




# When danger strikes: A linguistic tool for tracking America's collective response to threats

Virginia K. Choi<sup>a,1</sup>, Snehesh Shrestha<sup>b</sup>, Xinyue Pan<sup>a</sup>, and Michele J. Gelfand<sup>c,1</sup> 

<sup>a</sup>Department of Psychology, University of Maryland, College Park, MD 20742; <sup>b</sup>Department of Computer Science, University of Maryland, College Park, MD 20742; and <sup>c</sup>Graduate School of Business, Stanford University, Stanford, CA 94305

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected in 2021.

Contributed by Michele J. Gelfand; received July 29, 2021; accepted November 22, 2021; reviewed by Daniel Newman and Thomas Talhelm

**In today's vast digital landscape, people are constantly exposed to threatening language, which attracts attention and activates the human brain's fear circuitry. However, to date, we have lacked the tools needed to identify threatening language and track its impact on human groups. To fill this gap, we developed a threat dictionary, a computationally derived linguistic tool that indexes threat levels from mass communication channels. We demonstrate this measure's convergent validity with objective threats in American history, including violent conflicts, natural disasters, and pathogen outbreaks such as the COVID-19 pandemic. Moreover, the dictionary offers predictive insights on US society's shifting cultural norms, political attitudes, and macroeconomic activities. Using data from newspapers that span over 100 years, we found change in threats to be associated with tighter social norms and collectivistic values, stronger approval of sitting US presidents, greater ethnocentrism and conservatism, lower stock prices, and less innovation. The data also showed that threatening language is contagious. In all, the language of threats is a powerful tool that can inform researchers and policy makers on the public's daily exposure to threatening language and make visible interesting societal patterns across American history.**

collective threats | socioecology | language | historical change | mass communication

In today's world, ominous warnings about imminent threats are prevalent in advertisements, political rhetoric, and newscasts (1–3). Especially within the vast digital landscape, we absorb a torrent of online content on platforms, apps, and news feeds designed to elicit our fear of potential threats (4). By raising the alarm on our impending doom, threat-related words and messages can instantly attract our attention—activating the fear circuitry in the human brain that collates and accentuates valuable survival information (5, 6). The rise of social media, in particular, has fueled research on the negative effects of fear mongering and the mass circulation of misinformation (7–10). Whether the goal is to inform or exploit, claims about threats are pervasive in public discourse, and there is an urgent need for both researchers and policy makers to understand its social consequences. However, to date, our ability to detect threatening language in texts is limited by the lack of accessible, validated linguistic dictionaries.

To understand the implications of threats broadcasted through mass communication channels, such as the news or social media, the present study created a computational linguistic tool that indexes threat levels from texts. This research specifically aims to 1) develop and validate a computationally derived measure to identify threat-related content using natural language processing (NLP) methods; 2) track how fluctuations in threats have been changing over the past 100 years of US history; and 3) examine how these changing threat levels offer predictive insights into America's shifting cultural norms, political attitudes, and macroeconomic activity. To develop a semantic measure that tracks the form, frequency, and magnitude of

communicated threats, we used a dictionary-based approach (11). This involves scanning documents for the presence of select keywords representative of a construct of interest. We develop a threat dictionary to advance theory on the societal impact of mass communicated threats and to enable the analysis of these effects historically and in real time.

Early efforts to develop dictionaries relied heavily on human judges for brainstorming and finalizing keywords (12). By contrast, to generate the threat dictionary, we applied a technique known as word-embedding models (WEMs). Far exceeding human vocabulary retrieval abilities, these models intake a large collection of textual data and encode millions of fine-grained semantic connections found between the input words (13). WEMs first vectorize each word from a set of given texts and map them on to a multidimensional Euclidean space, arranging words that share semantic meaning closer together, typically based on their co-occurrence probability (14). By applying this technique, we were able to efficiently extract hundreds of English terms relevant to the mass communication of threats.

