UNIVERSITÀ DI SIENA 1240

QUADERNI DEL DIPARTIMENTO DI ECONOMIA POLITICA E STATISTICA

Francesco Sarracino
Talita Greyling
Kelsey O'Connor
Chiara Peroni
Stephanie Rossouw

Trust predicts compliance to Covid-19 containment policies: evidence from ten countries using big data

n. 858 – Giugno 2021

STATEC RESEARCH

Luxembourg National Research Fund

Trust predicts compliance to Covid-19 containment policies: evidence from ten countries using big data*

Francesco Sarracino[a], Talita Greyling[b], Kelsey O'Connor[c], Chiara Peroni[d], and Stephanie Rossouw[e]

June 30, 2021

Abstract

Previous evidence indicates that trust is an important correlate of compliance with Covid-19 containment policies. However, this conclusion hinges on two crucial assumptions: first, that compliance does not change over time, and second, that mobility and self-reported measures are good proxies for compliance. We demonstrate that compliance changes over the period March 2020 to January 2021, in ten mostly European countries, and that increasing (decreasing) trust in others predicts increasing (decreasing) compliance. We develop the first time-varying measure of compliance, which is calculated as the association between containment policies and people's mobility behavior using data from Oxford Policy Tracker and Google. We also develop new measures of both trust in others and national institutions by applying sentiment analysis to Twitter data. We test the predictive role of trust using a variety of dynamic panel regression techniques. This evidence indicates compliance should not be taken for granted and confirms the importance of cultivating social trust.

*The study is supported by the Luxembourg National Research Fund (grant number FNR-14878312).

[a]Francesco Sarracino, STATEC Research a.s.b.l., Francesco.Sarracino@statec.etat.lu. Telephone: +352 247 88 469. Address: 13, rue Erasme, L-2013, Luxembourg.

[b]Talita Greyling, University of Johannesburg, talitag@uj.ac.za.

[c]Kelsey J. O'Connor, STATEC Research a.s.b.l., IZA, Kelsey.OConnor@statec.etat.lu.

[d]Chiara Peroni, Institute of Statistics and Economics Studies (STATEC) Chiara.Peroni@statec.etat.lu.

[e]Stephanie Rossouw, Auckland University of Technology, stephanie.rossouw@aut.ac.nz.

JEL codes: D91, I18, H12.

Keywords: compliance; covid-19; trust; big data; Twitter.

1 Introduction

The effectiveness of public policies to contain epidemics hinges on people's adherence to the prescribed behaviors, i.e. compliance (Brailovskaia and Margraf, 2020). Low rates of compliance are undesirable because they can hamper the efficacy of public measures to limit contagion; they can lead to increased health care costs because of inefficient and potentially wasteful uses of resources; and induce substantial delays in which time viruses can mutate (Kyngäs et al., 2000; Wright et al., 2021; Lewnard and Lo, 2020; Chen et al., 2020). Hence, compliance with health policies is crucial. Evidence from empirical analysis and mathematical models indicates that measures such as self-isolation, quarantine and social distancing are effective strategies to limit the spread of epidemics (Kucharski et al., 2020; Anderson et al., 2020; O'Connor, 2020b). However, these measures require the cooperation of a large share of the population, they are individually and socially costly, and their psychological toll increases with stringency and length of containment policies (O'Hara et al., 2020; Kim and Jung, 2021; Zhao et al., 2020). Hence, the incentives for not complying are high. Moreover, enforcing containment measures on a recalcitrant population entails high social and economic costs at the expense of people's freedom, confidence in institutions, and sense of social cohesion.

A broad and inter-disciplinary literature on compliance indicates that people vary considerably in their willingness to adapt their behaviors to containment measures, such as those to limit the spread of Covid-19 (Levi and Stoker, 2000; Luttmer and Singhal, 2014; Fan et al., 2020). This observation motivated a number of studies investigating the correlates of compliance, whereby compliance means people's adherence to government's prescriptions. In this paper we analyze the role of trust in promoting compliance with Covid-19 containment policies. A large literature agrees that trust, either in national institutions or in others, facilitates compliance. We contribute to this literature both methodologically and empirically. While most papers use self-reported or mobility data as measures of compliance, we measure compliance as the association between containment policy stringency and individuals' time spent at home. Empirically, we make use of daily data to study the relationship between compliance and trust over time.

Compliance and its correlates are the subjects of a wide range of disciplines, including medicine, nursing, psychology, and health economics. It is, therefore, unsurprising that there is no commonly accepted definition of compliance (Kyngäs et al., 2000). More or less explicitly, however, the majority of empirical investigations on compliance adopted a cognitive-motivational framework which emphasises the relationship between attitudes and intentions towards recommended or prescribed healthcare measures (Cameron, 1996). This framework recognizes the importance of individual and social characteristics in shaping people's willingness to comply and, importantly, it clarifies that compliance should not be seen as a duality between compliance and non-compliance. Previous research documents that individuals' perceptions of the costs and risks associated with non-compliance predicts their behavior (Bish and Michie, 2010; Williams et al., 2015). These perceptions differ among people and over time, especially as the duration of treatment increases. In the case of Covid-19, containment measures lasted more than one year and compliance with containment measures varied widely across time, countries, the number of fatalities, and the strictness of lockdown measures (Becher et al., 2020; Bargain and Aminjonov, 2020).

In medical disciplines, there are three main methods to measure compliance: direct observation of practice (number of episodes over number of opportunities); self-reports, i.e. individual's declaration about their own behavior, usually collected through surveys; and indirect calculation based, for instance, on drug consumption (Haas and Larson, 2007). Direct observation is generally regarded as the preferred measure of compliance, as the alternative methods are prone to more errors. For instance, Dresselhaus et al. (2000) administered a survey to analyze the preventive care measures adopted by physicians in their daily routines. The authors found evidence of overshooting, that is physicians tend to perform more preventive care initiatives than they declare in medical records. Other scholars cautioned that direct survey questions are likely to suffer from measurement errors due to social desirability bias (Barari et al., 2020; Daoust et al., 2020). This is consistent with the observation that surveys in different countries found that self-reported compliance with recommended health prescriptions is high (Perrotta et al., 2021; Brouard et al., 2020). Another method is through survey experiments (Becher et al., 2020). This solution consists in administering a list of activities to respondents and to see how many tasks the respondent declared to have performed over the previous week. Becher et al. (2020) used survey experiments to examine the prevalence of non-compliance with social distancing in nine countries[1]. The authors split the respondents into two groups,

[1]Countries are: Australia, Austria, France, Germany, Italy, New Zealand, Sweden, United Kingdom, and United States.

one of which received a list including an item related to violating the social distancing norm (meeting people when not allowed). Results indicate that on average 25.8% of respondents did not adhere to social distancing guidelines. This share is higher than what is computed using estimates based on direct questions from surveys administered in the same countries, at approximately the same time (Becher et al., 2020; Perrotta et al., 2021).

Unfortunately, survey experiments and direct measures of compliance are not widely available. This is why the vast majority of studies on compliance with Covid-19 containment measures and its determinants used either self-reported declarations of compliance, or mobility data sourced from big data, such as cellphone location or Google Mobility Data. Wright et al. (2021) interviewed a longitudinal sample of 51,000 British adults every week over a period of three months of lockdown (April, 1 to June, 22 2020). Their measure of adherence is based on the following question: "Are you following the recommendations from authorities to prevent spread of Covid-19?". Respondents could reply using a seven-point Likert scale, where 1 stands for "Not at all" and 7 - "Very much so". The availability of repeated, individual-level observations, allowed the team of researchers to account for a wide list of possible correlates. They found that trust in government is the only factor predicting future compliance, whereas other factors such as mental health, confidence in the health care system, social experiences and awareness of Covid-19 were not statistically related to future changes in compliance.

There are various reasons to expect that trust, one of the key components of social capital, facilitates compliance. For instance, trust enables the co-operation necessary to achieve common goals, such as limiting the diffusion of a virus (Fukuyama, 1995; Putnam, 2000). Experimental evidence showed that people's propensity to cooperate increases if they expect that others will do the same (Fischbacher et al., 2001; Shinada and Yamagishi, 2007). Thus, trust alleviates the incentive to free-ride when i. containment measures are individually costly, and ii. individual behaviour has a negligible impact on containing the virus – a typical social dilemma (Ostrom, 2000). Additionally, available evidence shows trust can change in a relatively short time span (Sarracino and Mikucka, 2017; Mikucka et al., 2017). The evidence above suggests these changes can alter the extent to which people choose to comply with containment policies.

