economics* 0.0 (2022), pp. 1–18. doi: 10.1080/00036846.2022.2061903.

13. Wouter van der Wielen and Salvador Barrios. “Economic sentiment during the COVID pandemic: Evidence from search behaviour in the EU”. In: *Journal of Economics and Business* 115 (2021). COVID-19 – Economic and Financial Effects, p. 105970. issn: 0148-6195. doi: https://doi.org/10.1016/j.jeconbus.2020. 105970.

14. Scott Baker et al. *COVID-Induced Economic Uncertainty*. Tech. rep. 2020. doi:10.3386/w26983. url: https://doi.org/10.3386/w26983.
15. Nalini Chintalapudi, Gopi Battineni, and Francesco Amenta. “Sentimental Anal- ysis of COVID-19 Tweets Using Deep Learning Models”. In: *Infectious Disease Reports* 13.2 (2021), pp. 329–339. issn: 2036-7449. doi: 10. 3390/ idr13020032. url: https://www.mdpi.com/2036-7449/13/2/32.
16. Faheem Aslam et al. “Sentiments and emotions evoked by news headlines of coro- navirus disease (COVID-19) outbreak”. In: *Humanities and Social Sciences Com- munications* 7.1 (2020). doi: 10.1057/s41599-020-0523-3.
17. May Oo Lwin et al. “Global Sentiments Surrounding the COVID-19 Pandemic on Twitter”. In: *JMIR Public Health Surveill* 6.2 (2020), e19447. issn: 2369-2960. doi: 10.2196/19447.
18. Sebastian Jungkunz. “Political Polarization During the COVID-19 Pandemic”. In: *Frontiers in Political Science* 3 (2021). issn: 2673-3145. doi: 10.3389/fpos.2021. 622512.
19. Julie Jiang et al. “Political polarization drives online conversations about COVID- 19 in the United States”. In: *Human Behavior and Emerging Technologies* 2.3 (2020), pp. 200–211. doi: https://doi.org/10.1002/hbe2.202.
20. Yao-Tai Li, Man-Lin Chen, and Hsuan-Wei Lee. “Health communication on social media at the early stage of the pandemic”. In: *Social Science & Medicine* 347 (2024), p. 116748. issn: 0277-9536. doi: https://doi.org/10.1016/j.socscimed.2024. 116748.
21. Danilo Cereda et al. “The early phase of the COVID-19 epidemic in Lombardy, Italy”. In: *Epidemics* 37 (2021), p. 100528. issn: 1755-4365. doi: https://doi. org/10.1016/j.epidem.2021.100528.
22. Marco Polignano et al. “Alberto: Italian BERT language understanding model for NLP challenging tasks based on tweets”. In: *6th Italian Conference on Computa- tional Linguistics, CLiC-it 2019*. Vol. 2481. CEUR. 2019, pp. 1–6. url:

https://api.semanticscholar.org/CorpusID:204914950.

23. Jacob Devlin et al. “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Lan- guage Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 4171–4186. doi: 10 . 18653/v1/N19-1423.

24. Scott R. Baker, Nicholas Bloom, and Steven J. Davis. “Measuring Economic Policy Uncertainty”. In: *The Quarterly Journal of Economics* 131.4 (2016), pp. 1593–1636. doi: 10.1093/qje/qjw024.

25. Adam Hale Shapiro, Moritz Sudhof, and Daniel J. Wilson. “Measuring news senti- ment”. In: *Journal of Econometrics* 228.2 (2022), pp. 221–243. doi: 10.1016/j. Jeconom.2020.07.053.

26. Guillermo Blanco and An´alia Louren¸co. “Optimism and pessimism analysis using deep learning on COVID-19 related twitter conversations”. In: *Information Pro- cessing and Management* 59.3 (2022), p. 102918. issn: 0306-4573. doi: https ://doi.org/10.1016/j.ipm.2022.102918.

27. Tashko Pavlov and Georgina Mirceva. “COVID-19 Fake News Detection by Using BERT and RoBERTa models”. In: *2022 45th Jubilee International Convention on Information, Communication and Electronic Technology (MIPRO)*. 2022, pp. 312– 316. doi: 10.23919/MIPRO55190.2022.9803414.

