CESifo WORKING PAPERS

10803
2023
November 2023

Refueling a Quiet Fire: Old Truthers and New Discontent in the Wake of Covid-19

Gabriele Beccari, Matilde Giaccherini, Joanna Kopinska, Gabriele Rovigatti

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editor: Clemens Fuest
https://www.cesifo.org/en/wp
An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: https://www.cesifo.org/en/wp

CESifo *Working Paper No. 10803*

Refueling a Quiet Fire: Old Truthers and New Discontent in the Wake of Covid-19

Abstract

This paper investigates the factors that contributed to the proliferation of online COVID skepticism on Twitter across Italian municipalities. We demonstrate that socio-demographic factors are likely to mitigate the emergence of skepticism, while populist political leanings were more likely to foster it. Furthermore, we find that the presence of pre-COVID anti-vax sentiment, represented by old "truthers" on Twitter, amplifies online COVID skepticism in local communities. Additionally, exploiting the spatial variation in economic restrictive policies with severe implications for suspended workers belonging to non-essential economic sectors, we find that COVID skepticism spreads more in municipalities significantly affected by this economic lockdown. Finally, the diffusion of COVID skepticism is positively associated with COVID vaccine hesitancy.

JEL-Codes: I120, I180, J280, J600, C800, L820.

Keywords: Twitter, scepticism, public health, media, vaccines, Covid-19.

Gabriele Beccari
Scuola Normale Superiore / Italy
gabriele.beccari@sns.it

Matilde Giaccherini
CEIS, University of Rome "Tor Vergata"
Rome / Italy
matilde.giaccherini@uniroma2.it

*Joanna Kopinska**
DISSE, Sapienza University of Rome / Italy
joanna.kopinska@uniroma1.it

Gabriele Rovigatti
Bank of Italy
gabriele.rovigatti@bancaditalia.it

*corresponding author

We would like to thank Sauro Mocetti and Lorenzo Rovigatti for reviewing previous versions of the manuscript and providing useful suggestions. The views expressed here are our own and do not necessarily reflect those of the Bank of Italy.
Declarations of interest: none.

Introduction

The COVID-19 pandemic has presented unparalleled challenges to global health, demanding swift and well-coordinated responses. In this critical context, a significant threat to public health actions is the dissemination of inaccurate or deceptive information, which has the potential to erode trust and undermine adherence to crucial health protective measures Del Vicario et al. (2016), West and Bergstrom (2021).

Since the onset of the pandemic, there has been a considerable degree of skepticism concerning both the existence and the severity of the novel coronavirus. Some individuals have advanced theories positing SARS-CoV-2 as a weaponized pathogen developed in a laboratory or as a result of radiation emitted by 5G towers. It is worth noting that, apart from the extreme cases of conspiracy theories, this misinformation has also cultivated more moderate forms of skepticism regarding the suitability and efficacy of public health policies Krohnert et al. (2023). The extent to which such skepticism has proliferated may also be influenced by cognitive biases associated with various socio-demographic characteristics Bavel et al. (2020). Importantly, empirical research has demonstrated a positive association between individuals who have previously embraced a conspiracy theory and their inclination to endorse other unrelated ones (Goertzel, 1994, Wood et al., 2012). Consequently, it is likely that individuals who were affiliated with the anti-vaccination movement prior to the COVID-19 pandemic were more predisposed to harbor skepticism regarding COVID-19 and the corresponding vaccination efforts.

In this study, we investigate the origins and repercussions of COVID-19 skepticism in Italy, a nation that faced early and severe impacts during the pandemic. The unique timing of the pandemic onset in Italy enables us to trace sentiments surrounding COVID-19 from its earliest origins, when the virus was relatively unknown and its existence and management were minimally politicized. This renders Italy a valuable case study shedding light on the dynamics of public opinion during a significant public health crisis. Our analysis specifically explores the influence of public health interventions, socio-demographic factors, and pre-existing anti-vaccination sentiments (referred to as 'old truthers') in predicting skepticism and disbelief. We leverage Twitter data to gauge skepticism and assess sentiment towards public policies. During the pre-pandemic period, our focus is on tracing the propagation of anti-vaccination sentiment, while in the context of COVID-19, we examine tweets expressing skepticism about the pandemic. Unraveling the underlying factors contributing to this skepticism, especially when individual beliefs diverge from public health endeavors, can provide valuable guidance to policymakers seeking to devise more effective strategies for mitigating the divisive consequences of unpopular policy choices. While our study is rooted in the Italian context, its findings may bear relevance for other countries with similar socio-economic principles, levels of development, and political dynamics, providing insights into the generalizable aspects of online skepticism during a public health crisis.

Our focus is on the first wave of the COVID-19 pandemic, which extended from February to July 2020. This temporal choice allows us to explore COVID skepticism in a context where public sentiment towards the pandemic was in its nascent stages, and the progression of the epidemiological situation exhibited quasi-random patterns. Moreover, public health policies were consistently implemented at the national level, helping to mitigate potential selection biases.

First, we analyze the primary predictors of COVID-related skepticism at the municipal level, with a specific focus on socio-demographic, economic, and political factors. By providing a comprehensive description and examination of inequality in skepticism during the COVID-19 pandemic in terms of education and income, we offer a framework for understanding potential disparities in compliance with public health policy measures which are also likely to explain the actual disparities in health outcomes. From this standpoint, our contribution aligns with the extensive and continually developing body of literature, which demonstrates that, all else being equal, higher educational achievement and income levels are associated with improved health and longevity. This relationship has been well-documented in influential seminal papers, including works by Baker et al. (2017), Jackson et al. (2022), Montez et al. (2019). Furthermore, we evaluate the influence of pre-existing anti-vaccination sentiments expressed on Twitter before the pandemic on the likelihood of harboring skepticism about COVID. Our findings reveal that COVID skepticism was more prevalent in municipalities where a higher propensity to endorse pediatric anti-vaccination beliefs existed. Under the assumption that the rate of skepticism can approximate the degree of mistrust in the government, our findings also speak to the literature on the impact of trust and social capital on compliance with public policy prescriptions (Ananyev et al., 2021, Bargain and Aminjonov, 2020, Brodeur et al., 2021, Durante et al., 2021).

Secondly, we measure the causal impact of economic restrictive policies on COVID skepticism, by leveraging the spatial variation in the effects of these measures. On March 22, the government enforced a nationwide lockdown, resulting in significant economic implications for all employees in *non-essential* economic sectors. The non-uniform distribution of non-essential economic sectors across the nation, coupled with the consistent enforcement of these restrictions, resulted in an unequal distribution of economic impacts, which, in turn, fostered varying levels of dissatisfaction and skepticism. This natural experiment allows us to gauge to what extent policy-induced discontent fuelled COVID mistrust. On a theoretical ground, personal discontent and political alienation are linked to an increased likelihood of supporting conspiracy theories (Abalakina-Paap et al., 1999, Bruder et al., 2013, Goertzel, 1994, Nyhan and Zeitzoff, 2018). Our research confirms that the economic difficulties resulting from COVID-19 containment measures played a substantial role in fostering COVID skepticism and eroding trust in public health policies. These effects were particularly pronounced in municipalities that not only experienced economic hardships but were also characterized by a higher prevalence of individuals with pre-existing anti-vaccination sentiments ("old truthers").

To demonstrate the tangible influence of skepticism on adherence to public health recommendations, we establish a robust correlation between the regional diffusion of COVID skepticism on Twitter and disparities in the initial acceptance of COVID vaccines. These results underscore the possible unintended consequences of rigorous containment measures on public compliance to health guidelines.

This research makes a significant contribution to multiple facets of the public health field and the effectiveness of public policies. Firstly, we demonstrate that skepticism towards public health measures is an additional factor that can amplify existing disparities in health based on education and income. Furthermore, we reveal that, even after accounting for socio-demographic determinants, a predisposition to online skepticism is likely to foster mistrust across various dimensions, potentially impacting compliance with health protection measures advocated by public health policymakers.