The relationship between compliance and trust, both in others and in institutions, has received particular attention in previous research (Pagliaro et al., 2021; Brodeur et al., 2020; Bargain and Aminjonov, 2020). Chan et al. (2020) show that people in European regions where confidence in the health care system is high tend to comply more. Moreover, they find that compliance depends on confidence in government, trust in media and belief in science, as well as moral support, social norms, peer pressure, and a mix of characteristics, including income, risk-taking behavior, personality features,and political orientation. Plohl and Musil (2021) further explored the role of trust in science and Covid-19 risk perception for compliance using a random sample of 525 respondents, and structural equation modelling. The authors found that trust in science and risk perception predict compliance, while the effects of other variables, such as political conservatism, religious orthodoxy, conspiracy ideation, and intellectual curiosity, are mediated by trust in science.

Civic capital, which usually correlates with trust, is another important correlate of compliance (Barrios et al., 2021). The authors analyse American individuals, American counties and European regions. Besides showing that civic capital correlates positively with compliance, the authors also found that social distancing was more likely to stay steady in high civic capital counties, even when it was not mandated by law. Durante et al. (2021) reach a similar conclusion using mobility data across Italian provinces in early 2020. The authors found that mobility decreased more in provinces where civic capital was higher. Predictions indicate that if all provinces shared the same level of civic capital as the top 25%, mortality in Italy would have been 60% lower (Durante et al., 2021).

Using survey data from 23 countries, Pagliaro et al. (2021) document that people's reported compliance relates positively to various forms of trust (in others, in the government, and in science), but not to the number of infections. According to the authors, feelings of fairness and care are at the origin of these relations. This conclusion finds partial support in a study by Nofal et al. (2020) on a sample of about 8,500 Japanese people. The authors observed compliance using a battery of self-reported behaviours in which respondents selected the extent to which they adopted a specific policy (on a scale from 1 to 5, where higher scores indicate more adoption of the behaviour). Results indicate personality traits predict people's compliance. For instance, people high in conscientiousness, openness to experience and agreeableness – which positively correlate with feelings of fairness and care – were more likely to adopt Covid-19 transmission mitigation behavioural guidelines than others. On the contrary, people high in extraversion were less likely than others to comply. In a longitudinal study on a small sample of UK residents, Stevenson et al. (2020) find that community identification, a proxy of the local network of relations and trust, predicted compliance, as measured by the extent to which respondents adhered to Covid-19 containment policies (on a scale from 1 to 5, where higher scores indicate higher compliance).

Additional studies using mobility data support the hypothesis that trust predicts compliance. For instance, Brodeur et al. (2020) merged US cell phone data from Unacast with data on social capital and trust from the General Social Survey (prevalence by county). They found that counties with higher shares of people trusting others showed, on average, a greater decrease in mobility once a containment policy was introduced. Bargain and Aminjonov (2020) reach the same conclusion using Google Mobility Data. The authors use a double difference approach to estimate the impact of trust (at the regional level in Europe) on decreases in mobility around the time of lockdown: high-trust regions decreased their non-neccessary mobility more than low-trust regions.

The main limitation of this burgeoning literature is that the adopted measures of compliance are subject to measurement error. On one hand, self-reports can be up-ward or down-ward biased, depending on the context and respondents. Mobility data, on the other hand, provide an accurate picture of people's movements, but not in relation to containment policies, i.e. compliance. People may choose to stay home even if it is not mandated by law. The few studies exploiting the association between changes in mobility and changes in policies overcome this problem, but at the cost of focusing on a specific point in time, whereas compliance and its correlates may change over time, especially in the case of long-lasting events.

We overcome limitations in previous studies by developing and analysing a time-varying measure of compliance based on the association between mobility and policy stringency. Our measure of mobility is drawn from Google Mobility Data, which provides a direct observation of people's distancing behavior. We extract information about containment policies from the Oxford Policy Tracker. Our measure of compliance reflects how individuals' distancing behavior changed in correspondence with given levels of policy stringency at any point in time. Importantly, this measure is comparable across countries. Our measure is not free from limitations. For instance, it is estimated at the national level, and therefore cannot be used to study subnational heterogeneity.

The paper is organized as follows. The next section introduces our measure of compliance as well as the additional data used in the study. A discussion of the validity of our measure of compliance is available in Appendix C. Section 3 describes the empirical strategy adopted in present study, whereas section 4 illustrates how compliance differs across countries and its changes over time, before discussing the role of trust for compliance. Section 5 concludes.

2 Data

We define compliance as the degree of association between people's behaviours and Covid-19 containment policies. We imagine compliance as a continuum ranging from no compliance at all to perfect association between what is mandated and people's behaviors. In accordance with previous literature, we acknowledge that the degree of compliance changes over time as a consequence of changes of its determinants and, in particular, of trust in others and institutions.

We observe compliance at country level to overcome the measurement limitations of previous studies which relied mainly on self-reported compliance, or changes in mobility (thus observing people's behavior without accounting for the policy). The focus on countries allows us to measure compliance as the association between stringency of containment policies (left-hand side y-axis of figure 1) and the increase in time that people spend at home (right-hand y-axis of figure 1). The changes of the two variables are fairly consistent over the considered period (from March 2020 to January 2021). For instance, figure 1 shows a sudden increase in stringency and the relative share of time people spend at home at the very beginning of the pandemic (March to May 2020). However, the association between the two measures is not constant over time. The end of December 2020 provides a good example: we observe a much stronger increase in the relative share of people staying at home than in policy stringency. In other words, people stayed at home much more than what was mandated. In this case, it is clear that the reason is Christmas, but this helps us illustrating the point that compliance changes over time as a consequence of specific conditions.

Data on policies are the Stringency Index provided by Oxford's Covid-19 Government Response Tracker (Hale et al., 2020)[2], whereas information about the time spent at home are sourced from Google Mobility Reports (Google, 2020). Google provides information based on users' location history, and informs about daily mobility/visitation to various places by geographic location. We focus on the time spent at home as this variable requires less assumptions about people's movements. Data are expressed as relative visits' numbers compared to the number of visits during a period of reference, i.e. January 3 to February 6, 2020. Data on policy stringency and on time spent at home are provided daily throughout the pandemic. We measure compliance as the association between the two lines reported in figure 1, and

[2]The Tracker monitors 18 indicators of policy response to the pandemic. The Stringency Index uses the following nine indicators: school closing, workplace closing, cancel events, restriction of gatherings, close public transport, stay at home requirements, restrictions on internal movement, international travel controls, and public information campaigns.

Figure 1: Average policy stringency and time spent at home across countries over time.

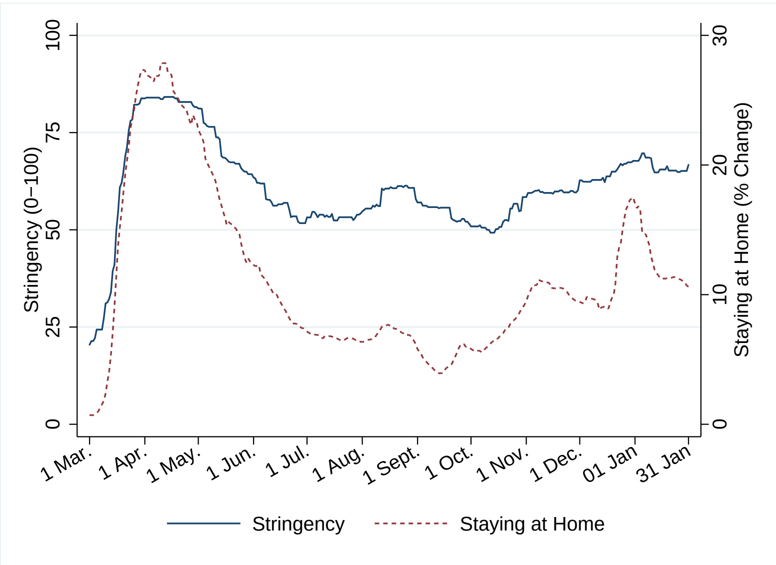

Note: Staying at home data are presented using a seven-day (centered) moving average.

Source: data are sourced from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

it is computed as follows:

$$res_{ct} = \alpha + \beta_{cm} \cdot Country_c \cdot Month_m \cdot Policy_{ct} + \delta_c \cdot Country_c + \lambda_{ws} + \varepsilon_{ct} \quad (1)$$

where res_{ct} is residential mobility in country c in day t; $Country$ is a vector of dummies for each country included in the dataset; $Month$ is a vector of dummies for the m months from February 2020 to January 2021. We focus on this period because prior to February 2020 data on mobility are not available, while containment policies started being in place by the end of February; *lambda* is a vector of dummies for each combination of week w and hemisphere s, to account for the different seasons and evolution of the pandemic among Northern and Southern hemisphere. The coefficient β_{cm} is our measure of compliance. It provides the correlation between policy stringency and mobility by country and month. In principle, we could have estimated compliance on a weekly or daily basis, but policy measures do not change frequently enough, and shorter time periods carry the risk of introducing noise in our correlation. Section C in the Appendix provides evidence supporting the validity of our measure of compliance.

