

NEUROSCIENCE

Interconnectedness and (in)coherence as a signature of conspiracy worldviews

Alessandro Miani^{1*}, Thomas Hills^{2,3}, Adrian Bangert¹

Conspiracy theories may arise out of an overarching conspiracy worldview that identifies common elements of subterfuge across unrelated or even contradictory explanations, leading to networks of self-reinforcing beliefs. We test this conjecture by analyzing a large natural language database of conspiracy and nonconspiracy texts for the same events, thus linking theory-driven psychological research with data-driven computational approaches. We find that, relative to nonconspiracy texts, conspiracy texts are more interconnected, more topically heterogeneous, and more similar to one another, revealing lower cohesion within texts but higher cohesion between texts and providing strong empirical support for an overarching conspiracy worldview. Our results provide inroads for classification algorithms and further exploration into individual differences in belief structures.

INTRODUCTION

Conspiracy theories (CTs) propose alternative explanations of publicly relevant events [e.g., vaccination, coronavirus disease 2019 (COVID-19), climate change, and Princess Diana's death], evoking secret plots by malevolent and powerful groups who act at the expense of an unwitting population (1). CTs are popular. In a nationally representative sample of American adults, half of respondents believed in at least one medical CT, such as that the Food and Drug Administration deliberately prevents access to effective natural cures for cancer and that fluoride is a dangerous by-product that the government allows phosphate mines to dump into the public water supply (2). In a cross-national survey in 17 nations, more than half of respondents believed in CTs associated with the 9/11 terrorist attack (3). Furthermore, CTs may emerge and propagate during times of crisis (4) such as epidemics, deaths of public figures, or even natural disasters, which, in turn, may increase exposure to other CTs (5).

Belief in CTs is associated with numerous negative outcomes. Belief in medical CTs correlates with reduced likelihood of influenza vaccination or annual checkups (2). Belief in HIV/AIDS CTs reduces intentions to use condoms (6). Exposure to CTs reduces intentions to limit carbon footprints (7). Belief in CTs is also associated with political extremism and violence (8, 9) and increased intentions to engage in everyday crime (10). At the societal level, these phenomena may lead to loss of human lives, waste of public funds, and disruption of social order (11–13). The ubiquity and negative impact of CTs warrant increasing efforts to understand them.

One pathway to understanding CTs is the observation that people who believe in one CT tend to believe in others, irrespective of how unrelated or contradictory they may seem (14–16). For example, believing that the AIDS virus was deliberately engineered in a government laboratory is associated with believing that the Federal Bureau of Investigation was involved in Martin Luther King's assassination (16). Conspiracy believers tend to identify meaningful relationships among randomly co-occurring events (17, 18), confuse

aspects of reality such as believing that prayers have the capacity to heal (19), and believe in the paranormal (20).

This collection of events, however, can devolve into contradiction. Individuals who felt that it was more likely that Osama Bin Laden was already dead before the U.S. military forces arrived at his compound in Pakistan were also more likely to believe that he is still alive (15). As the authors of that work put it, "...the specifics of a conspiracy theory do not matter as much as the fact that it is a conspiracy theory at all" (p. 5). This was further reiterated by Lewandowsky and colleagues (21), "The incoherence does not matter to the person rejecting the official account because it is resolved at a higher level of abstraction" (p. 179).

Thus, belief in multiple conspiracies may be self-supporting: The interconnectivity among conspiracy beliefs is supported by a meta-belief that resolves the apparent contradictions at the lower level. Hence, CTs may thus constitute a mutually reinforcing network of beliefs, creating self-sustaining evidence for a world dominated by deceptive agents (15). These networks would constitute a top-down conspiracy worldview, which coerces unconnected or even contradictory observations into support for a global conspiracy (21).

If the criterion for being correct is not about the veridicality or mutual compatibility of individual events but rather their support for the existence of deceptive agents, then a conspiracy worldview can even sustain entirely fictitious beliefs. For example, in one study (22), the more participants believed in popular CTs (e.g., about 9/11), the more they perceived as real a set of fictitious CTs created ad hoc for the study. This study suggests that the belief in some CTs increases the chances that an individual will accept evidence for novel CTs.

The combination of incompatible (or, more generally, incoherent) lower-level explanations that get resolved by mutual compatibility with a higher-level belief that authorities are deceptive has recently been claimed to be the hallmark of conspiracy worldviews (21). These patterns of beliefs might facilitate the creation and endorsement of the unusual patterns seen in conspiracy narratives (23). CTs serve sense-making functions, as they allow one to reduce the complexity of real-world events (24, 25), dismissing them as evidence for a grand conspiracy. Thus, the narrative structure of CTs may be important in supporting a CT worldview. The conjecture that conspiracy worldviews have local incoherence (supporting evidence is drawn from numerous unrelated and sometimes contradictory sources) but global coherence (appealing to an overarching belief in the deceptive

Copyright © 2022
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

¹Institute of Work and Organizational Psychology, University of Neuchâtel, Rue Emile-Argand 11, 2000 Neuchâtel, Switzerland. ²Department of Psychology, University of Warwick, University Road, Coventry CV47AL, UK. ³The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, UK.

*Corresponding author. Email: alessandro.miani@unine.ch

nature of authorities) constitutes an important theoretical step forward in the psychology of conspiracy beliefs. However, it remains virtually untested outside of a handful of studies focusing on specific CTs (15, 21, 26).

In the current study, we exploit the abundant instantiations of CT beliefs in naturally occurring narrative text. We perform a large-scale text analysis to test the conjectures that conspiracy narratives are characterized by a multitude of interconnected ideas (27) that manifest in patterns of lower local (within-text) cohesion but higher global (between-text) cohesion (15, 21). If a conspiracy worldview coerces unconnected observations to support for an overarching belief in the deceptive nature of authorities, then a network of potential conspiracy-related topics will be more tightly interconnected in conspiracy documents than in nonconspiracy documents [hypothesis 1 (H1)]. Similarly, if conspiracy narratives focus less on individual topics than nonconspiracy narratives, then topic specificity will be lower for conspiracy than nonconspiracy documents (H2a). Furthermore, text cohesion between paragraphs should be less internally cohesive for conspiracy documents than nonconspiracy documents (H2b). Last, if conspiracy narratives reference similar worldviews, then similarity between documents should be higher for conspiracy documents (as a group) than nonconspiracy documents (H3).

Our analyses rely on the largest corpus of CTs available today, language of conspiracy (LOCO) (28), an 88 million-word corpus composed of topic-matched conspiracy ($N = 23,937$) and nonconspiracy ($N = 72,806$) text documents (i.e., webpages) harvested from 150 websites. LOCO provides two types of semantic indexes for each document. One is represented by seeds ($N = 39$), keywords used via a Google search to retrieve documents associated with events that have generated CTs (e.g., 9/11 terrorist attacks, the death of Princess Diana, and COVID-19). The other one is represented by topics, extracted from documents with latent Dirichlet allocation (LDA) (29). Topics are expressed as distribution of probabilities, while seeds are either present or not in the document. In terms of content, topics differ from seeds because they represent out-of-domain themes extracted from the corpus a posteriori (versus seeds that were defined, by us, a priori). LDA topics differ from seeds in terms of granularity: We provide three sets of topics that represent the semantic resolution (at 100, 200, and 300 topics) of our corpus. These differences between seeds and topics allow us to evaluate within- and cross-theme similarities as they differ between conspiracy and nonconspiracy documents.