The timeline of the events

During the first wave of the pandemic, Italy declared the state of national emergency on January 31, 2020, and on February 23 implemented the first “red zone,” restricting access to and exit from eleven municipalities with high rates of infection - see Figure 1.[1] Following a sharp rise in the number of infections and hospitalizations due to COVID-19, the Italian government

[1]The quarantined municipalities were Bertonico, Casalpusterlengo, Castelgerundo, Castiglione D’Adda, Codogno, Fombio, Maleo, San Fiorano, Somaglia, Terranova dei Passerini, Vo.

implemented national restrictions on cultural events, sports contests, and schooling through a Decree of the President of the Council of Ministers (DPCM) on March 9. On March 22, the government imposed a nationwide lockdown and reduced the population's freedom of movement. The latter was only allowed for those working in the *essential* sectors, given that the law also mandated an immediate halt to all commercial and industrial activities, but for those sectors considered essential for keeping the economy operative and sustaining the population during the pandemic. Crucially, the distribution of essential sector workers across Italian municipalities is rather heterogeneous. Hence, depending on the pre-determined structure of economic activities, the economic impact of the lockdown varied widely across the country, in direct relationship with the number and the share of essential workers resident. The distinction between *essential* and *non-essential* (*suspended*) economic activities is provided in Appendix A.[2]

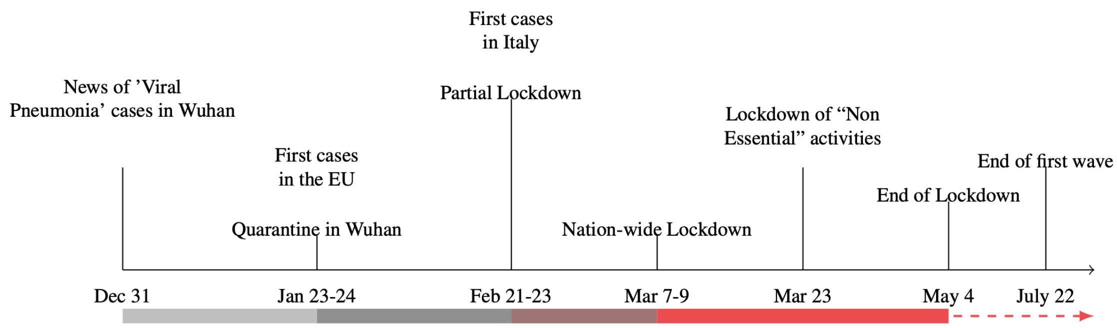

Fig. 1. The evolution of COVID-19 pandemic in the first wave of contagion (Italy, 2020). Red areas correspond to partial/national lockdown periods.

As Italy was the first European country to be significantly affected by the COVID-19 pandemic, its experience provides insights into the challenges and consequences that may arise when dealing with a novel and highly contagious virus. The initial state-of-emergency declarations in Italy faced skepticism from the public and policymakers, with some Italian politicians even engaged in public handshaking to downplay economic concerns due to the virus. The absence of clear and transparent communication was also observed in other countries, where the decision-making process at many levels of government was slow, fragmented, and inconsistent. Therefore, lessons learned from Italy can be valuable for other countries in crafting strategies in response to a rapidly spreading pandemic in a complex and diverse society.

[2] We use the updated essential sectors' list published on a Decree published on March 25, instead of the original, three-day lived list contained in the March 22 DPCM.

Data and measures

This study draws upon a variety of data sources, encompassing Twitter data regarding skepticism, socio-demographic features of Italian municipalities, political election data, economic sector information related to essential activities, COVID mortality, mobility, and vaccination rates, and other supplementary variables. The subsequent section provides a comprehensive account of these data sources and outlines the methodology for constructing the principal variables employed in our analysis.

COVID Skepticism on Twitter. We used the Twitter Academic Application Programming Interface (API) to collect all available tweets in Italian containing COVID-related keywords (such as "`coronavirus`" and "`covid`") for the year 2020, for a total of 17,318,849 tweets. According to Kim (2022), Twitter data are particularly suitable to measure and track public attitudes towards policy-relevant topics over time and across different locations.

Each retrieved tweet encompasses the tweet's plain text and distinctive identification details, which comprise the tweet ID, creation date, and counts of replies, likes, mentions, retweets, hashtags, and multimedia content. Where applicable, additional tweet-specific information, including location and user details such as the user ID, Twitter handle, display name, brief bio, verification status, Twitter join date, and location were also incorporated. Furthermore, user metrics such as the count of friends, followers, and posted tweets were included in the dataset. A sample of the objects obtained for each tweet is available in the appendix (refer to Figure 7).

To geo-localize the tweets, we use three types of information. First, we gather the geotag information of each tweet.[3] Second, for the tweets with missing geotags, we exploit the user geotag.[4] Finally, for any remaining tweets without geotag information, we use place name dictionaries found in the users' bio[5] and pin their tweets to the appropriate municipality based on Italy's administrative boundaries. Out of the initial 17,318,849 tweets, we geolocalize a total of 6,758,177 tweets, with 3,209,030 being original content (not retweets).

To label the tweets according to their COVID skeptic content, we use a Natural Language Processing (NLP) transfer learning model. NLP models allow us to build task-specific models on top of large pre-trained neural networks. This enables us to leverage the power of models

[3] It is known as the "place field" object, and includes the latitude and longitude from the bounding box.
[4] This variable is known as the "location" object, and may include information such as "Rome - Italy".
[5] These may contain the "description" object, such as 'I live in Rome"

trained on billions of documents and to provide the "last mile" of data to fine-tune the model for our specific task. We train the Italian version of Google's BERT (Bidirectional Encoder Representations from Transformers) model, as in AlBERTo (Polignano et al., 2019), and we fine-tune it to label the COVID-skeptic tweets. To classify tweets according to COVID-skepticism, we first train the neural network on a set of tweets labeled manually as either COVID-skeptic ($l_i = 1$) or regular COVID tweets ($l_i = 0$). As our goal is to identify skeptical tweets that express lend ground to scientifically unfounded viewpoints, as opposed to legitimate concerns about the pandemic and its associated phenomena, we construct our training dataset by selecting tweets from established fake news outlets, in line with the methodology employed by Pierri et al. (2020). Lastly, we apply the trained model to the remaining tweets, enabling us to predict the level of leniency $l_i \in 0, 1$ for each individual tweet i. Technical details on the creation of the training and test sets can be found in Appendix C.

Following the approach of Cinelli et al. (2020) and Giaccherini et al. (2022), we aggregate COVID-related tweets to determine rates of COVID skepticism prevalence at the municipality-week level, as the average leaning of all COVID-related tweets authored by users situated within the municipality's boundaries. The stances on COVID are assumed to be embedded in a one-dimensional space $CovidS_{mw} \in [0, 1]$. Following the same procedure, as also done in Giaccherini et al. (2022), we additionally retrieve the average municipality-level Twitter stance concerning skepticism towards pediatric vaccines in the pre-COVID period ($VaxS_i \in [0, 1]$). The measure is constructed through a process similar to that of COVID-related tweets. However, in this case, the algorithm is fine-tuned to identify skepticism related to vaccines. The descriptive statistics presented in Table 1 shed light on the degrees of skepticism associated with COVID-19 and vaccination (Vax skepticism) across municipalities, along with Twitter engagement metrics during the pandemic. In *Panel a*, the median COVID skepticism is 0.51, signifying a moderate level of skepticism, accompanied by a mean of 0.55 and a noteworthy range spanning from 0.33 to 0.89. In contrast, Vax skepticism in 2018 presents a lower mean of 0.36, indicating a relatively reduced skepticism level in the pre-pandemic period. Moving to *Panel b*, the Twitter metrics for COVID tweets unveil substantial variations in user engagement. Metrics like User, Retweet, Replies, Likes, and Quotes exhibit diverse activity levels, with a notably elevated mean Retweet value of 27.20, highlighting significant information dissemination. Follower and Friend counts accentuate the diverse levels of influence across municipalities, with median and mean values reflecting varying Twitter network sizes. Importantly, all figures are weighted by municipality population size, ensuring a proportional representation in the data.

	Median	Mean	SD	Min	Max	Obs.
Panel a: Main variables						
COVID skepticism	0.51	0.55	0.18	0.33	0.89	56,430
Vax skepticism (2018)	0.34	0.36	0.19	0.00	1.00	2,313
Panel b: Twitter metrics of COVID tweets						
User	1,263.50	5,487.86	18,543.99	1.00	664,377.00	56,430
Retweet	0.43	27.20	248.89	0.00	9,332.00	56,430
Replies	0.00	0.18	2.95	0.00	353.00	56,430
Likes	0.00	1.54	44.98	0.00	17,405.33	56,430
Quotes	0.00	0.04	1.28	0.00	534.00	56,430
Followers	62.56	404.54	5,131.25	0.00	2,031,074.00	56,430
Friends	48.56	256.68	1,473.71	0.00	125,886.00	56,430

Table 1. Descriptive statistics of Twitter data. Panel (a) COVID skepticism refers to 6,758,177 geo-located tweets, across 1640 municipalities and 52 weeks, and Vax skepticism is the municipality-level measure observed in 2018. Panel (b) refers to the average number of COVID tweets. All numbers are weighted by the municipality population size.

COVID policy stringency and contagion. We establish the degree of local stringency in COVID mitigation policies by focusing on the economic lockdown measures. Specifically, following the approach in Borri et al. (2021) and Di Porto et al. (2022), we construct a measure of the local impact of the COVID economic lockdown, by considering the heterogeneity in the distribution of essential workers across various economic sectors within Italian municipalities.