Our measure of compliance is available for nearly every country in the world. However, in present analysis we restrict our attention to ten countries, namely Australia, Belgium, France, Germany, Italy, Luxembourg, New Zealand, South Africa, Spain, and United Kingdom. The reason is that our main explanatory variables, trust in others and in national institutions, are available only for these countries. Figure 2 shows the average monthly levels of compliance for European countries in panel 2a, and for Australia, New Zealand and South Africa, henceforth A-NZ-SA, in panel 2b (dashed lines).

The dashed line for European countries is always positive, indicating a positive association between changes in policy stringency and distancing. In particular, compliance follows a U-shaped trajectory, which reaches a minimum in the month of August, the period when also policy stringency was at its minimum (if we exclude the pre-pandemic period). In A-NZ-SA, compliance declines throughout 2020, more so after August, and it is negative after December 2020. This is a period characterized by few positive cases, and a gradual relaxation of containment measures. However, residential mobility did not closely follow containment policies, as is the case in August 2020 and, in particular, in January 2021. In both months, policy stringency increased, but people did not reduce their mobility accordingly. We focus on the changes of trust in others and in national institutions to explain the differences in compliance across countries and over time.

The measures of trust are compiled by extracting tweets from Twitter, a voluntary online social platform. Tweets are analyzed using sentiment

Figure 2: Compliance over time in two country groups.

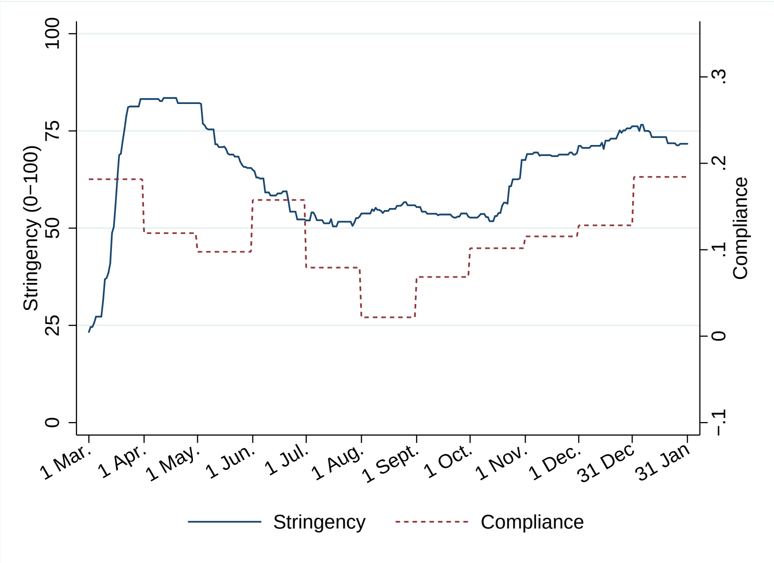

(a) Average compliance in seven European countries.

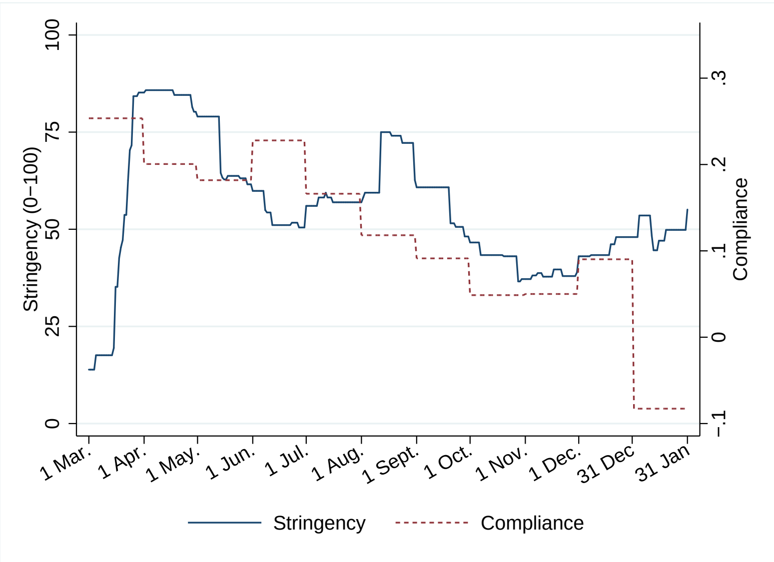

(b) Average compliance in Australia, New Zealand, and South Africa.

Source: own computation on data from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

analysis, an automated process to extract the emotional content of a written text (Hailong et al., 2014). This technique has already been used extensively in various fields of social sciences (Eichstaedt et al., 2015; Caldarelli et al., 2014; Gayo-Avello, 2013; Bollen et al., 2011; Asur and Huberman, 2010). The daily score of trust extracted from sentiment analysis is the average of the trust expressed in all the tweets of a given day. [3] We derive a measure of trust in national institutions by extracting the trust content of tweets that included selected keywords.[4] Sarracino et al. (2021) showed that the measures of trust in others and in national institutions extracted from Twitter correlate meaningfully with corresponding measures of trust issued from survey data. Therefore, the two variables can be regarded as valid measures of trust.

We use sentiment analysis applied to Twitter data to extract also three control variables, namely anticipation, fear and economic fear. Anticipation and fear are extracted from Twitter using the same procedure illustrated for trust. Economic fear results from extracting the fear content from tweets that included keywords related to the economic situation.[5]

Finally, we include in the analysis the daily number of new positive cases of Covid-19 (new confirmed cases per million in population) to account for the evolution of the pandemic. This variable is extracted from Our World In Data (Roser et al., 2020) and it is transformed using an inverse hyperbolic sine function. This transformation is similar to a logarithmic one, but it is identified for zeros.

Table 1 provides summary statistics for the variables included in our study.

[3]For more details about the construction of this variable, please, see Greyling et al. (2021); Sarracino et al. (2021).

[4]The keywords used are: government, parliament, ministry, minister, senator, MPs, legislator, political, politics, prime minister.

[5]The list of keywords includes: jobs, economy, saving, work, wages, income, inflation, stock market, investment, unemployment, unemployed, employment rate, tech start-up, venture capital.

Table 1: Descriptive statistics. Average monthly values.

Variable	Obs	Mean	Std. Dev.	Min	Max
Compliance	100	0.108	0.082	-0.209	0.328
Residential mobility	100	11.232	7.142	0.645	33.067
Stringency	100	62.106	16.542	22.220	95.127
IHS New Cases	100	4.112	2.094	0.165	7.533
Lag Trust	90	6.658	0.822	5.560	9.023
Lag Nat. Trust	81	8.917	1.837	6.769	12.068
Lag GNH	90	6.999	0.426	5.992	7.640
Lag Anticipation	90	5.921	0.722	4.917	7.559
Lag Fear	90	3.620	0.526	2.927	5.090
Lag Economic Fear	81	4.073	1.090	2.803	6.273

Note: Monthly Data April 2020 - Jan. 2021. Compliance begins in March, the first month for which mobility data cover the whole month, and March is dropped to include lags. Fewer observations are available for national trust and economic fear, because they could not be collected for Luxembourg due to data limitations.

Source: all sources are described in the text. They are omitted for brevity.

3 Methodology

The correlation between compliance and trust over time is likely affected by other variables, such as the severity of the pandemic, economic situation, weather conditions, or seasons. To account for the possible confounding effects of these variables, we resort to various econometric techniques.

In principle, we would like to estimate the following equation:

$$Compl_{cm} = \alpha + \rho Compl_{cm-1} + \beta Trust_{cm-1} + \delta IHS(newcases)_{cm} + \lambda_m + \mu_c + \varepsilon_{cm} \quad (2)$$

where $Compl_{cm}$ is compliance as defined in the data section for country c in month m; $Trust_{cm-1}$ is the average level of trust for country c in the month $m-1$. $IHS(newcases)_{cm}$ is the transformed number of new daily cases per one million residents (on average over the month) in country c. λ_m are month effects, while μ_c are country effects. Additional controls are added in robustness tests, discussed below.

This specification addresses many potential sources of bias. Reverse causality is reduced by lagging trust and including lagged compliance as a control. Omitted variables are addressed in part through the inclusion of fixed effects, which account for all fixed-characteristics, both observed and unobserved. Lagged compliance likewise captures anything that affects both current and lagged compliance, both observed and unobserved. Current cases capture the country specific time-varying evolution of the pandemic. Cases in the current period, one period after lagged trust, also capture people's expectations of the pandemic in their country. Common effects across countries, such as seasonal effects and global trends in the pandemic are captured using month controls.