RESULTS

Interconnectedness (H1)

To test H1—conspiracy-related topics will be more tightly interconnected in conspiracy documents than in nonconspiracy documents—we ran network analyses on the co-occurrences of seeds and topics extracted from the conspiracy and nonconspiracy subcorpora. We calculated the average degree of connectedness in the networks by extracting how many edges were connected to each node (either seeds or topics).

Seed interconnectedness

We first tested whether conspiracy documents, on average, contained more different seeds than nonconspiracy documents. That is, we calculated thematic richness by counting how many seeds were present in each document. We fitted a linear mixed-effects regression model predicting the number of seeds contained in documents by subcorpus (either conspiracy or nonconspiracy), adding word

count as covariate and nesting documents within websites. On average, conspiracy documents contained more seeds than nonconspiracy documents: $\beta = 0.85$, $SE = 0.129$, $t_{119.44} = 6.59$, $P < 0.001$, $R^2_{m/c}$ (marginal/conditional) = 0.093/0.454 (conspiracy: $M = 1.293$, $SD = 0.741$, range: 1 to 13; nonconspiracy: $M = 1.073$, $SD = 0.305$, range: 1 to 6).

We then tested our main hypothesis (H1), which is how the degree of interconnectedness differed between conspiracy and nonconspiracy networks. We fitted a linear mixed-effects regression model predicting interconnectedness (i.e., number of edges per node) by subcorpus (either conspiracy or nonconspiracy) while nesting nodes within nodes' names (similar to a paired test that tracks differences within a seed, e.g., "lady_diana" between conspiracy and nonconspiracy networks). The conspiracy network was more interconnected compared to the nonconspiracy network: $\beta = 0.97$, $SE = 0.121$, $t_{38} = 8.03$, $P < 0.001$, $R^2_{m/c} = 0.235/0.720$. Figure 1 shows the two networks.

In table S1, we report an additional set of analyses that further confirm the above network differences. CT narratives formed a larger giant component than the nonconspiracy network, with fewer distinct subnetworks. Entropy was higher in the conspiracy (versus nonconspiracy) network, suggesting that seeds are, to a larger extent, agglomerated in a random, nonsystematic way. Conspiracy nodes (compared to nonconspiracy nodes) were more similar to each other in terms of connection patterns. The clustering coefficient (the probability that the adjacent vertices of a node are interconnected) was higher in the conspiracy network. The average shortest path length through nodes, i.e., distance, was lower in the conspiracy network, suggesting again more interconnectedness in conspiracy network. Last, density, the ratio of the number of edges to the number of possible edges, was also higher in the conspiracy, compared to the nonconspiracy, network.

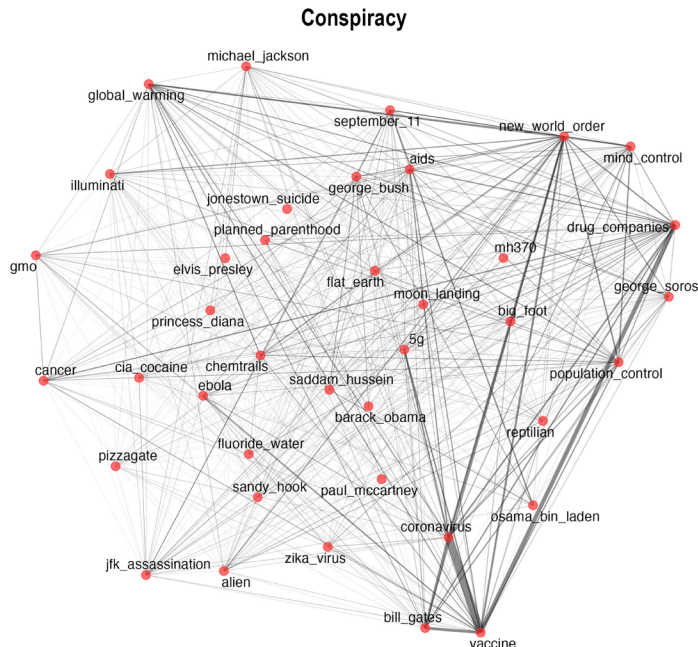
Topic interconnectedness

Because seeds represent specific mentions of themes that have generated conspiracies (e.g., 9/11 and Princess Diana's death), we further investigated H1 using a more general pattern of word co-occurrences associated with LDA topics that were extracted in an unsupervised fashion from the corpus. We created topic networks, where nodes are topics and edges are the degree of correlation between topics. Relying on LOCO's three sets of LDA topics (that contain 100, 200, and 300 topics; henceforth, LDA100, LDA200, and LDA300), we then compared the average degree of interconnectedness between conspiracy and nonconspiracy networks. In all three sets, the conspiracy networks were more interconnected than the nonconspiracy networks, LDA100: $\beta = 0.30$, $SE = 0.089$, $t_{99} = 3.37$, $P = 0.001$, $R^2_{m/c} = 0.022/0.607$; LDA200: $\beta = 0.301$, $SE = 0.068$, $t_{199} = 4.42$, $P < 0.001$, $R^2_{m/c} = 0.023/0.538$; and LDA300: $\beta = 0.33$, $SE = 0.06$, $t_{299} = 5.55$, $P < 0.001$, $R^2_{m/c} = 0.027/0.470$. Similar to the results obtained for seed networks, entropy was higher in the conspiracy networks. Conspiracy nodes were more similar to each other in terms of connection patterns. Clustering coefficients were higher in the conspiracy network. Distance was lower in conspiracy networks. In addition, density was higher in conspiracy networks. In table S2, we report the properties of each of the six networks (conspiracy and nonconspiracy networks for the three LDA topic matrices).

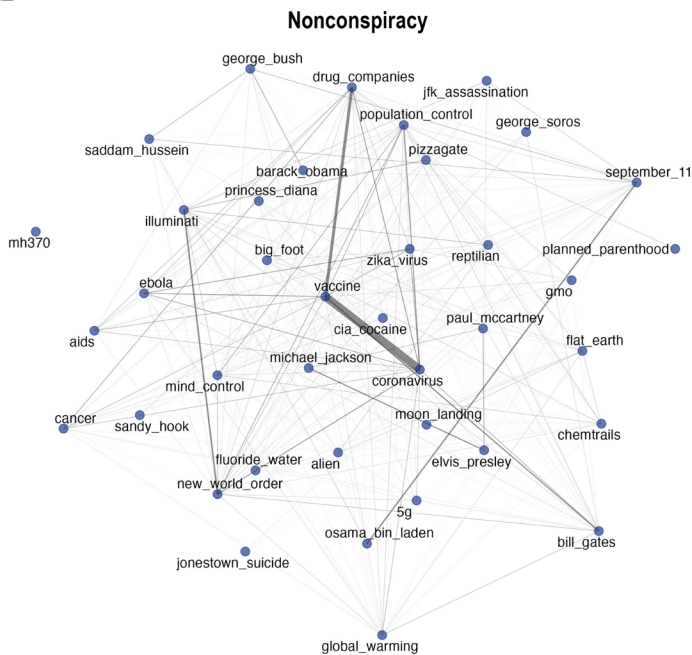
Local cohesion (H2)