Specifically, we calculate the municipality-level proportion of workers impacted by the economic lockdown (referred to as "suspended workers") by combining data on the number of workers at the NACE three-digit sector level from the 2018 "Registro Statistico delle Unità locali" and the list of sectors affected by the lockdown, as outlined in the March 25 Decree.[6] Subsequently, we compute the share of the affected population by dividing the number of suspended workers (aged 15 and above) by the total municipality population. Additionally, we construct a binary indicator variable, taking value one for municipalities where the share of workers affected by the economic lockdown exceeded the national average, weighted by the population size of each municipality. Table 2 shows the metrics on economic and health dimensions during the pandemic. The share of suspended workers reveals a mean of 0.47, highlighting a significant portion of the workforce affected by economic disruptions, with a range from 0.04

[6] We rely on the 2018 data from the "Registro Statistico delle Unità locali" (ASIA), which represents the most recent dataset available for assessing the number of workers in each municipality in Italy across various NACE sectors. This data remains suitable for our purposes as no substantial changes transpired within the Italian economy during the subsequent year in 2019.

to 0.91. Excess mortality indicates relatively low overall mortality rates (mean value 0.10), but with considerable variability from -1.48 to 5.49. Mobility, computed at the province level between January 25 and June 17, 2020, exhibits a mean of 1.82, signifying diverse levels of mobility restrictions across regions during that period. These statistics, all weighted by municipality population size, provide nuanced insights into the multifaceted impacts of COVID policies.

	Median	Mean	SD	Min	Max	Obs.
Share of suspended workers	0.46	0.47	0.09	0.04	0.91	3,152
Excess mortality	0.01	0.10	0.46	-1.48	5.49	56,430
Mobility	1.79	1.82	1.30	0.14	5.89	27,093

Table 2. Descriptive statistics of COVID policy stringency data. Worker-related variables refer to the municipality-level measure observed in 2018. *Mobility* is computed at the province level between January 25 and June 17, 2020. All numbers are weighted by the municipality population size.

COVID mobility and excess mortality. Following Pepe et al. (2020), To measure the extent of individual mobility from January to June, we utilize the weekly users' average radius of gyration by province, made available through the Facebook – Data for Good program.

Additionally, we generate excess mortality data during the COVID pandemic period using the official dataset of deaths registered at the municipality level, sourced from the Italian National Statistics Office (ISTAT), spanning from January 1 to June 30, 2020. With this data, we calculate the difference in the weekly number of deaths at the municipality level in 2020 compared to the historical average for the years 2015 to 2019.

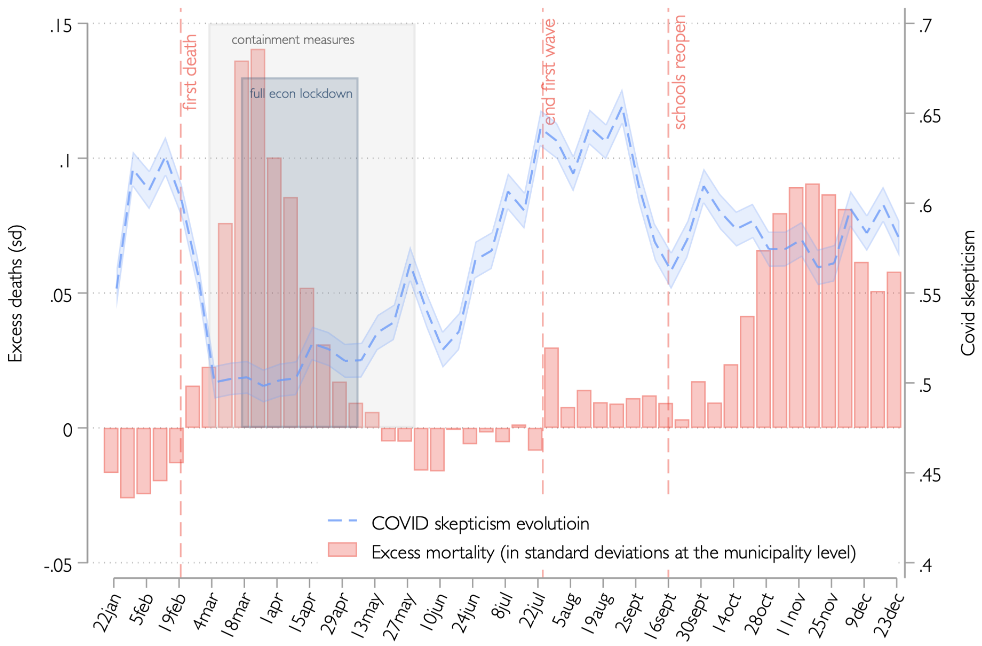

Fig. 2. The evolution of Twitter COVID skepticism and excess mortality (Italy, 2020). The blue line shows the OLS coefficient estimates of Twitter COVID skepticism, along with the 95% confidence intervals, weighted by population size (on the right-hand axis). The model is constructed as an event study, regressing Twitter COVID skepticism at the municipality/week level on 49 week-specific dummy variables starting from the last week of January, when the first COVID cases in Italy were detected, and ending in the last week of 2020. The estimation sample is an unbalanced panel of 1,639 municipalities for 56,397 observations. The red bars plot their weekly excess mortality in 2020 with respect to the average of the same week in the period 2015-2019 (in standard deviations, left-hand axis). The clear and dark gray areas represent the periods of i) containment measures (March 4 to June 2) and ii) full economic lockdown (March 22 to May 4), respectively.

Figure 2 presents the evolution of the pandemic and Twitter skepticism in 2020. The red bars show the excess mortality (expressed in standard deviations) in the subset of municipalities defined by the Twitter data, measured with respect to the average of the same week in the period 2015-2019 on the left-hand axis. The blue line displays the trend in the average Twitter COVID skepticism along with the 95% confidence interval obtained through an event study Ordinary Least Squares (OLS) estimation on the right-hand axis. The clear and dark gray areas represent the periods of containment measures (March 4 to June 2) and full economic lockdown (March 22 to May 4), respectively. Figure 2 illustrates the countercyclical evolution of skepticism in relation to the epidemic. Initial mistrust about COVID decreased as the first casualties emerged but reemerged after the first wave of infections ended in the summer 2020, before dropping again with the arrival of the second wave.

Socio-demographic characteristics Table 3 offers insights into the demographic and political variables across municipalities. We collect data on the share of the population with secondary educational attainment at the municipality level from the official ISTAT Census of 2011 (mean value 0.43), which represents the most recent detailed source available for such information at the municipal level. Furthermore, we use municipality-level per-capita annual income data from the Ministry of Economy and Finance for the year 2019. Per capita income, measured in thousands, showcases the economic diversity within municipalities, ranging from 2.72 to 22.67.

	Median	Mean	SD	Min	Max	Obs.
Share of secondary education	0.42	0.43	0.08	0.18	0.74	3,136
Per capita income (in thousands)	7.84	7.47	2.14	2.72	22.67	3,152
Populist (governing)	0.15	0.18	0.10	0.00	0.55	3,152
Populist (not governing)	0.30	0.31	0.11	0.05	0.69	3,152

Table 3. Descriptive statistics of other data. All numbers are weighted by the municipality population size.

Political variables We gather information regarding the proportion of votes cast for Italian political parties at the municipal level during the 2019 European Parliament election from the Ministry of Interior. We specifically define two variables that reflect the proportion of votes given to populist parties that are in government (Movimento 5 Stelle) and out of government (Lega Nord), respectively (Aassve et al., 2018). In Table 3, the summary statistics reveal that the governing party has a mean of 0.18. Similarly, the non-governing counterpart exhibits a mean of 0.31.

COVID vaccination rates and other relevant data We obtain COVID vaccination coverage rates for the first year of the vaccination rollout, spanning from January to December 2021, from the Italian regional vaccination census known as the "Anagrafe Vaccinale." Throughout this period, vaccination efforts effectively covered nearly all citizens in Italy, accumulating valuable information and individual experiences pertaining to COVID infections and vaccines.

We assess coverage at two distinct time points: one after the first 6 months of the vaccination campaign and another after 12 months. Our primary focus is on the first doses administered since they most directly reflect the decision to pursue immunization. In Table 4, descriptive statistics provide insights into various vaccine-related variables across 20 regions, shedding light on the vaccination landscape. Notably, the mean vaccine coverage at 6 months stands at 0.58, indicating a considerable proportion of the population receiving vaccination during the early stages. This coverage increases to 0.79 at 12 months, reflecting a significant uptick in vaccination rates over time. Concurrently, COVID skepticism, with a mean of 0.58, demonstrates a nuanced public sentiment towards vaccination efforts. The Vax skepticism variable, with a mean of 0.37, suggests a discernible level of skepticism within the population.

	Mean	SD	Min	Max
Vaccine coverage (6 months)	0.58	0.03	0.49	0.61
Vaccine coverage (12 months)	0.79	0.02	0.74	0.81
COVID skepticism	0.58	0.04	0.53	0.64
Vax skepticism	0.37	0.02	0.28	0.41
Public health exp. (pc in thousands)	1.84	0.08	1.67	2.07
Share of LHAs with prevention units	0.99	0.05	0.50	1.00
Share of C-sections	0.29	0.05	0.21	0.55
Elderly index	1.71	0.25	1.30	2.57

Table 4. Descriptive statistics of vaccine data. All numbers are weighted by the region population size. The sample refers to 20 regions.

We also gather additional relevant regional variables related to the supply side of the Italian healthcare system. These include the regional per-capita public health expenditure , as recorded in 2018 (mean value 1.84), the proportion of public Local Health Authorities (LHAs) with a dedicated prevention department (mean value of 0.99 with a variability from 0.50 to 1.0), and the percentage of Cesarean sections (C-sections) performed in 2018 (mean value 0.29). The C-sections serve as a proxy for assessing healthcare system appropriateness since high and unwarranted C-section rates are often associated with financial incentives, concerns about malpractice, or shortages of nursing staff.

