Lockdowns and Conspiracy Theories: Evidence from European Policy Variation

Ciarán Murphy[a,1]

[a] *University College Dublin Dublin Ireland*

Abstract

When European governments locked down their populations during Covid-19, they inadvertently increased demand for conspiracy theories. This paper provides demand-side evidence on what drives online conspiracy narratives, exploiting differential variation in lockdown stringency across 31 European countries. Stricter lockdowns significantly raised conspiracy search interest during 2020–2021. Evidence points to social isolation as the operative mechanism: the sustained level of restriction predicts conspiracy search while week-on-week changes do not, and restrictions that physically confine people drive the result while mask mandates do not. The findings suggest that social capital disruption is a distinct driver of conspiratorial information-seeking.

Keywords: conspiracy theories, lockdowns, Covid-19, Google Trends, social isolation, social capital

JEL Classification: D83, I18, Z13.

*Corresponding author

[1] Email: ciaran.murphy10@ucdconnect.ie. Corresponding author.

1. Introduction

The growth and persistence of online conspiracy theories have emerged as a central concern for policymakers. Conspiracy narratives questioning the legitimacy of political institutions and scientific expertise have become increasingly salient in public discourse and are associated with concrete behavioural outcomes, including voting behaviour, compliance with public health guidelines, and criminal activity (Allcott et al., 2020; Van Es et al., 2025). Despite these observable behavioural consequences, there is little causal evidence on the economic and institutional conditions under which demand for conspiracy narratives increases.

Although conspiracy theories predate the Covid-19 pandemic, their sudden rise in prominence and reach coincided with a period of unprecedented lockdowns, during which daily routines and peer interactions were sharply disrupted. Such changes plausibly altered how individuals allocated attention across information sources and narratives. In particular, reduced social interaction and weakened institutional anchoring may have increased the appeal of conspiratorial explanations that offer simple accounts of complex events.

This paper provides the first demand-side evidence on what drives online conspiracy narratives, exploiting Covid-19 lockdown policies as an exogenous source of variation in social isolation. Using differential changes in lockdown stringency across 31 European countries over time, combined with Google Trends data on searches for prominent conspiracy topics, I identify the effect of lockdown policies, which sharply curtailed offline social interaction, on online conspiracy search behaviour. The central hypothesis is that lockdown policies shifted attention toward conspiratorial narratives that provide identity-consistent explanations for uncertainty and institutional distrust.

The paper makes two main contributions. First, it provides demand-side evidence on digital radicalisation: using within-country variation in policy intensity and week fixed effects, I estimate how social-capital shocks shift searches for conspiracy-related content. Second, the paper provides evidence that these effects occur through the mechanism of social isolation. Examining individual policy components reveals that restrictions which physically confine people (workplace closures, stay-at-home orders, and school closures) significantly increase conspiracy search, while measures that require compliance without isolating individuals, such as facial covering mandates, do not. The effect operates through sustained levels of restriction rather than transient shocks: week-on-week changes in stringency have no independent effect, inconsistent with news-salience and policy-backlash interpretations. Controlling for an index of government economic support throughout all specifications is consistent with the interpretation that the effect is not driven by economic hardship induced by lockdowns: the isolation channel operates independently of any economic anxiety mechanism. Similarly, controlling for lagged Covid-19 cases and deaths argues against pandemic severity as an alternative explanation, since the relationship between stringency and conspiracy search holds conditional on the epidemiological environment. These findings point to an ongoing social isolation channel distinct from pandemic severity, policy

backlash, or economic anxiety.

2. Literature

The paper broadly relates to three literatures: the economics of belief and identity formation, the political economy of shocks and radicalisation, and the effects of Covid-19 (particularly a recent literature which has used Google Trends to examine these effects). A rich economics literature examines the role of the media in shaping beliefs, political attitudes, and trust in institutions. Canonical studies examine the effect of access to different media sources on voting behaviour in the U.S. (DellaVigna & Kaplan, 2007) and Russia (Enikolopov et al., 2011). These studies establish that supply-side shocks in information can persuade individuals and shift ideologies, exploiting natural experiments in media exposure to provide clean causal evidence of persuasion. Social media and online platforms also affect identity and preferences: for example, Müller & Schwarz (2021) show that anti-refugee sentiment on social media predicts real-life hate crimes in Germany. Theories of identity formation (Akerlof & Kranton, 2000) suggest that beliefs are sustained through ongoing social interaction: engagement with peers and institutions provides continuous reinforcement of socially mainstream interpretations of events. Forced isolation severs these anchoring ties, potentially increasing susceptibility to conspiratorial narratives that offer alternative, identity-consistent frameworks for understanding uncertainty. The literature to date has focused heavily on the media supply of such content, whereas this paper examines the demand-side conditions (specifically, the disruption of social capital) that make individuals receptive to it.

A second strand of literature related to this work is the large political economy research linking economic and institutional shocks to extreme or radical outcomes. Adverse economic events such as the China trade shock cause voters to shift to the right (Autor et al., 2020). Similar effects have been documented in Western Europe recently (Colantone & Stanig, 2019; Fetzer, 2019) and throughout modern history (Funke et al., 2016). These contributions emphasise the "grievance" channel: when people face economic pain or uncertainty, they may turn to radical ideologies or anti-establishment narratives. However, the shocks studied in this literature are usually financial in nature (trade, fiscal cuts, and crises) and the outcomes are generally electoral (votes for extremist parties and policies). This study adds a new dimension by examining an institutional social shock (lockdowns that forced social isolation) and a high-frequency outcome which more accurately reflects the demand for conspiratorial narratives. The hypothesis follows from existing findings (economic hardship fosters extremism) but focuses on social-capital hardship. Lockdowns are conceptually similar to an economic shock in that they deprive people of a key "asset" (peer interaction and structure), potentially generating frustration or identity-seeking behaviour.

Finally, the paper builds on a growing body of research about the Covid-19 pandemic's wide-ranging social effects. A subset of these papers has specifically analysed how pandemic

lockdowns affected online search behaviour. For example, Brodeur et al. (2021) use Google Trends and difference-in-differences to show that the first lockdowns in Europe resulted in substantial increases in searches for boredom and loneliness, but decreases in searches for stress, suicide and divorce. Other designs use similar data but are mainly correlational; for example, Berger et al. (2021) estimate potential effects of lockdowns on fertility by examining trends in searches related to marriage and contraception. Other studies use Google Trends to examine the association between lockdowns and mental health searches (Rosa et al., 2022) and general interests (Li et al., 2025). Beyond Google Trends, qualitative interdisciplinary research also documents a rise in online extremist communities during the pandemic. Wedel & Coufal (2025) examine a large "incel" forum and find that during the initial Covid lockdown, user activity and extreme rhetoric intensified. However, this evidence lacks a counterfactual and is confined to specific online groups. Despite this suggestive evidence, no studies provide a causal quantitative estimate of the effect of lockdown policies on demand for conspiracy or extremist content. The paper also adds to a literature on unintended effects of Covid-19 lockdowns (for example, the effect on domestic violence (Leslie & Wilson, 2020)).

3. Data and Descriptive Statistics

To assess the effect of lockdown policies on demand for conspiracy content, I assemble a novel dataset merging online search activity indicators, pandemic policy measures, and covariates, at the national level. The core data components are described below.

3.1. Google Trends

The primary outcome data come from Google Trends, which provides measures of relative web search interest[2] for specific topics over time and across geographic units. The unit of observation in the analysis is a country-week-topic cell (or, in the case of the United States, a state-week-topic cell). For each topic, Google Trends reports a normalised index that ranges from 0 to 100, capturing search interest within a given region and week. Let Q_{gct} denote the number of Google searches in region c during week t that Google assigns to topic g, and let $Q_{\cdot ct}$ denote the total number of Google searches in the same region-week. The underlying object of interest is the topic's share of all searches:

$$\text{Share}_{gct} \equiv \frac{Q_{gct}}{Q_{\cdot ct}}. \tag{1}$$

Google Trends does not report Share_{gct} directly. Instead, the series is normalised within each region so that the maximum observed share for a given topic over the selected time window equals 100:

[2]Data were separately collected on YouTube searches, but were not used due to a lack of variation.

$$GT_{gct} = 100 \times \frac{\text{Share}_{gct}}{\max_{\tau \in T} \text{Share}_{gc\tau}}, \qquad GT_{gct} \in [0, 100], \tag{2}$$

where T denotes the full 2020-2022 download window (Appendix A provides a numerical example). As a result, the Google Trends index captures changes in interest in topic g relative to all search activity in the same region and week, abstracting from fluctuations in overall online activity or search intensity. This feature mitigates concerns that increases in general online activity during lockdowns mechanically drive the results.

Using search behaviour to proxy for underlying attitudes and interests is an established empirical strategy, particularly where survey measurement is costly, infrequent, or subject to social desirability bias. Stephens-Davidowitz (2014) uses Google search volume for a racially charged term to measure racial animus at a geographic granularity unavailable through survey data, validating the measure against real-world political outcomes. Da et al. (2011) show that search volume is a direct, quantifiable measure of attention, predictive of subsequent economic behaviour. Böhme et al. (2020) use geo-referenced online search data to measure migration intentions, showing that the resulting measure partly reflects genuine intentions and outperforms conventional predictors of actual migration flows. The present analysis follows this tradition, treating search interest in conspiracy-related topics as a revealed-preference measure of demand for conspiratorial content.

I collect weekly Google Trends data for a broad set of topics related to conspiracy theories, across the 27 European Union member states, as well as the United Kingdom, Norway, Switzerland, and Iceland. In Google Trends, a "topic" represents a semantic grouping of related search queries identified by Google, aggregating searches that refer to the same underlying concept across different spellings, phrasings, and languages. Topics are preferred to individual search terms because they provide broader coverage and more stable time series, particularly at high frequencies. In a European context, topics are particularly useful as they aggregate semantically equivalent searches across different languages, ensuring comparability across countries.

A broad set of candidate topics was compiled by enumerating prominent online conspiracy narratives and identifying those available as semantic topics in Google Trends. However, many topics exhibit extremely sparse data at the region-week level due to Google's suppression of low search volume, which is reported as zero rather than as small positive values. To ensure that the outcomes contain usable variation, topics were screened and included only if they satisfied three criteria. First, a topic was required to exhibit sufficient non-zero variation, defined as having positive search interest in at least 30% of region-week observations over the sample period. Second, topics were required to display broad geographic coverage, defined as positive search interest in at least 80% of geographic units at some point during the download sample period (January 2020 to December 2022). Third, topics were required to exhibit non-zero search interest prior to the onset of the Covid-19 pandemic (defined for the purpose of this screening process as March 1st, 2020). The full set of considered topics and screening metrics are shown in Appendix

B, while results without screening are reported in Appendix C.[3]

The main analysis focuses on eight conspiracy-related topics listed in Table 1: "QAnon", "Propaganda", "New World Order", "Brainwashing", "Fake news", "Chemtrails", "David Icke", and "Flat Earth". These are the topics that passed all three screening criteria. They span long-standing conspiracy theories ("Flat Earth", "New World Order"), prominent conspiracy figures ("David Icke"), and disinformation narratives ("Fake news", "Propaganda"). The pre-pandemic search criterion excludes topics that emerged specifically in the Covid-19 context, meaning the analysis captures a broader shift in demand for conspiracy narratives rather than the more mechanical result that pandemic-specific conspiracy theories increased during the pandemic.

Table 1: Topic screening metrics for conspiracy topics included in the European analysis.

Topic	Zero	Countries	Pre-Covid	Peak period
Propaganda	0.294	1.000	Yes	Feb 2022
Fake news	0.392	1.000	Yes	Mar 2020
QAnon	0.515	0.968	Yes	Jan 2021
Flat Earth	0.568	0.935	Yes	Feb 2020
New world order	0.582	0.871	Yes	Mar 2020
David Icke	0.682	0.968	Yes	Apr 2020
Chemtrails	0.684	0.903	Yes	Apr 2020
Brainwashing	0.687	0.903	Yes	Dec 2020

Note: Zero is the share of country-week observations with zero search interest. Countries is the share of European countries with positive search interest at any point in the sample period. Pre-Covid indicates whether the topic had positive interest prior to March 2020. Peak period is the month in which search interest was highest in 2020-2022.

To address the heavy mass at zero in the Google Trends data while retaining all observations, I use the inverse hyperbolic sine (IHS) transformation of the Google Trends index as the primary outcome measure. This transformation preserves zero values while approximating a logarithmic transformation for positive values. The analysis primarily focuses on a composite conspiracy index, constructed in two steps. First, each topic's IHS series is standardised using the mean and standard deviation of that topic computed over the 2020–2021 estimation sample. Second, the standardised series are averaged within each country-week to form a raw composite, which is then

[3]The result is not sensitive to these specific cutoffs: varying the 30% non-zero and 80% geographic-coverage thresholds individually (e.g. to 25%/35% and 70%/90%) leaves the composite-index stringency coefficient in the range 0.0044–0.0050, significant at the 5% level throughout.

re-standardised using its own 2020–2021 mean and standard deviation. The resulting index has mean zero and unit standard deviation in the regression sample. This index provides a summary measure of overall engagement with conspiracy-related content and reduces noise associated with individual topic-level series. Results for individual topics are reported separately. Alternative outcome specifications are reported in Appendix D; results are insensitive to the scale of the raw index prior to the IHS transformation (Appendix E).

3.2. Covid-19 Policy Data

Information on Covid-19 containment policies is drawn from the Oxford Covid-19 Government Response Tracker (OxCGRT), which provides daily measures of policy responses at the national level. The primary policy variable is the Oxford Stringency Index, a continuous composite measure summarising the overall severity of Covid-19 lockdown policies. The index aggregates information on multiple policy dimensions, including school closures, workplace closures, restrictions on movement and public gatherings, and stay-at-home orders, and is scaled from 0 to 100, with higher values indicating more stringent lockdowns. In addition to the composite index, the OxCGRT provides ordinal indicators for each individual policy component, which I use to examine which specific types of restriction drive the aggregate result.

The daily OxCGRT data are aggregated to the weekly level to align with the Google Trends outcomes, with each week beginning on Sunday and ending on Saturday. For the continuous Stringency Index, daily values are averaged within each week. For the ordinal policy component indicators, a week is assigned the maximum value observed on any day within that week. The main analysis is conducted at the country-week level. Table 2 displays descriptive statistics for the key policy variables, with the United States included for comparison.