However, this equation cannot be estimated by ordinary least squares (OLS) without bias. Nickell Bias (Nickell, 1981) arises when including both lagged trust and fixed effects.[6] However, equation 2 can be estimated when excluding either the fixed effects or lagged compliance. For this reason, we use fixed-effects (FE) estimation when excluding lagged compliance, and dynamic OLS (DOLS) when including lagged compliance but excluding the fixed effects.[7] Estimates from the two methods should also bound the true

[6]Fixed-effects models are typically estimated by subtracting from each variable its mean value over time, and in the case of a dynamic panel, the mean of the lagged dependent variable is correlated with the mean error term. In other words, de-meaning introduces a source of endogeneity.

[7]Similar approaches were used in (Flèche and Layard, 2017; Krekel et al., 2020; O'Connor, 2020a).

estimate (Angrist and Pischke, 2008, pag. 184).[8]

Under certain conditions it is possible to use an alternative approach to account for both fixed effects and dynamics (Anderson and Hsiao, 1981). To account for fixed effects, the authors apply first differences to equation 2. First differences, however, cause the lagged differenced dependent variable to be related to the differenced error term. To overcome this problem, lagged differenced compliance is predicted in a first stage and then used in a two-stage instrumental variable approach. Equation 3 presents the second-stage specification:

$$\Delta Compl_{cm} = \rho \Delta \widehat{Compl_{cm-1}} + \beta \Delta \widehat{Trust_{cm-1}}) + \delta \Delta IHS(\widehat{newcases_{cm}}) + \Delta\lambda_m + \Delta\varepsilon_{cm} \tag{3}$$

where $\Delta Compl_{cm} = Compl_{cm} - Compl_{cm-1}$ and $\Delta \widehat{Compl_{cm-1}}$ is the predicted value for $\Delta Compl_{cm-1}$. Anderson and Hsiao (1981) suggest $Compl_{cm-2}$ as a valid instrument for $\Delta Compl_{cm-1}$. $Compl_{cm-2}$ is relevant because it is correlated with $\Delta Compl_{cm-1}$ and it is excludable (not correlated with $\Delta\varepsilon_{cm}$) if there is no autocorrelation in the level equation 2 (i.e., $cov(\varepsilon_{cm}, \varepsilon_{cm-1}) = 0$). $\Delta Trust_{cm-1}$ and $\Delta IHS(newcases)_{cm}$ are also allowed to be endogenous and predicted in the same way. We extended this approach by including the additional instruments: $\Delta Compl_{cm-2}$, $\Delta Trust_{cm-2}$, and $\Delta IHS(newcases)_{cm-1}$, which improved the first stage predictions.

This IV approach has limitations. The assumption is that there is no autocorrelation in the level equation, but we know there is serial correlation in each of the variables: compliance, trust, and new cases. Typically, an overidentification test would be used to assess whether the instruments are excludable, but in this case, we have too few clusters for the Hansen test. It is for this reason that we do not emphasize the IV results alone.

There is an additional approach which builds on the work by Anderson and Hsiao, described by Arellano and Bond (1991). Arellano and Bond recognized that further lags could be used as additional instruments. For example, both $Compl_{c1}$ and $Compl_{c2}$ are valid instruments for $\Delta Compl_{c3}$. For $\Delta Compl_{c4}$ even more instruments are available, specifically: $Compl_{c1}$, $Compl_{c2}$, and $Compl_{c3}$. In this way, an additional instrument is added for each period. Estimation of this structure is then performed using generalized methods of moments (see Arellano and Bond, 1991, for more details). The method proposed by Arellano and Bond has the same limitations of the

[8]If fixed effects represent the true data generating process, but a dynamic model is used, then the resulting estimate will be biased downward. However, if the true data generating process is dynamic but fixed effects are used, then the estimate will be biased upward.

method proposed by Anderson and Hsiao, but improves efficiency with the additional instruments. To test for instrument validity, one assesses the degree of autocorrelation in the predicted residuals. It is expected that there is first-order autocorrelation in the predicted residuals from equation 3 due to the mechanical relation between $\Delta Compl_{it-1}$ and $\Delta Compl_{it-2}$, but there should be no second-order autocorrelation for the instruments to be valid. A further limitation of the approach is that the results are often unstable in small samples. In any case, we also include an Arellano and Bond estimation as a robustness test.

We assess statistical significance using Wild Cluster Bootstrap methods. Clustering standard errors at the country level is necessary to allow for serial correlation within countries. Bootstrap methods are necessary because a small number of clusters leads to rejecting the null hypothesis relatively more frequently, in some cases at more than double the critical value (Bertrand et al., 2004). To address this problem Wild Cluster Bootstrap methods are used (using 999 replications). The limitation is that only p-values from the bootstrap distribution can be obtained. Standard errors cannot be estimated using this method because it includes asymptotic refinement (sample estimates approach the population values at a faster rate), which can only be performed on statistics that do not depend on unknown parameters. For this reason, the bootstrapped p-values are reported in the tables. For a further explanation of Wild Cluster Bootstrap methods see Cameron and Miller (2015); when using instrumental variables, see Davidson and MacKinnon (2010); and for implementation using STATA, see Roodman et al. (2019).

4 Results

Figure 3 shows the changes of trust and compliance over time in Europe (panel 3a) and in A-NZ-SA (panel 3b). Compliance follows a U-shaped pattern in Europe which is partly matched by the changes of trust. The first three months of pandemic are characterized by declining compliance and raising trust in others. In June, trust reaches a peak and from that moment onward the changes of trust accompany compliance rather well. Compliance declines from March 2020 to January 2021 in A-NZ-SA. In this case the association with trust is unclear. However, many factors may confound the association between our two measures. We turn to regression analysis in an attempt to single out the net effect of trust on compliance.

Table 2 shows the results of our baseline model. The first column shows the results of the most basic model in which compliance is regressed over the number of new positive cases of Covid-19, and the lagged value of trust

Figure 3: Compliance and trust across countries

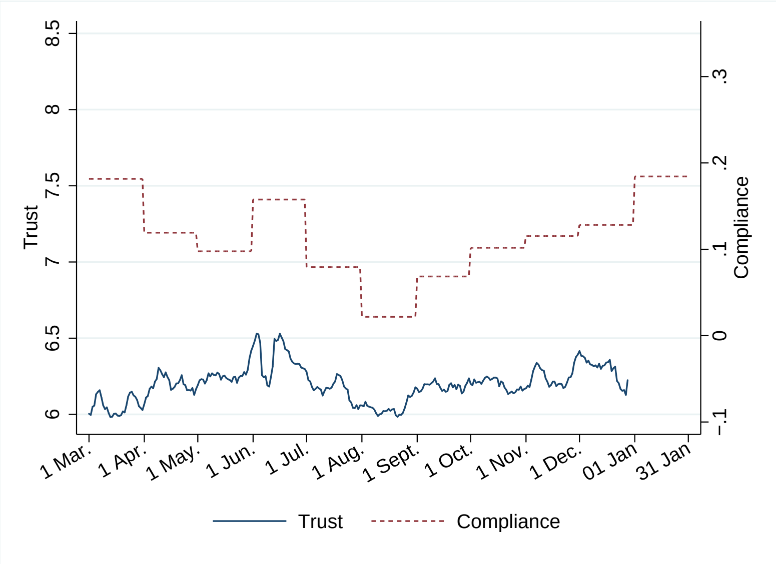

(a) Average compliance in seven European countries.

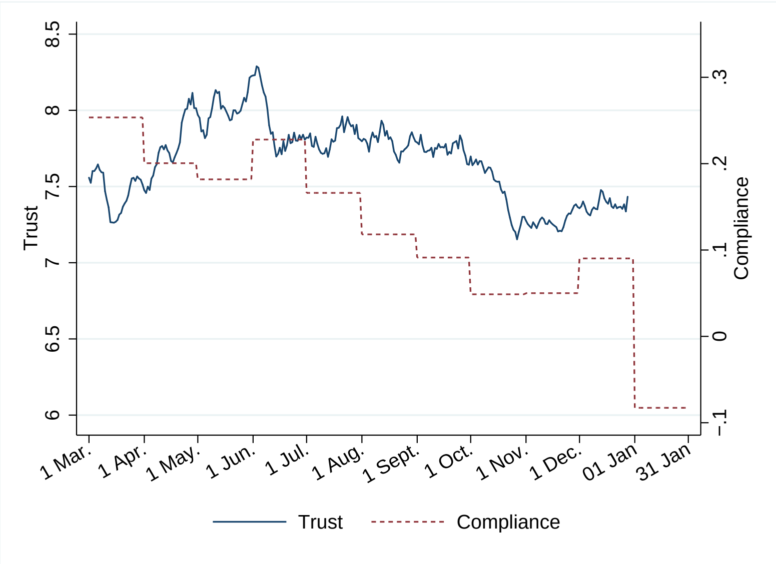

(b) Average compliance in Australia, New Zealand, and South Africa.