Furthermore, we create an elderly index, calculated as the ratio of the resident population aged 65 or older to the young population (0-14). This index, with a mean of 1.71, reflects the demographic composition, underscoring the significance of age-related considerations in public health planning. All values are weighted by the region's population size.

Methodology

We commence our analysis by investigating the relationship between COVID skepticism sentiment during the first wave of the pandemic (February to July 2020) and several municipality-level attributes.[7]

[7] In our context, the "first wave" extends from the declaration of the state of emergency in the last week of January 2020 to the end of July 2020. Others, such as Boschi et al. (2021), define it as the period from February to

Among these characteristics, we consider various socio-demographic factors like educational attainment and per-capita income. Additionally, we empirically investigate whether municipalities with a higher prevalence of anti-vaccine sentiment in the pre-pandemic era exhibited greater skepticism toward COVID. Lastly, we incorporate political affiliation and several COVID-related variables, including the local impact of containment measures, excess mortality, and mobility patterns. As a result, we estimate the following equation:

$$\begin{gathered} CovidS_{FW_m} = Demo'_m\alpha + \beta VaxS_m + Pol'_m\beta + \delta EconLockShare_m + \\ Covid'_{FWm}\zeta + Twitter'_m\sigma + \rho_a + \varepsilon_m \end{gathered} \tag{1}$$

where $CovidS_{FW_m}$ is the degree of COVID skepticism as measured by the average COVID-related Twitter stance in the first wave (FW) of the pandemic at the municipality m level (ranging between 0 and 1), while $Demo'_m$ is a vector of municipality socio-economic characteristics, namely share of the population with secondary educational attainment and per-capita income (linear and square), $VaxS_m$ is a municipality-level measure of skepticism related to vaccines observed in 2018 (also ranging between 0-1), Pol'_m reports the share of votes obtained by populist parties, either governing or at the opposition, $EconLockShare_m$ is the share of suspended workers due to the economic lockdown introduced after March 22, 2020, $Covid'_{FWm}$ is a vector of variables describing the local COVID burden in the first wave of the pandemic, including excess mortality and mobility. Finally, $Twitter'_m$ controls for a number of Twitter metrics, including the average number of Twitter users, followers, and followings (friends), retweets, replies, likes, and quotes, while ρ_a are geographic macro-area fixed effects accounting for systematical differences between the North-East, North-West, Center, South, and Islands. Finally, the error term is clustered at the local labor market level (LLM), while the estimates are weighted by municipal population size. For notational convenience, we suppress subscripts a for the municipality's macro-area (5 in total) and s for the municipality's LLM (610 in total).[8]

In the second part of the empirical analysis, we shift our focus to investigate the causal impact of the economic lockdown on COVID skepticism. The implementation of the national mandate to suspend non-essential economic activities varied among municipalities, even within the same province. This variation in policy enforcement provides an ideal setting to discern the

April.

[8]LLMs are sub-regional geographical areas where the bulk of the labor force lives and works—they are developed through an allocation process that creates the LLMs based on the analysis of commuting patterns (https://www.istat.it/en/labour-market-areas).

causal impact of contentious policy measures on skepticism toward public policy and, consequently, the likelihood of compliance with public health measures. In our pursuit of identifying the causal effect, we employ a distributed lags difference-in-differences design. We thus compare local average COVID skepticism with the share of the suspended population due to the economic lockdown above and below the national average before and after the introduction of the lockdown. For this purpose, we estimate the following model:

$$CovidS_{mw} = EconLock_m \times \sum_{w=-5|w\neq-1}^{20} \beta_w I_w + Covid'_{mw}\zeta + \\ Twitter'_{mw}\sigma + \theta_m + \delta_w + \phi_r \times w + \varepsilon_{mw} \tag{2}$$

where $CovidS_{mw}$ is the average COVID skepticism stance of municipality m in a week w, $EconLock_m$ is an indicator dummy variable taking value one for municipalities where the share of workers affected by the economic lockdown was greater than the national average. It is interacted with the indicator variables $I(w = -5, 20)$, i.e., 5 anticipatory effects (leads) and 20 post-treatment effects (lags). We also control for mobility and COVID weekly excess mortality in the $Covid'_{mw}$ vector, for time-varying Twitter metrics at the municipal level in $Twitter'_{mw}$, and for weekly trends at the regional level, $\phi_r \times w$. Again, we suppress subscript r for the municipality's region (20 in total) for notational convenience. Finally, we include week and municipality fixed effects (θ_m δ_w, respectively). Again, the idiosyncratic error term is clustered at the LLM level, and the estimates are weighted by municipal population size.

Additionally, we also test if municipalities showing a stronger anti-vax sentiment pre-COVID were more COVID-skeptical, complementing the analysis with an indicator variable ($VaxSm$) in the interaction term that takes value 1 for municipalities where the anti-Vax views were greater than 0.5, and 0 otherwise.

Finally, in order to examine the implications of online skepticism on compliance with public health recommendations, we estimate the correlation between COVID skepticism and COVID vaccine coverage. In carrying out this exercise, we are limited by the availability of COVID vaccination data, accessible only at the regional level. This represents an important limitation as the Italina regions are only 20. We nevertheless aim to understand if the COVID skepticism registered in 2020 and the general vaccine skepticism measured in 2018 are likely to explain differences in the share of the population immunized by the novel vaccines after 6(12) months from the beginning of the campaign. We thus estimate the following model:

$$VaccineC_r^t = \beta^t CovidS_r + \delta^t VaxS_r + HC_r' \zeta^t + \gamma^t X_r + Twitter_r' \sigma^t + \varepsilon_r^t \quad (3)$$

where $VaccineC_r^t$ is COVID vaccine coverage in the region $r = 1, 20$ in the first t months from the vaccination introduction ($t = 6, 12$), $CovidS_r$ is the regional average COVID skepticism in 2020, $VaxS_r$ is the regional level vaccine skepticism referring to pediatric and elderly vaccines in 2018, HC_r' is a vector of the regional healthcare system characteristics measured in 2018, such as per-capita public health expenditure, the share of public LHAs with prevention units, and the share of C-sections. Finally, X_r is the 2018 regional elderly index, and $Twitter_r'$ controls for the Twitter average number of followers, friends, and replies, while ε_r is the idiosyncratic error term, and the estimates are weighted by regional population size.

Results

Drivers of COVID skepticism. In the first part of the analysis, we examine the most important predictors of COVID skepticism at the municipality level. The estimates are shown in Figure 3. The dependent variable and most of the covariates (unless otherwise specified) range between 0 and 1, and the estimates show how much the dependent variable changes when the predictor changes by 1 unit, conditional on other covariates.

First of all, we find the rate of excess mortality - measured in hundred deaths - with respect to the same week in the previous year and the extent of mobility feature a negative association with the rate of skepticism (-7.8 and -3.8 pp, respectively). These coefficient estimates align with the idea that municipalities more directly affected by the pandemic toll were more likely to comply with mobility restrictions and less likely to spread skepticism.

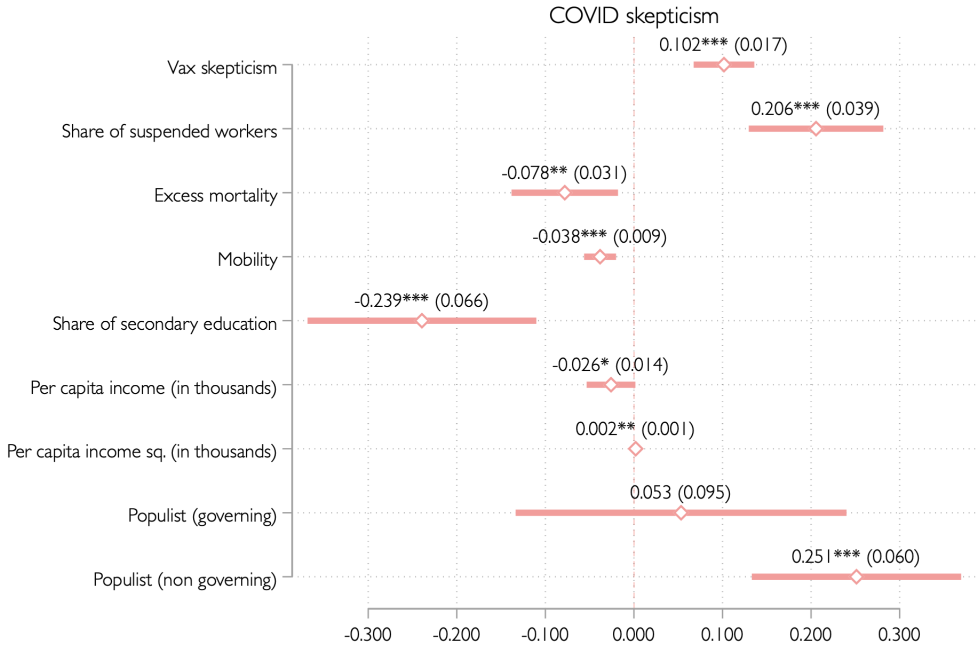

Fig. 3. **COVID skepticism in the first wave and municipality-level characteristics.** COVID skepticism is the municipality-level average over the first wave of the pandemic (Feb-Jul). Also, the other time-varying covariates, such as excess mortality, mobility, and Twitter metrics, are measured over the first wave of the pandemic. The share of suspended workers is defined for sectors enlisted in the March 22 D.L. n. 6/2020 decree on the economic lockdown in function of the distribution of workers by economic sector in 2018 provided by the National Statistical Office. Education and per capita income are sourced from the 2011 municipality-level census and 2019 municipality-level income data, respectively. The municipality-level support for populist parties in power (the 5-star movement) and out-of-government parties is based on the proportion of votes received at the 2019 European Parliament election. Using OLS regression, we estimate the association between COVID skepticism and these covariates, with the coefficients weighted by the municipal population size. Standard errors are clustered at the Labor Market Area level.