28. Hanyi Min et al. “Using Machine Learning to Investigate the Public’s Emotional Responses to Work from Home During the COVID-19 Pandemic”. In: (2021). doi: 10.31234/osf.io/xk5hz.

29. Quyen G To et al. “Applying Machine Learning to Identify Anti-Vaccination Tweets during the COVID-19 Pandemic”. In: *International Journal of Environmental Re- search and Public Health* 18.8 (2021). issn: 1660-4601. doi: 10.3390/ijerph18084069.

30. Ioana-Elena Oana, Alessandro Pellegata, and Chendi Wang. “A cure

worse than the disease? Exploring the health-economy trade-off during COVID-19". In: *West European Politics* 44.5-6 (2021), pp. 1232–1257. doi: 10.1080/01402382.2021. 1933310.

31. Daniel Devine et al. "Trust and the Coronavirus Pandemic: What are the Conse- quences of and for Trust? An Early Review of the Literature". In: *Political Studies Review* 19.2 (2020), pp. 274–285. doi: 10.1177/1478929920948684.

32. Dominik Schraff. "Political trust during the Covid-19 pandemic: Rally around the flag or lockdown effects?" In: *European journal of political research* 60.4 (2021), pp. 1007–1017. url: https://doi.org/10.1111/1475-6765.12425.

33. Sven Hegewald and Dominik Schraff. "Who rallies around the flag? Evidence from panel data during the Covid-19 pandemic". In: *Journal of Elections, Public Opinion and Parties* (2022), pp. 1–22. doi: 10.1080/17457289.2022.2120886.

34. Martin Baekgaard et al. ""Rallying Around the Flag in Times of COVID-19". In: *Journal of Behavioral Public Administration* 3 (2020), pp. 1–28. doi: https://doi.org/10.30636/jbpa.32.172.

35. Catherine E. De Vries et al. "Crisis signaling: how Italy's coronavirus lockdown affected incumbent support in other European countries". In: *Political Science Re- search and Methods* 9.3 (2021), pp. 451–467. doi: 10.1017/psrm.2021.6.

36. Joost Oude Groeniger et al. "Dutch COVID-19 lockdown measures increased trust in government and trust in science". In: *Social Science & Medicine* 275 (2021), p. 113819. doi: 10.1016/j.socscimed.2021.113819.

37. Diane Alexander and Ezra Karger. "Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior". In: *Review of Economics and Statistics* (2023), pp. 1–11. doi: 10.1162/rest_a_01108.

38. Sabina De Rosis et al. "The early weeks of the Italian Covid-19 outbreak: sentiment insights from a Twitter analysis". In: *Health*

Policy 125.8 (2021), pp. 987–994. issn: 0168-8510. doi: https://doi.org/10.1016/j.healthpol.2021.06.006.