To test H2a—topic specificity should be lower for conspiracy than nonconspiracy documents—we evaluated the inequality of within-document topic distributions using the Gini coefficient. Because each topic has a probability associated with each document, a more

A1



A2



B

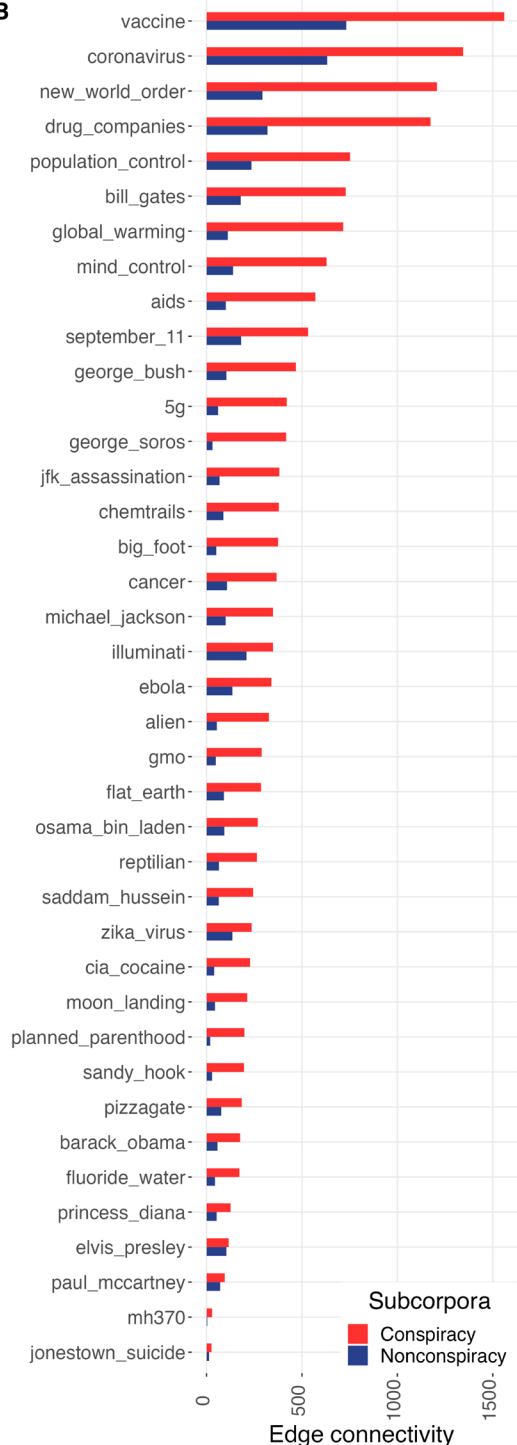


Fig. 1. Network interconnectedness from seeds. Nodes represent seeds (documents’ keywords associated with events that have generated CTs), and edges represent the co-occurrence of seeds in documents. Thicker edges indicate higher co-occurrence. Left: Network plots extracted from the conspiracy (A1) (red) and nonconspiracy (A2) (blue) subcorpora. Right: Numbers of links held by each node (edge connectivity) as a measure of seed interconnectedness (B).

unequal topic distribution (i.e., higher Gini coefficient) indicates that the document is focused on fewer topics than a document with a more equal topic distribution (fig. S1). Results of the linear mixed-effects models for predicting topic specificity (controlling for word count and nested within websites) show that conspiracy

documents had lower topic specificity compared to nonconspiracy documents: $\beta = -0.29$, $SE = 0.063$, $t_{145.09} = -4.66$, $P < 0.001$, $R^2_{m/c} = 0.019/0.143$. We then tested H2b—within-document lexical cohesion should be lower for conspiracy documents than nonconspiracy documents—by analyzing how well within-document

paragraphs are semantically connected to each other (e.g., lexical cohesion across paragraphs). The measure we use correlates with perceived text coherence (30). Conspiracy documents showed lower lexical cohesion than nonconspiracy documents: $\beta = -0.68$, $SE = 0.08$, $t_{146,54} = -7.98$, $P < 0.001$, $R^2_{m/c} = 0.080/0.302$. Note that lexical cohesion and topic specificity are two different constructs (they correlate poorly; $r_{96,741} = 0.19$, $P < 0.001$; see table S3): While topic specificity measures how many topics account for a document's content, lexical cohesion measures the lexical overlap of paragraphs.

Global cohesion (H3)

To test H3—similarity between documents should be higher for conspiracy documents than nonconspiracy documents—we reasoned that, at the global level, shared lexical patterns would make conspiracy documents more similar to each other than nonconspiracy documents. We computed a measure of between-document lexical overlap, the cosine similarity (CS) scores between documents within subcorpora (either conspiracy or nonconspiracy), obtaining a value of similarity for each document with all other documents within the subcorpus. Pairwise CS was computed between each document and the remaining documents within the same subcorpus. For each computation, we excluded documents with the same seed and those gathered from the same website to avoid same-topic and same-author lexicons inflating document similarity. Compared to nonconspiracy documents, conspiracy documents were more similar to each other: $\beta = 0.96$, $SE = 0.064$, $t_{148,14} = 14.87$, $P < 0.001$, $R^2_{m/c} = 0.422/0.559$. Our results (see table S4 and fig. S2) are robust and replicate across six different subsets of LOCO (in which we artificially created subcorpora perfectly matched for topics and word count; tables S5 to S7 and fig. S3).

DISCUSSION

Popular belief in conspiracies is a riddle. Belief in one CT leads to believe in other CTs (14–16). Irrespective of how related events are, an accumulation of CTs reflect a view of a world dominated by deception, resulting in a self-sustaining network of supportive beliefs (15). This conjecture was previously tested by measuring the extent to which participants simultaneously endorsed researcher-designed potentially contradictory conspiratorial items. Here, moving beyond individual beliefs and beyond single-case studies of specific CTs, we performed a large-scale text and network analyses on the abundant naturally occurring instantiations of CTs, providing strong empirical support for an overarching conspiracy worldview in conspiracy narratives.

Our results show that conspiracy texts exhibit a pattern of strong interconnectedness with each other, linking multiple ideas that result in a dense and highly interconnected network (H1). Individual conspiracy documents are built from multiple sources and are, on average, less locally (within-document) coherent than corresponding nonconspiracy documents (H2). They nevertheless exhibit higher global (between-document) cohesion, being more lexically similar to each other than nonconspiratorial documents (H3).

Compared to nonconspiracy, conspiracy narratives are more interconnected via a dense and unstructured network of shared themes. These properties emerge not only from themes associated with events that have generated CTs (seeds) but also from out-of-domain themes (LDA topics). Such a high topical interconnectedness mirrors the psychological need to reduce uncertainty and gain

control by explaining and finding order in real-world events that might otherwise seem random. Individuals who believe in CTs have an overarching conspiracy mentality that makes them more likely to draw implausible causal connections between random or unrelated events (17, 18).