More importantly, we find that the share of individuals with completed secondary education at the municipality level is negatively associated with COVID skepticism (-23.9 pp). Interestingly, we also identify a statistically significant U-shaped relationship between COVID skepticism and the average per-capita income at the municipality level. When we account for other covariates, we observe that municipalities with the lowest and highest per-capita income exhibit pronounced levels of COVID skepticism. Considering the income distribution among the municipalities, the coefficient estimates indicate that, after controlling for other characteristics, skepticism reaches its minimum point around the average value of per capita income for the Italian municipalities. Beyond that threshold, skepticism tends to increase for municipalities with higher income levels. Our evidence might present an important contribution to the literature that provides empirical evidence on the beneficial role of both education and income for health outcomes (e.g. Baker et al., 2017, Jackson et al., 2022, Montez et al., 2019). We show that in the case of attitudes towards public health policies, the gradient in terms of socio-demographics characteristics is not monotonic, which is also consistent with several papers that study preventive behavior in terms of vaccination or compliance with stay-at-home orders Drummond and Fischhoff (2017), MacDonald et al. (2015).

It is important to stress that our measure of skepticism by construction addresses the extent to which scientifically unfounded posts propagated by famous fake news outlets were locally shared and distributed and does not refer to posts regarding a healthy debate about the benefits and costs of COVID-related policy measures, also expressed in the economic literature (e.g. Allen, 2022, Mader and Rüttenauer, 2022). From this perspective, the results of our study should be read in line with the literature on misinformation and vaccine hesitancy. Moffitt et al. (2022) show that individuals who embrace conspirational forms of skepticism are frequently characterized by a well-defined cognitive profile that makes them susceptible to such contents. In fact, our evidence supports the hypothesis that pre-existing anti-vaccine sentiment creates a fertile ground for the emergence of COVID skepticism (Hornsey et al., 2022, Romer and Jamieson, 2020), as indicated by a positive coefficient estimate of 10.2 percentage points (pp).

Additionally, we show that the political affiliation of the municipalities matters, as the share of municipality votes for the non-governing populist political parties exhibited a positive and statistically significant relationship with skepticism ($+25.1$ pp). On the other hand, the share of votes for the ruling populist party (i.e., the 5-star movement) was not associated with skepticism, which is likely to reflect the support for the Government decisions. The politicization of the COVID-related debate has also been confirmed by Rathje et al. (2022), who show that vaccine hesitancy is associated with a systematic tendency to interact with low-quality social-media content as well as conservative-leaning in the US. In a similar fashion, Gollwitzer et al. (2020) find a strong relationship between pro-Trump voting US counties and lower compliance with a number of public health recommendations and orders, such as social distancing, stay-at-home orders, and subsequently worse health outcomes in terms of infection and mortality.

Importantly, we also show that the unequal burden of the economic lockdown significantly contributed to the spread of conspiracy theories, with a positive coefficient of 20.6 pp. This evidence is in line with de Figueiredo et al. (2021), who examine the effects of the introduction of vaccine passports for domestic use and to facilitate international travel for UK residents. Their analysis shows that passports are likely to have had a polarizing effect on the attitude towards routine immunization programs, discouraging those individuals with low overall trust in authorities and the Government due to the curtailing effect on individual freedom. They suggest that by alienating critical minorities, it may be actually possible to attain an overall decrease in inclination to vaccinate.

The impact of economic lockdown. To investigate the impact of restrictions on COVID skepticism in a causal setting, we leverage a difference-in-differences strategy (DiD), comparing the change in local average COVID skepticism between municipalities with low ($treatment\ group$) and high ($control\ group$) proportions of essential workers. The uniform implementation of the restriction on non-essential economic activities across the country resulted in an unequal distribution of economic impacts, which was directly related to the percentage of the population working in essential sectors. As a result, neighboring municipalities that were ex-ante identical faced somewhat different economic consequences in an unpredictable way (i.e., as good as random), leading to heterogeneous levels of discontent and disbelief.[9]

Figure 4-(a) shows COVID skepticism trends across treated and control municipalities. After an initial decreasing trend in both groups, the two series start to diverge—even though not significantly in the period between the first nationwide restrictions and the full economic lockdown. During the six weeks of economic lockdown, the treated municipalities immediately feature a steady, increasing trend that is not matched by the dynamics in the control counterpart, which is relatively flat. Interestingly, after the end of the most rigorous emergency measures, skepticism plummet in control municipalities while remaining relatively high in treated locations. From July onwards, when the number of infections, hospitalizations, and deaths reached the lowest level since March 2020, the differential between the two series vanish. In Figure 4-(b), we plot the estimated DiD coefficients and their 95% confidence intervals contrasted with the average excess mortality. In line with the ballpark graphical evidence, the analysis confirms that imposing the economic lockdown caused an increment in COVID skepticism between 3 and 5 pp in treated municipalities, that such effect lasted until July 2020, and that it was not directly linked to the dynamics of excess deaths.

[9]The use of DiD offers the additional advantage of netting out from the estimation any effects on skepticism that are independent of the treatment. If the imposition of stringent policies generated a strong response in skepticism across e.g. younger generations, such an increase would not affect the parameters as long as there is no correlation between the average age and the proportion of essential workers across provinces.

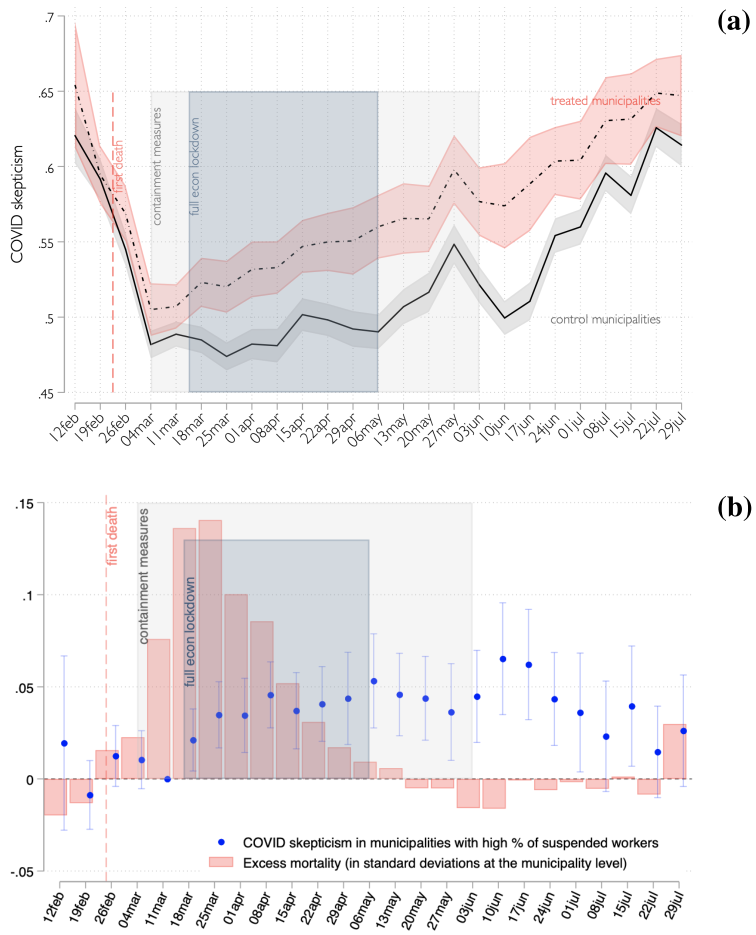

Fig. 4. The impact of economic lockdown on COVID skepticism (a) reports predicted values of COVID skepticism based on a distributed lags DiD design. The graph compares local average COVID skepticism stances coefficient estimates for treated (dash-dot line) and control municipalities (solid line). The treatment is defined as unity for municipalities where the share of suspended workers was greater than the national average. Blue plot (**b**) reports the DiD coefficients between treated and control municipalities compared with the excess mortality expressed in standard deviations at the municipality level. The regression controls for Twitter metrics at the municipal level (average number of users, followers, followings (friends), retweets, replies, likes, and quotes), weekly trends at the regional level, and week and municipality fixed effects. The error term is clustered at the Labor Market Area level, and the estimates are weighted by municipal population size. The estimation sample refers to 1,639 municipalities over 25 weeks.