Note: Trust data are presented using a seven-day (centered) moving average.

Source: own computation on data from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

($month\, m - 1$). We find positive and statistically significant coefficients for both variables: the higher the number of new positive cases, and the higher the lagged trust, the higher is compliance. The positive association between new cases and compliance suggests that the more severe people perceive the pandemic, the more they tend to follow the required behaviours. This result is based on the analysis of ten countries over a period of ten months, for a total of a hundred observations. The adjusted R-squared is 0.29, indicating that large part of the total variance remains unexplained.

In columns 2 to 7, we run dymanic and fixed effects model for trust and trust in institutions separately. The inclusion of trust in institutions reduces the number of observations to 90, because data for Luxembourg are not available. Despite this decrease, the results qualitatively change very little. In the dynamic models, the coefficient of trust is positive but generally not statistically significant. In the fixed effects models, the coefficient on trust is positive, large and statistically significant. The larger coefficients in each of the FE regressions is consistent with expectations, according to which the true effect of trust should lie between the DOLS and FE estimates (see footnote 8). The coefficient of trust in institutions is small and largely not significant; it turns negative in the models with fixed effects. The coefficient of lagged compliance is large, positive, and statistically significant. The dynamic models explain the largest share of total variance (about 67%).

The last two models (columns 8 and 9) include at the same time trust and trust in institutions. The results confirm that the latter attracts a negative and not statistically significant coefficient, whereas lagged trust, lagged compliance and new cases attract positive and statistically significant coefficients in both the dynamic and fixed effects models (except new cases in the fixed effect model).

This evidence confirms that trust correlates with compliance over time above and beyond the effect of new positive cases, previous compliance, or country fixed effects. However, this does not completely exclude the hypothesis of a spurious relationship driven by a third (time-changing) variable. To address this issue, we run an additional set of dynamic and fixed effects OLS to check the robustness of our finding to the inclusion of controls for economic fear, anticipation and fear (all variables are lagged by one month). Previous results do not change significantly. As reported in table 3, trust retains its positive and statistically significant coefficient in most specifications, as well as lagged compliance and new positive cases (not significant in case of fixed effects models). In addition, we find that the experience of fear or economic fear in the previous month tend to reduce compliance in the dynamic models, whereas lagged anticipation does not attract a statistically significant coefficient.

Table 2: Baseline regression results on the association between correlation and trust in others and in national institutions. Results are from OLS, dynamic OLS and OLS with fixed effects.

	(1) OLS	(2) DOLS	(3) FE	(4) DOLS	(5) FE	(6) DOLS	(7) FE	(8) DOLS	(9) FE
Lag Trust	0.055	0.013	0.115			0.014	0.175	0.026	0.201
	[0.088]	[0.220]	[0.039]			[0.246]	[0.006]	[0.062]	[0.007]
Lag Nat. Trust				0.001	-0.013			-0.007	-0.053
				[0.840]	[0.483]			[0.141]	[0.127]
Lag Compliance		0.795		0.841		0.797		0.798	
		[0.000]		[0.001]		[0.000]		[0.000]	
IHS New Cases	0.029	0.014	0.017	0.012	0.027	0.016	0.016	0.015	0.018
	[0.016]	[0.009]	[0.201]	[0.013]	[0.056]	[0.017]	[0.298]	[0.005]	[0.237]
Constant	-0.398	-0.172	-0.761	-0.089	0.098	-0.184	-1.154	-0.205	-0.878
N	100	100	100	90	90	90	90	90	90
# of Countries	10	10	10	9	9	9	9	9	9
Adj. R Sq.	0.294	0.658	0.38	0.662	0.252	0.67	0.424	0.671	0.454
Bootstrapped p-values in brackets									

Note: dummy variables for months are included, but omitted for brevity.

Source: own computation on data from the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

The joint reading of results from dynamic and fixed effects OLS provides evidence supporting the hypothesis that trust predicts compliance across countries and over time. To check the robustness of our findings, we run an additional model that allows for dyanamics and fixed effects, i.e., we use the two-stage instrumental variable approach described in the Methodology section.

The first three columns of table 4 show the results of the first stage, column 4 reports the results of the second stage, and column 5 shows the results using the alternative technique provided by Arellano and Bond (1991). The number of observations is 80 because we use twice lagged values of the independent variables to predict the lagged, differenced endogenous variables. 2SLS and Arellano and Bond provide fairly consistent results. Lagged differenced compliance remains statistically significant in both specifications. New positive cases remains statistically significant, but the sign changes depending on the regression method used (negative in case of two-stages least squares, positive otherwise). Only trust continues to attract a positive and statistically significant coefficient independently from the model. Additionally, as expected, the magnitude of the A&B coefficient is between the lower and upper boundaries set by the coefficients of the dynamic and fixed effects OLS (see columns 2 and 3 of table 2) and the 2SLS coefficient is near the upper bound.

Table 3: Dynamic and fixed effects models with additional control variables.

	(1) DOLS	(2) FE	(3) DOLS	(4) FE	(5) DOLS	(6) FE
Lag Trust	0.032	0.174	0.018	0.193	0.024	0.073
	[0.004]	[0.006]	[0.350]	[0.053]	[0.003]	[0.040]
Lag Economic Fear	-0.017	0.004				
	[0.003]	[0.877]				
Lag Anticipation			-0.006	-0.124		
			[0.679]	[0.177]		
Lag Fear					-0.026	0.079
					[0.020]	[0.159]
Lag Compliance	0.805		0.783		0.836	
	[0.000]		[0.000]		[0.000]	
IHS New Cases	0.015	0.015	0.014	0.014	0.013	0.017
	[0.008]	[0.387]	[0.010]	[0.339]	[0.023]	[0.215]
Constant	-0.236	-1.16	-0.167	-0.534	-0.152	-0.771
N	90	90	100	100	100	100
# of Countries	9	9	10	10	10	10
Adj. R Sq.	0.675	0.417	0.655	0.404	0.66	0.4
Bootstrapped p-values in brackets						

Note: dummy variables for months are included, but omitted for brevity.

Source: own computation on data from the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

In sum, various regression methods and specifications provide evidence supporting the hypothesis that trust is an important, if not the single most important, predictor of compliance.

Table 4: Two-stages least squares and Arellano & Bond estimates to account for endogeneity in lagged differenced compliance.

	(1) First Stage Lag Δ Compl.	(2) First Stage Lag Δ Trust	(3) First Stage Δ Cases	(4) 2SLS Δ Compl.	(5) A&B Δ Compl.
Lag2 Compliance	-0.337	0.187	1.275		
	(0.05)	(0.365)	(1.172)		
Lag2 Δ Compl.	-0.13	-0.966	-2.184		
	(0.084)	(0.97)	(1.665)		
Lag2 Trust	0.021	-0.114	-0.651		
	(0.008)	(0.074)	(0.095)		
Lag2 Δ Trust	-0.005	-0.215	-0.409		
	(0.006)	(0.102)	(0.245)		
Lag IHS New Cases	0.014	-0.013	-0.289		
	(0.003)	(0.022)	(0.021)		
Lag Δ IHS New Cases	-0.002	-0.006	0.556		
	(0.003)	(0.025)	(0.096)		
Lag Δ Compl.				0.621	0.625
				(0.22)	(0.1)
Lag Δ Trust				0.149	0.062
				(0.037)	(0.026)
Δ IHS New Cases				-0.03	0.015
				(0.018)	(0.004)
Constant	-0.17	0.777	5.561	0.026	
	(0.064)	(0.523)	(0.65)	(0.011)	
N	80	80	80	80	80
# of Countries	10	10	10	10	
Adj. R Sq.	0.627	0.086	0.377	-0.391	
Kleib. F Stat				6.587	
AR1p					0.015
AR2p					0.112
Clustered standard errors in parentheses					

Note: Dummy variables for months are included, but omitted for brevity.