Table 2: Descriptive Statistics: Covid-19 Containment Policies

Region	Variable	Mean	SD	Min	Max	N
Europe	Stringency index	39.848	24.208	0	96.30	4867
Europe	School closures (binary)	0.324	0.468	0	1.00	4867
Europe	Workplace closures (binary)	0.424	0.494	0	1.00	4867
US	Stringency index	37.832	20.107	0	93.52	7850
US	School closures (binary)	0.375	0.484	0	1.00	7850
US	Workplace closures (binary)	0.206	0.404	0	1.00	7850

Note: School and workplace closures are binary indicators equal to one if mandatory closures were in place. The stringency index ranges from 0 to 100, with higher values indicating more stringent lockdown policies. All variables are measured at the unit-week level (state-week for the US, country-week for Europe). Statistics are reported for the full 2020-2022 period.

Figure 1 documents substantial variation in Covid-19 lockdown policies across regions and over time. Panels A and B plot the share of geographic units with strict school and workplace closures, respectively, for European countries and, for comparison, U.S. states. Both types of closure were nearly universal during the initial phase of the pandemic in spring 2020, but reopening trajectories diverged substantially thereafter, with considerable cross-jurisdictional variation in timing and pace throughout 2020 and 2021. Panel C plots the weekly mean and interquartile range of the Oxford Stringency Index for U.S. states and European countries. In both regions, average stringency peaks early in the pandemic and declines over time, but there remains substantial cross-unit dispersion throughout the sample period. By early 2022, stringency had fallen to near-zero across virtually all European countries, marking the effective end of the active restriction period.

Figure 2 plots the average IHS-transformed conspiracy index across European countries over time, with the interquartile range and the rescaled mean stringency index overlaid. Three distinct spikes in conspiracy search are visible. The first, in March-April 2020, coincides clearly with the onset of lockdowns across Europe. The second spike, in October–November 2020, is consistent with the European second wave of restrictions, though it also coincides with the U.S. presidential election — a global media event that may have independently increased searches for conspiracy-related content even in Europe. The third spike, in December 2020-January 2021, aligns with the Alpha variant wave and renewed European restrictions, although its peak also aligns with the U.S. Capitol riot and presidential inauguration. After the peak in early 2021, the index declines through the first half of 2021 before a further spike around March 2022, slightly after the Omicron wave, and coinciding with the Russian invasion of Ukraine, which

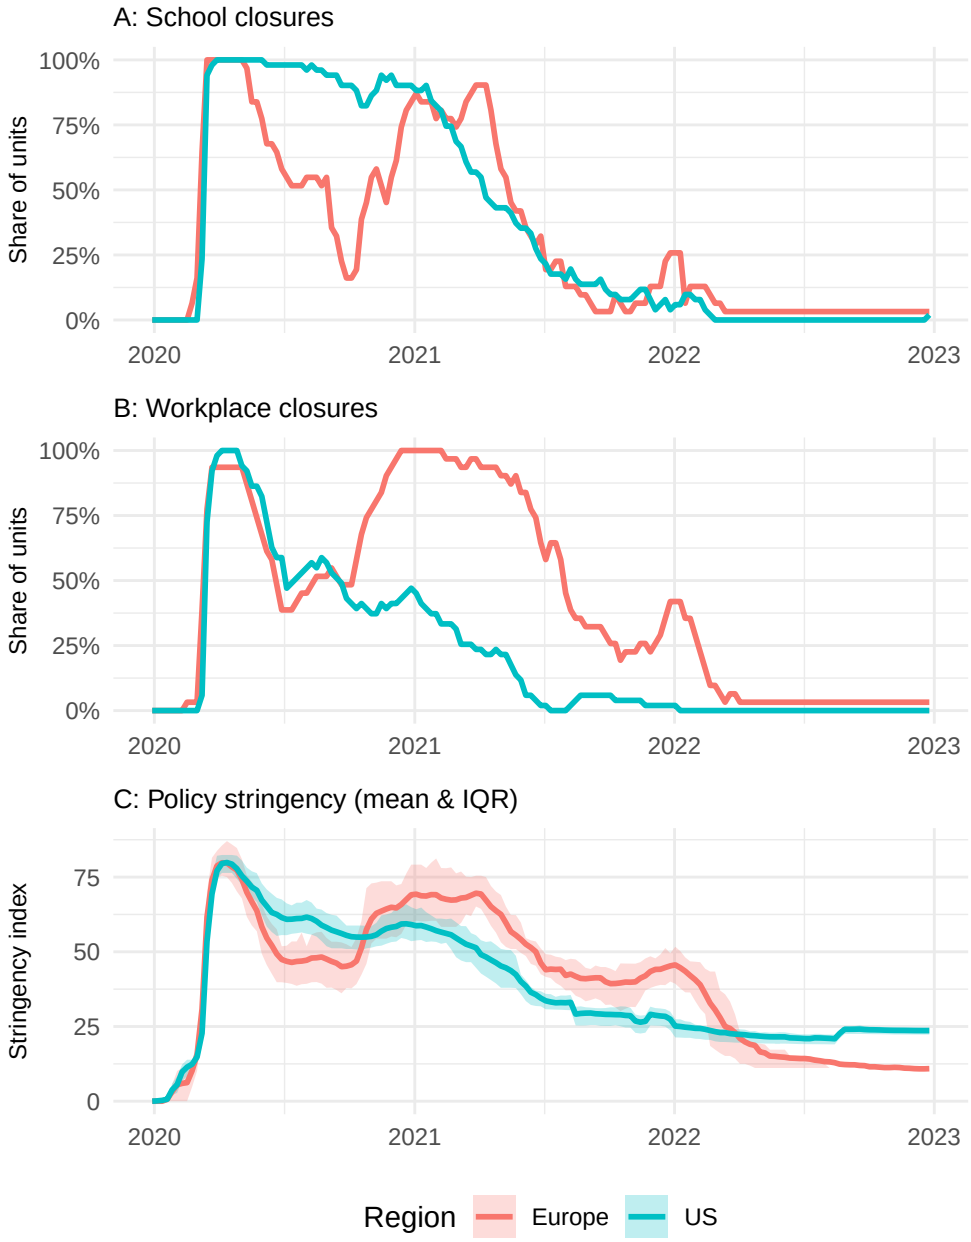


Figure 1: School closures, workplace closures, and lockdown stringency across U.S. states and European countries.

likely amplified searches for propaganda-related content. The index then declines to lower levels through the remainder of 2022.

The first spike is most cleanly attributable to lockdown onset; the second and third spikes are consistent with the lockdown hypothesis but also coincide with major global political events. The raw time series therefore cannot distinguish whether conspiracy search rose because of lockdown-induced isolation, because of pandemic severity, or because of unrelated global media shocks. The main empirical challenge is to separate the effect of lockdown policies themselves from these common time-varying influences. The identification strategy addresses this directly. By including week fixed effects, the analysis absorbs all shocks common to European countries in a given week, including aggregate pandemic dynamics, global media events, and international

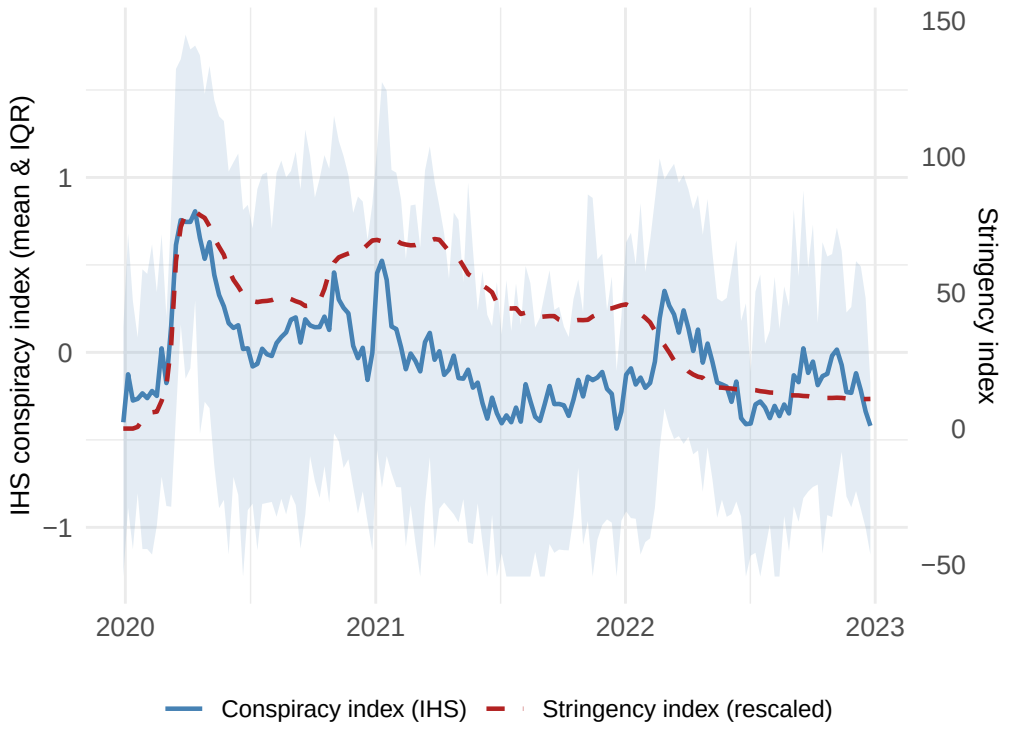


Figure 2: Average conspiracy index (mean and IQR) and rescaled mean stringency index across European countries.

political crises. Identification therefore comes exclusively from differential variation in lockdown stringency across countries within the same week. In effect, the design compares European countries exposed to the same global shocks in the same week, asking whether those that were more stringently locked down exhibited higher demand for conspiratorial narratives.

3.3. Covariates and heterogeneity

To account for time-varying pandemic severity, I control for weekly new cases and deaths. These data are sourced from Our World in Data country-level daily data aggregated to the weekly frequency (or, for the supplementary United States analysis, the Centers for Disease Control and Prevention aggregated to the state-week level). I also include the OxCGRT Economic Support Index, a composite measure of government financial support to households and businesses (income support, debt relief, and related interventions), scaled from 0 to 100 and drawn from the same OxCGRT dataset as the policy variables. I utilise a number of supplementary data sources to examine cross-country heterogeneity. These data are sourced from the World Bank (for internet penetration, urbanisation rate, and GDP per capita) and Reporters Without Borders (for press freedom) and are defined in 2019, prior to the onset of the pandemic.

4. Methodology

The empirical approach estimates reduced-form effects of Covid-19 containment policies on demand for online conspiracy-related content, as proxied by Google Trends search activity. The main analysis exploits differential changes in lockdown stringency across 31 European countries

over time. Specifically, I estimate regressions of the form:

$$y_{ct} = \beta\,\text{Stringency}_{ct} + \boldsymbol{\delta}\mathbf{X}_{c,t-1} + \theta Z_{ct} + \alpha_c + \gamma_t + \varepsilon_{ct}, \tag{3}$$

where y_{ct} denotes a Google Trends outcome for country c in week t, Stringency_{ct} is the Oxford Stringency Index measuring the overall severity of lockdown policies, $\mathbf{X}_{c,t-1}$ is a vector of controls containing new Covid-19 cases and deaths in country c in week $t-1$, Z_{ct} is the OxCGRT Economic Support Index for country c in week t, α_c are country fixed effects, and γ_t are time fixed effects. Standard errors are clustered at the country level. To address the small number of clusters ($G = 31$), inference uses critical values from a $t(G-1)$ distribution rather than the standard normal, following Donald & Lang (2007): the t-statistic from the cluster-robust variance estimator is compared to a $t(30)$ distribution, which places more probability mass in the tails and yields conservative inference relative to the asymptotic approximation. Appendix F reproduces the main results under wild cluster bootstrap inference (Cameron & Miller, 2015); the conclusions are unchanged.

Country fixed effects absorb all time-invariant differences between countries, including baseline levels of conspiracy interest, media ecosystems, and institutional characteristics. Time fixed effects control for shocks common to all countries, such as aggregate pandemic dynamics, widespread media coverage, and policy announcements at the European level. Identification therefore comes from differential changes in containment policies across countries within the same time period.

The identification assumption is that there are no time-varying country-specific factors that are correlated with lockdown policy and which affect demand for conspiracy content. The main identification concern is that online conspiracy-related search activity may respond to contemporaneous pandemic conditions; in particular, rising infection or mortality rates may prompt governments to tighten restrictions while independently affecting public attention, anxiety, or information-seeking behaviour. To mitigate this concern, I control for Covid-19 cases and deaths at the country level. In all specifications, Covid-19 cases and deaths enter with a one-week lag to avoid conditioning on an outcome. Results are robust to using contemporaneous controls.

A second identification concern is that lockdown policies cause economic disruption, and economic hardship may independently drive demand for conspiracy narratives. A substantial political economy literature documents that adverse economic shocks increase support for radical and anti-establishment movements (Autor et al., 2020; Colantone & Stanig, 2019; Funke et al., 2016), and the same mechanism could in principle generate conspiratorial information-seeking. If countries that imposed stricter lockdowns also experienced greater economic disruption, the estimated effect of stringency on conspiracy search could reflect economic grievance rather than social isolation. To address this concern, I control for the OxCGRT Economic Support Index, a measure of the financial assistance governments provided during the pandemic.

A third identification concern is reverse causality: elevated conspiracy search interest could

generate political pressure to ease restrictions, introducing a negative association between search and stringency within countries over time. Time-invariant sources of such pressure, such as persistent institutional distrust and political culture, are absorbed by country fixed effects. Within a given week, reverse causality is unlikely: stringency reflects decisions taken in prior periods, and governments do not typically adjust codified policy in response to shifts in public sentiment within days. Any such reverse channel would in any case bias the estimated coefficient toward zero, so the estimates represent a lower bound on the true effect.

Week fixed effects are used throughout. Given that the outcome is observed weekly, this is the most demanding identification test: it requires that the relationship between stringency and conspiracy search holds after controlling for all common weekly fluctuations, relying only on differential within-week variation in policy intensity across countries. The main analysis restricts the sample to January 2020 to December 2021, the period of active lockdown restrictions across European countries. Data are collected over the full 2020–2022 period, which is used for descriptive statistics and topic screening. The estimates should be interpreted as reduced-form policy effects. Specifically, the coefficient β captures whether countries that experienced more stringent lockdowns exhibited systematically different patterns of relative search interest in conspiracy-related content compared to other countries during the same period.