An important measurement development goal of this project was to create a threat dictionary that generalized across different linguistic contexts. For instance, word choice on social media is typically more conversational than language found in

## Significance

**People are constantly exposed to threatening language in mass communication channels, yet we lack tools to identify language about threats and track its impact on human groups. We developed a threat dictionary, a computationally derived linguistic tool that indexes threat levels from texts with high temporal resolution, across media platforms, and for different levels of analysis. The dictionary shows convergent validity with objective threats in American history, including violent conflicts, natural disasters, and pathogen outbreaks. Moreover, fluctuations in threat levels from the past 100 years coincide with America's shifting cultural norms, political attitudes, and macroeconomic activity, demonstrating how this linguistic tool can be applied to understand the collective shifts associated with mass communicated threats.**

Author contributions: V.K.C. and M.J.G. designed research; V.K.C., S.S., X.P., and M.J.G. performed research; V.K.C. and X.P. analyzed data; and V.K.C., M.J.G., and X.P. wrote the paper.

Reviewers: D.N., University of Illinois; and T.T., University of Chicago.

The authors declare no competing interest.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<sup>1</sup>To whom correspondence may be addressed. Email: vavkc@terpmail.umd.edu or gelfand1@stanford.edu.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2113891119/-/DCSupplemental>.

Published January 24, 2022.


official write-ups (15, 16). Consequently, only relying on a single WEM would limit our sampling of words. To ensure we identified threat words expressed in multiple communication settings, we used an ensemble of WEMs, each separately pre-trained on a unique corpus: 1) Wikipedia articles, 2) Twitter posts, and 3) Common Crawl's randomized sample of web pages (GloVe) (14). Wikipedia provides encyclopedic content, whereas Twitter posts are real-time interactions on social media. The third model trained on Common Crawl's metadata of eclectic web pages represents language applied in a broad spectrum of contexts. Drawing from all three WEMs, we extended our lexical sampling and capacity to glean threat terms common across a variety of communication channels.

*SI Appendix, Table S1* provides the pseudocode (algorithmic flow) of our dictionary development's 10-step process. First, after loading all three of the pretrained WEMs, we identified words proximal to "threat" as well as its synonyms within each model. From this word generation, we filtered out common words, such as "the," "of," and "all." This produced the first list of words  $\chi_1, \dots, \chi_n$  semantically associated with threat. Before vetting this list, we applied spectral clustering from each word's vector coordinates to evaluate how they were interrelated (17, 18). The spectral clustering process involved computing a matrix based on closeness as a measure of similarity  $S_{ij} \geq 0$  between all possible word pairings  $\chi_i$  and  $\chi_j$  from the list. An optimal number of segmentation points was determined from this spatial distance information to cluster together similar words and divide up dissimilar words. The resulting clusters brought disambiguation and clarity to the extracted words, thereby providing more contextual information.

Following these steps, based on a predefined exclusion criterion, inapposite clusters of words were removed if full interrater agreement was achieved between the study's authors. For example, we excluded clusters predominantly made up of foreign words, numbers, or named entities (Europe, Gaza, Balkan). We repeated these steps with the two other WEMs. The final threat dictionary was composed of terms ( $n = 240$ ) that converged across all three model outputs. A sample of these final threat terms includes "attack," "crisis," "destroy," "disaster," "fear," "frightening," "injuries," "lethal," "looming," "meltdown," "outbreak," "suffer," "tension," "toxic," "unrest," "unstable," and "violent." *SI Appendix, Table S2* provides the full dictionary.

Analysis of textual documents using the threat dictionary can be performed here: <https://bit.ly/3zp2cYi>. We next set out to examine the convergent validity of the threat dictionary by showing how patterns of change in the use of threat words over time correspond to real moments in US history when Americans faced grave threats. We applied our threat dictionary to time-stamped news articles from 1900 to 2020 via <https://www.Newspapers.com> (19). As the largest online repository of historical and contemporary newspapers, this publicly available data source contains over 600 million pages of digitized news content. We tallied the occurrence rates of threat dictionary terms in news articles at monthly and annual intervals at both the state and national levels. Since these estimates can be influenced by the number of newspapers published within a given time, we adjusted these total counts by the coinciding approximate number of article pages published. This produced our primary time series dataset, which we used to track variance in threat levels over the past century of US history.