39. Damien Bol et al. "The effect of COVID-19 lockdowns on political support: Some good news for democracy?" In: *European Journal of Political Research* 60.2 (2020), pp. 497–505. doi: 10.1111/1475-6765.12401. url: https://doi.org/10.1111/ 1475-6765.12401.
40. *Gazzetta Ufficiale* n.45 (Feb. 23, 2020): https://www.gazzettaufficiale.it/eli/id/2020/02/23/20A01228/sg
41. *Gazzetta Ufficiale* (Mar. 8, 2020): https://www.gazzettaufficiale.it/eli/id/2020/03/08/20A01522/sg
42. *Gazzetta Ufficiale* (Mar. 9, 2020): https://www.gazzettaufficiale.it/eli/id/2020/03/ 09/20A01558/sg
43. *La Stampa* (Feb. 20, 2022): https://www.lastampa.it/cronaca/2022/02/20/news/codogno_20_febbraio_il_tampone_0_che_scopri_il_covid_in_italia_la_dottoressa_annalisa_malara_ricordo_ancora_la_paura_or-2859352/.
44. Branson, Zach. 2021. "Randomization Tests to Assess Covariate Balance When Designing and Analyzing Matched Datasets." Observational Studies 7:1–36. https://doi.org/10.1353/obs.2021.0031.
45. ISTAT. "CODICI STATISTICI DELLE UNITA` AMMINISTRATIVE TERRITORIALI: COMUNI, CITTA` METROPOLITANE, PROVINCE E REGIONI". In: *Istituto nazionale di statistica e informatica* (2022-2023), (accessed on July 6, 2022).url: https://www.istat.it/it/archivio/6789.
46. ISTAT. "Censimenti Permanenti: Censimento della Popolazione e delle Abitazioni". In: *Istituto nazionale di statistica e informatica* (2019), (acessed on July 6, 2022). url: http://dati-censimentipermanenti.istat.it/.
47. ISTAT. "Dataset analitico con i dati comunali nei settori attivi e sospesi". In: *Istituto nazionale di statistica e informatica* (2020), (accessed on September 8, 2022). url: https://www.istat.it/it/archivio/241341.
48. ISTAT. "POPOLAZIONE RESIDENTE AL 1° GENNAIO 2020". In: *Istituto nazionale di statistica e informatica* (2020), (accessed on July 8, 2022). url: http://dati.istat.it/index.aspx?datasetcode=dcis_popres1.
49. ISTAT. "Tavola decessi per 6.866 comuni al 31marzo2020". In: *Istituto nazionale di statistica e informatica* (2020), (accessed on July 9, 2022). url: https://www.istat.it/it/archivio/240401.

50. Ali Shariq Imran et al. “Cross-Cultural Polarity and Emotion Detection Using Sentiment Analysis and Deep Learning on COVID-19 Related Tweets”. In: *IEEE Access* 8 (2020), pp. 181074–181090. doi: 10.1109/ACCESS.2020.3027350.

51. Yixian Zhang et al. *COVID-19 Public Opinion and Emotion Monitoring System Based on Time Series Thermal New Word Mining*. 2020. doi: 10.48550/ARXIV. 2005.11458.

52. Metzler, Pellert, and Garcia. “Using social media data to capture emotions be- fore and during COVID-19”. In: *World Happiness Report*. New York: Sustainable Development Solutions Network., 2022, pp. 77–104.

53. Xin Zheng, Jialong Han, and Aixin Sun. “A Survey of Location Prediction on Twitter”. In: *IEEE Transactions on Knowledge and Data Engineering* 30.9 (2018), pp. 1652–1671. doi: 10.1109/TKDE.2018.2807840.

54. ISTAT. “CONFINI DELLE AMMINISTRATIVE A FINI STATISTICIUNITA` AL 1° GENNAIO 2022”. In: *Istituto nazionale di statistica e informatica* (2022), (accessed on July 6, 2022). url: https://www.istat.it/it/archivio/222527.

55. Kung-Yee Liang and Scott L. Zeger. “Longitudinal data analysis using generalized linear models”. In: *Biometrika* 73.1 (1986), pp. 13–22. issn: 0006-3444. doi: 10. 1093/biomet/73.1.13.

56. Travis Adams et al. *More than Words*. Federal Reserve Board Working Paper 2023-034. Board of Governors of the Federal Reserve System, 2023. doi: doi.org/10. 17016/FEDS.2023.034.

57. Chris G. Sibley et al. “Effects of the COVID-19 pandemic and nationwide lockdown on trust, attitudes toward government, and well-being.” In: *American Psychologist* 75.5 (2020), pp. 618–630. doi: 10.1037/amp0000662.

58. Ashesh Rambachan and Jonathan Roth. “A More Credible Approach to Parallel Trends”. In: *The Review of Economic Studies* (2023). doi: 10 . 1093 / restud/ rdad018.

59. Kyle Butts. *Difference-in-Differences Estimation with Spatial Spillovers*. Papers. arXiv.org, 2023. url: https:// EconPapers. repec. org/ RePEc: arx: papers: 2105.03737.

60. Yoav Benjamini and Yosef Hochberg. “Controlling The False Discovery Rate - A Practical And Powerful Approach To Multiple Testing”. In: *J. Royal Statist. Soc., Series B* 57 (1995), pp. 289–300. doi: 10.2307/2346101.