5. Results

5.1. Stringency index

Table 3: Lockdown Stringency and Conspiracy Search

Model:	(1)	(2)	(3)
Variables			
Stringency index	0.0046***	0.0047***	0.0050***
	(0.0015)	(0.0015)	(0.0016)
Covid cases (t-1, per 100k)		0.0585**	0.0498*
		(0.0251)	(0.0266)
Covid deaths (t-1, per 1k)		-0.0176	-0.0158
		(0.0121)	(0.0123)
Economic support index			-0.0017**
			(0.0007)
Fixed-effects			
Country	Yes	Yes	Yes
Week	Yes	Yes	Yes
Fit statistics			
Observations	3,255	3,224	3,224
Within R^2	0.01396	0.01754	0.02535

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 3 presents the main results. The primary outcome is the inverse hyperbolic sine (IHS) transformation of the composite conspiracy index. The sample covers the two full calendar years during which lockdowns were in place (January 2020 to December 2021).

Column 1 reports the overall Oxford Stringency Index with country and week fixed effects but no additional controls. Column 2 adds lagged Covid-19 cases and deaths per capita. Column 3 further adds the OxCGRT Economic Support Index. The stringency coefficient in the fully specified model is 0.0050, significant at the 1% level. An increase in the stringency index of 21.6 points (one standard deviation over the 2020-21 period) is associated with an increase in conspiracy demand of 0.11 standard deviations. The coefficient is not attenuated by including measures of pandemic severity.

The Economic Support Index (column 3) enters with a negative and significant coefficient: higher government financial support is associated with reduced conspiracy search, consistent with economic hardship independently contributing to conspiracy demand. Crucially, the stringency coefficient is not attenuated by its inclusion, confirming that the isolation and economic anxiety channels are distinct.

Figure 3 presents a residualised binscatter of the within-country, within-week relationship be-

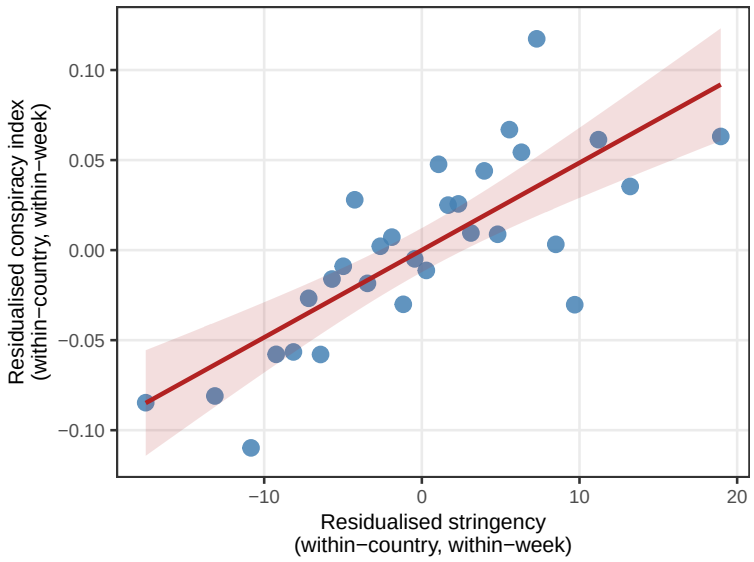


Figure 3: Residualised binscatter of conspiracy search on lockdown stringency. Both variables are residualised on country fixed effects, week fixed effects, lagged epidemiological controls, and the OxCGRT Economic Support Index. Each point is the mean within one of 30 equal-frequency bins of the residualised stringency variable.

tween lockdown stringency and conspiracy search. Both variables are residualised against country fixed effects and week fixed effects. The residuals therefore represent pure within-country, within-week variation in each variable, net of all common shocks and country-specific levels. Each point corresponds to the mean of the residualised outcome within one of 30 equal-frequency bins of the residualised stringency variable. The positive, approximately linear relationship is consistent with the regression estimates.

I also examine whether the effect varies systematically with baseline country characteristics by interacting stringency with pre-pandemic GDP per capita, internet penetration, urbanisation, and press freedom. These interaction terms are generally small and statistically insignificant (Appendix G), suggesting that the relationship is not strongly moderated by observable cross-country structural differences. The estimate is similarly stable to excluding any individual country from the sample (Appendix G).

5.2. Policy components

To examine which specific restrictions drive the stringency result, each OxCGRT policy component is estimated separately on the composite conspiracy index, replacing the ordinal indicator with a full set of level dummies. This approach directly identifies the severity threshold at which each policy begins to affect conspiracy search. Figure 4 displays the estimates, with level 0 (no restriction) as the reference category in each case.

For school closures, levels run from recommended closure (1) through required closure at some grades or ages (2) to required closure at all levels (3); the estimated effect increases across levels and is largest when all schools are required to close. For workplace closures, levels range from recommending that workers stay home (1) through required closure for some sectors (2) to

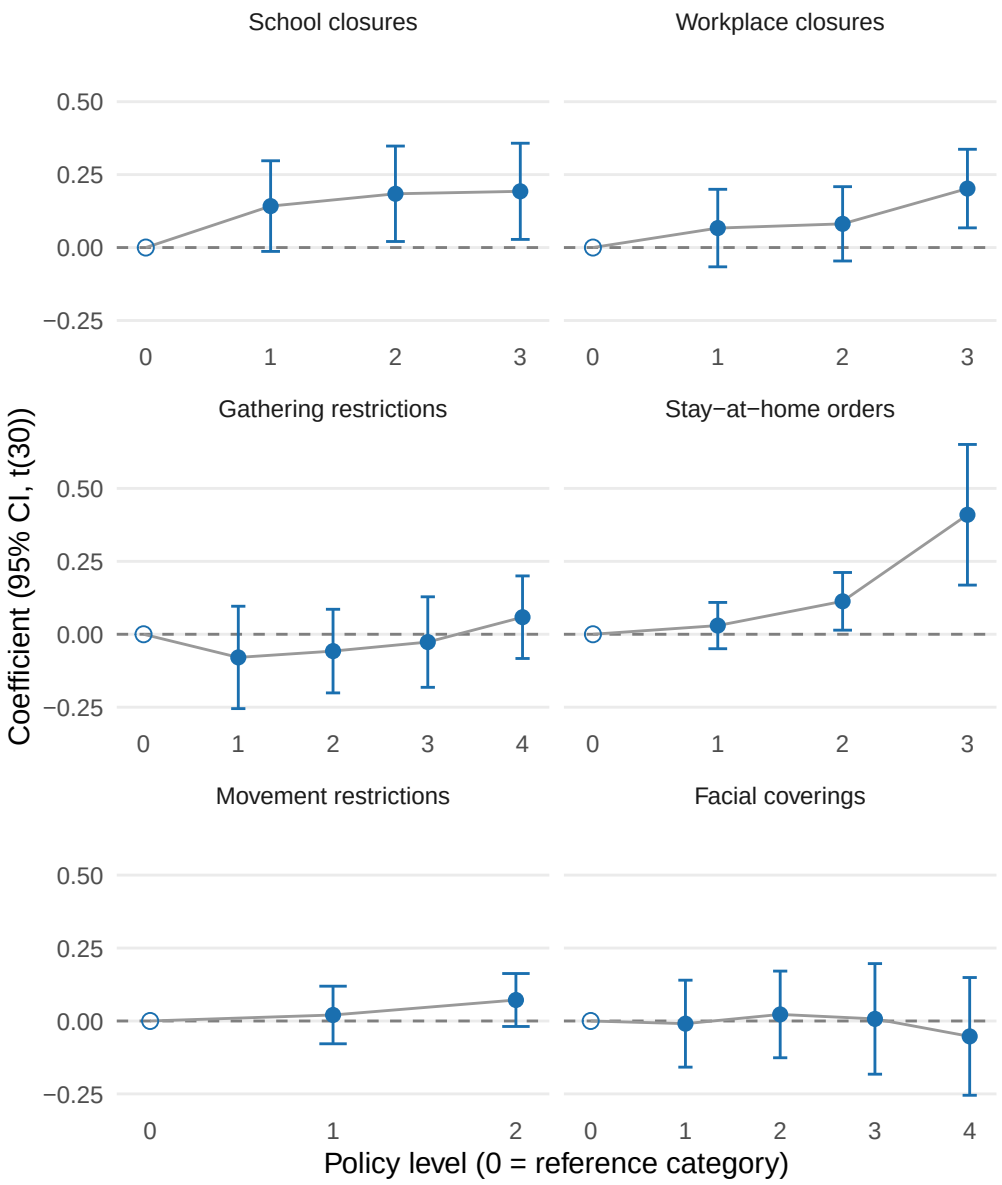


Figure 4: Dose–response coefficients for individual lockdown policy components. Each panel shows the estimated effect of each restriction level relative to level 0 (no restriction), from a regression replacing the ordinal variable with a full set of level dummies. Country and week fixed effects, lagged epidemiological controls, and the OxCGRT Economic Support Index are included throughout. Error bars are 95% confidence intervals using $t(30)$ critical values.

required closure for all but essential workers (3); the effect is concentrated at the most restrictive level, when the great majority of workers are confined to their homes. Gathering restrictions are indexed by group size: level 1 restricts gatherings above 1,000 people, level 2 above 100, level 3 above 10, and level 4 restricts gatherings of 10 or fewer; no effect is detectable at any level. For stay-at-home orders, level 1 is a recommendation not to leave the house, level 2 requires remaining at home with broad exceptions such as grocery shopping and exercise, and level 3 requires remaining at home with minimal exceptions; the effect increases sharply at the most restrictive level, when residents face binding confinement. Internal movement restrictions run from a recommendation not to travel between regions (1) to an enforced requirement (2).

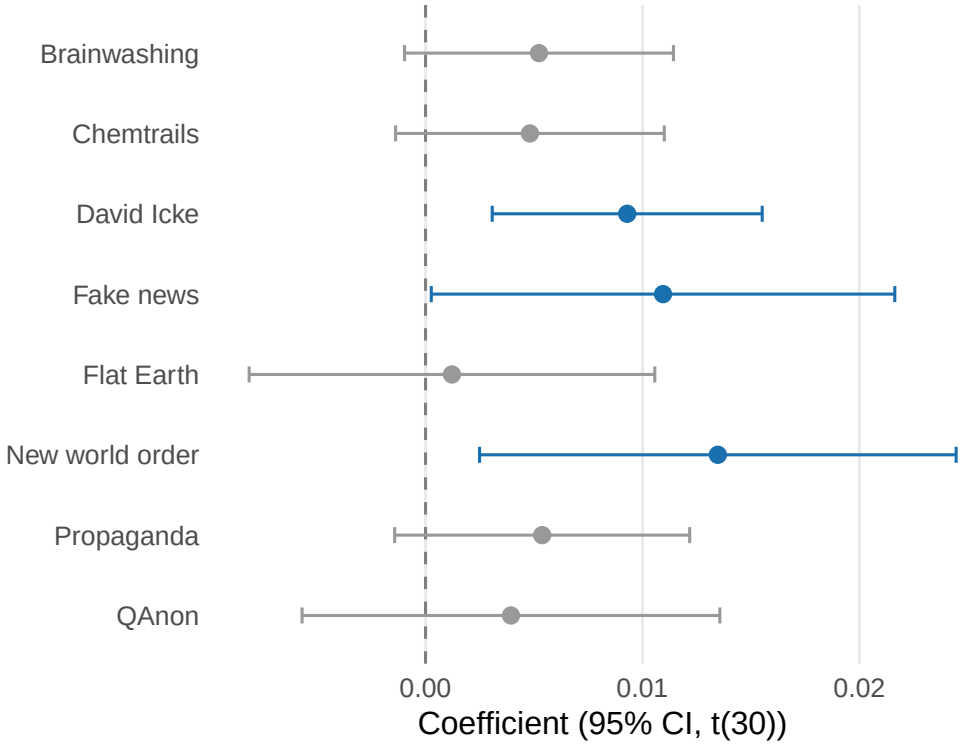


Figure 5: Effect of lockdown stringency on individual conspiracy topics. Each point is the coefficient on the stringency index from a separate regression of the IHS topic search index on stringency, with country and week fixed effects, lagged epidemiological controls, and the OxCGRT Economic Support Index. Error bars are 95% confidence intervals using $t(30)$ critical values. Blue points are significant at the 5% level.

Facial covering mandates span from recommended (1) through required in some public spaces (2) and all public spaces (3) to required outside the home at all times (4), and show no effect at any level. The dose–response profiles are broadly monotone for policies that physically confine people, consistent with social isolation as the operative mechanism, and are consistent with the degree of physical confinement, rather than the intrusiveness of government policy in general, driving the effect. Full regression results underlying Figure 4 are reported in Appendix H.

5.3. Individual conspiracy topics

Figure 5 examines heterogeneity across individual conspiracy topics. When examined individually, three topics show statistically significant responses to lockdown stringency at the 5% level: "David Icke", "Fake news", and "New World Order". A one standard deviation (21.6-point) increase in stringency is associated with demand for these topics increasing by approximately 20–29%. The coefficients for the remaining topics are all positive but not individually significant at the 5% level. The significant topics span prominent conspiracy figures ("David Icke"), disinformation narratives ("Fake news"), and anti-establishment ideology ("New World Order"), suggesting the composite index captures a broad increase in conspiracy-related search rather than a single narrative.