Qualitative research has also suggested that conspiracy narratives rely on an accumulation of ideas in support of their claims (27). Popular figures, organizations, technologies, and even states (e.g., China or the United States) are often cited and interact, giving rise to the unusual narrative patterns emerging in CTs (23, 31). For example, according to one CT, the COVID-19 pandemic is a pretext for distributing harmful vaccines, activated by 5G radiation, which lead to a mass depopulation, all being commanded by George Soros and Bill Gates (31). In our study, the co-occurrence of these themes (operationalized by seeds: e.g., covid, 5g, soros, gates, and vaccines) represents the narrative richness of CTs. We have quantified this richness, showing that, on average, conspiracy narratives have higher number of seeds and lower topic specificity (i.e., more topics) in comparison to nonconspiracy narratives. These results confirm previous observations (27, 31). Although displaying high thematic richness (less local cohesion), conspiracy documents are more lexically similar to each other (more global cohesion) compared to nonconspiracy documents. One could think that this conspiratorial (in)coherence is counterintuitive because more topical richness should maximize within-document lexical diversity, hence potentially decreasing between-document lexical similarity. In our sample, the reverse is true, and we stress that this phenomenon is evidence for top-down thematic coercion, where themes fit an overarching conspiracy worldview. In this way, each theme is translated into a conspiracy by adding a set of recurrent lexical patterns involving language of deception, questioning, social identification, and negative emotions (28). These lexical patterns can be reused in any conspiratorial context and are shared across conspiracy narratives (28).

Despite being less internally cohesive, conspiracy narratives may appear coherent to believers because the believers' mental representations are based on world knowledge that is not explicitly represented in the text. Proneness to hallucinations and delusional ideations, traits that are shared with belief in CTs (32), might help connecting the dots (17, 18), so as to reduce uncertainty and gain control. People with delusional ideations (33), similar to conspiracy believers (34), tend to draw conclusions quicker and tend to be based on less evidence than people without delusions. Quickly connecting poorly related ideas gives rise to a highly chaotic and randomly connected network of ideas. This might explain why the conspiracy networks we built from narratives are denser and more unstructured in comparison to nonconspiracy networks.

Our findings resemble phenomena in thought of individuals on the schizophrenia spectrum, which includes the subclinical and milder form schizotypy. Schizotypy is a trait that correlates with belief in CTs (32). Individual with schizotypal personality tend to jump to conclusions quicker and make decisions based on less evidence compared to people without schizotypy (35). Schizotypy and schizophrenia overlap substantially and are characterized, yet on different levels, by an impairment in thought and perception that lead to psychotic symptoms (36). This impairment is manifested in language production (37, 38). Patients with schizophrenia show disruptions in speech production at the level of causal-motivational and thematic coherence (39) and in structural cohesive markers (40). Moreover, patients with schizophrenia also show disorganized

semantic networks (41, 42). These studies indicate a certain degree of overlap between the schizophrenia spectrum and belief in CTs in regard to semantic processing, suggesting that further research should pay attention to this overlap.

Our findings help advance research on fighting misinformation. Both schizotypal personality and belief in CTs are linked to vulnerability to misinformation (43, 44). The fact that conspiracy narratives are characterized by high global cohesion despite low local cohesion may help develop classification algorithms to detect conspiratorial language either online or offline. This could be achieved by extracting the lexical patterns that are shared across individual conspiracies such as language of deception, questioning, and social identification (e.g., “Are they lying to us?”). Moreover, future computational endeavors in natural language processing could move further, helping detect contradictory statements in texts, hence replicating seminal findings (15). This move might benefit not only misinformation and conspiracy research but also cognitive science and clinical research in general.

Our study has some limitations. One limitation is that, despite the number of controls we introduced in our analyses, our results could be, in part, driven by other factors. The variance explained by some of our models (especially topic interconnectedness and topic specificity) was modest, suggesting that other factors might be in play to explain differences in local and global (in)coherence between conspiracy and nonconspiracy narratives. Further research might explore other indicators of cohesion to test the robustness of these findings.

This study leaves some questions open for future research. To what extent is individuals’ conspiratorial (in)coherence specifically related to their belief in CTs? Are there any other individual differences that affect such tolerance to incoherence? For example, because worldviews are essential for making sense of the world, disconfirming a worldview would affect the very sense of an individual’s reality (45). To avoid this, people enable defensive mechanisms such as confirmation bias (28, 46) that allows them to preserve a worldview by seeking confirmation while avoiding challenges. Similar mechanisms could be used to protect any type of worldview. For example, irrespective of how related they are, ideas could be simultaneously endorsed to fit a coherent political or religious worldview. Furthermore, individual characteristics such as education or holistic thinking style might increase the tolerance to accept incoherent ideas. Conversely, analytic thinking style, negatively associated with belief in CTs (47), decreases the endorsement of contradictory statements (48).

To summarize, our findings contribute to a better understanding of the textual structure of CTs, linking theory-driven psychological research on CT beliefs measured in individuals (1) with data-driven, computational approaches to CT narratives measured in texts (23, 49). This move links CT research to a larger body of research on computational approaches to fake news and misinformation (50, 51) and offers inroads to develop classification algorithms and design debunking campaigns and institutional communication to counteract the spread of CTs online.

MATERIALS AND METHODS

Material

Our text material is gathered from the LOCO corpus (28). LOCO is a freely available, multilevel, topic-specific, ~88 million-token corpus of documents extracted from ~100,000 webpages. LOCO is composed

of both conspiracy ($N = 23,937$) and nonconspiracy ($N = 72,806$) documents nested within websites ($N = 150$).

LDA topic extraction

LOCO’s topics were extracted with LDA (29). LDA is an unsupervised probabilistic machine learning model capable of identifying co-occurring word patterns and extracting the underlying topic distribution for each text document. By setting a priori the number of topics desired from a given corpus, LDA computes, for each document in a corpus, the probabilities for all topics to be represented in the document. Each word of the corpus has a probability to be part of a topic. That is, a word x has probability β of being part of topic k ; a topic k has probability γ of being part of document n . The sum of all the word probabilities within one topic is 1, and the sum of all the topic probabilities within one document is 1.

Before topic extraction, texts were preprocessed: Texts were converted to American Standard Code for Information Interchange (ASCII) characters; lower cased; cleaned by Uniform Resource Locators (URLs), punctuation, numbers, symbols, separators, and stop words [for the full list, see SM5 in the supplementary materials of (28)]; and stemmed. We then generated a document term matrix, from which we extracted only the most frequent 10,000 words. Topic extraction was performed with the topicmodels R package (52), using Gibbs sampling. The other LDA parameters were set as default. LOCO is provided with three LDA topic matrices, which contain 100, 200, and 300 topics (LDA100, LDA200, and LDA300, respectively).

Seed extraction

In LOCO, seeds are keywords used to retrieve webpages via Google during the corpus construction. A document can be associated with more than one seed. This is because a single webpage can be returned by a Google search using different keywords. For example, if a document relates to Lady Diana’s death because of an Illuminati plot, then this document would be returned twice for both “lady_diana_death” and “illuminati” seeds.

Seeds differ from LDA-identified topics. Seeds are a set of keywords (related to straightforwardly identifiable themes such as the Sandy Hook school shooting or AIDS) built a priori to retrieve documents; LDA topics are extracted a posteriori from the given set of documents in an unsupervised fashion. Although they sometimes overlap, they constitute two methodologically different approaches. A webpage is returned by Google if the seed is present in the webpage (but note, not necessarily in the main text) at least once. However, the seed presence in the webpage does not necessarily indicate that the seed reflects the main topic of the document’s text because the seed can be contained in boilerplate texts or in the comment section of the webpage.