61. juan Palomino, Juan G Rodriguez, and Raquel Sebastian. “Wage inequality and poverty effects of lockdown and social distancing in Europe”. In: *European economic review* 129 (2020), p. 103564.

62. Jianlei Yang and Chunpeng Yang. “Economic policy uncertainty, COVID-19 lock- down, and firm-level volatility: Evidence from China”. In: *Pacific-Basin Finance Journal* 68 (2021), p. 101597. issn: 0927-538X. doi: doi.org/10.1016/j.pacfin. 2021.101597.

63. Vishala Mishra and Joseph Dexter. “Comparison of Readability of Official Public Health Information About COVID-19 on Websites of International Agencies and the Governments of 15 Countries”. In: *JAMA Network Open* 3 (2020), e2018033. doi: 10.1001/jamanetworkopen.2020.18033.

Supplementary Files

Figures

Red zone (lockdown of February 23, 2020 in Lombardy, Italy) and the surrounding selected control municipalities included in the analysis (in orange). We found tweets from 8 out of 10 cities in the red zone. The control area includes the municipalities surrounding Codogno in a circle with a radius of 42 km. The area we considered is homogeneous in terms of demographics, socioeconomics and virus exposure.

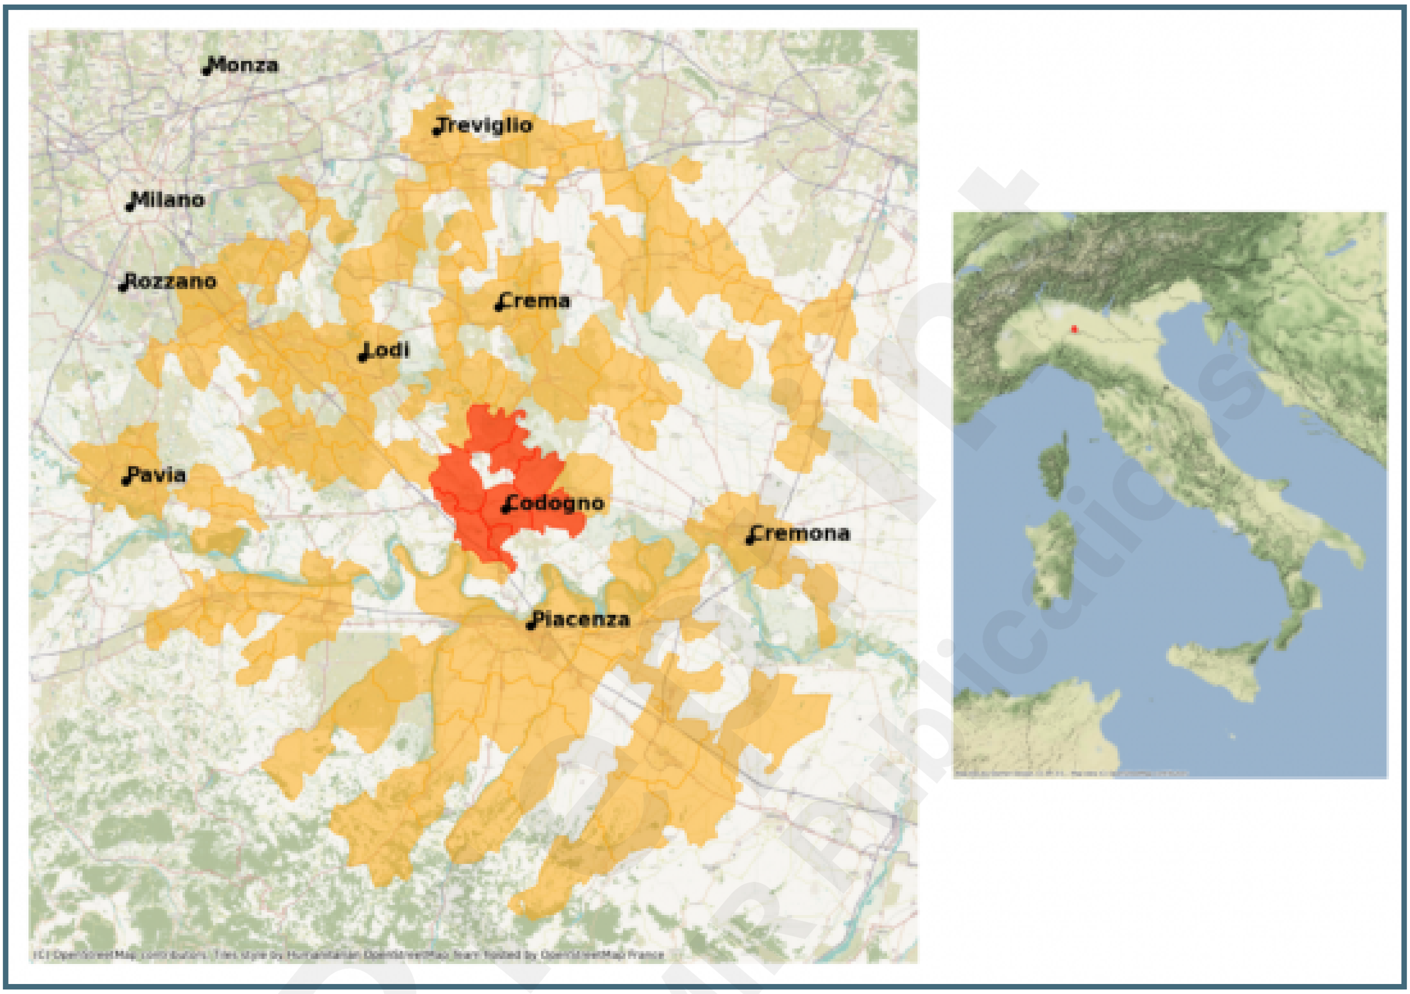

Standardized covariate mean difference along with 7.5% and 92.5% complete randomization quantiles using 2000 permutations. We check for covariate balance among the red zone and control municipalities of the orange zone matching the user location of tweets featured in the main analysis (8 red zone and 118 orange zone municipalities). We considered Istat data [45-49] on social, economic, and demographic characteristics: number of residents at January 1rst, 2020, January, February and March average number of deaths over 2015-2019, share of residents in occupational categories, age cohorts, and education levels. We also collect information on industry and services from 2017, including total output and value added (in euros), number of employees, total staff, and local units. We impute with the group mean missing values on industry and services for two treated cities (Castelgerundo, 1,473 residents, and Bertonico, 1,059 residents) and missing entries on monthly average deaths for two controls (Calendasco, 2,409 residents, and Cerro all'Ambro, 5\,149 residents) and one treated city (Bertonico, 1,059 residents). Before testing, we standardize the covariates on mean mortality, industry, services, and total residents.