A natural concern is that some topics in the index could have multiple interpretations: one might argue, for example, that "Fake news" partly reflects interest in counter-disinformation

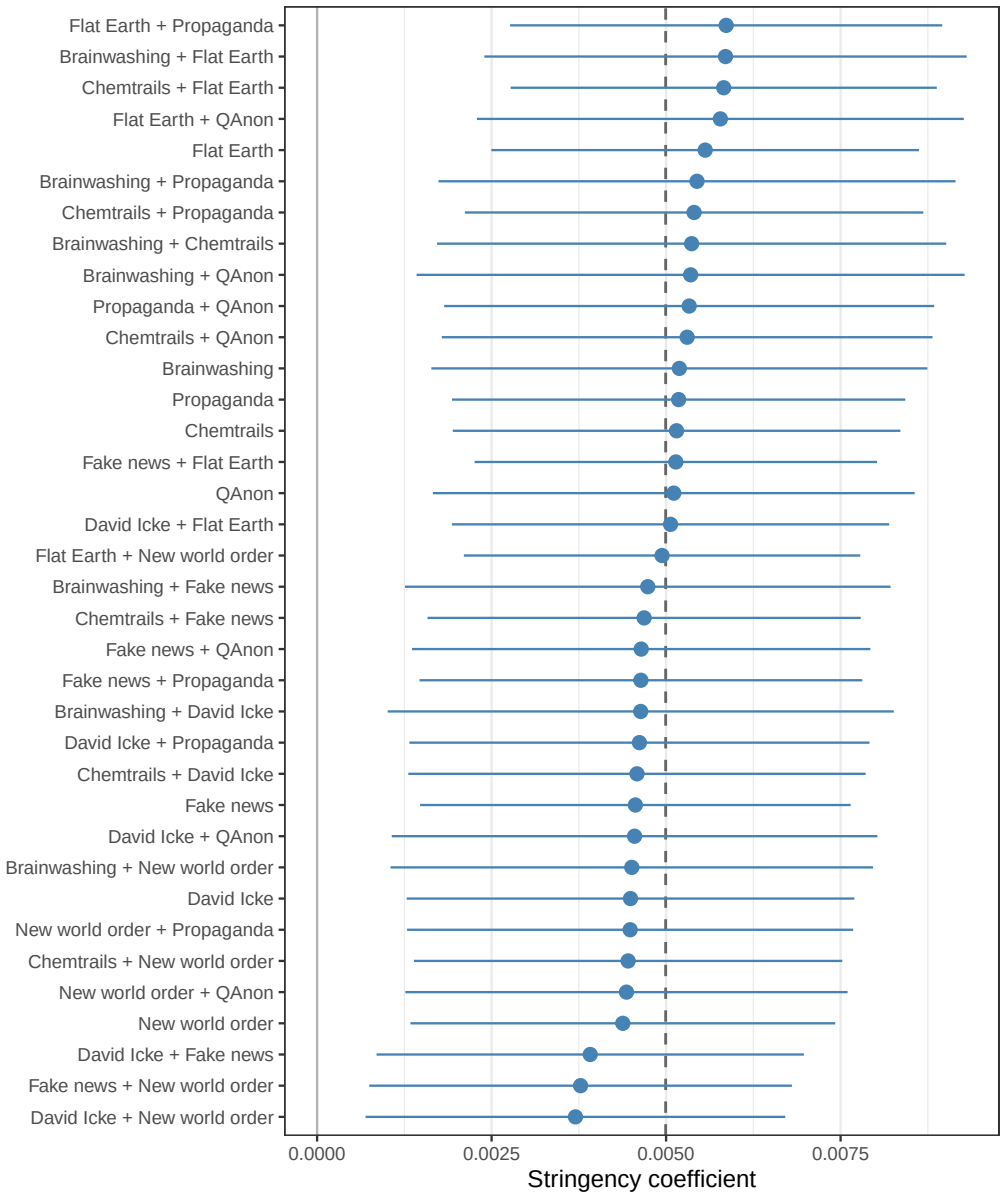


Figure 6: Index composition robustness: stringency coefficients from all leave-one-out (8 specs) and leave-two-out (28 specs) sub-indices, with 95% confidence intervals using $t(30)$ critical values. Dashed line: full eight-topic estimate. Specifications are sorted by coefficient.

rather than conspiratorial content. To address this, Figure 6 reports the stringency coefficient on the composite index recomputed after dropping every topic individually (8 specifications) and every pair of topics simultaneously (28 specifications), for a total of 36 sub-indices. The dashed line marks the full eight-topic estimate. Across all 36 specifications the coefficient remains positive and statistically significant, ranging from 0.0037 to 0.0060: the result does not depend on any single topic or pair of topics.

Full topic-level regression results are reported in Appendix I.

6. Heterogeneity and Mechanisms

6.1. Persistence

Figure 7 examines the temporal structure of the stringency–conspiracy relationship. If conspiracy search responds primarily to the act of imposing or tightening restrictions, two mechanisms are plausible. First, news salience: media coverage of policy changes may temporarily increase searches for conspiracy-related content. Second, a policy-backlash mechanism: the imposition of new restrictions may trigger resentment that manifests in conspiracy search. Under either mechanism, week-on-week changes in stringency should predict the outcome, and the effect should dissipate once the policy change is no longer newsworthy. If, on the other hand, the mechanism operates through sustained social isolation, then it is the level of stringency that should matter, and transient changes should have no independent effect.

The left panel replaces the level of stringency with three change specifications: the symmetric first difference (ΔStringency), and its tightening (Δ^+, positive changes only) and loosening (Δ^-, negative changes only) components. All three are close to zero and statistically insignificant at the 5% level. The results are therefore inconsistent with both the news-salience and policy-backlash interpretations: it is not the imposition or relaxation of restrictions that shifts conspiracy search, but the sustained state of restriction itself.

The right panel reports regressions on the contemporaneous level of stringency and its first through fourth weekly lags, together with a four-week moving average. All specifications yield positive and individually significant coefficients. This is unsurprising given that stringency is highly autocorrelated across weeks; each lag largely proxies for the current level of restriction. The four-week moving average coefficient is comparable in magnitude to the contemporaneous specification, consistent with the effect reflecting sustained rather than transient exposure. Appendix J reports further specifications, including a model that enters four weekly lags jointly and one that combines the current level with change terms in a single regression; the conclusions are unchanged.

A related concern is that this persistence pattern is also consistent with a simpler time-use channel: confinement leaves people with more idle time to fill, which could increase search across a broad range of undemanding content, conspiracy-related or otherwise, without reflecting weakened institutional anchoring specifically. Brodeur et al. (2021) document a sharp rise in boredom-related search terms under the same lockdowns, consistent with this channel. The share normalisation underlying the Google Trends index (Section 3.1) addresses one version of this concern, since it nets out any mechanical increase in conspiracy search driven purely by higher overall online activity during lockdowns. It does not address a compositional shift in which a larger share of search activity is reallocated toward low-stakes or fringe content in general, of which conspiracy topics are one instance; the present design cannot separate this channel from social isolation directly.

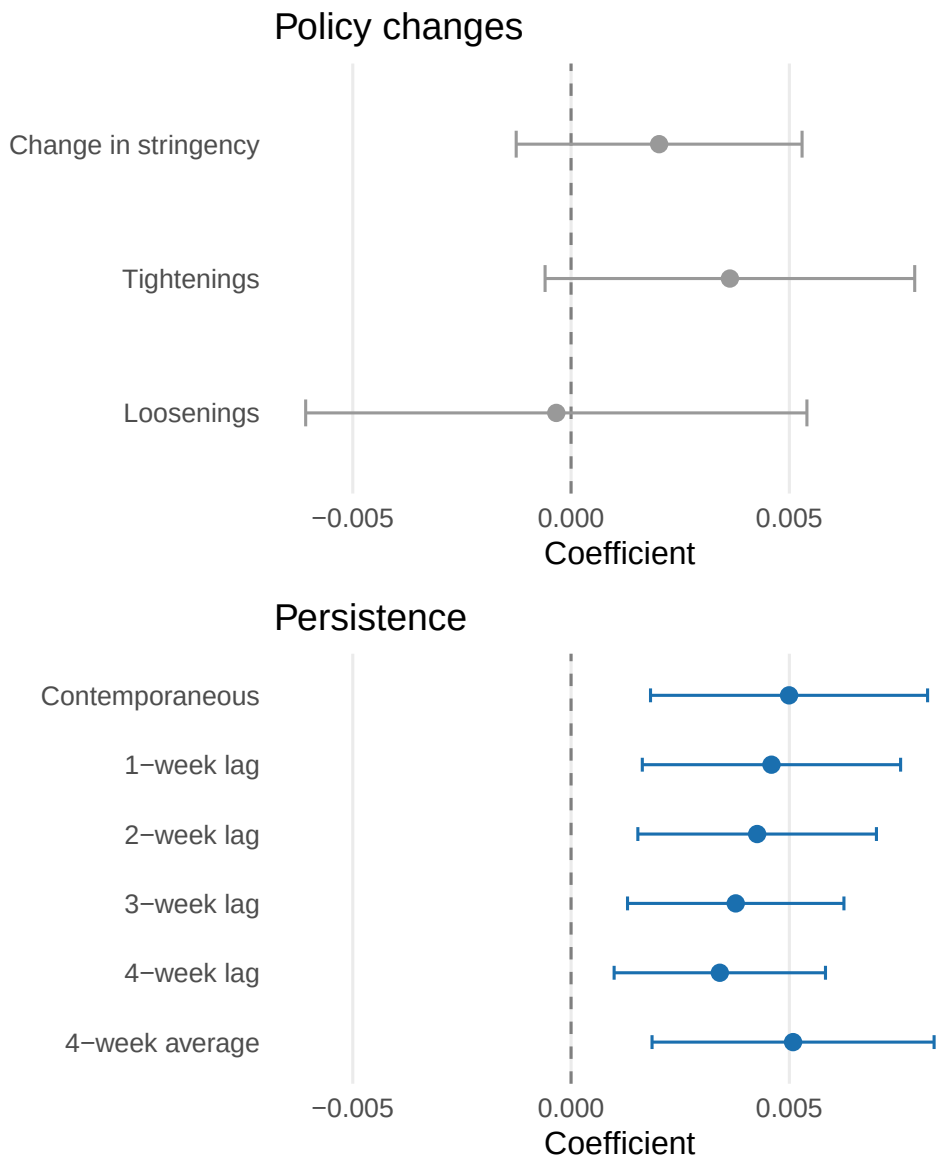


Figure 7: Lag structure of lockdown stringency on the conspiracy index. Left panel: coefficient on the first difference of stringency and its tightening/loosening decomposition, each estimated in a separate regression. Right panel: coefficient on the contemporaneous level, individual weekly lags, and a four-week moving average, each estimated in a separate regression. Both panels share the same x-axis scale. Error bars are 95% confidence intervals using $t(30)$ critical values.

6.2. The United States

An analogous analysis for the United States, exploiting within-state variation in lockdown intensity with week fixed effects, yields a null result (reported in Appendix K).[4] The coefficient on stringency is 0.0008 (SE = 0.0024) for the composite index in the fully specified model, and no individual topic shows a significant positive response. The U.S. conspiracy landscape during 2020-2021 was dominated by partisan identity; vaccine scepticism was strongly aligned

[4]The U.S. topics to pass the screening process were: QAnon, Propaganda, Brainwashing, Fake news, New world order, and Cover-up.

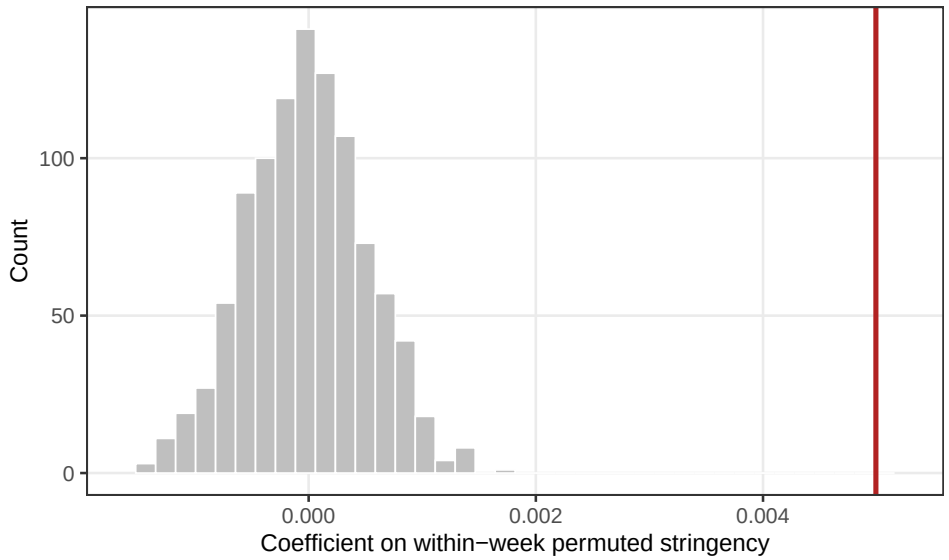


Figure 8: Distribution of 1,000 permuted coefficients from the within-week randomisation placebo. The red line marks the true estimated coefficient (0.0050).

with partisan affiliation (Allcott et al., 2020), as were prominent conspiracy theories such as "QAnon" and "Stop the Steal".[5] Conspiracy search behaviour may therefore have been saturated by political identity rather than responsive to lockdown-induced isolation. A further candidate explanation is that the U.S. screening process yields only six topics compared to the EU's eight, which may reduce the composite index's power to detect a comparable effect; this cannot be fully distinguished from the partisan-saturation explanation with the current data. The U.S. null also provides a useful contrast: if the European result were driven by a spurious time-series correlation between lockdowns and conspiracy interest, one would expect to observe a similar pattern in the United States.

7. Robustness

7.1. Randomisation Placebo

To assess whether the estimated coefficient could arise mechanically from the panel structure rather than genuine cross-country variation in stringency, I conduct a within-week randomisation placebo test. In each of 1,000 replications, country-level stringency values are randomly permuted across countries within each calendar week, and the main TWFE regression is re-estimated on the permuted dataset. Under this permutation, the within-week identifying variation is pure noise: any cross-country differences in stringency within a given week are by construction unrelated to anything that could affect conspiracy search. The distribution of permuted coefficients provides an empirical null distribution — the distribution of estimates that would arise if the within-week cross-country assignment of stringency to countries were random.

[5] See, for example, https://www.pewresearch.org/journalism/2020/09/16/political-divides-conspiracy-theories-and-divergent-news-sources-heading-into-2020-election/.

Figure 8 reports the distribution of permuted coefficients. The true coefficient (0.0050) lies 9.6 standard deviations above the mean of the permuted distribution (mean = −0.000019, SD = 0.000522). None of the 1,000 permuted regressions yielded a coefficient as large in absolute value as the true estimate (permutation p-value < 0.001). This indicates that the estimated effect is driven by genuine cross-country variation in stringency within weeks, and is not an artefact of common time trends, seasonal patterns, or the panel structure.