We first examined longitudinal descriptive trends in threat levels over the last 100 years of US history based on our threat index computed from <https://www.Newspapers.com>. Extant theories and data-driven works by numerous scholars (20–23) have argued that threat levels are in historical decline. Here, we examined whether threat levels do appear to be decreasing across US newspapers over time by modeling the monthly



**Fig. 1.** Tracking threat language in the United States over the past century. The y axis contains the monthly values of national threat levels, computed from the relative frequency of threat terms found in US newspapers. (*Upper*) The red line depicts the line of best linear fit, which follows a downward linear trend. (*Lower*) The 20-year forecast of threat levels is based on an ARIMA model, with the mean projection of threat levels from 2020 to 2040 returning a slightly upward trend. Note that the shaded areas represent 80 and 95% prediction intervals.

increases with time as a predictor. As portrayed in Fig. 1, discussion of threats in American newspapers has been in steady decline ( $B = -0.02$ , 95% CI  $[-0.024, -0.022]$ ,  $P < 0.001$ ), with this linear temporal trend accounting for 65% of the variance in the series. Since an ordinary least squares estimation with time series data is prone to self-correlating residuals, we fitted a model with optimal autoregressive integrated moving average (ARIMA) errors (24–26). The gradual decrease in US threat levels from 1900 to 2020 remained significant ( $B = -0.002$ , 95% CI  $[-0.003, -0.001]$ ,  $P < 0.001$ ), accounting for 98% of the ARIMA-adjusted series variance. ARIMA models are also beneficial for making forecasts since they capture both the autocorrelation and nonlinear trends in the series (27). Applying the same ARIMA parameters determined earlier, we forecasted the threat series two decades into the future (2020 to 2040). Unlike a linear approximation of predicted values, the mean predictions from the ARIMA-based forecasts indicate threat levels may rise in the coming decades, reflecting the positive trend in threat levels from the past eight years (Fig. 1). However, the wide prediction bounds indicate that this trend is uncertain and should be interpreted with caution, especially because irregular and unpredictable future events can affect such forecasts.

We also sought to validate the threat dictionary by demonstrating its convergence with actual life-threatening events. We focused on three domains of socioecological threats that have indiscriminately endangered human life throughout history: violent conflicts, natural disasters, and pathogen outbreaks. When these collective dangers increase at a certain time or region in US history, more words related to threats are expected to increase in mass communication channels as well. The data taken as our ground truth were times the United States became involved in militaristic conflicts since 1900, Federal Emergency Management Agency (FEMA) reports on severe natural disaster cases from 1953 to 2020, and data on regional mortality rates due to major infectious diseases from the Institute for Health Metrics and Evaluation (IHME) from 1980 to 2014. Additionally, in the early months of 2020 as the