The propensity for skepticism, mistrust, and conspiratorial thinking is likely to be a consistent individual trait, as indicated by studies like Hornsey et al. (2022), Moffitt et al. (2022), Romer and Jamieson (2020). As further outlined in Hornsey et al. (2022), at the individual level, various cognitive, clinical, motivational, personality, and developmental factors may predispose individuals to engage in conspiratorial thinking. Conversely, at the aggregate level, economic, political, cultural, and socio-historical contexts can also systematically influence individuals' inclination to embrace different forms of skepticism and mistrust. Hence, it is plausible that the distribution of anti-vaccination sentiment among municipalities at the outset of the lockdown serves as a significant source of heterogeneity in the impact of policy-induced discontent.

To capture the heterogeneous effect on the new discontent in the treated municipalities across different pre-exisitng levels of skepticism towards pediatric vaccines, we define an indicator variable taking value one for municipalities with pre-COVID anti-vaccination sentiment above the median. We then augment our DiD setting with the interaction term between the new binary variable and the treatment status for municipalities with a high share of suspended workers. Figure 5 reports the estimated parameters of interest alongside their standard errors. As expected, the pre-existing anti-vaccination sentiment determines a higher level of COVID skepticism (+2.8 pp). We observe a similar treatment effect in municipalities that were particularly exposed to the economic lockdown following its implementation, resulting in a +2.8 pp increase in skepticism. More interestingly, the treatment effect nearly doubles in municipalities with initially high anti-vax sentiment. In these areas, where the economic lockdown was experienced to a more significant extent, the coefficient estimate indicates an additional 2.4 pp increase in the local skepticism rate.

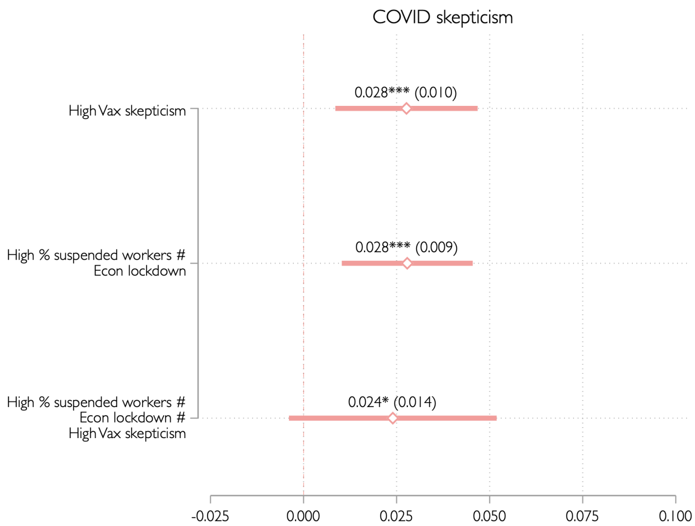

Fig. 5. The impact of economic lockdown and anti-Vax views pre-COVID on COVID skepticism. COVID skepticism is the municipality-level average over the first wave of the pandemic (Feb-Jul). The treatment is defined as unity for municipalities where the share of suspended workers was greater than the national average. The share of suspended workers is defined for sectors enlisted in the March 22 D.L. n. 6/2020 decree on the economic lockdown in function of the distribution of workers by economic sector in 2018 provided by the National Statistical Office. High vax skepticism is a dummy variable that takes value 1 for municipalities where the anti-vax skepticism (pre-COVID) was greater than the median value. Dots are DiD coefficients weighted by the municipal population size with 95% confidence intervals clustered at the Labor Market Area level.

Vaccine hesitancy. From the first phases of the initial phase of the COVID vaccine rollout, significant public policy efforts were dedicated to promoting their safety and effectiveness, aiming to maximize acceptance. However, even with these extensive public outreach campaigns and mandates in place, a noteworthy segment of the population continued to opt from the immunization.

In this context, our objective is to investigate whether the online COVID skepticism, as measured in our analysis, is associated with variations in the uptake of COVID vaccines during the early stages of their introduction. Within an OLS framework, we establish correlations between COVID vaccine coverage observed in the initial months of the campaign and several key factors, including the COVID skepticism indicator, pre-pandemic anti-vax sentiment, and additional variables designed to control for the characteristics of the local healthcare system that might influence the administration of COVID vaccinations. Specifically, these control variables comprise regional-level data encompassing local healthcare expenditure, the proportion of Local Health Authorities (LHAs) equipped with prevention units, the share of Caesarean sections (C-sections) in relation to total births, and the elderly index. The regional level represents the most detailed level available for COVID vaccine coverage data.

Figure 6 illustrates that a higher level of COVID skepticism is associated with a reduction in vaccine coverage. Specifically, a 1 pp increase in COVID skepticism corresponds to a 0.4 pp (0.3 pp) decrease in the coverage of the first dose of anti-COVID vaccination during the first 6(12) months of the campaign. Notably, pre-COVID vaccine skepticism exhibits an even stronger correlation, albeit noisier, with COVID vaccine uptake in the initial 6 months of the campaign, resulting in a decrease of 0.7 pp. However, this correlation loses statistical significance after 12 months of the campaign.

Moreover, COVID vaccine coverage rates demonstrate a positive association with local public healthcare expenditure and the presence of prevention units in Local Health Authorities (LHAs). Conversely, coverage decreases as the level of healthcare quality, as measured by the share of C-sections, and the proportion of elderly residents in the local population increase.

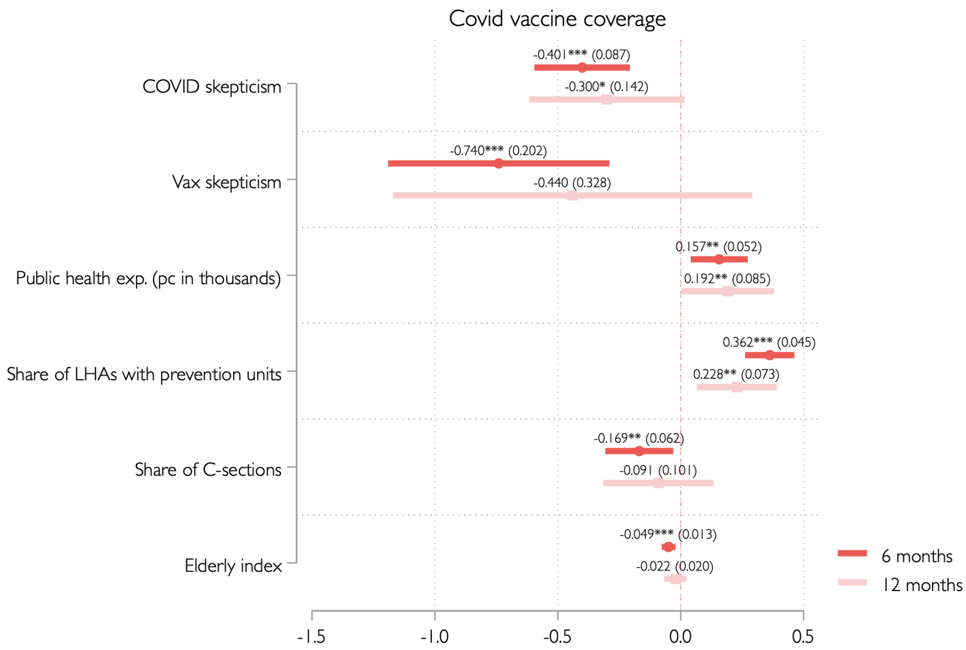

Fig. 6. COVID vaccine coverage, COVID skepticism, vaccine skepticism, and healthcare provision The figure reports OLS coefficient estimates (weighted by the regional population size) together with the 95% confidence intervals. The red gradient refers to estimates carried out on the coverage measured in the first 6 (vivid) and 12 (bland) months of the vaccination campaign. All regressions control for Twitter metrics at the regional level measured in 2020. COVID skepticism is the 2020 regional average. The correlation sample consists of all 20 regions.

It is important to emphasize that our results do not indicate an actual detrimental effect of skepticism on health. Instead, what we are measuring is the impact of skepticism on adherence to policy recommendations. It's worth noting that while the vaccination of vulnerable groups likely reduced the incidence of death and severe illness within those populations, it's possible that achieving high coverage among high-risk individuals led low-risk individuals to make a rational decision to opt out of vaccination campaigns. In this scenario, lower coverage rates may not necessarily indicate inappropriate immunization but could reflect a rational decision-making process. In fact, as highlighted by de Figueiredo et al. (2021), vaccination mandates may have contributed to mistrust, resulting in coverage being too low among some vulnerable groups and too high in the general population.

Conclusions

The COVID-19 pandemic provides a fertile ground for the proliferation of misinformation and diffrent forms of skepticism as individuals attempt to comprehend the complex and often unclear events of the world. Old and new "truthers" flood the social media platforms to share their speculations about the origin of the coronavirus, global plots to restrict freedom, and the reasons behind government-imposed lockdowns, mask mandates, and vaccine requirements. The increased usage of social media during the stay-at-home orders further facilitated the spread of COVID skepticism and related theories.