7.2. Difference-in-Differences

The main analysis exploits continuous variation in policy stringency across countries and over time. While this approach captures the effects of sustained differences in policy intensity, it does not isolate the discrete impact of entering lockdown during the first wave of the pandemic. To examine this initial transition more directly, and to provide corroborating evidence using a qualitatively different source of identification, I implement an alternative difference-in-differences design following Brodeur et al. (2021), which uses the same calendar period in 2019 as an explicit counterfactual rather than relying on cross-country stringency variation.[6]

For each country c, I define a country-specific lockdown date as the first week in 2020 in which the OxCGRT composite stringency index exceeds 30, corresponding approximately to the onset of meaningful first-wave restrictions such as school closures and public event cancellations.[7] I then construct a binary indicator Post_{ct} equal to one whenever the week of observation t is greater than or equal to the lockdown week of country c. The sample covers January through April in both 2019 and 2020,[8] yielding approximately 10–11 pre-lockdown weeks (January to early March) and 4–6 post-lockdown weeks (mid-March to April) for most countries. This window is chosen because all major first-wave lockdowns had occurred by April 2020, making it the natural endpoint for the first-wave design. In 2020, $\text{Post}_{ct} = 1$ identifies the period of active lockdown restrictions; in 2019, it flags the same calendar weeks, which experienced no lockdown and serve as the counterfactual seasonal baseline.

An important feature of the first-wave timing distribution is that 29 of the 31 European countries crossed the stringency-30 threshold within a single 11-day window (8–18 March 2020), with only Italy and France entering meaningfully earlier (22 February and 29 February, respec-

[6] Brodeur et al. (2021) use daily Google Trends data, which allows finer identification of onset timing and a cleaner separation of announcement from implementation effects. The series used here are only available at weekly frequency; daily disaggregation would also compound the zero-search-interest problem, since sparsely searched topics that record positive weekly counts would register zeros on most individual days.

[7] The DiD estimate is not sensitive to this specific cutoff: redefining the lockdown date using thresholds of 20, 25, 35, or 40 instead of 30 yields pooled coefficients in the range 0.534–0.550, all significant at the 1% level.

[8] The DiD analysis requires a separate Google Trends download covering January 2019 to April 2020. Google Trends normalises each download independently to its own internal peak, so a single download spanning January 2019 to April 2020 is required to ensure the 2019 and 2020 series are on a common scale. The period covers January through April, providing a sufficient pre-treatment window while encompassing all major first-wave lockdowns, which had occurred by April 2020. Brodeur et al. (2021) use the same January–April window in their original design, for the same scaling reason.

tively) (Appendix L). The variation in treatment timing is therefore largely clustered rather than staggered across the sample, which limits the scope for the heterogeneous-treatment-effect bias that arises in staggered difference-in-differences designs.

The estimating equation is:

$$y_{ct} = \alpha \left(\text{Post}_{ct} \times \mathbf{1}[\text{Year} = 2020]_t \right) + \beta \, \text{Post}_{ct} + \gamma \, \mathbf{1}[\text{Year} = 2020]_t + \delta_c + \delta_\tau + \varepsilon_{ct} \tag{4}$$

where y_{ct} is the IHS-transformed conspiracy search index for country c in week t, δ_c are country fixed effects, and δ_τ are week-of-year fixed effects that absorb any common seasonal pattern in search behaviour shared across years (distinct from the calendar-week fixed effects in the main specification, which absorb shocks specific to a given week in the sample). Standard errors are clustered by country.[9] No additional controls are included; the lagged case counts, death counts, and stringency values used in the main specification are all zero in 2019, rendering them uninformative in a pooled 2019–2020 regression. The coefficient of interest is α, the change in conspiracy search interest in the post-lockdown calendar period in 2020, relative to the change in the same calendar period in 2019. The key identifying assumption is that, absent lockdown restrictions, the seasonal evolution of conspiracy search interest in 2020 would have followed the same within-year pattern as in 2019.

[9]Standard errors are clustered at the country level, applying $t(G-1) = t(30)$ critical values as in the main analysis. Country-level clustering is preferred because identification in the DiD rests on country-specific lockdown timing, making the country the natural unit for within-cluster correlation. For comparison, clustering by time period yields a standard error on the interaction-term coefficient of SE = 0.0522; the larger standard error from country-level clustering indicates that the preferred specification is the more conservative of the two. Appendix L reports both specifications.

Table 4: Difference-in-Differences: Effect of Lockdown on Conspiracy Index.

Model:	(1)
Variables	
Post-lockdown (same calendar period)	-0.1769**
	(0.0738)
Year 2020	-0.0534
	(0.0391)
DiD: Post × 2020	0.5379***
	(0.0919)
Fixed-effects	
Country	Yes
Week	Yes
Fit statistics	
Observations	2,170
R^2	0.87716

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Post = 1 from first week OxCGRT Stringency > 30.
Applied identically in 2019 (placebo) and 2020 (treatment year).
Sample: January 2019–April 2020, 31 European countries.

Table 4 reports the estimates. The DiD coefficient is $\hat{\alpha} = 0.5379$ (SE = 0.0919), indicating that conspiracy search interest rose by approximately 0.54 standard deviations in the post-lockdown period of 2020 relative to what the 2019 seasonal baseline would predict. The onset of the first lockdown corresponded on average to an increase in the stringency index of approximately 80 points; scaling the main TWFE coefficient of 0.0050 by this magnitude implies an effect of approximately 0.40 standard deviations, broadly consistent with the difference-in-differences estimate of 0.54. The convergence of two approaches using qualitatively different sources of identification suggests that the result is robust to the choice of identification strategy.

A limitation of this design is that lockdown restrictions coincided with pandemic onset in spring 2020, so the estimate may capture the combined effect of lockdowns and contemporaneous pandemic-related factors, such as mortality salience and news coverage, rather than the lockdown-only effect in isolation. The DiD estimate is therefore best interpreted as corroborating the directional finding of the main TWFE specification.

8. Conclusion

This paper provides evidence that Covid-19 lockdown policies increased demand for online conspiracy narratives in Europe. Using within-country variation in lockdown stringency and week fixed effects, I find that European countries with more stringent lockdowns exhibited higher relative search interest in conspiracy-related content during the 2020–2021 lockdown period.

Several pieces of evidence point to social isolation as the main mechanism. First, the effect operates through sustained levels of restriction rather than transient policy shocks: lagged stringency predicts conspiracy search up to four weeks later, while week-on-week changes in stringency (whether tightenings or loosenings) have no independent effect. This is consistent with the ongoing deprivation of peer interaction and daily routine, not the initial shock of entering lockdown, shifting attention toward conspiratorial content.

Second, decomposing the stringency index into individual policy components reveals that restrictions that physically isolate people (workplace closures, school closures, and stay-at-home orders) are the specific drivers of the effect. Facial covering mandates, which require behavioural compliance but do not isolate individuals from social contact, have no effect. This distinction between isolation and mere policy intrusiveness provides direct evidence that the mechanism operates through the disruption of social interaction rather than through generalised resentment of government intervention. This pattern is consistent with theories of identity and belief formation (Akerlof & Kranton, 2000), which emphasise that individuals seek narratives that reinforce identity-consistent interpretations of events. Social isolation may weaken institutional anchoring and peer moderation, increasing susceptibility to identity-reinforcing conspiracy narratives.

Third, controlling for the OxCGRT Economic Support Index does not attenuate the stringency coefficient. Economic support itself enters negatively, consistent with economic hardship independently contributing to conspiracy demand. However, the stringency coefficient is unaffected by this control, indicating that the isolation channel operates distinctly from any economic anxiety mechanism. While the political economy literature has emphasised economic grievance as a driver of radicalisation (Autor et al., 2020; Colantone & Stanig, 2019; Fetzer, 2019; Funke et al., 2016), the present result suggests that non-economic institutional shocks, specifically disruptions to social capital, can generate similar demand-side shifts toward anti-establishment narratives.

The composite null result in the United States is consistent with a partisan saturation story. Where conspiracy narratives are already tightly bound to partisan identity, the marginal effect of lockdown-induced isolation may be zero: individuals' information-seeking behaviour may be driven more by political affiliation than by the social environment. The European context, where conspiracy narratives are less tightly aligned with partisan identity, allows the environmental channel to be identified.

From a policy perspective, these findings highlight an unintended informational cost of severe lockdown policies. They do not imply that lockdowns were unwarranted, but suggest that policy-

makers should consider complementary interventions which limit social isolation and strengthen communication during periods of enforced isolation. The finding that isolation-inducing restrictions specifically drive the result, while non-isolating measures such as mask mandates do not, suggests that the informational costs of lockdowns are concentrated in the most socially disruptive policies.

Whereas the literature on belief formation and political attitudes (Allcott et al., 2020; DellaVigna & Kaplan, 2007; Enikolopov et al., 2011) has focused largely on media shocks, the findings here suggest that conspiracy narratives should be understood not only as products of online platforms, algorithms, and supply of misinformation, but also as endogenous responses to social and institutional shocks. These dynamics are economically and socially consequential: demand for conspiratorial content is closely linked to compliance with public policies, institutional trust, political participation, and other outcomes of social and political consequence. Future work could examine whether similar patterns arise following other large-scale disruptions that reduce social interaction (such as natural disasters), and whether the effects documented here translate into downstream outcomes such as vaccine hesitancy, voting behaviour, or hate crime.

Acknowledgements

I thank Kevin Denny, Sarah Parlane, Nora Strecker, and particularly Paul Devereux, as well as seminar participants at University College Dublin, for helpful comments on earlier versions of this paper.

Data Availability

All data used in this paper are publicly available: Google Trends search data; the Oxford Covid-19 Government Response Tracker (OxCGRT); Our World in Data's Covid-19 dataset; the U.S. Centers for Disease Control and Prevention's Covid-19 surveillance data; World Bank data on internet penetration, urbanisation, and GDP per capita; the Reporters Without Borders Press Freedom Index; and the EF English Proficiency Index. A full replication package, including code and processed data, can be provided upon request.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Competing Interests

The author declares that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, *115*(3), 715–753. https://doi.org/10.1162/003355300554881

Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., & Yang, D. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. *Journal of Public Economics*, *191*, 104254. https://doi.org/10.1016/j.jpubeco.2020.104254

Autor, D., Dorn, D., Hanson, G., & Majlesi, K. (2020). Importing political polarization? The electoral consequences of rising trade exposure. *American Economic Review*, *110*(10), 3139–3183. https://doi.org/10.1257/aer.20170011

Berger, L. M., Ferrari, G., Leturcq, M., Panico, L., & Solaz, A. (2021). COVID-19 lockdowns and demographically-relevant Google Trends: A cross-national analysis. *PLOS ONE*, *16*(3), e0248072. https://doi.org/10.1371/journal.pone.0248072

Böhme, M. H., Gröger, A., & Stöhr, T. (2020). Searching for a better life: Predicting international migration with online search keywords. *Journal of Development Economics*, *142*, 102347. https://doi.org/10.1016/j.jdeveco.2019.04.002

Brodeur, A., Clark, A. E., Fleche, S., & Powdthavee, N. (2021). COVID-19, lockdowns and well-being: Evidence from Google Trends. *Journal of Public Economics*, *193*, 104346. https://doi.org/10.1016/j.jpubeco.2020.104346

Cameron, C. A., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*(2), 317–372. https://doi.org/10.3368/jhr.50.2.317

Colantone, I., & Stanig, P. (2019). The surge of economic nationalism in Western Europe. *Journal of Economic Perspectives*, *33*(4), 128–151. https://doi.org/10.1257/jep.33.4.128

Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *The Journal of Finance*, *66*(5), 1461–1499. https://doi.org/10.1111/j.1540-6261.2011.01679.x

DellaVigna, S., & Kaplan, E. (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, *122*(3), 1187–1234. https://doi.org/10.1162/qjec.122.3.1187

Donald, S. G., & Lang, K. (2007). Inference with difference-in-differences and other panel data. *The Review of Economics and Statistics*, *89*(2), 221–233. https://doi.org/10.1162/rest.89.2.221

Enikolopov, R., Petrova, M., & Zhuravskaya, E. (2011). Media and political persuasion: Evidence from Russia. *American Economic Review*, *101*(7), 3253–3285. https://doi.org/10.1257/aer.101.7.3253

Fetzer, T. (2019). Did austerity cause Brexit? *American Economic Review*, *109*(11), 3849–3886. https://doi.org/10.1257/aer.20181164

Funke, M., Schularick, M., & Trebesch, C. (2016). Going to extremes: Politics after financial

crises, 1870–2014. *European Economic Review*, *88*, 227–260. https://doi.org/10.1016/j.euroecorev.2016.03.006

Leslie, E., & Wilson, R. (2020). Sheltering in place and domestic violence: Evidence from calls for service during COVID-19. *Journal of Public Economics*, *189*, 104241. https://doi.org/10.1016/j.jpubeco.2020.104241

Li, H., Zhang, N., Ma, X., Wang, Y., Yang, F., Wang, W., Huang, Y., Xie, Y., & Du, Y. (2025). Google Trends applications for COVID-19 pandemic: A bibliometric analysis. *Digital Health*, *11*, 20552076241310055. https://doi.org/10.1177/20552076241310055

Müller, K., & Schwarz, C. (2021). Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*, *19*(4), 2131–2167. https://doi.org/10.1093/jeea/jvaa045

Rosa, P. A. de la, Cowden, R. G., Filippis, R. de, Jerotic, S., Nahidi, M., Ori, D., Orsolini, L., Nagendrappa, S., Pinto da Costa, M., Ransing, R., Saeed, F., Shoib, S., Turan, S., Ullah, I., Vadivel, R., & Ramalho, R. (2022). Associations of lockdown stringency and duration with google searches for mental health terms during the COVID-19 pandemic: A nine-country study. *Journal of Psychiatric Research*, *150*, 237–245. https://doi.org/10.1016/j.jpsychires.2022.03.026

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics*, *118*, 26–40. https://doi.org/10.1016/j.jpubeco.2014.04.010

Van Es, L., Van Prooijen, J.-W., Van Buuren, J., & Van Veldhuizen, A. P. (2025). From 'Great Awakening' to sending death threats: A study of police-recorded crimes driven by conspiracy theories in the Netherlands. *European Journal of Criminology*, 24. https://doi.org/10.1177/14773708251351728

Wedel, L., & Coufal, L. (2025). Radicalization or relief: Divergent impacts of the COVID-19 lockdown on incels by seniority. *New Media & Society*, *0*(0), 14614448251336432. https://doi.org/10.1177/14614448251336432

A. Google Trends normalisation: a numerical example

The following example illustrates how the Google Trends index is constructed. Suppose that in a given week t, Germany records a total of 1,000 Google searches. Of these, 40 are attributed to the topic "QAnon", 20 to "Propaganda", and 5 to "Flat Earth". The corresponding search shares are 0.040, 0.020, and 0.005. Suppose further that across the entire sample period, the highest observed share for each topic in Germany was 0.080 for QAnon, 0.020 for Propaganda, and 0.040 for Flat Earth.