Source: own computation on data from the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

5 Conclusions

Compliance with Covid-19 containment policies is crucial for the effectiveness of public policies to: limit contagion, reduce the opportunites for virus mutations, and limit waste. At the same time it entails heavy individual and social costs. For government policies to be effective, a large share of the population must cooperate by adjusting their individual behaviors to match govenment mandates. Unsurprisingly, previous studies documented that compliance with Covid-19 containment policies varies considerably across countries (Becher et al., 2020; Margraf et al., 2020). To this end, a large number of recent studies have tried to uncover which factors increase compliance (Fan et al., 2020). Many factors have been investigated using a variety of data sources and methods, and the available evidence agrees that trust, either in others or in institutions, enables compliance (Pagliaro et al., 2021; Brodeur et al., 2020; Bargain and Aminjonov, 2020). This conclusion, however, hinges on two hypothesis: first, that self-reports or mobility data are valid measures of compliance; and second, that compliance does not change over time.

Previous research clarified that both hypothesis are fragile. Self-reported measures of compliance suffer from up- or down-ward bias depending on the target population and the kind of survey. Mobility measures are not measures of compliance, as they reveal individual behaviors, but not their correspondence to policies. Finally, the literature on compliance illustrates that people comply to different degrees over time (Bish and Michie, 2010; Williams et al., 2015). To what extent does the violation of these two hypotheses challenge what we know about the role of trust for compliance? We overcome these limitations in the present work.

Our contribution is both methodological and empirical: first, we measure compliance throughout 2020 as the correlation between the stringency of containment policy and the time people spend at home; second, we use trust data sourced from sentiment analysis of Twitter data to check whether it predicts compliance over time and across countries.

The availability of daily data from the Oxford Policy Tracker and Google allowed us to build a time-series of observations on compliance ranging from early 2020 to the end of January 2021 for ten countries. Our list of countries includes: Australia, Australia, Belgium, France, Germany, Italy, Luxembourg, New Zealand, South Africa, Spain, and United Kingdom. We then enriched these data with information on the pandemic, such as the number of new positive cases, and with a set of variables obtained using sentiment analysis applied to Twitter data. Following Greyling et al. (2021), we used sentiment analysis to extract the emotional content of Tweets to derive information about trust, fear, and anticipation. In addition, we used a set of keywords to retrieve information about users' fear about the economy, and users' confidence in national institutions. As far as we know, this is the first time that this kind of variables is retrieved using sentiment analysis applied to Twitter data. The result is a rich dataset of monthly observations about compliance and a number of possible explanatory factors observed from March 2020 to the end of January 2021. Also, as novel data, we assessed the validity of each variable, finding that our measure of compliance correlates meaningfully and significantly with experimental data, such as those provided by Becher et al. (2020). Concerning the sentiment-derived variables, various tests were performed in an earlier paper that provide evidence supporting their validity (Sarracino et al., 2021).

To analyse the relationship between trust and compliance over time and across countries, we used dynamic Ordinary Least Squares, and fixed effects panel regressions, both using with Wild Clusered Bootstrapped error terms. Bootstrapping allows us to account for the small number of countries included in the analysis. Both methods have specific advantages that, jointly, allow us to confirm that trust is a robust correlate of compliance over time and across countries. This relationship also holds after accounting for the role of new positive cases. We further check the robustness of our findings using Two-Stages Least Squares and the Arellano and Bond dynamic panel approach. Both techniques allow us to account for the endogeneity in lagged compliance and trust. We found that trust positively and significantly contributes to compliance across all models. However, the coefficient of trust in national institutions is not statistically significant. In other words, we found that it is the relationships among people that seem to play a major role for compliance followed, to a lesser extent, by the number of new positive cases.

Our study confirms the results of similar previous studies while overcoming some of their limitations; it is not free from drawbacks, however. First, the small number of countries and short period reduces the number of observations, which reduces statistical power and the possible set of control variables. Nonetheless, we feel that our specifications account for many of the possible sources of bias. Our approach focuses necessarily on countries: it does not allow us to study heterogeneity within countries. Finally, sentiment analysis applied to Twitter data is likely susceptible to the number of Twitter users in a country. For instance, in small countries, such as Luxembourg, emotions linked to specific keywords do not provide enough Tweets to derive reliable information.

Acknowledgements

We would like to thank the Luxembourg National Research Fund (FNR) for the generous funding of the project. Additionally, we thank AFSTEREO for the I.T. service provided.

A Policy stringency and time spent at home

Figure 4: Average changes of policy stringency and time spent at home by groups of countries.

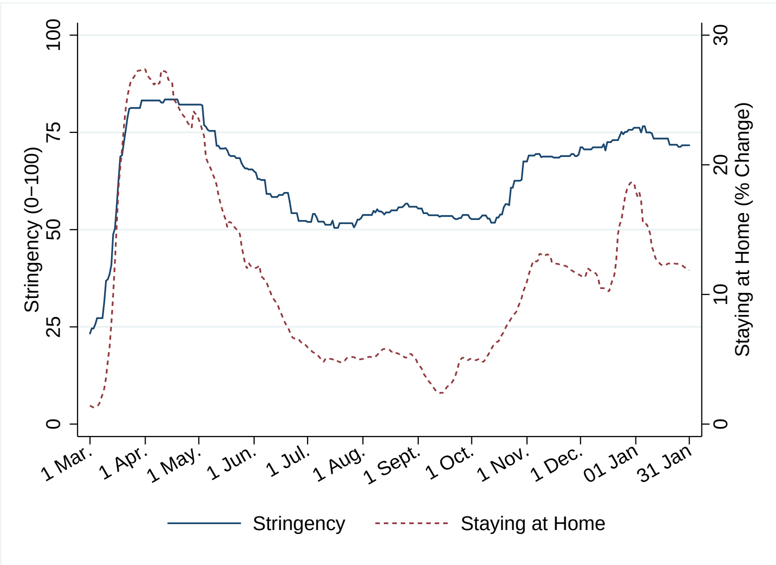

(a) Average changes in seven European countries.

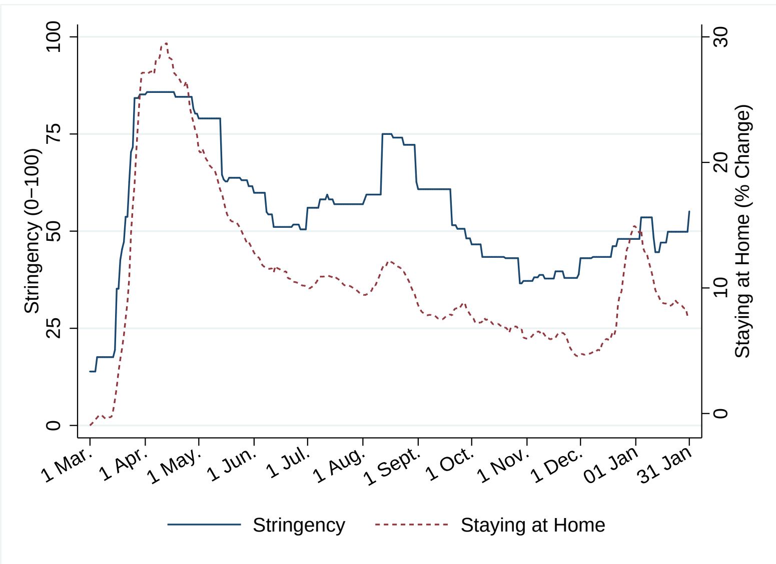

(b) Average changes in Australia, New Zealand, and South Africa.

Note: Staying at home data are presented using a seven-day (centered) moving average.

Source: data are sourced from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

B Compliance by country

Figure 5: Compliance and its two components by country over time.

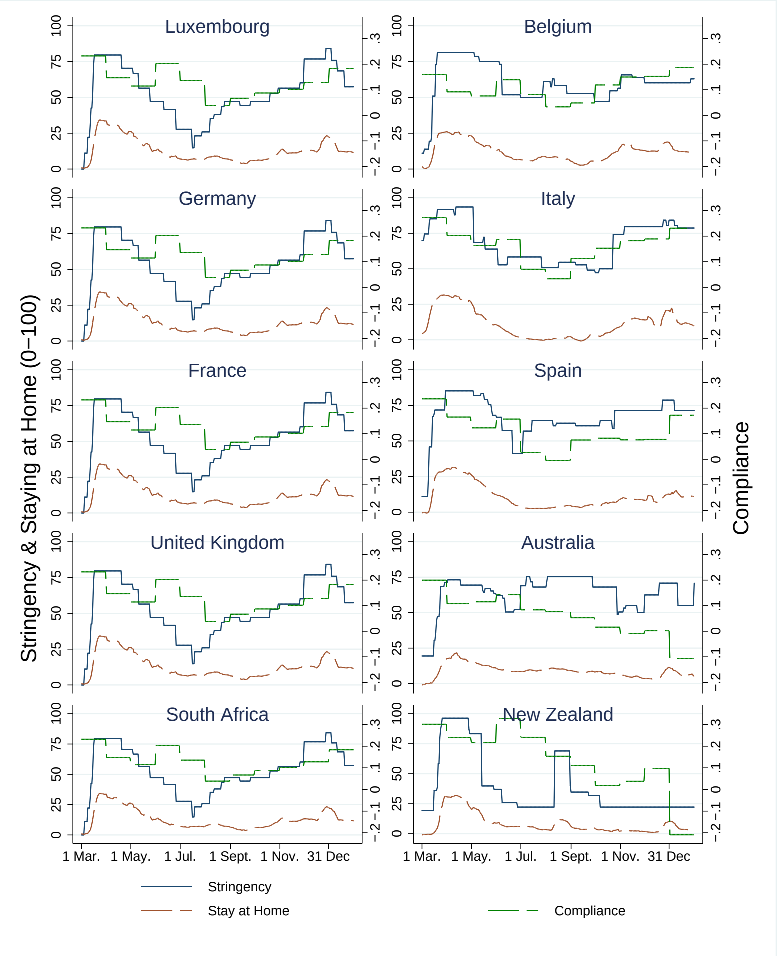

Note: compliance is the association between policy stringency and time spent at home. It is computed following equation 1.