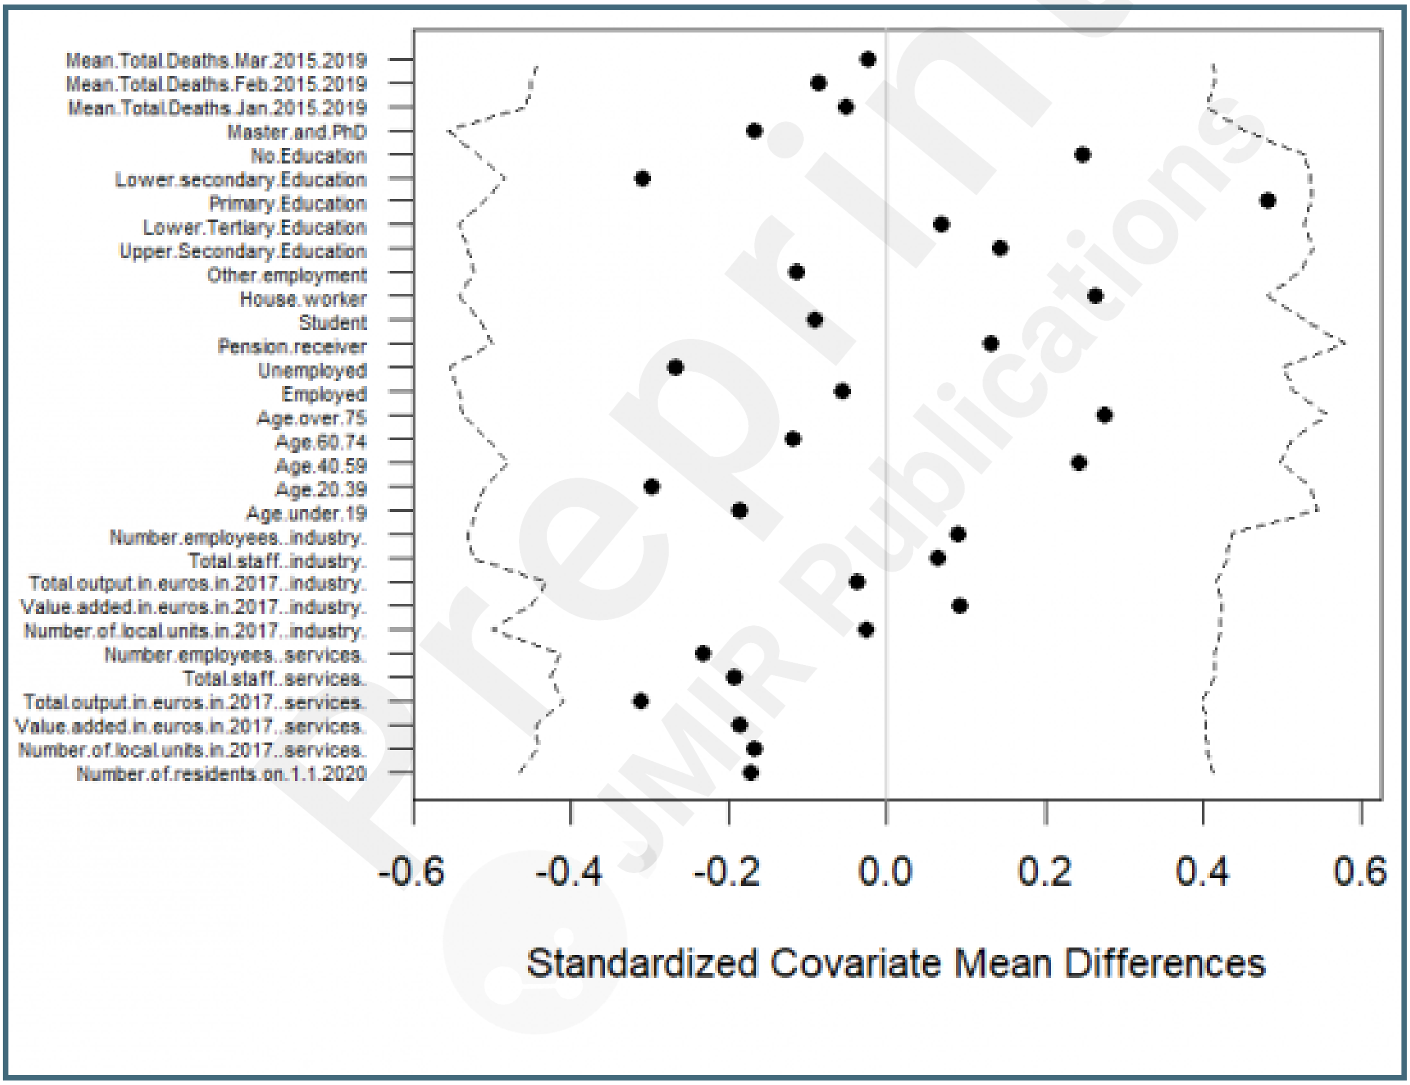

Word clouds of the first 100 most frequent terms in tweets, classified as aggre- gated (a) uncertainty and (b) negative sentiment. In terms of our thematic categories, the most prominent terms in Figure (a) relate to health (e.g. coronavirus, virus, cases (casi), deaths (morti), patient (paziente), asintomatic (asintomatico)) and politics (e.g. Europe (europa), parties (partito)), and the lockdown policy (supplies/food (viveri necessari), closures (chiuderlo)). In Figure (b), the most frequently occurring words belong to the category health (coronavirus, deaths (morti), ambulances (ambulanze)), and to the cat- egory lockdown policy (closures (chiuderlo), provisions/food supplies (viveri necessari)). The complete set of words belonging to each thematic category is listed and translated in Tables 1 and 2 of the Supplementary Material.