The Google Trends index for each topic in week t is then:

Topic	Searches	Share	Max share	GT index	IHS(GT)
QAnon	40	0.040	0.080	$100 \times 0.040/0.080 = 50$	4.61
Propaganda	20	0.020	0.020	$100 \times 0.020/0.020 = 100$	5.30
Flat Earth	5	0.005	0.040	$100 \times 0.005/0.040 = 12.5$	3.22

This construction exhibits several notable features. First, each topic is normalised to its own sample maximum, so a value of 100 does not mean a topic was the most-searched in that week — it means that week was the peak for that topic in that country. In the example above, Propaganda achieves GT = 100 because week t happens to be its own peak, even though it has fewer raw searches than QAnon. Second, the normalisation is region-specific: the same search share may correspond to different GT values in different countries if the country-level peaks differ. Third, the index captures search share rather than search volume, so aggregate increases in overall online activity during lockdowns do not mechanically inflate the index. Fourth, the IHS transformation $\text{IHS}(x) = \log\left(x + \sqrt{x^2+1}\right)$ is applied to the GT index before entering the regression; for GT values above roughly 5 this closely approximates $\log(2x)$, providing a proportional-change interpretation for positive observations while preserving the zero values that arise when a topic records no searches in a given region-week.

B. Full list of variables considered

Table B.1: Topic screening metrics for all candidate European conspiracy topics.

Topic	Topic ID	Zero	Countries	Pre-Covid	Peak period
Propaganda	/m/05s0v	0.294	1.000	Yes	Feb 2022
Fake news	/g/1210rwkh	0.392	1.000	Yes	Mar 2020
QAnon	/g/11hc_gvb_r	0.515	0.968	Yes	Jan 2021
Flat Earth	/m/032dv	0.568	0.935	Yes	Feb 2020
New world order	/m/02vhjq	0.582	0.871	Yes	Mar 2020
David Icke	/m/0hvwn	0.682	0.968	Yes	Apr 2020
Chemtrails	/m/02ll82	0.684	0.903	Yes	Apr 2020
Brainwashing	/m/058bq	0.687	0.903	Yes	Dec 2020
Deep state	/m/04gkqtf	0.759	0.968	Yes	Oct 2022
Political corruption	/m/0hj49	0.880	0.871	Yes	Jan 2021
Cover-up	/m/014zmn	0.922	0.742	Yes	Apr 2020
False flag	/m/05cqz4	0.930	0.806	Yes	Nov 2020
Big pharma conspiracy theory	/m/010gq78_	0.936	0.613	Yes	Jul 2021
Plandemic	/g/11fv8jj_r0	0.937	0.742	Yes	Dec 2020
Illuminati	/g/121ths1c	0.948	0.548	Yes	Nov 2020
Great Replacement	/g/11bbsws11q	0.953	0.677	Yes	Oct 2021
Mainstream media	/m/0b6gkfv	0.958	0.613	Yes	Jan 2021
David Icke (alt)	/g/11fxfhbk3w	0.965	0.613	Yes	Mar 2022
Reptilian conspiracy theory	/g/121qfqbk	0.970	0.581	Yes	Oct 2021
COVID-19 lab leak theory	/g/11nxpm2x3x	0.985	0.548	No	Aug 2022
Great Replacement conspiracy	/g/11t15k2ps1	0.995	0.484	No	Apr 2021
Plandemic (alt)	/g/11rmt1q2b4	0.996	0.484	Yes	Dec 2022
Anti-vaccine activism	/g/11vf3mqfjc	0.997	0.516	No	Feb 2022
Russell Brand	/g/11j2nsd6x9	0.997	0.452	Yes	Nov 2021
Tucker Carlson	/g/11vkz0h1bm	0.997	0.419	No	Oct 2020
False flag (alt)	/g/11gfn91vdr	0.997	0.387	Yes	Jul 2021
Great Replacement (alt)	/g/11sk4prhjv	0.997	0.387	No	Jan 2021

Note: Zero is the share of country-week observations with zero search interest. Countries is the share of European countries with positive search interest at any point in the sample period. Pre-Covid indicates whether the topic had positive interest prior to March 2020. Peak period shows the month-year of the highest observed search interest across all country-weeks. Topic ID is the Google Trends internal semantic topic identifier. Where a topic has two Google Trends semantic topic identifiers, both rows are included; (alt) denotes the second.

C. Unscreened conspiracy index

The main analysis restricts the outcome to the eight conspiracy topics that pass all three screening criteria (fewer than 70% zero observations, positive search interest in at least 80% of European countries, and positive pre-pandemic interest). Table C.1 replicates the main specification using an index constructed from all candidate conspiracy topics, including those excluded by the screening criteria. The unscreened index is noisier: many excluded topics have very high zero rates, which mechanically attenuates the estimated effect. Nevertheless, the coefficient on the unscreened index in column 2 remains positive and significant. Column 1 reports the main result for comparison.

Table C.1: Robustness: Unscreened Index (All 27 Topics)

Model:	(1)	(2)
Variables		
Stringency index	0.0050***	0.0027**
	(0.0016)	(0.0013)
Covid cases (t-1, per 100k)	0.0498*	0.0582
	(0.0266)	(0.0637)
Covid deaths (t-1, per 1k)	-0.0158	0.0558**
	(0.0123)	(0.0235)
Economic support index	-0.0017**	-0.0009
	(0.0007)	(0.0007)
Fixed-effects		
Country	Yes	Yes
Week	Yes	Yes
Fit statistics		
Observations	3,224	3,224
Within R^2	0.02535	0.01940

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

D. Alternative outcomes and index constructions

Table D.1: Alternative Outcome Specifications and Index Constructions

Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
Stringency index	0.0050***	0.0051***	0.0015***	0.0013	0.0119***	0.0068***	0.0102***	0.0545**
	(0.0016)	(0.0015)	(0.0005)	(0.0013)	(0.0041)	(0.0021)	(0.0032)	(0.0257)
Covid cases (t-1, per 100k)	0.0498*	0.0526*	0.0330***	0.0604***	0.2639***	0.0664*	0.1008*	-0.6248*
	(0.0266)	(0.0274)	(0.0102)	(0.0178)	(0.0815)	(0.0353)	(0.0541)	(0.3448)
Covid deaths (t-1, per 1k)	-0.0158	-0.0173	-0.0212***	-0.0416***	-0.1695***	-0.0206	-0.0310	0.7100**
	(0.0123)	(0.0125)	(0.0045)	(0.0107)	(0.0358)	(0.0167)	(0.0245)	(0.3027)
Economic support index	-0.0017**	-0.0017**	-0.0005**	-0.0009	-0.0041**	-0.0023**	-0.0033**	-0.0217*
	(0.0007)	(0.0007)	(0.0002)	(0.0007)	(0.0017)	(0.0010)	(0.0015)	(0.0127)
Fixed-effects								
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Observations	3,224	3,224	3,224	3,224	3,224	3,224	3,224	3,224
Within R^2	0.02535	0.02627	0.03430	0.01299	0.03430	0.02479	0.02520	0.01871

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table D.1 reports robustness checks for the main specification across eight alternative outcome and index measures. Column 1 is the IHS transformation used in the main analysis. Column 2 winsorises the IHS outcome at the 1st and 99th percentiles to reduce the influence of extreme values. Column 3 is the extensive margin (share of conspiracy topics with any positive search activity in a country-week). Column 4 is a binary above-median indicator. Column 5 is the count of topics with positive search (an integer between 0 and 8). Column 6 uses a simple (unweighted) average of raw IHS values across topics, without z-score standardisation. Column 7 uses the first principal component of the topic-level IHS matrix, a data-driven linear combination that maximises common variation across topics. Column 8 uses the simple average of raw Google Trends hit counts (0-100) without any transformation. The stringency coefficient is positive and significant across seven of the eight specifications, indicating that the result is not sensitive to the choice of outcome measure or index aggregation method.

E. IHS scaling sensitivity

Table E.1 replicates the main specification with the raw Google Trends index scaled by factors of 0.1 and 10 before the IHS transformation. Columns 1–3 correspond to scales of $\times 1$ (baseline), $\div 10$, and $\times 10$ respectively. The stringency coefficient is stable across all three scales (0.0047–0.0055, all significant at the 1% level), confirming that the result is not sensitive to the unit of measurement.

Table E.1: Scaling Sensitivity: IHS at Different Raw Index Scales

Model:	(1)	(2)	(3)
Variables			
Stringency index	0.0050***	0.0055***	0.0047***
	(0.0016)	(0.0018)	(0.0015)
Covid cases (t-1, per 100k)	0.0499*	-0.0010	0.0659**
	(0.0266)	(0.0271)	(0.0270)
Covid deaths (t-1, per 1k)	-0.0159	0.0255	-0.0322**
	(0.0123)	(0.0156)	(0.0119)
Economic support index	-0.0017**	-0.0019**	-0.0016**
	(0.0007)	(0.0009)	(0.0007)
Fixed-effects			
Country	Yes	Yes	Yes
Week	Yes	Yes	Yes
Fit statistics			
Observations	3,224	3,224	3,224
Within R^2	0.02536	0.02372	0.02741

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

F. Wild Cluster Bootstrap

Table F.1 reproduces the main sequential-controls table (Table 3) using wild cluster bootstrap (WCB) inference with Webb weights ($B = 9{,}999$ draws) in place of $t(G-1)$ critical values. Cluster-robust standard errors are shown in parentheses and WCB p-values in brackets; significance stars are based on the WCB p-values. All three columns retain significance at the 1% level, confirming that the main result is not an artefact of the $t(G-1)$ asymptotic approximation.

Figures F.1 and F.2 reproduce the topic-coefficient and policy-component figures with WCB confidence intervals in place of $t(30)$ critical values. The pattern of results is qualitatively identical.

Table F.1: Wild Cluster Bootstrap: Lockdown Stringency and Conspiracy Search

	(1)	(2)	(3)
Stringency index	0.004607***	0.004660***	0.004994***
	(0.001531)	(0.001518)	(0.001554)
	[p = 0.006]	[p = 0.005]	[p = 0.004]
Covid cases ($t-1$, per 100k)		0.05848**	0.04985**
		(0.02515)	(0.02656)
		[p = 0.016]	[p = 0.021]
Covid deaths ($t-1$, per 1k)		-0.01765	-0.01585
		(0.01206)	(0.01231)
		[p = 0.164]	[p = 0.178]
Economic support index			-0.001708**
			(0.0007291)
			[p = 0.049]
Country FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Observations	3,255	3,224	3,224

Cluster-robust (country) standard errors in parentheses; WCB p-values (Webb weights, $B = 9{,}999$) in brackets.
Significance stars based on WCB p-values: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
Reproduces Table 3 with wild cluster bootstrap inference.

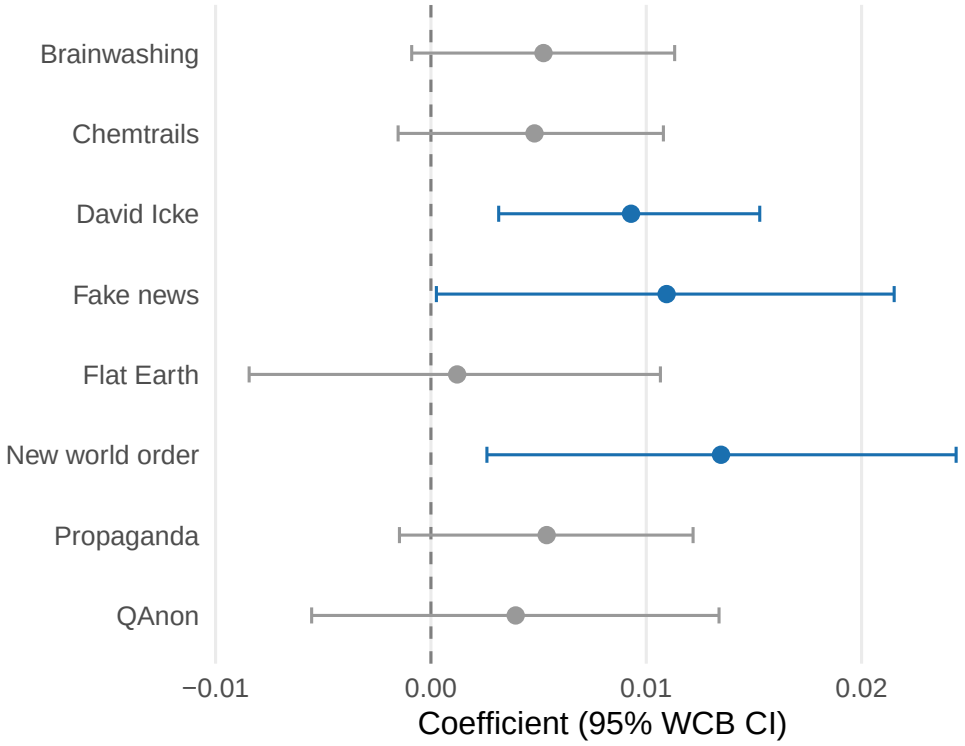


Figure F.1: Effect of lockdown stringency on individual conspiracy topics: wild cluster bootstrap. Reproduces Figure 5 with 95% WCB confidence intervals (Webb weights, $B = 9{,}999$). Blue points are significant at the 5% level under WCB.