We tested to what extent seeds reflect text content. We started by searching for the words that compose seeds in each document (e.g., “climate” and “change” for documents associated with the seed “climate_change”). We then we tested the agreement between seeds and text content. For all seeds, the mean of accuracy and precision were 0.909 (SD = 0.126, range: 0.300 to 0.997) and 0.993 (SD = 0.005, range: 0.982 to 0.999), with a sensitivity of 0.911 (SD = 0.132, range: 0.268 to 1.00) and a specificity of 0.788 (SD = 0.141, range: 0.304 to 0.938), respectively. These results show that there is a substantial overlap between seeds and the content of texts; hence, seeds are useful for indexing the semantic content of documents.

During LOCO’s construction, some seeds were entered with synonyms to accommodate different spellings [e.g., “new_world_order” and

“NWO” as well as “climate_change” and “global_warming”; see table 4 in (28)]. Before our analyses, here, we aggregated the synonym seeds, reducing the seed pool from 47 to 39. A list of the 39 seeds is visible in Fig. 1 and fig. S2.

Networks from LDA topics and seeds

Interconnectedness—how documents are connected to each other via seeds or LDA topics—was tested on the networks resulting from the co-occurrences of seeds and LDA topics. We tested interconnectedness as edge connectivity, which is the number of nodes interconnected to each other.

Extracting networks from topics

To extract co-occurrences of LDA topics, we created network objects from the three LDA gamma values matrices provided with LOCO [LDA100, LDA200, and LDA300, whose dimensions are N documents (rows) by k topic (columns)], which contain 100, 200, and 300 topics, respectively. This was done by creating correlation matrices from the LDA topics matrices and then converting those matrices into graph objects, i.e., networks, using the *igraph* R package (53). In these networks, nodes are represented by topics, while edges are represented by their co-occurrences. We assessed interconnectedness by computing the edge connectivity, which is the number of edges associated with each node.

We started by computing, for each LDA matrix, the between-topic correlation matrix, extracting the Pearson r coefficient from the correlation between topic i and topic j within each matrix. To convert these correlation matrices into co-occurrence matrices, we needed a threshold of correlation values above which we consider a co-occurrence of topics. This is because if no threshold is provided, then all topics co-occur with each other (i.e., all topics whose $|r| > 0$; highest degree of connectivity that is equal to k , the number of topics). Conversely, if the threshold is too high, then no topic will co-occur (i.e., the degree of connectivity is equal to zero).

We explored how different $|r|$ thresholds would return a degree of connectivity different from zero and different from k (i.e., all topics co-occurring). To this purpose, we created a vector of $|r|$ values for each of the three correlation matrices (ranging from 0 to max r). For each value in the $|r|$ vector, we created a network object and extracted the degree of connectivity. As a threshold, for each of the three sets of networks, we selected the mean of all absolute correlation values from both conspiracy and nonconspiracy networks (see fig. S4). For LDA100, the mean of the absolute correlation values, above which we considered a topic co-occurrence, was $r = 0.034$ (range: 0 to 0.313). For LDA200, the mean was $r = 0.023$ (range: 0 to 0.400). For LDA300, the mean was $r = 0.018$ (range: 0 to 0.430).

Extracting networks from seeds

To create the networks of seeds for each subcorpus (both conspiracy and nonconspiracy), we created two co-occurrence matrices using the *fcm* function from the *quantda* R package (54). We then created the graph networks using the R package *igraph* (53). The nodes of this network represent seeds, and the edges represent the co-occurrences of seeds within each matrix.

Topic specificity

Topic specificity measures the extent to which documents contain more or less topics. Here, we computed topic specificity by extracting documents' topics using LDA and by computing the inequality of topic distribution within each document using the Gini coefficient. Topic specificity can be thought in terms of inequality: The more

unequal a distribution of topic is, the more a document is well represented by the highest value. As a measure of inequality, we used the unbiased Gini coefficient, which ranges from 0 to 1, where lower values indicate equal distribution. Thus, documents with higher Gini coefficients are better represented by a single LDA topic, whereas documents with lower Gini coefficients are more equally represented by a large number of topics. To extract the Gini coefficient, we used the function *Gini* from the R package *DescTools* that relies on the following equation

$$Gini = \frac{2 \sum_{i=1}^n i y_i}{n \sum_{i=1}^n y_i} - \frac{n+1}{n}$$

The Gini coefficient was computed (for each document) on the top 10 topics with the highest gamma value. The choice of top 10 topics was justified for two main reasons. First, most of a document's content can be summarized within a handful of topics. Second, we visually explored the distributions of the documents where the Gini coefficient was either the highest or the lowest value, and we assessed that, by visual inspection, most of the variation in topic distributions occurs within the top 10 highest gamma values. To put differently, the top 10 highest gamma topics account for most of the document's content. This is visible in fig. S1, where we show the gamma values (black lines) and their cumulative sum (red lines) for the top 500 documents with the highest Gini (left) and top 500 documents with the lowest Gini (right) coefficient. The figure shows that in high-Gini documents, the cumulative topic probability (i.e., gamma, on the y axis) reaches around 0.70 to 0.80 proportion of all topic distributions with less than five documents. Differently, in low-Gini documents, the cumulative topic probability reaches around 0.70 to 0.80 proportion of all topic distributions with about 25 documents. This shows that in documents with high Gini coefficient, fewer topics are needed to account for a large part of the document's semantic content, suggesting therefore higher topic specificity.

Before testing our hypotheses on topic specificity, we removed all documents shorter than six paragraphs to provide a sufficient amount of text to evaluate topic distribution. Note that results do not change in a meaningful way when all documents (including those having less than six paragraphs) are included (see table S4).

Lexical cohesion

Texts provide the means to objectively measure lexical cohesion. Cohesive devices in text include word substitution, pronominal reference, conjunctions, and lexical repetition. Here, we computed lexical cohesion features using the Tool for the Automatic Analysis of Cohesion (TAACO) (30), a freely available standalone application that allows batch processing of text files. In TAACO, cohesion measures are extracted using several methods such as computing type/token ratios, as well as lexical and semantic overlaps for different part-of-speech categories. For our purpose, investigating semantic cohesion (i.e., different topics within a text across paragraphs), we used measures of semantic overlap. This is computed in TAACO with three computational models: latent semantic analysis (LSA) (55), LDA (29), and Word2vec (56). Different from other word-counting tools such as the Linguistic Inquiry and Word Count (LIWC) (57), these probabilistic models are capable of extracting the underlying semantic relations in texts. These models are based on unsupervised

machine learning algorithms, meaning that human biases (e.g., associating a word to a category) are minimized. Last, TAACO outperforms similar model-based tools such as Coh-Matrix (58) because its semantic space is built on a larger corpus (~219 million words), and correlations with human rating of coherence were stronger than those with Coh-Matrix (30). TAACO uses these models to provide measures of lexical cohesion for segments within a text, i.e., adjacent paragraphs. For the LSA and the Word2vec models, TAACO computes similarity scores by the CS between segments (ranging between 0 and 1). LDA scores are computed using the Jensen-Shannon divergence between the normalized summed vector weights for the words in each segment (ranging between 0 and 1).