Source: own computation on data from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

C Reliability of compliance

As far as we know, this is the first work that conceptualizes compliance as the association between policy and behaviors, and that uses big data to this purpose. It is therefore legitimate to question the validity of our measure, that is its ability to correctly measure the extent to which people abide by the rules.

A major difficulty to assess the validity of compliance is to find a yardstick to compare to. Unfortunately, other measures of compliance based on the association between policy and behaviors over time are not available. We therefore resorted to comparing yearly compliance with data provided by other sources.

Figure 6 shows the average yearly compliance computed using our data. For instance, if containment policies become more stringent by ten points (0-100) then the time Italians spend at home increases by 3 percent. Assuming base time at home is 12 hours, $0.03*12 = 0.36$ hours $= 21.6$ minutes. In fact stringency went from 0 to more than 80, which would lead to a $80*0.3 = 24$ percent increase in time at home, on 12 hours that's $0.24*12 = 2.88$ hours, on average.

An external source of reliable data on compliance is provided for 9 countries by Becher et al. (2020). The authors used experimental evidence to infer information about people's *non*-compliance to social distancing guidelines during Covid-19 pandemic. We computed the correlation between our measure of compliance and the one by Becher and colleagues. The two measures should be negatively related to each other. The result is reported in figure 7. The two datasets have only six countries in common, and the Pearson's correlation coefficient is -0.88 (significant at 5%, with N = 6 observations). If we increase the number of observations by adding figures about Spain from Margraf et al. (2020), who use self-reported measures of adherence to containment behaviours, the correlation coefficient does not change, but statistical significance improves (p-value = 0.0090, with N = 7 observations).

This evidence provides some support in favor of the validity of our measure of compliance, although only in levels. We were not able to find other measures of compliance over time to test the validity of our measure.

Figure 6: Compliance across countries

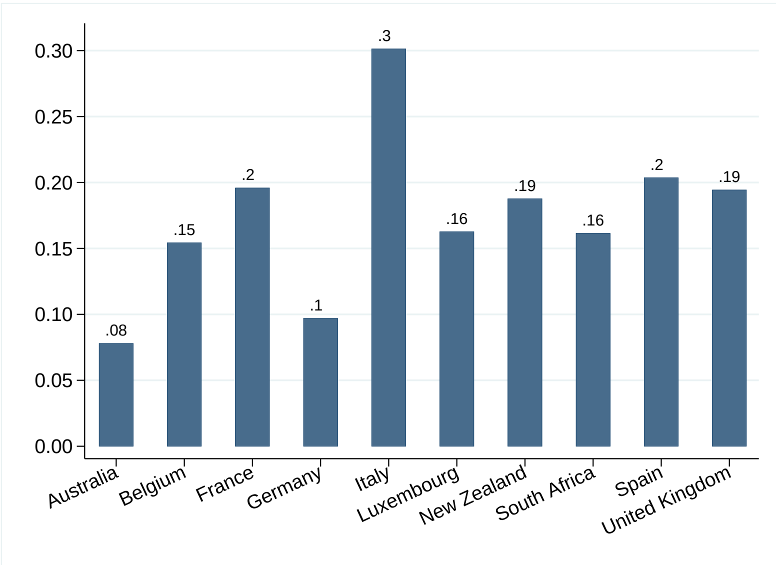

Note: the average country levels of compliance are computed using a modified version of equation 1 whereby the β coefficient is estimated by country without the interaction by month.

Source: own computation on data from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT.

Figure 7: Correlation between two measures of compliance.

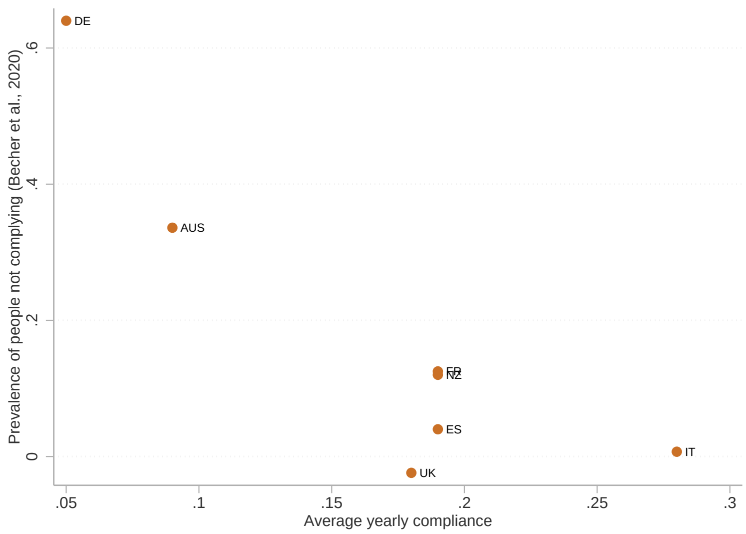

Note: the experimental data from Becher and colleagues have been expanded to add Spain to the set of considered countries using data from Margraf and colleagues.

Source: compliance is computed using data from Oxford Policy Tracker and Google Mobility Data within the framework of the project "Preferences Through Twitter" with the support of FNR, UJ and AUT. Non-compliance is issued from the paper by Becher et al. (2020), and by Margraf et al. (2020).

References

Anderson, R. M., Heesterbeek, H., Klinkenberg, D., and Hollingsworth, T. D. (2020). How will country-based mitigation measures influence the course of the COVID-19 epidemic? *The lancet*, 395(10228):931–934.

Anderson, T. W. and Hsiao, C. (1981). Estimation of dynamic models with error components. *Journal of the American statistical Association*, 76(375):598–606.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297.

Asur, S. and Huberman, B. A. (2010). Predicting the future with social media. In *2010 IEEE/WIC/ACM international conference on web intelligence and intelligent agent technology*, volume 1, pages 492–499. IEEE.

Barari, S., Caria, S., Davola, A., Falco, P., Fetzer, T., Fiorin, S., Hensel, L., Ivchenko, A., Jachimowicz, J., King, G., et al. (2020). Evaluating covid-19 public health messaging in Italy: Self-reported compliance and growing mental health concerns. *MedRxiv*.

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

Barrios, J. M., Benmelech, E., Hochberg, Y. V., Sapienza, P., and Zingales, L. (2021). Civic capital and social distancing during the Covid-19 pandemic. *Journal of Public Economics*, 193:104310.

Becher, M., Stegmueller, D., Brouard, S., and Kerrouche, E. (2020). Comparative experimental evidence on compliance with social distancing during the COVID-19 pandemic. *Available at SSRN 3652543*.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*, 119(1):249–275.

Bish, A. and Michie, S. (2010). Demographic and attitudinal determinants of protective behaviours during a pandemic: A review. *British journal of health psychology*, 15(4):797–824.

Bollen, J., Mao, H., and Zeng, X. (2011). Twitter mood predicts the stock market. *Journal of computational science*, 2(1):1–8.

Brailovskaia, J. and Margraf, J. (2020). Predicting adaptive and maladaptive responses to the Coronavirus (COVID-19) outbreak: A prospective longitudinal study. *International Journal of Clinical and Health Psychology*, 20(3):183–191.

Brodeur, A., Grigoryeva, I., and Kattan, L. (2020). Stay-at-home orders, social distancing and trust. IZA Discussion Paper 13234, IZA.

Brouard, S., Vasilopoulos, P., and Becher, M. (2020). Sociodemographic and psychological correlates of compliance with the Covid-19 public health measures in France. *Canadian Journal of Political Science/Revue canadienne de science politique*, 53(2):253–258.