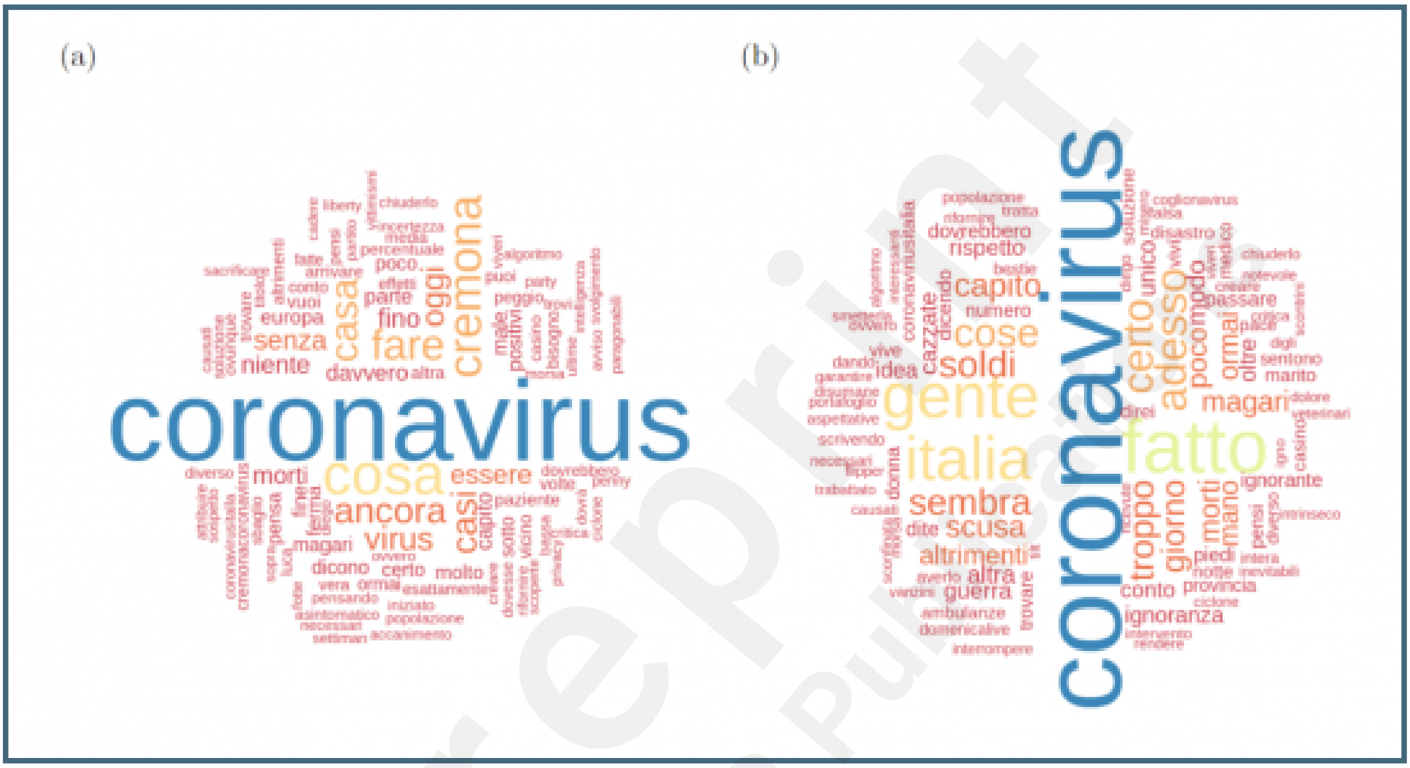

Multimedia Appendixes

Supplementary Material 1.
URL: http://asset.jmir.pub/assets/738b7fa1191ebb477466fa5cf4c4a16b.pdf