Using a dataset of Italian Twitter posts paired with municipality characteristics, our research suggests that socio-demographic factors are likely to mitigate the emergence of skepticism, while pre-existing anti-vaccine sentiment or populist political leniency were more likely to foster it. Additionally, our difference-in-differences analysis shows that the unequal distribution of economic consequences resulting from lockdown measures during the COVID-19 pandemic played a significant role in fueling skepticism. Restrictive policy measures implemented during the COVID-19 pandemic generated discontent, which was exacerbated in areas where pre-existing skepticism towards pediatric vaccines existed. These findings align with previous research on the non-linear effects of controversial issues in promoting social media polarization (de Figueiredo et al., 2021, Giaccherini et al., 2022).

Research has shown that individuals who opt out of COVID vaccination tend to be younger, female, and have low trust in government and science Callaghan et al. (2021), Galasso et al. (2021), Khubchandani and Macias (2021), Lazarus et al. (2021), Schwarzinger et al. (2021). In the pre-COVID era, Broniatowski et al. (2018) argued that the spread of vaccine misinformation on social media fuels the vaccination debate, fosters negative opinions about vaccines, and may even contribute to the resurgence of diseases like measles (Burki, 2019, Wilson and Wiysonge, 2020). Others have asserted that online misinformation can pose significant health risks and contribute to vaccine hesitancy (Gallotti et al., 2020, Loomba et al., 2021). However, the precise relationship between vaccine skepticism and COVID-19 vaccine uptake rates is not entirely clear. In a recent analysis, (Pierri et al., 2022) demonstrated that lower COVID-19 vaccine uptake rates in the United States may actually be driven by online misinformation on platforms like Twitter and Facebook.

We contribute to the debate within a novel empirical framework, leveraging a state-of-the-art natural language processing model and detailed socio-economic data. In this context, we investigate shifts in skepticism and their relationship with COVID vaccine uptake in Italy over the initial 6 and 12 months of the campaign, at the regional level. Our findings reveal a strong negative correlation between vaccine coverage and COVID skepticism, as well as with pre-existing anti-vax sentiment. While COVID coverage rates may not solely stem from inadequate coverage among high-risk individuals but rather reflect a rational choice made by lower-risk segments of the population, they serve as a proxy for compliance with public health policy recommendations.

Our results contribute to the literature investigating the impact of the digital revolution on health-related outcomes. While numerous studies have emphasized the positive effects of broadband access, mobile phones, and the internet on reproductive health and protective behaviors (e.g. Bellou, 2015, Billari et al., 2020, Guldi and Herbst, 2017, Pesando, 2022, Rotondi et al., 2020), our research highlights that social media, acting as an amplifier of intentions, may also exacerbate a wide range of potentially undesirable choices from a policy perspective and promote mistrust within the population. Finally, our findings indicate that coercion in policy implementation may inadvertently lead to reduced overall compliance with public health recommendations.

References

Aassve, A., G. Daniele, and M. Le Moglie (2018). Never forget the first time: The persistent effects of corruption and the rise of populism in italy. *BAFFI CAREFIN Centre Research Paper* (2018-96).

Abalakina-Paap, M., W. G. Stephan, T. Craig, and W. L. Gregory (1999). Beliefs in conspiracies. *Political Psychology 20*(3), 637–647.

Allen, D. W. (2022). Covid-19 lockdown cost/benefits: A critical assessment of the literature. *International Journal of the Economics of Business 29*(1), 1–32.

Ananyev, M., M. Poyker, and Y. Tian (2021). The safest time to fly: Pandemic response in the era of fox news. *Journal of Population Economics 34*, 775–802.

Baker, D. P., W. C. Smith, I. G. Muñoz, H. Jeon, T. Fu, J. Leon, D. Salinas, and R. Horvatek (2017). The population education transition curve: Education gradients across population exposure to new health risks. *Demography 54*(5), 1873–1895.

Bargain, O. and U. Aminjonov (2020). Trust and compliance to public health policies in times of covid-19. *Journal of public economics 192*, 104316.

Bavel, J. J. V., K. Baicker, P. S. Boggio, V. Capraro, A. Cichocka, M. Cikara, M. J. Crockett, A. J. Crum, K. M. Douglas, J. N. Druckman, et al. (2020). Using social and behavioural science to support covid-19 pandemic response. *Nature human behaviour 4*(5), 460–471.

Bellou, A. (2015). The impact of internet diffusion on marriage rates: Evidence from the broadband market. *Journal of Population Economics 28*, 265–297.

Billari, F. C., V. Rotondi, and J. Trinitapoli (2020). Mobile phones, digital inequality, and fertility. *Demographic Research 42*, 1057–1096.

Borri, N., F. Drago, C. Santantonio, and F. Sobbrio (2021). The "great lockdown": Inactive workers and mortality by covid-19. *Health economics 30*(10), 2367–2382.

Boschi, T., J. Di Iorio, L. Testa, M. A. Cremona, and F. Chiaromonte (2021). Functional data analysis characterizes the shapes of the first covid-19 epidemic wave in italy. *Scientific reports 11*(1), 1–15.

Brodeur, A., I. Grigoryeva, and L. Kattan (2021). Stay-at-home orders, social distancing, and trust. *Journal of Population Economics 34*(4), 1321–1354.

Broniatowski, D. A., A. M. Jamison, S. Qi, L. AlKulaib, T. Chen, A. Benton, S. C. Quinn, and M. Dredze (2018). Weaponized health communication: Twitter bots and russian trolls amplify the vaccine debate. *American journal of public health 108*(10), 1378–1384.

Bruder, M., P. Haffke, N. Neave, N. Nouripanah, and R. Imhoff (2013). Measuring individual differences in generic beliefs in conspiracy theories across cultures: Conspiracy mentality questionnaire. *Frontiers in psychology 4*, 225.

Burki, T. (2019). Vaccine misinformation and social media. *The Lancet Digital Health 1*(6), e258–e259.

Callaghan, T., A. Moghtaderi, J. A. Lueck, P. Hotez, U. Strych, A. Dor, E. F. Fowler, and M. Motta (2021). Correlates and disparities of intention to vaccinate against covid-19. *Social science & medicine (1982) 272*, 113638.

Cinelli, M., G. D. F. Morales, A. Galeazzi, W. Quattrociocchi, and M. Starnini (2020). Echo chambers on social media: A comparative analysis. *arXiv preprint arXiv:2004.09603*.

de Figueiredo, A., H. J. Larson, and S. D. Reicher (2021). The potential impact of vaccine passports on inclination to accept covid-19 vaccinations in the united kingdom: Evidence from a large cross-sectional survey and modeling study. *EClinicalMedicine 40*.

Del Vicario, M., A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, and W. Quattrociocchi (2016). The spreading of misinformation online. *Proceedings of the national academy of Sciences 113*(3), 554–559.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Di Porto, E., P. Naticchioni, and V. Scrutinio (2022). Lockdown, essential sectors, and covid-19: Lessons from italy. *Journal of Health Economics 81*, 102572.

Drummond, C. and B. Fischhoff (2017). Individuals with greater science literacy and education have more polarized beliefs on controversial science topics. *Proceedings of the National Academy of Sciences 114*(36), 9587–9592.

Durante, R., L. Guiso, and G. Gulino (2021). Asocial capital: Civic culture and social distancing during covid-19. *Journal of public economics 194*, 104342.

Galasso, V., P. Profeta, M. Foucault, and V. Pons (2021). Covid-19 vaccine's gender paradox. *medRxiv*.

Gallotti, R., F. Valle, N. Castaldo, P. Sacco, and M. De Domenico (2020). Assessing the risks of 'infodemics' in response to covid-19 epidemics. *Nature human behaviour 4*(12), 1285–1293.

Giaccherini, M., J. Kopinska, and G. Rovigatti (2022). Vax populi: the social costs of online vaccine skepticism. *CESifo Working Paper 10184*.

Goertzel, T. (1994). Belief in conspiracy theories. *Political psychology*, 731–742.

Gollwitzer, A., C. Martel, W. J. Brady, P. Pärnamets, I. G. Freedman, E. D. Knowles, and J. J. Van Bavel (2020). Partisan differences in physical distancing are linked to health outcomes during the covid-19 pandemic. *Nature human behaviour 4*(11), 1186–1197.

Guldi, M. and C. M. Herbst (2017). Offline effects of online connecting: the impact of broadband diffusion on teen fertility decisions. *Journal of Population Economics 30*, 69–91.

Hornsey, M. J., K. Bierwiaczonek, K. Sassenberg, and K. M. Douglas (2022). Individual, intergroup and nation-level influences on belief in conspiracy theories. *Nature Reviews Psychology*, 1–13.

Jackson, M. I., E. Rauscher, and A. Burns (2022). Social spending and educational gaps in infant health in the united states, 1998–2017. *Demography 59*(5), 1873–1909.

Johnson, K. and D. Goldwasser (2016, November). Identifying stance by analyzing political discourse on Twitter. In *Proceedings of the First Workshop on NLP and Computational Social Science*, Austin, Texas, pp. 66–75. Association for Computational Linguistics.