Caldarelli, G., Chessa, A., Pammolli, F., Pompa, G., Puliga, M., Riccaboni, M., and Riotta, G. (2014). A multi-level geographical study of Italian political elections from Twitter data. *PloS one*, 9(5):e95809.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of human resources*, 50(2):317–372.

Cameron, C. (1996). Patient compliance: recognition of factors involved and suggestions for promoting compliance with therapeutic regimens. *Journal of advanced nursing*, 24(2):244–250.

Chan, H. F., Brumpton, M., Macintyre, A., Arapoc, J., Savage, D. A., Skali, A., Stadelmann, D., and Torgler, B. (2020). How confidence in health care systems affects mobility and compliance during the COVID-19 pandemic. *PloS one*, 15(10):e0240644.

Chen, S., Yang, J., Yang, W., Wang, C., and Bärnighausen, T. (2020). COVID-19 control in China during mass population movements at New Year. *The Lancet*, 395(10226):764–766.

Daoust, J.-F., Nadeau, R., Dassonneville, R., Lachapelle, E., Bélanger, É., Savoie, J., and van der Linden, C. (2020). How to survey citizens' compliance with COVID-19 public health measures: Evidence from three survey experiments. *Journal of Experimental Political Science*, pages 1–8.

Davidson, R. and MacKinnon, J. G. (2010). Wild bootstrap tests for IV regression. *Journal of Business & Economic Statistics*, 28(1):128–144.

Dresselhaus, T. R., Peabody, J. W., Lee, M., Wang, M. M., and Luck, J. (2000). Measuring compliance with preventive care guidelines. *Journal of General Internal Medicine*, 15(11):782–788.

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

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., and Sap, M. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological science*, 26(2):159–169.

Fan, Y., Orhun, A. Y., and Turjeman, D. (2020). Heterogeneous actions, beliefs, constraints and risk tolerance during the covid-19 pandemic. NBER Working paper series 27211, National Bureau of Economic Research.

Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics letters*, 71(3):397–404.

Flèche, S. and Layard, R. (2017). Do more of those in misery suffer from poverty, unemployment or mental illness? *Kyklos*, 70(1):27–41.

Fukuyama, F. (1995). *Trust: The social virtues and the creation of prosperity*, volume 99. Free press New York.

Gayo-Avello, D. (2013). A meta-analysis of state-of-the-art electoral prediction from Twitter data. *Social Science Computer Review*, 31(6):649–679.

Google (2020). Google COVID-19 Community Mobility Reports. Accessed: 7 December 2020.

Greyling, T., Rossouw, S., and Adhikari, T. (2021). The good, the bad and the ugly of lockdowns during Covid-19. *PloS one*, 16(1):e0245546.

Haas, J. and Larson, E. (2007). Measurement of compliance with hand hygiene. *Journal of hospital infection*, 66(1):6–14.

Hailong, Z., Wenyan, G., and Bo, J. (2014). Machine learning and lexicon based methods for sentiment classification: A survey. In *2014 11th web information system and application conference*, pages 262–265. IEEE.

Hale, T., Angrist, N., Cameron-Blake, E., Hallas, L., Kira, B., Majumdar, S., Petherick, A., Phillips, T., Tatlow, H., and Webster, S. (2020). Oxford COVID-19 government response tracker.

Kim, H. H.-s. and Jung, J. H. (2021). Social isolation and psychological distress during the COVID-19 pandemic: A cross-national analysis. *The Gerontologist*, 61(1):103–113.

Krekel, C., Swanke, S., De Neve, J.-E., and Fancourt, D. (2020). Are happier people more compliant? Global evidence from three large-scale surveys during Covid-19 lockdowns. IZA Discussion Paper 13690, IZA.

Kucharski, A. J., Russell, T. W., Diamond, C., Liu, Y., Edmunds, J., Funk, S., Eggo, R. M., Sun, F., Jit, M., Munday, J. D., et al. (2020). Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The lancet infectious diseases*, 20(5):553–558.

Kyngäs, H., Duffy, M. E., and Kroll, T. (2000). Conceptual analysis of compliance. *Journal of clinical nursing*, 9(1):5–12.

Levi, M. and Stoker, L. (2000). Political trust and trustworthiness. *Annual review of political science*, 3(1):475–507.

Lewnard, J. A. and Lo, N. C. (2020). Scientific and ethical basis for social-distancing interventions against COVID-19. *The Lancet Infectious Diseases*, 20(6):631–633.

Luttmer, E. F. and Singhal, M. (2014). Tax morale. *Journal of economic perspectives*, 28(4):149–68.

Margraf, J., Brailovskaia, J., and Schneider, S. (2020). Behavioral measures to fight covid-19: An 8-country study of perceived usefulness, adherence and their predictors. *Plos one*, 15(12):e0243523.

Mikucka, M., Sarracino, F., and Dubrow, J. K. (2017). When does economic growth improve life satisfaction? Multilevel analysis of the roles of social trust and income inequality in 46 countries, 1981–2012. *World Development*, 93:447–459.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica: Journal of the econometric society*, pages 1417–1426.

Nofal, A. M., Cacciotti, G., and Lee, N. (2020). Who complies with COVID-19 transmission mitigation behavioral guidelines? *PloS one*, 15(10):e0240396.

O'Connor, K. J. (2020a). Life satisfaction and noncognitive skills: Effects on the likelihood of unemployment. *Kyklos*, 73(4):568–604.

O'Connor, K. J. (2020b). Physical distancing worked to combat the spread of Covid-19 using Google Mobility data. Économie et Statistiques: Working Papers du STATEC 116, STATEC.

O'Hara, L., Rahim, H. F. A., and Shi, Z. (2020). Gender and trust in government modify the association between mental health and stringency of social distancing related public health measures to reduce COVID-19: a global online survey. *medRxiv*.

Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of economic perspectives*, 14(3):137–158.

Pagliaro, S., Sacchi, S., Pacilli, M. G., Brambilla, M., Lionetti, F., Bettache, K., Bianchi, M., Biella, M., Bonnot, V., Boza, M., et al. (2021). Trust predicts COVID-19 prescribed and discretionary behavioral intentions in 23 countries. *PloS one*, 16(3):e0248334.

Perrotta, D., Grow, A., Rampazzo, F., Cimentada, J., Del Fava, E., Gil-Clavel, S., and Zagheni, E. (2021). Behaviours and attitudes in response to the COVID-19 pandemic: insights from a cross-national Facebook survey. *EPJ Data Science*, 10(1):1–13.

Plohl, N. and Musil, B. (2021). Modeling compliance with COVID-19 prevention guidelines: the critical role of trust in science. *Psychology, Health & Medicine*, 26(1):1–12. PMID: 32479113.

Putnam, R. (2000). *Bowling alone: the collapse and revival of American community*. Simon and Schuster.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal*, 19(1):4–60.

Roser, M., Ritchie, H., Ortiz-Ospina, E., and Hasell, J. (2020). Coronavirus pandemic (COVID-19). *Our world in data*.

Sarracino, F., Greyling, T., O'Connor, K. J., Peroni, C., and Rossouw, S. (2021). A year of pandemic: levels, changes, and validity of well-being data from twitter. evidence from ten countries. Munich Personal Repec Archive 1111111111111, Global Labor Organization (GLO).

Sarracino, F. and Mikucka, M. (2017). Social capital in Europe from 1990 to 2012: Trends and convergence. *Social Indicators Research*, 131(1):407–432.

Shinada, M. and Yamagishi, T. (2007). Punishing free riders: direct and indirect promotion of cooperation. *Evolution and Human Behavior*, 28(5):330–339.

Stevenson, C., Wakefield, J., Drury, J., and Felsner, I. (2020). Collectively coping with coronavirus: local community identification predicts giving support and lockdown adherence during the COVID-19 pandemic.

Williams, L., Rasmussen, S., Kleczkowski, A., Maharaj, S., and Cairns, N. (2015). Protection motivation theory and social distancing behaviour in response to a simulated infectious disease epidemic. *Psychology, health & medicine*, 20(7):832–837.

Wright, L., Steptoe, A., and Fancourt, D. (2021). Predictors of self-reported adherence to COVID-19 guidelines. A longitudinal observational study of 51,000 UK adults. *The Lancet Regional Health - Europe*, (4):100061.

Zhao, S. Z., Wong, J. Y. H., Wu, Y., Choi, E. P. H., Wang, M. P., and Lam, T. H. (2020). Social distancing compliance under Covid-19 pandemic and mental health impacts: A population-based study. *International journal of environmental research and public health*, 17(18):6692.