Extracting cohesion from documents

To obtain an output from TAACO, we first needed to feed it with a batch of documents. For this purpose, we first exported all documents as text files. Note that we did not perform any text preprocessing before this step (e.g., removing stop words or stemming) because TAACO performs analysis on the parsed text that needs to be syntactically valid. From the TAACO output, we extracted the three sets of measures computed with LDA, LSA, and Word2vec models. Specifically, we extracted the measures that computed dis/similarity between all adjacent paragraphs. While LSA and Word2vec outputs are in the form of similarity (LSA CS and Word2vec similarity scores, respectively), LDA output was computed as divergence; hence, LDA scores were reversed (i.e., by subtracting them from one). To obtain a single score of similarity for each document, we aggregated all three measures by computing the mean. Before testing our hypotheses on lexical cohesion, we removed all webpage documents whose length was less than six paragraphs, for reasons described above (note that, using all documents, results do not change in a meaningful way). Cohesion scores are assigned to each document and range from 0 (low cohesion) to 1 (high cohesion).

Testing cohesion metrics

To test whether TAACO measures the within-document lexical cohesion, we generated a sample of documents whose internal cohesion was artificially lowered. To this aim, we created a sample of “synthetic” documents composed of scrambled paragraphs randomly obtained from LOCO and tested whether cohesion was lower compared to a sample of “natural” documents. Our reasoning is that we cannot simply use TAACO to test differences between conspiracy and nonconspiracy documents because we do not know (i) whether TAACO is capable of detecting cohesion differences and (ii) whether there are real differences in lexical cohesion between conspiracy and nonconspiracy documents. Therefore, we first tested TAACO on documents that we know a priori are different, namely, by creating two groups of documents in which there are true differences in cohesion. It follows that if we find differences between these two groups (synthetic versus natural), then TAACO is capable of detecting cohesion differences.

To build our two test corpora, we first selected a random sample of 1000 documents from LOCO (both nonconspiracy and conspiracy) and created a bag of paragraphs ($N = 19,528$). We then selected a random sample of 500 nonconspiracy documents from LOCO and kept those that had at least six paragraphs, obtaining 385 nonconspiracy documents, whose length was 18.32 paragraphs on average ($SD = 14.56$). This set of high-cohesion, natural documents is used as a control group for the low-cohesion scrambled, i.e., synthetic, group. To build low-cohesion documents, we extracted the exact number of paragraphs from each of the high-cohesion document

and generated a matched-by-length scrambled version. This resulted in a set of scrambled documents as large as the set of natural documents with the exact number of paragraphs (hence, two groups of $N = 385$ documents). For example, if a high-cohesion document is composed of 18 paragraphs, then we create its low-cohesion version by taking 18 paragraphs randomly selected from the bag of paragraphs. The two sets of documents did not differ in word count: $t_{753} = 0.176$, $P = 0.86$, $d = 0.01$. Using TAACO, we extracted between-paragraphs cohesion metrics and aggregated them in a unique score for each document. A t test between the two groups showed that the synthetic documents had lower cohesion than the natural ones: $t_{651,697} = 42.05$, $P < 0.001$, $d = 3.03$ (synthetic: $M = 0.363$, $SD = 0.028$; natural: $M = 0.477$, $SD = 0.045$). We conclude that the cohesion metrics (and their aggregated value) are reliable in capturing differences in text cohesion.

Document similarity

To test the between-document similarity, we used the documents' pairwise CS with other documents in the same subcorpora (i.e., conspiracy and nonconspiracy). We used CS instead of other measures, for example, Jaccard similarity, because while the latter relies on unique word overlaps, CS is more sensitive to repetitions [CS and Jaccard similarity are highly correlated (in LOCO: $r_{96741} = 0.82$, $P < 0.001$)].

One could argue that a similarity score that is tested on all documents within a subcorpus might be inflated if documents rely on the same topic, simply because of word overlap. In addition, similarity between documents extracted from the same website might also be inflated by authors that might copy and paste pieces of narratives across webpages. We control for these confounds by computing the pairwise similarity of each document with the remaining documents in the subcorpus that (i) were not extracted from the same website and (ii) did not have the same seed. Texts were pre-processed following the same method used to extract LDA topics. Documents' similarity scores were computed using the `textstat_simil` function from the R package `quanteda` (54). Values range from 0 to 1, indicating either no overlap (0) or a perfect overlap (1) of terms. The returned output length of the CS for each document was a vector whose length was equal to the number of documents against which the similarity was tested. We therefore averaged this vector, obtaining a unique value for each document.

Statistical analyses (testing H1, H2, and H3)

To test H1, i.e., conspiracy networks have a higher interconnectedness than nonconspiracy networks, we extracted the degree of interconnectedness by counting the number of edges for each node in the network. To statistically test whether interconnectedness was higher in conspiracy compared to nonconspiracy networks, we ran linear mixed-effects models using the `lme4` and the `lmerTest` R packages (59, 60). In each model, we predicted the number of edges by the subcorpus (i.e., conspiracy and nonconspiracy), clustering observations within nodes. Note that this is similar to running a paired t test (β coefficients from these models are equal to the Cohen's d obtained from t tests). We preferred to rely on these multilevel models—instead of t tests—for consistency, so results are expressed in the same format throughout the paper.

For each network, we also provide descriptive statistics. In particular, we measured (i) entropy [with the function `graph.entropy` from the R package `statGraph` (61)], related to the extent to which

nodes in a network are interconnected in a random, nonsystematic way; (ii) similarity [with the function *similarity* from the R package *igraph* (53)], which measures of how similar are connection patterns within a network; (iii) clustering (with the function *transitivity* from the package *igraph*), which calculates the probability that adjacent vertices of a node are interconnected; (iv) distance (with the function *mean_distance* from the package *igraph*), which extracts the average shortest path length through nodes; and (v) density (with the function *edge_density* from the package *igraph*), which computes the ratio of the number of edges to the number of possible edges.

To test H2 and H3, for each dependent variable (topic specificity, lexical cohesion, and CS), we ran a series of linear mixed-effects models using the *lme4* and the *lmerTest* R packages. In each model, we specified as fixed effects the dichotomous subcorpus variable (i.e., whether the document is conspiracy or nonconspiracy) and added document word count as a covariate. As random intercept, we specified the websites from which documents were extracted. Theoretically, it is reasonable to assume that longer documents have space to accommodate more topics than shorter documents, which, consequently, decreases topic specificity and lexical cohesion. Likewise, larger documents, with a potential larger vocabulary, have higher chances to resemble the whole subcorpus vocabulary, hence resulting in high CS scores. Conspiracy documents are longer than nonconspiracy documents in word count: $t_{32,452} = 47.11$, $P < 0.001$, $d = 0.35$ (conspiracy: $M = 1236$, $SD = 1307$; nonconspiracy: $M = 806$, $SD = 939$). Second, word count correlates with our dependent variables (see tables S3 and S7). Thus, we include word count as a covariate in our analyses.