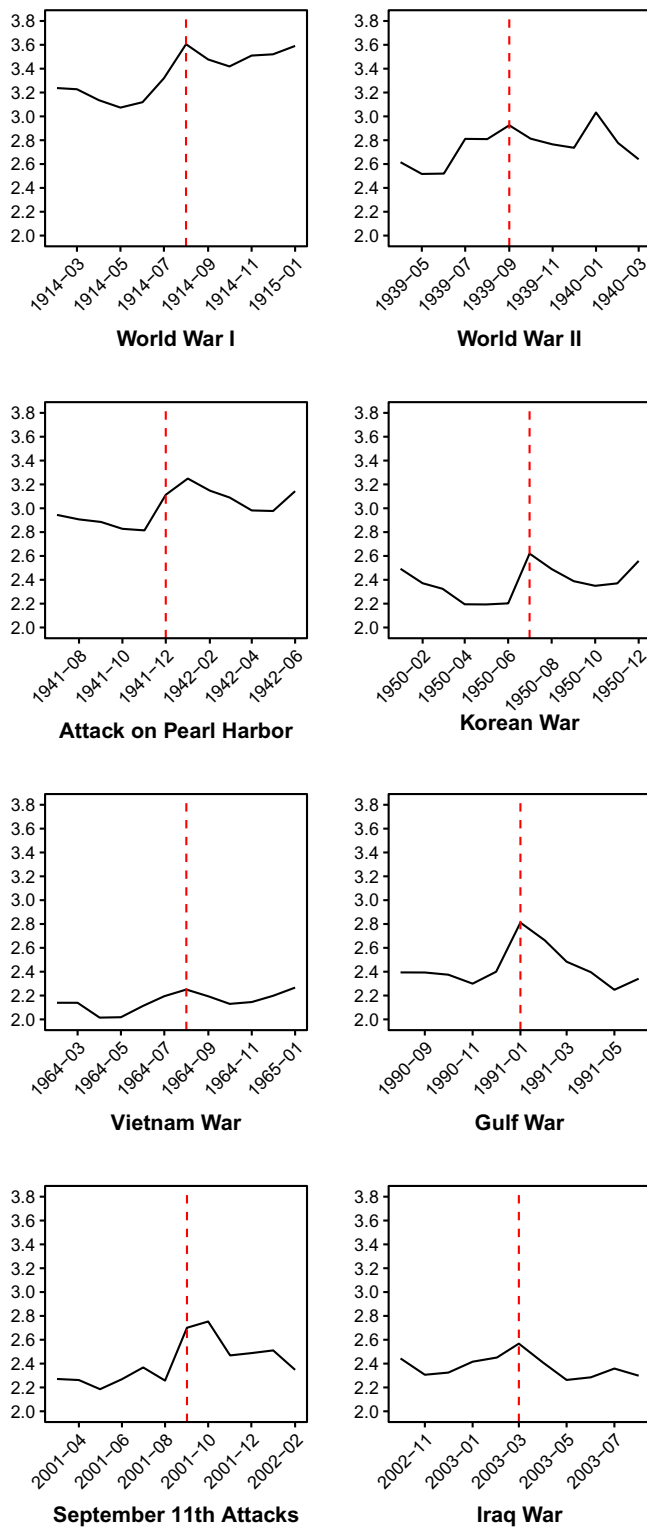
severity of the COVID-19 pandemic escalated, we analyzed 0.24 million Twitter posts sampled from each US state daily for 56 days. This dataset enabled us to examine the real-time threat dynamics relevant to the unfolding COVID-19 pandemic. Testing a variety of convergent indicators was critical to assessing our measure's sensitivity to multiple types of major collective threats and its application across different platforms.

Lastly, our 100-year plus threat data indexed from US newspapers enabled our historical analysis of national correlates with threats. In times of great danger, both preindustrial and contemporary societies have galvanized their populations with more structure and cooperation to better withstand collective threats (28–30). Thus, heightened periods of threat in US history are expected to coincide with shifts in America's increased preference for order—stronger norms, group orientation, and conservatism—yet lower levels of openness and innovation. Specifically, we examine how changes in threat levels correspond to an array of changes in cultural norms (cultural tightness and collectivism), political shifts (approval of sitting US presidents, Republican identification, and anti-immigrant attitudes), and macroeconomic activities (changes in the US stock market and innovation rates).

**Results**

**War and Conflicts.** Although US news coverage of threats is in gradual decline, we expected to find momentary increases in threat words during times of foreign conflict, such as major wars and attacks on US soil. We considered numerous wars from the past century that jointly provided sufficient coverage of America's major militaristic embroilments, including US involvement in World War I (WWI), World War II (WWII), the attack on Pearl Harbor, the Korean War, the Vietnam War, the Gulf War, the September 11th terrorist attacks, and the Iraq War. Discontinuous growth models (DGMs) (31) enabled the estimation of these upticks in threat levels when major conflicts occurred. DGMs test the change in trajectory of a given variable measured over time and its interrupted response pattern due to a specified shock, relative to the times directly preceding the shock. We conducted a DGM analysis per event, with each base model including three time-related covariates that reflect threat frequencies 5 mo before (prethreat), during (onset), and 5 mo after (postthreat) the advent of a conflict. *SI Appendix, Table S4* illustrates an example of how these time predictors are coded. In our DGMs, the onset points were designated based on the official dates of these ordeals (*SI Appendix, section C* has details). This onset vector, also referred to as the transition or time change variable, tests how the intercept has changed right after an expected event. For our research purposes, we anticipated that this onset time estimate for threat levels would reflect a significant discontinuity (i.e., an increase) after each foreign conflict, relative to the expected change pattern prethreat.