Khubchandani, J. and Y. Macias (2021). Covid-19 vaccination hesitancy in hispanics and african-americans: A review and recommendations for practice. *Brain, behavior, & immunity-health 15*, 100277.

Kim, T. (2022). Measuring police performance: Public attitudes expressed in twitter. In *AEA Papers and Proceedings*, Volume 112, pp. 184–87.

Krohnert, K., A. Haslam, T. B. Hoeg, and V. Prasad (2023). Statistical and numerical errors made by the us centers for disease control and prevention during the covid-19 pandemic. *Available at SSRN 4381627*.

Lazarus, J. V., S. C. Ratzan, A. Palayew, L. O. Gostin, H. J. Larson, K. Rabin, S. Kimball, and A. El-Mohandes (2021). A global survey of potential acceptance of a covid-19 vaccine. *Nature medicine 27*(2), 225–228.

Loomba, S., A. de Figueiredo, S. J. Piatek, K. de Graaf, and H. J. Larson (2021). Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa. *Nature human behaviour 5*(3), 337–348.

MacDonald, N. E. et al. (2015). Vaccine hesitancy: Definition, scope and determinants. *Vaccine 33*(34), 4161–4164.

Mader, S. and T. Rüttenauer (2022). The effects of non-pharmaceutical interventions on covid-19 mortality: A generalized synthetic control approach across 169 countries. *Frontiers in Public Health 10*, 820642.

Moffitt, T. E., A. Caspi, A. Ambler, K. Bourassa, H. Harrington, S. Hogan, R. Houts, S. Ramrakha, S. L. Wood, and R. Poulton (2022). Deep-seated psychological histories of covid-19 vaccine hesitance and resistance. *PNAS nexus 1*(2), pgac034.

Montez, J. K., A. Zajacova, M. D. Hayward, S. H. Woolf, D. Chapman, and J. Beckfield (2019). Educational disparities in adult mortality across us states: How do they differ, and have they changed since the mid-1980s? *Demography 56*(2), 621–644.

Nyhan, B. and T. Zeitzoff (2018). Conspiracy and misperception belief in the middle east and north africa. *The Journal of Politics 80*(4), 1400–1404.

Pepe, E., P. Bajardi, L. Gauvin, F. Privitera, B. Lake, C. Cattuto, and M. Tizzoni (2020). Covid-19 outbreak response, a dataset to assess mobility changes in italy following national lockdown. *Scientific data 7*(1), 1–7.

Pesando, L. M. (2022). Safer if connected? mobile technology and intimate partner violence. *Demography 59*(2), 653–684.

Pierri, F., A. Artoni, and S. Ceri (2020). Investigating italian disinformation spreading on twitter in the context of 2019 european elections. *PloS one 15*(1), e0227821.

Pierri, F., B. L. Perry, M. R. DeVerna, K.-C. Yang, A. Flammini, F. Menczer, and J. Bryden (2022). Online misinformation is linked to early covid-19 vaccination hesitancy and refusal. *Scientific reports 12*(1), 1–7.

Polignano, M., P. Basile, M. De Gemmis, G. Semeraro, and V. Basile (2019). Alberto: Italian bert language understanding model for nlp challenging tasks based on tweets. In *6th Italian Conference on Computational Linguistics, CLiC-it 2019*, Volume 2481, pp. 1–6. CEUR.

Rathje, S., J. K. He, J. Roozenbeek, J. J. Van Bavel, and S. van der Linden (2022). Social media behavior is associated with vaccine hesitancy. *PNAS nexus 1*(4), pgac207.

Romer, D. and K. H. Jamieson (2020). Conspiracy theories as barriers to controlling the spread of covid-19 in the us. *Social science & medicine 263*, 113356.

Rotondi, V., R. Kashyap, L. M. Pesando, S. Spinelli, and F. C. Billari (2020). Leveraging mobile phones to attain sustainable development. *Proceedings of the National Academy of Sciences 117*(24), 13413–13420.

Schwarzinger, M., V. Watson, P. Arwidson, F. Alla, and S. Luchini (2021). Covid-19 vaccine hesitancy in a representative working-age population in france: a survey experiment based on vaccine characteristics. *The Lancet Public Health 6*(4), e210–e221.

See, A., P. Liu, and C. Manning (2017). Get to the point: Summarization with pointer-generator networks. In *Association for Computational Linguistics*.

West, J. D. and C. T. Bergstrom (2021). Misinformation in and about science. *Proceedings of the National Academy of Sciences 118*(15), e1912444117.

Wilson, S. L. and C. Wiysonge (2020). Social media and vaccine hesitancy. *BMJ global health 5*(10), e004206.

Wood, M. J., K. M. Douglas, and R. M. Sutton (2012). Dead and alive: Beliefs in contradictory conspiracy theories. *Social psychological and personality science 3*(6), 767–773.

A Supplementary Material

Data on share of non/essential workers

In order to determine the share of *suspended* workers in each municipality, we sourced the data on the number of workers at the NACE three-digit sector level from the "Registro Statistico delle Unità Locali" (ASIA-UL).[10] To classify workers as either active or suspended, we use the list of sectors allowed to operate (essential sectors) specified in the Appendix 1 of the March 25 DPCM, which introduced the economic lockdown. Along with the restrictions currently in place under the March 9 DPCM, such as the lockdown of bars and restaurants, shops, schools and universities, sport infrastructure, cultural and leisure activities, and a national stay-at-home order, the March 25 DPCM provided for a general lockdown of all industrial and commercial activities, with some exceptions (essential sectors), such as:

- health;
- agro-food;
- logistic and energy activities;
- financial;
- insurance;
- maintenance and vigilance services;
- activities aimed at ensuring the continuity of the supply chain of the essential sectors' activities;
- activities of plants with a continuous production cycle;
- activities of the aerospace and defense industry and strategically relevant activities for the national economy;
- the production, transport, commercialization and delivery of medical products;
- healthcare technology and medical devices and any agricultural and food supplies;
- any activity useful to manage the emergency.

[10] The most recent information on the number of workers in each Italian municipality is from ASIA-UL data for 2018. The Italian economy did not significantly change between 2018 and the COVID epidemic in 2020, making this data a trustworthy indicator.

B Additional Figures

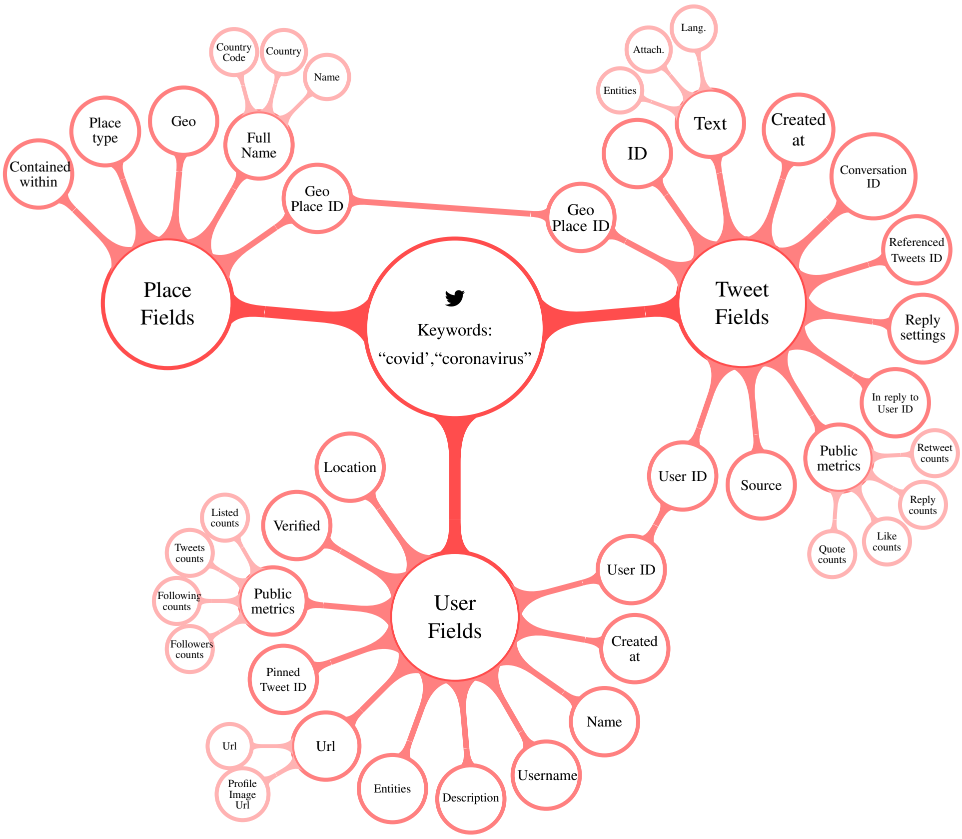

Fig. 7. Twitter objects. The figure shows the structure of tweet objects returned by Twitter API. In particular, they are constituted by three linked entities - featuring several “children” each: i) Tweet, ii) User, and iii) Place fields for each retrieved tweet related to the keywords “covid” or “coronavirus” written in Italian during the 2020 period.