For multilevel models, as measures of effect sizes, we report the standardized regression coefficients beta (β) for predictors of interest and measures of fit such as R^2 (62). We report both marginal and conditional R^2 ($R^2_{m/c}$) associated with the variance explained by the fixed effects (marginal) and the variance explained by the entire model that includes both fixed and random effects (conditional), respectively. As a measure of effect size for t tests, we use Cohen's d . Note that because of the large samples we used in our main analyses, most P values are significant at $P < 0.001$. Although we report all P values from our analyses following the APA (American Psychological Association) style, we suggest readers focus more on the effect sizes and variance explained by the models, which vary greatly between analyses.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <https://science.org/doi/10.1126/sciadv.abq3668>

REFERENCES AND NOTES

- K. M. Douglas, J. E. Uscinski, R. M. Sutton, A. Cichocka, T. Nefes, C. S. Ang, F. Deravi, Understanding conspiracy theories. *Polit. Psychol.* **40**, 3–35 (2019).
- J. E. Oliver, T. Wood, Medical conspiracy theories and health behaviors in the United States. *JAMA Intern. Med.* **174**, 817–818 (2014).
- J. Allen, No consensus on who was behind Sept 11: Global poll (Reuters, 2008); www.reuters.com/article/us-sept11-qaeda-poll-idUSN1035876620080910.
- J.-W. van Prooijen, K. M. Douglas, Conspiracy theories as part of history: The role of societal crisis situations. *Mem. Stud.* **10**, 323–333 (2017).
- K. M. Douglas, COVID-19 conspiracy theories. *Group Process. Intergroup Relat.* **24**, 270–275 (2021).
- L. M. Bogart, S. Thorburn, Are HIV/AIDS conspiracy beliefs a barrier to HIV prevention among African Americans? *J. Acquir. Immune Defic. Syndr.* **38**, 213–218 (2005).
- D. Jolley, K. M. Douglas, The social consequences of conspiracism: Exposure to conspiracy theories decreases intentions to engage in politics and to reduce one's carbon footprint. *Br. J. Psychol.* **105**, 35–56 (2014).
- R. Imhoff, L. Dieterle, P. Lamberty, Resolving the puzzle of conspiracy worldview and political activism: Belief in secret plots decreases normative but increases nonnormative political engagement. *Soc. Psychol. Personal. Sci.* **12**, 71–79 (2021).
- D. Jolley, J. L. Paterson, Pylons ablaze: Examining the role of 5G COVID-19 conspiracy beliefs and support for violence. *Br. J. Soc. Psychol.* **59**, 628–640 (2020).
- D. Jolley, K. M. Douglas, A. C. Leite, T. Schrader, Belief in conspiracy theories and intentions to engage in everyday crime. *Br. J. Soc. Psychol.* **58**, 534–549 (2019).
- P. Davies, Antivaccination activists on the world wide web. *Arch. Dis. Child.* **87**, 22–25 (2002).
- J. Tollefson, Tracking QAnon: How Trump turned conspiracy-theory research upside down. *Nature* **590**, 192–193 (2021).
- P. Ball, Anti-vaccine movement could undermine efforts to end coronavirus pandemic, researchers warn. *Nature* **581**, 251–251 (2020).
- V. Swami, T. Chamorro-Premuzic, A. Furnham, Unanswered questions: A preliminary investigation of personality and individual difference predictors of 9/11 conspiracist beliefs. *Appl. Cogn. Psychol.* **24**, 749–761 (2010).
- M. J. Wood, K. M. Douglas, R. M. Sutton, Dead and alive. *Soc. Psychol. Personal. Sci.* **3**, 767–773 (2012).
- T. Goertzel, Belief in conspiracy theories. *Polit. Psychol.* **15**, 731 (1994).
- R. C. van der Wal, R. M. Sutton, J. Lange, J. P. N. Braga, Suspicious binds: Conspiracy thinking and tenuous perceptions of causal connections between co-occurring and spuriously correlated events. *Eur. J. Soc. Psychol.* **48**, 970–989 (2018).
- J.-W. van Prooijen, K. M. Douglas, C. De Inocencio, Connecting the dots: Illusory pattern perception predicts belief in conspiracies and the supernatural. *Eur. J. Soc. Psychol.* **48**, 320–335 (2018).
- E. Lobato, J. Mendoza, V. Sims, M. Chin, Examining the relationship between conspiracy theories, paranormal beliefs, and pseudoscience acceptance among a university population. *Appl. Cogn. Psychol.* **28**, 617–625 (2014).
- H. Darwin, N. Neave, J. Holmes, Belief in conspiracy theories. The role of paranormal belief, paranoid ideation and schizotypy. *Pers. Individ. Differ.* **50**, 1289–1293 (2011).
- S. Lewandowsky, J. Cook, E. Lloyd, The 'Alice in Wonderland' mechanics of the rejection of (climate) science: Simulating coherence by conspiracism. *Synthese* **195**, 175–196 (2018).
- V. Swami, R. Coles, S. Stieger, J. Pietschnig, A. Furnham, S. Rehim, M. Voracek, Conspiracist ideation in Britain and Austria: Evidence of a monological belief system and associations between individual psychological differences and real-world and fictitious conspiracy theories. *Br. J. Psychol.* **102**, 443–463 (2011).
- T. R. Tangherlini, S. Shahsavari, B. Shahbazi, E. Ebrahimzadeh, V. Roychowdhury, An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, Pizzagate and storytelling on the web. *PLOS ONE* **15**, e0233879 (2020).
- B. Franks, A. Bangerter, M. W. Bauer, Conspiracy theories as quasi-religious mentality: An integrated account from cognitive science, social representations theory, and frame theory. *Front. Psychol.* **4**, 424 (2013).
- A. Bangerter, P. Wagner-Egger, S. Delouvée, How conspiracy theories spread, in *Routledge Handbook of Conspiracy Theories*, M. Butter, P. Knight, Eds. (Routledge, 2020), pp. 206–218.
- P. Lukić, I. Žeželj, B. Stanković, How (ir)rational is it to believe in contradictory conspiracy theories? *Eur. J. Psychol.* **15**, 94–107 (2019).
- S. Oswald, Conspiracy and bias: Argumentative features and persuasiveness of conspiracy theories. *OSSA Conf. Arch.* **168**, 1–16 (2016).
- A. Miani, T. Hills, A. Bangerter, LOCO: The 88-million-word language of conspiracy corpus. *Behav. Res. Methods* **54**, 1794–1817 (2022).
- D. M. Blei, A. Y. Ng, M. I. Jordan, Latent Dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (2003).
- S. A. Crossley, K. Kyle, M. Dascalu, The Tool for the Automatic Analysis of Cohesion 2.0: Integrating semantic similarity and text overlap. *Behav. Res. Methods* **51**, 14–27 (2019).
- A. Bruns, S. Harrington, E. Hurcombe, 'Corona? 5G? or both?': The dynamics of COVID-19/5G conspiracy theories on Facebook. *Media Int. Aust.* **177**, 12–29 (2020).
- N. Dagnall, K. Drinkwater, A. Parker, A. Denovan, M. Parton, Conspiracy theory and cognitive style: A worldview. *Front. Psychol.* **6**, 206 (2015).
- M. Rodier, M. Prévost, L. Renoult, C. Lionnet, Y. Kwann, E. Dionne-Dostie, I. Chapleau, J. B. Debrulle, Healthy people with delusional ideation change their mind with conviction. *Psychiatry Res.* **189**, 433–439 (2011).
- R. Moulding, S. Nix-Carnell, A. Schnabel, M. Nedeljkovic, E. E. Burnside, A. F. Lentini, N. Mehzabin, Better the devil you know than a world you don't? Intolerance of uncertainty and worldview explanations for belief in conspiracy theories. *Pers. Individ. Differ.* **98**, 345–354 (2016).
- V. Juárez-Ramos, J. L. Rubio, C. Delpero, G. Mioni, F. Stablum, E. Gómez-Milán, Jumping to conclusions bias, BADE and feedback sensitivity in schizophrenia and schizotypy. *Conscious. Cogn.* **26**, 133–144 (2014).
- U. Ettinger, I. Meyhöfer, M. Steffens, M. Wagner, N. Koutsouleris, Genetics, cognition, and neurobiology of schizotypal personality: A review of the overlap with schizophrenia. *Front. Psychiatry.* **5**, 18 (2014).