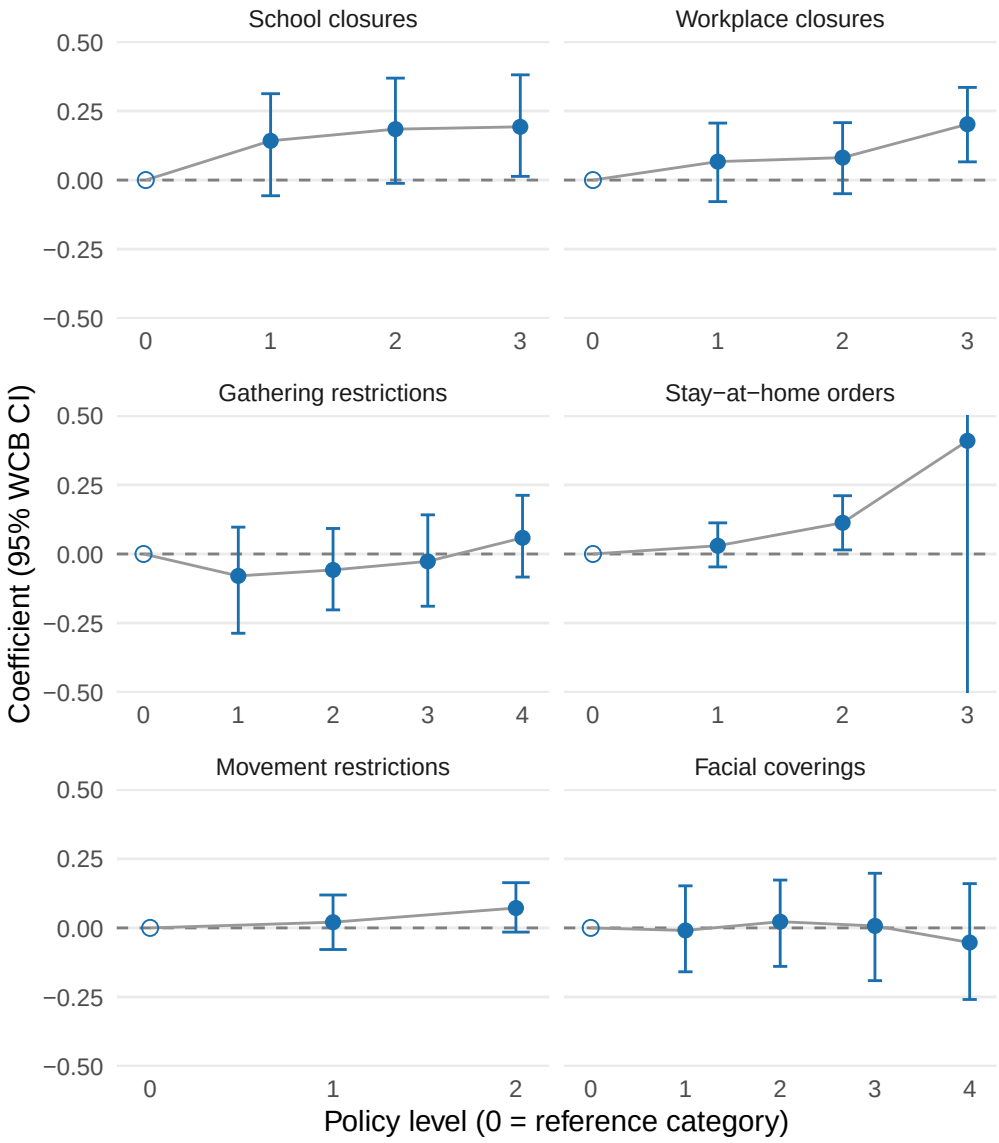


Figure F.2: Dose–response coefficients for individual lockdown policy components: wild cluster bootstrap. Reproduces Figure 4 with 95% WCB confidence intervals (Webb weights, $B = 9{,}999$). The y-axis is capped for readability; the confidence interval for stay-at-home orders at level 3 extends well beyond the axis because very few countries implemented that restriction level, leaving the bootstrap with very limited cross-country variation.

G. Cross-country heterogeneity and sensitivity

To test whether the lockdown–conspiracy relationship varies with country-level characteristics, I interact the stringency index with measures of internet penetration, GDP per capita (PPP), and urbanisation rate (all World Bank, 2019 data), press freedom (Reporters Without Borders Press Freedom Index, 2019), and English proficiency.[10] Continuous moderators are z-standardised (mean 0, SD 1); English proficiency is a binary indicator. Each column of Table G.1 reports a separate model including one interaction term. Three of the five interaction terms are statistically insignificant, and the main stringency coefficient remains positive and significant across all five specifications (ranging from 0.0039 to 0.0088), suggesting that the isolation mechanism operates broadly across European countries, albeit with some variation in magnitude across moderators.

Two interactions are marginally significant. The interaction term Stringency × Internet penetration is negative, indicating that countries with higher internet penetration exhibit a smaller response of conspiracy search to lockdown stringency. This is counterintuitive if one expects broader internet access to amplify exposure to conspiracy content. One possible interpretation is that in highly connected countries, individuals already rely heavily on digital information sources, so lockdowns induce smaller marginal changes in information environments and online attention patterns.

The interaction term Stringency × English proficiency is positive and marginally significant, indicating that countries with high English proficiency exhibit a somewhat stronger conspiracy search response to lockdown stringency; this result is sensitive to the classification of Austria and Belgium as high-proficiency countries, and loses significance if either is excluded. The most obvious interpretation is that high-proficiency countries are more integrated into English-language online ecosystems, where conspiracy content is more abundant and readily accessible, so lockdown-induced increases in online attention translate more readily into elevated conspiracy search.

Table G.2 reports leave-one-country-out estimates for the main specification. Each row drops one of the 31 European countries and re-estimates the model on the remaining 30. The estimated coefficient is stable across all 31 specifications, indicating that the main result is not driven by any single country.

[10] Countries classified as high-proficiency are those rated "Very High" on the EF English Proficiency Index (2023 European rankings), plus the United Kingdom and Ireland as native-speaker countries: Netherlands, Sweden, Norway, Denmark, Finland, Austria, Belgium, United Kingdom, and Ireland.

Table G.1: Cross-Country Heterogeneity: Stringency Interactions with Country Characteristics

Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Stringency index	0.0054***	0.0088***	0.0057***	0.0052***	0.0039**
	(0.0019)	(0.0027)	(0.0018)	(0.0015)	(0.0016)
Covid cases (t-1, per 100k)	0.0496*	0.0472*	0.0484*	0.0504*	0.0537**
	(0.0265)	(0.0272)	(0.0269)	(0.0268)	(0.0254)
Covid deaths (t-1, per 1k)	-0.0163	-0.0175	-0.0138	-0.0163	-0.0163
	(0.0123)	(0.0128)	(0.0123)	(0.0123)	(0.0116)
Economic support index	-0.0017**	-0.0016**	-0.0017**	-0.0017**	-0.0018**
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0007)
Stringency × GDP p.c.	-0.0003				
	(0.0011)				
Stringency × Internet		-0.0039*			
		(0.0022)			
Stringency × Urbanisation			-0.0015		
			(0.0016)		
Stringency × Press freedom				-0.0004	
				(0.0006)	
Stringency × English proficiency					0.0030*
					(0.0016)
Fixed-effects					
Country	Yes	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	3,224	3,224	3,224	3,224	3,224
Within R^2	0.02564	0.03032	0.02787	0.02591	0.03284

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table G.2: Leave-One-Country-Out Robustness: European Stringency on Conspiracy Index

Dropped country	Coefficient	SE	p-value	N	G	Sig
Austria	0.00508	0.00158	0.0031	3120	30	***
Belgium	0.00471	0.00158	0.0059	3120	30	***
Bulgaria	0.00514	0.00158	0.0029	3120	30	***
Switzerland	0.00528	0.00154	0.0018	3120	30	***
Cyprus	0.00537	0.00154	0.0016	3120	30	***
Czechia	0.00471	0.00158	0.0057	3120	30	***
Germany	0.00508	0.00162	0.0038	3120	30	***
Denmark	0.00500	0.00162	0.0045	3120	30	***
Estonia	0.00495	0.00159	0.0042	3120	30	***
Spain	0.00494	0.00160	0.0043	3120	30	***
Finland	0.00502	0.00161	0.0040	3120	30	***
France	0.00512	0.00159	0.0032	3120	30	***
United Kingdom	0.00490	0.00162	0.0051	3120	30	***
Greece	0.00511	0.00159	0.0032	3120	30	***
Croatia	0.00420	0.00143	0.0066	3120	30	***
Hungary	0.00499	0.00161	0.0042	3120	30	***
Ireland	0.00441	0.00154	0.0078	3120	30	***
Iceland	0.00474	0.00160	0.0059	3120	30	***
Italy	0.00505	0.00162	0.0041	3120	30	***
Lithuania	0.00559	0.00149	0.0008	3120	30	***
Luxembourg	0.00502	0.00159	0.0038	3120	30	***
Latvia	0.00457	0.00152	0.0054	3120	30	***
Malta	0.00520	0.00159	0.0028	3120	30	***
Netherlands	0.00497	0.00159	0.0041	3120	30	***
Norway	0.00549	0.00150	0.0010	3120	30	***
Poland	0.00509	0.00162	0.0038	3120	30	***
Portugal	0.00479	0.00159	0.0053	3120	30	***
Romania	0.00534	0.00157	0.0020	3120	30	***
Sweden	0.00466	0.00164	0.0079	3120	30	***
Slovenia	0.00525	0.00159	0.0026	3120	30	***
Slovakia	0.00492	0.00160	0.0045	3120	30	***

Note: Each row drops one country and re-estimates the week FE specification on the 2020–2021 sample. p-values use t(G-1) correction.

H. Ordinal indicators as level dummies

Tables H.1 and H.2 report the full regression results underlying Figure 4 of the main text.

Table H.1: Ordinal Policy Indicators as Level Dummies: School Closures, Workplace Closures, Gathering Restrictions

Model:	(1)	(2)	(3)	(4)
Variables				
Stringency index (continuous)	0.0050***			
	(0.0016)			
Covid cases (t-1, per 100k)	0.0498*	0.0479	0.0516*	0.0551*
	(0.0266)	(0.0284)	(0.0292)	(0.0280)
Covid deaths (t-1, per 1k)	-0.0158	-0.0083	-0.0213	-0.0088
	(0.0123)	(0.0122)	(0.0141)	(0.0128)
Economic support index	-0.0017**	-0.0015*	-0.0018**	-0.0017**
	(0.0007)	(0.0007)	(0.0008)	(0.0008)
School: level 1		0.1420*		
		(0.0761)		
School: level 2		0.1843**		
		(0.0801)		
School: level 3		0.1927**		
		(0.0807)		
Workplace: level 1			0.0668	
			(0.0650)	
Workplace: level 2			0.0813	
			(0.0624)	
Workplace: level 3			0.2021***	
			(0.0659)	
Gatherings: level 1				-0.0794
				(0.0859)
Gatherings: level 2				-0.0577
				(0.0702)
Gatherings: level 3				-0.0267
				(0.0760)
Gatherings: level 4				0.0585
				(0.0693)
Fixed-effects				
Country	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes
Fit statistics				
Observations	3,224	3,224	3,224	3,224
Within R^2	0.02535	0.01893	0.02289	0.02033

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table H.2: Ordinal Policy Indicators as Level Dummies: Stay-at-Home Orders, Movement Restrictions, Facial Coverings

Model:	(1)	(2)	(3)
Variables			
Stay-at-home: level 1	0.0298		
	(0.0389)		
Stay-at-home: level 2	0.1130**		
	(0.0484)		
Stay-at-home: level 3	0.4096***		
	(0.1180)		
Covid cases (t-1, per 100k)	0.0501*	0.0531*	0.0524*
	(0.0273)	(0.0269)	(0.0294)
Covid deaths (t-1, per 1k)	-0.0142	-0.0112	-0.0072
	(0.0130)	(0.0121)	(0.0124)
Economic support index	-0.0017**	-0.0015*	-0.0014*
	(0.0008)	(0.0007)	(0.0008)
Movement: level 1		0.0206	
		(0.0483)	
Movement: level 2		0.0720	
		(0.0445)	
Masks: level 1			-0.0092
			(0.0730)
Masks: level 2			0.0225
			(0.0727)
Masks: level 3			0.0072
			(0.0928)
Masks: level 4			-0.0528
			(0.0988)
Fixed-effects			
Country	Yes	Yes	Yes
Week	Yes	Yes	Yes
Fit statistics			
Observations	3,224	3,224	3,224
Within R^2	0.02024	0.01342	0.01203

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

I. Individual conspiracy topics

Table I.1 reports the full regression results underlying Figure 5 of the main text. Each column regresses the IHS-transformed topic-level search index on the stringency index with country and week fixed effects, lagged epidemiological controls, and the OxCGRT Economic Support Index.

Table I.1: Effect of Lockdown Stringency on Individual Topics

	Brainwashing	Chemtrails	David Icke	Fake news	Flat Earth	New world order	Propaganda	QAnon
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
Stringency index	0.0052*	0.0048	0.0093***	0.0109**	0.0012	0.0135**	0.0054	0.0039
	(0.0030)	(0.0030)	(0.0030)	(0.0052)	(0.0046)	(0.0054)	(0.0033)	(0.0047)
Covid cases (t-1, per 100k)	0.0477	-0.1228**	0.2255**	0.0520	0.1197*	0.1313**	0.0706*	0.0074
	(0.0346)	(0.0496)	(0.0880)	(0.0675)	(0.0587)	(0.0630)	(0.0403)	(0.0581)
Covid deaths (t-1, per 1k)	-0.0261	-0.0093	-0.0721*	0.0155	-0.0298	0.0030	-0.0301	-0.0158
	(0.0181)	(0.0299)	(0.0382)	(0.0305)	(0.0271)	(0.0416)	(0.0290)	(0.0421)
Economic support index	-0.0016	-0.0037***	-0.0022	-0.0002	-0.0011	-0.0039	-0.0013	-0.0044
	(0.0011)	(0.0013)	(0.0026)	(0.0022)	(0.0018)	(0.0026)	(0.0017)	(0.0031)
Fixed-effects								
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics								
Observations	3,224	3,224	3,224	3,224	3,224	3,224	3,224	3,224
Within R^2	0.00201	0.00509	0.01248	0.00873	0.00133	0.01176	0.00357	0.00480

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

J. Lag structure: full results

Table J.1 reports the full lag structure results. In addition to the specifications shown in Figure 7 of the main text, the table includes a "Level + shocks" column that enters the contemporaneous level alongside both tightening and loosening change terms simultaneously, and an "All lags" column that enters four weekly lags jointly. In the level-plus-shocks specification the level coefficient is unchanged at $\hat{\beta} = 0.0051$ while both shock terms remain negligible. In the all-lags specification the contemporaneous coefficient survives at $\hat{\beta} = 0.0047$, while lags 1 through 4 become small and statistically indistinguishable from zero, reflecting multicollinearity induced by the high week-to-week autocorrelation in stringency.