Our results, summarized in *SI Appendix, Table S6*, indicated that the intercepts for locally reported threat words in newspapers significantly increased as compared with their expected projections following news of WWI ( $P < 0.001$ ), WWII ( $P < 0.001$ ), the attack on Pearl Harbor ( $P < 0.001$ ), the Korean War ( $P < 0.001$ ), the Vietnam War ( $P < 0.001$ ), the Gulf War ( $P < 0.001$ ), the “war on terror” instigated by the September 11th terrorist attacks ( $P < 0.001$ ), and the Iraq War ( $P < 0.001$ ). For all conflicts, we found that the intercepts of threat levels among statewide newspapers significantly increased in threat words at the onset of the conflict relative to its prior trajectory (Fig. 2). To put the results in context, threat language in newspapers during the onset of WWI increased by 0.43 points from its expected levels prior to the start of WWI. At the transition time of WWII's start, threat increased by 0.15 points. The Pearl Harbor attack led



**Fig. 2.** Increase in threat words at the onset of major US conflicts. Plotted points represent the relative use of threat words found in US newspapers estimated at the national level (y axis) during the months that preceded and followed each major conflict (x axis). The red dotted lines demonstrate how threat levels spiked up at the time of each conflict's onset in comparison with its prior trajectory.

to a 0.16-point increase in threat. At the start of the Korean War, threat values jumped to 0.30 points. Threat rose by 0.18 points when Congress passed a resolution to increase

military presence in Vietnam and by 0.36 points when the Gulf War started. The September 11th terrorist attacks escalated threat levels by 0.16 points. At the onset of the Iraq War, threat levels increased by 0.21 points.

**Natural Disasters.** Next, we examined whether the threat dictionary is sensitive to the occurrence of objective natural disasters. The counts of major disaster declarations (MDDs) from FEMA reflect emergency response efforts instituted to help local governments combat severe natural events, such as tsunamis, earthquakes, tornadoes, hurricanes, flash floods, snowstorms, droughts, fires, and volcanic eruptions. We compiled FEMA's available monthly data on the number of MDDs enacted for each US state across the last 60 years and conducted a multilevel regression with data nested by states. This analysis was conducted at the state level, with MDDs likely varying based on each state's distinct ecological vulnerabilities. We found that greater instances of MDDs within each state were predictive of more threat words found in local statewide newspapers ( $B = 0.003$ ,  $P < 0.001$ , 95% CI [0.002, 0.004],  $R^2 = 0.24$ ).

**Pathogens.** Finally, we examined whether threat language in newspapers can be related to greater rates of pathogen-related deaths. Estimates of death rates from contagious diseases are documented in IHME's available state-level data spanning 35 years (1980 to 2014). Annual average death rates were computed per state by collapsing together mortality rates from all categories of major infectious diseases (hepatitis, HIV/AIDS, diarrheal diseases, lower respiratory infections, meningitis, and tuberculosis). The results showed a positive association with threat words and increases in mortality rates from infectious diseases ( $B = 0.10$ ,  $P < 0.001$ , 95% CI [0.07, 0.12],  $R^2 = 0.37$ ). [SI Appendix, Table S7](#) presents the results of the multilevel analysis for both infectious disease deaths and MDDs.

We also examined whether the threat dictionary can capture the growing severity of the COVID-19 pandemic. As more people in the United States were affected by COVID-19 in 2020, we tested how this growing public health crisis, measured by the number of cases and deaths, coincides with more threat words found in people's real-time tweets. We matched the time-stamped and geolocated tweets in our sample to its corresponding state's available pandemic statistics. Following the data processing detailed in [SI Appendix, section B2](#), we matched the remaining 105,209 tweets to the corresponding daily statewide number of COVID-19 cases and deaths from March to May of 2020. Since both cases and deaths were highly positively skewed, they were log transformed. We used negative binomial regression to account for the overdispersed distribution of threat words ( $M = 0.39$ ,  $SD = 0.68$ ) (32). The results of our analysis revealed how threat words found in tweets increased as both COVID-19 cases (incident rate ratio [IRR] = 1.02, 95% CI [1.01, 1.02],  $P < 0.001$ ) and deaths (IRR = 1.02, 95% CI [1.02, 1.03],  $P < 0.001$ ) grew in number. On average, 4% more threat words appeared per tweet with every 10-factor increase in positive cases, and 5% more threat words appeared with every 10-factor increase in deaths.