37. N. Rezaei, E. Walker, P. Wolff, A machine learning approach to predicting psychosis using semantic density and latent content analysis. *NPJ Schizophr.* **5**, 9 (2019).
38. B. Elvevåg, P. W. Foltz, D. R. Weinberger, T. E. Goldberg, Quantifying incoherence in speech: An automated methodology and novel application to schizophrenia. *Schizophr. Res.* **93**, 304–316 (2007).
39. M. C. Allé, J. Potheegadoo, C. Köber, P. Schneider, R. Coutelle, T. Habermas, J.-M. Danion, F. Berna, Impaired coherence of life narratives of patients with schizophrenia. *Sci. Rep.* **5**, 12934 (2015).
40. J. A. Willits, T. Rubin, M. N. Jones, K. S. Minor, P. H. Lysaker, Evidence of disturbances of deep levels of semantic cohesion within personal narratives in schizophrenia. *Schizophr. Res.* **197**, 365–369 (2018).
41. J. S. Paulsen, R. Romero, A. Chan, A. V. Davis, R. K. Heaton, D. V. Jeste, Impairment of the semantic network in schizophrenia. *Psychiatry Res.* **63**, 109–121 (1996).
42. M. S. Aloia, M. L. Gourovitch, D. R. Weinberger, T. E. Goldberg, An investigation of semantic space in patients with schizophrenia. *J. Int. Neuropsychol. Soc.* **2**, 267–273 (1996).
43. M. V. Bronstein, G. Pennycook, A. Bear, D. G. Rand, T. D. Cannon, Belief in fake news is associated with delusionality, dogmatism, religious fundamentalism, and reduced analytic thinking. *J. Appl. Res. Mem. Cogn.* **8**, 108–117 (2019).
44. A. Anthony, R. Moulding, Breaking the news: Belief in fake news and conspiracist beliefs. *Aust. J. Psychol.* **71**, 154–162 (2019).
45. M. E. Koltko-Rivera, The psychology of worldviews. *Rev. Gen. Psychol.* **8**, 3–58 (2004).
46. E. Brugnoli, M. Cinelli, W. Quattrociocchi, A. Scala, Recursive patterns in online echo chambers. *Sci. Rep.* **9**, 20118 (2019).
47. V. Swami, M. Voracek, S. Stieger, U. S. Tran, A. Furnham, Analytic thinking reduces belief in conspiracy theories. *Cognition* **133**, 572–585 (2014).
48. D. Santos, B. Requero, M. Martín-Fernández, Individual differences in thinking style and dealing with contradiction: The mediating role of mixed emotions. *PLOS ONE* **16**, e0257864 (2021).
49. S. Phadke, M. Samory, T. Mitra, What makes people join conspiracy communities?: Role of social factors in conspiracy engagement. *Proc. ACM Hum. Comput. Interact.* **4**, 223 (2020).
50. D. M. J. Lazer, M. A. Baum, Y. Benkler, A. J. Berinsky, K. M. Greenhill, F. Menczer, M. J. Metzger, B. Nyhan, G. Pennycook, D. Rothschild, M. Schudson, S. A. Sloman, C. R. Sunstein, E. A. Thorson, D. J. Watts, J. L. Zittrain, The science of fake news. *Science* **359**, 1094–1096 (2018).
51. S. Lewandowsky, U. K. H. Ecker, J. Cook, Beyond misinformation: Understanding and coping with the “post-truth” era. *J. Appl. Res. Mem. Cogn.* **6**, 353–369 (2017).
52. B. Grün, K. Hornik, topicmodels: An R package for fitting topic models. *J. Stat. Softw.* **40**, 1–30 (2011).
53. G. Csardi, T. Nepusz, The igraph software package for complex network research. *InterJournal Complex Syst.* **1695**, 1–9 (2006).
54. K. Benoit, K. Watanabe, H. Wang, P. Nulty, A. Obeng, S. Müller, A. Matsuo, quanteda: An R package for the quantitative analysis of textual data. *J. Open Source Softw.* **3**, 774 (2018).
55. T. K. Landauer, P. W. Foltz, D. Laham, An introduction to latent semantic analysis. *Discourse Process.* **25**, 259–284 (1998).
56. T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality. arXiv:1310.4546 [cs.CL] (16 October 2013).
57. Y. R. Tausczik, J. W. Pennebaker, The psychological meaning of words: LIWC and computerized text analysis methods. *J. Lang. Soc. Psychol.* **29**, 24–54 (2010).
58. A. C. Graesser, D. S. McNamara, M. M. Louwerse, Z. Cai, Coh-Metrix: Analysis of text on cohesion and language. *Behav. Res. Methods Instrum. Comput.* **36**, 193–202 (2004).
59. D. Bates, M. Mächler, B. Bolker, S. Walker, Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
60. A. Kuznetsova, P. B. Brockhoff, R. H. B. Christensen, lmerTest package: Tests in linear mixed effects models. *J. Stat. Softw.* **82**, 1–26 (2017).
61. D. R. da Costa, T. C. Ramos, G. E. C. Guzman, S. S. Santos, E. S. Lira, A. Fujita, statGraph: Statistical methods for graphs (2021); <https://CRAN.R-project.org/package=statGraph>.
62. J. Lohr, Effect size measures for multilevel models: Definition, interpretation, and TIMSS example. *Large Scale Assess. Educ.* **6**, 8 (2018).
63. D. Nguyen, M. Liakata, S. DeDeo, J. Eisenstein, D. Mimno, R. Tromble, J. Winters, How we do things with words: Analyzing text as social and cultural data. *Front. Artif. Intell.* **3**, 62 (2020).
64. A. Colin, J. Murdock, in *Dynamics Of Science: Computational Frontiers in History and Philosophy of Science*, G. Ramsey, A. de Block, Eds. (Pittsburgh Univ. Press, 2020).
65. M. Nikita, ldatuning: Tuning of the latent Dirichlet allocation models parameters (2020); <https://CRAN.R-project.org/package=ldatuning>.
66. B. Han, D. Duong, J. H. Sul, P. I. W. de Bakker, E. Eskin, S. Raychaudhuri, A general framework for meta-analyzing dependent studies with overlapping subjects in association mapping. *Hum. Mol. Genet.* **25**, 1857–1866 (2016).
67. W. Viechtbauer, Conducting meta-analyses in R with the metafor package. *J. Stat. Softw.* **36**, 1–48 (2010).

Acknowledgments

Funding: We acknowledge that we received no funding in support for this research.

Author contributions: Conceptualization: A.M., T.H., and A.B. Methodology: A.M. Investigation: A.M., T.H., and A.B. Visualization: A.M. Funding acquisition: None. Project administration: A.B. Supervision: A.B. and T.H. Writing—original draft: A.M. Writing—review and editing: A.M., T.H., and A.B. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Scripts for replication are available at <https://osf.io/aqnm>.

Submitted 4 April 2022

Accepted 6 September 2022

Published 26 October 2022

10.1126/sciadv.abq3668