Table J.1: Lag Structure of Lockdown Stringency on Conspiracy Index

Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Variables											
Stringency index	0.0050***									0.0051***	0.0047**
	(0.0016)									(0.0016)	(0.0020)
Covid cases (t-1, per 100k)	0.0498*	0.0516*	0.0524*	0.0531*	0.0524*	0.0514*	0.0488*	0.0487*	0.0496*	0.0500*	0.0487*
	(0.0266)	(0.0275)	(0.0286)	(0.0296)	(0.0300)	(0.0279)	(0.0283)	(0.0282)	(0.0286)	(0.0267)	(0.0275)
Covid deaths (t-1, per 1k)	-0.0158	-0.0155	-0.0157	-0.0152	-0.0147	-0.0172	-0.0083	-0.0079	-0.0084	-0.0155	-0.0178
	(0.0123)	(0.0125)	(0.0129)	(0.0132)	(0.0133)	(0.0128)	(0.0123)	(0.0123)	(0.0123)	(0.0123)	(0.0126)
Economic support index	-0.0017**	-0.0017**	-0.0017**	-0.0016**	-0.0016**	-0.0017**	-0.0015*	-0.0015*	-0.0015*	-0.0017**	-0.0017**
	(0.0007)	(0.0007)	(0.0007)	(0.0007)	(0.0008)	(0.0007)	(0.0008)	(0.0008)	(0.0008)	(0.0007)	(0.0007)
Stringency (t-1)		0.0046***									-0.0011
		(0.0014)									(0.0019)
Stringency (t-2)			0.0043***								0.0013
			(0.0013)								(0.0019)
Stringency (t-3)				0.0038***							-0.0011
				(0.0012)							(0.0014)
Stringency (t-4)					0.0034***						0.0015
					(0.0012)						(0.0014)
Stringency (4-week avg)						0.0051***					
						(0.0016)					
Δ Stringency							0.0020				
							(0.0016)				
Δ^+ Stringency								0.0036*		0.0011	
								(0.0021)		(0.0018)	
Δ^- Stringency									-0.0003	-0.0032	
									(0.0028)	(0.0028)	
Fixed-effects											
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics											
Observations	3,224	3,224	3,193	3,162	3,131	3,162	3,224	3,224	3,224	3,224	3,131
Within R^2	0.02535	0.02281	0.02062	0.01796	0.01607	0.02309	0.00986	0.01036	0.00932	0.02589	0.02428

Clustered (Country) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

K. United States results

All specifications use the 2020–2021 sample and include state and week fixed effects, with standard errors clustered by state. Inference uses $t(G-1) = t(49)$ critical values, as in the main analysis; with 50 clusters, these are virtually indistinguishable from the standard normal. The epidemiological controls (lagged cases and deaths per capita) are included; the Economic Support Index is not, as the OxCGRT subnational dataset does not provide state-level values for this variable.

Table K.1: Lockdown Stringency and Conspiracy Search: United States

Model:	(1)	(2)
Variables		
Stringency index	0.0011	0.0008
	(0.0023)	(0.0024)
Covid cases (t-1, per 100k)		0.0426
		(0.0529)
Covid deaths (t-1, per 1k)		0.0161
		(0.0249)
Fixed-effects		
State	Yes	Yes
Week	Yes	Yes
Fit statistics		
Observations	5,250	5,050

Clustered (State) standard errors in parentheses; $t(G-1)$ critical values
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

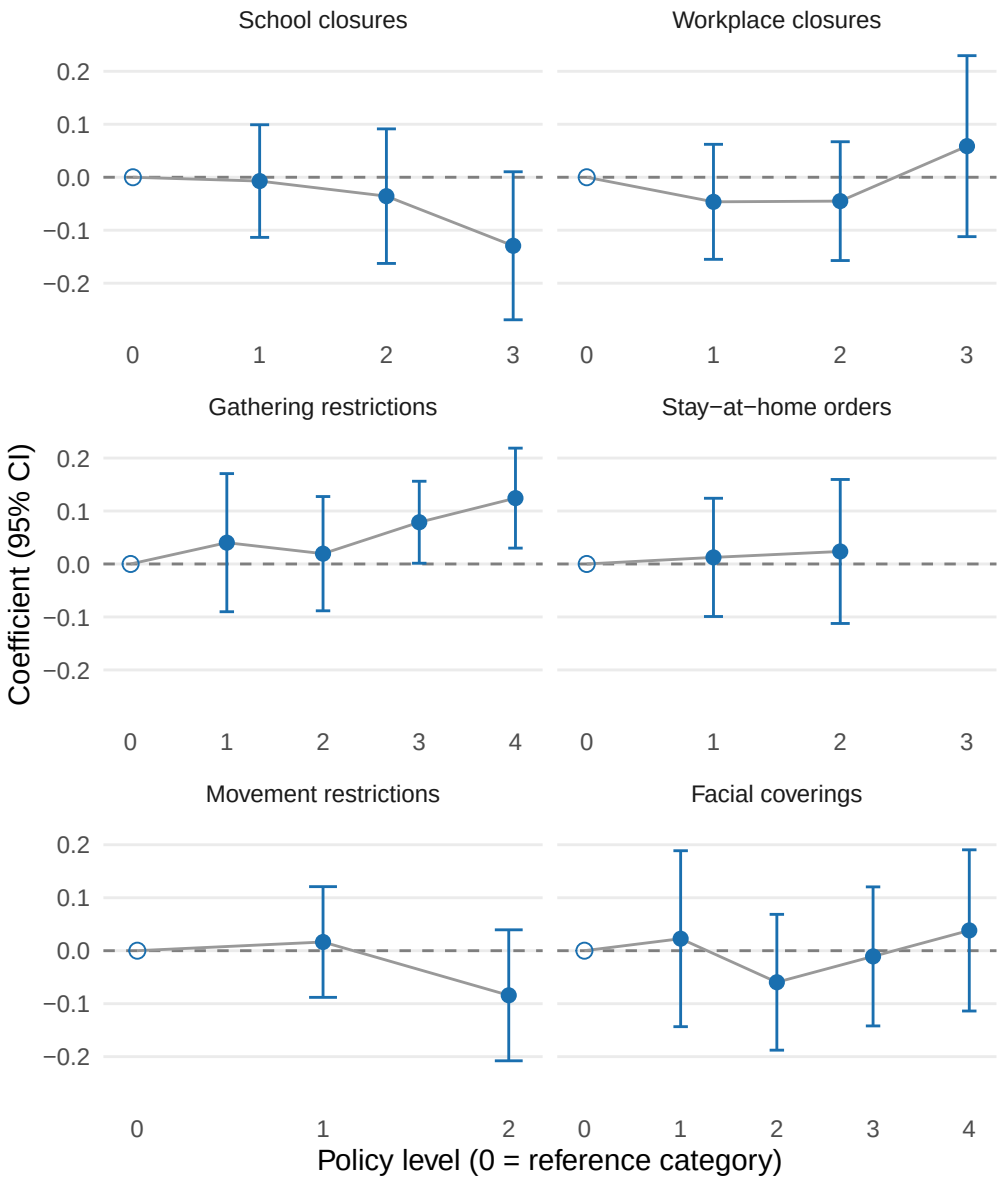


Figure K.1: Dose–response coefficients for individual lockdown policy components (United States). Each panel shows the estimated effect of each restriction level relative to level 0 (no restriction), from a regression replacing the ordinal variable with a full set of level dummies. State and week fixed effects and lagged epidemiological controls are included throughout. Error bars are 95% confidence intervals.

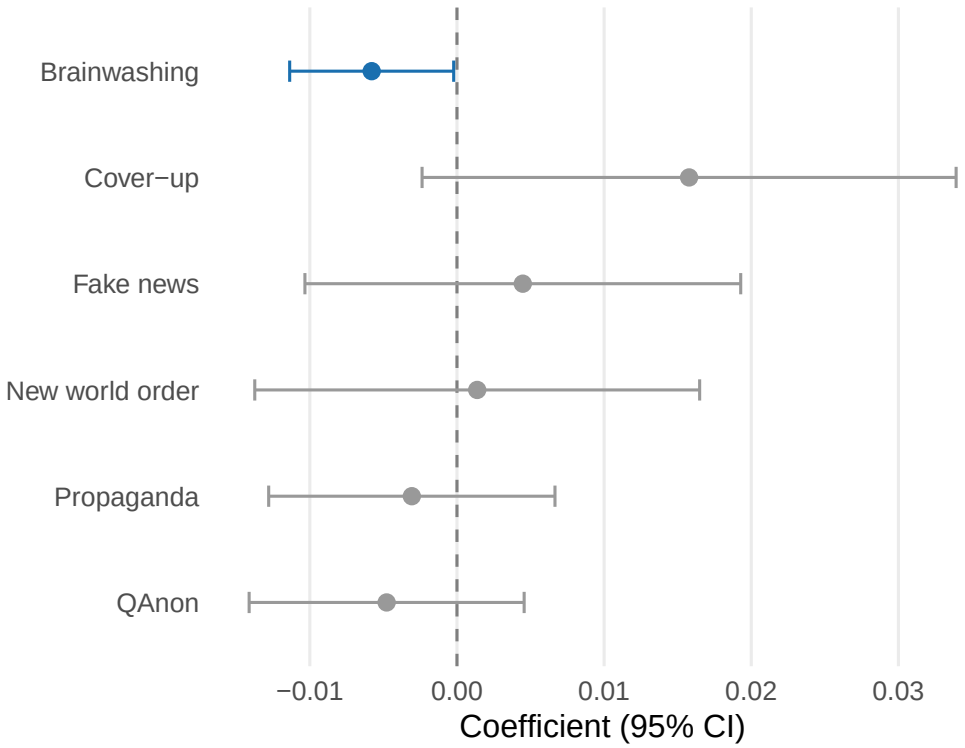


Figure K.2: Effect of lockdown stringency on individual conspiracy topics (United States). Each point is the coefficient on the stringency index from a separate regression with state and week fixed effects and lagged epidemiological controls. Error bars are 95% confidence intervals. Blue points are significant at the 5% level.

L. Difference-in-Differences Approach

Table L.1: Clustering Choice Comparison: DiD Estimates

Model:	(1)	(2)
Variables		
Post-lockdown (same calendar period)	-0.1769**	-0.1769**
	(0.0738)	(0.0702)
Year 2020	-0.0534	-0.0534
	(0.0391)	(0.0335)
DiD: Post × 2020	0.5379***	0.5379***
	(0.0919)	(0.0522)
Fixed-effects		
Country	Yes	Yes
Week of year	Yes	Yes
Fit statistics		
Observations	2,170	2,170
R^2	0.87716	0.87716

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Col. 1: clustered by country ($G = 31$). Col. 2: clustered by week of year.

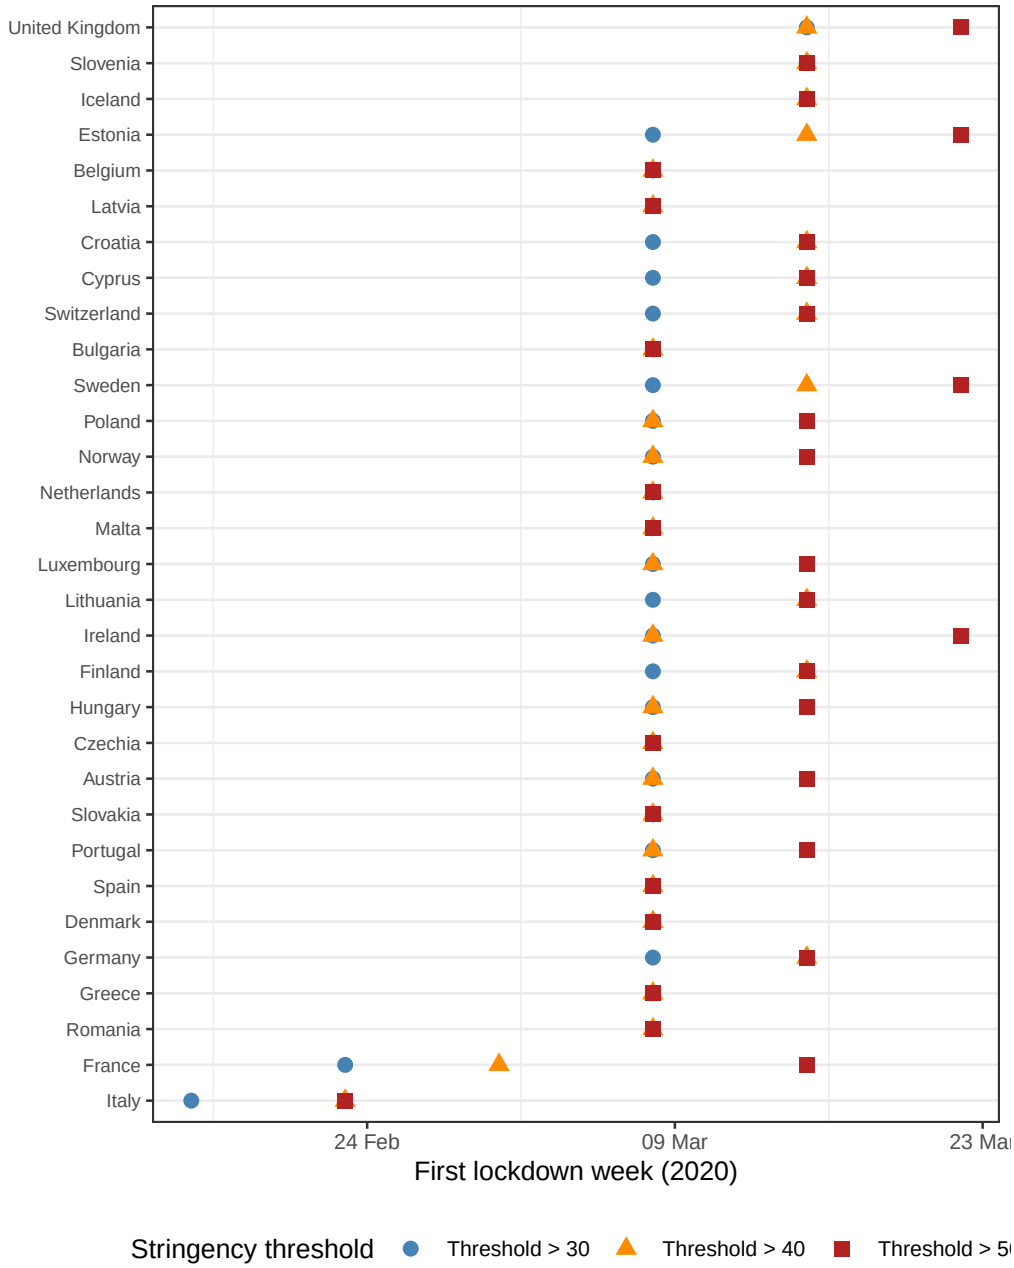


Figure L.1: Distribution of first-wave lockdown onset dates across 31 European countries, defined as the first week the OxCGRT Stringency Index exceeded 30, 40, or 50. 29 of 31 countries crossed the threshold-30 within a single 11-day window (8–18 March 2020).