For divergent validity, we performed a supervised machine learning classification test, wherein a random forest model was trained to identify tweets on the topic of COVID-19 ( $n = 105,348$ ) against tweets unrelated to COVID-19 ( $n = 133,313$ ). The model made these classifications based on how often words from different dictionaries were found in a particular tweet. This included the threat dictionary and other dictionaries with nomological overlap to threats, such as dictionaries on death, risk, negative emotion (12), and moral foundations (33). More detail on this analysis is described in [SI Appendix, SI Appendix, Fig. S1](#) shows how, as compared with other dictionaries, the

threat dictionary exhibited the highest feature importance for classification purposes, supporting its capacity to accurately distinguish threats in real-time broadcasted content.

We also examined whether threatening language is contagious. On Twitter, retweets involve people sharing a tweet to their own extended network (34). Based on previous research, we ran a negative binomial regression and tested whether the number of threat terms in a tweet is predictive of its retweet rate, while holding constant a number of twitter covariates as in ref. 32, including a user's number of followers, URLs found in the tweet, and the verified status of the user. We found that tweets on COVID-19 that expressed more threat terms accrued more retweets (IRR = 1.18, 95% CI [1.15, 1.21],  $P < 0.001$ ). On average, adding a single threat word to a tweet increased its expected retweet rate by 18%, indicative of the contagious properties of threat words on social media. Given that previous research (32) found that tweets high in moral-emotional words were more likely to be retweeted, we also tested whether our threat measure predicted retweets above and beyond the moral-emotional measure. We replicated the same analytical procedure with the same covariates found in previous research (32) and used measures of both threat and moral-emotional words as predictors in our model. After controlling for the effects of moral-emotional words, adding a single threat word to a tweet increased its expected retweet rate by 15% (IRR = 1.15, 95% CI [1.12, 1.19],  $P < 0.001$ ). These results, further described with additional robustness tests in [SI Appendix](#), suggest that linguistic cues that elicit a potential threat possess the rhetorical advantage of garnering people's attentional interest and can spread to more people. Particularly in times of public crises, the contagious nature of threat-laden online messages has important bearing on urgent issues regarding social media's role in amplifying misinformation campaigns and mass panic, a point we return to in [Discussion](#).

Having demonstrated how our threat measure coincides with actual threats in US history, we next examined its predictive power. Specifically, we applied our long-run data tracking national fluctuations in threat levels from newspapers to assess corresponding shifts in America's cultural, political, and economic standing over time.

For this analysis, we addressed common problematic features of time series data structures, including serial dependence, and lagged forecast errors that can result in spurious findings. To do this, we fitted ARIMA models to test the linear approximation of these relationships, with adjustments made to the error terms that address these temporal dependencies (26). Three main parameters ( $p$ ,  $d$ ,  $q$ ) are specified for the ARIMA errors. The  $p$  component denotes the number of lags that account for the autoregressive structure of the model, the  $d$  parameter refers to the order of differencing needed to stabilize the variance in the time series, and the  $q$  term represents the moving average value that explains the model's lagged random errors. For our data collected at monthly intervals, the same three parameters were extended to capture any seasonal influences, notated in capitalized form ( $P$ ,  $D$ ,  $Q$ ). Identification of parameters for our ARIMA models was obtained using an algorithm (24) that systematically searches for the optimal combination of parameters with the least fitting error based on the Akaike information criterion (AIC). The ARIMA model classification and results are presented in [Table 1](#).

**Cultural Shifts.** Research on the evolution of human culture demonstrates how groups adjust and adapt their attitudes and norms in accordance with changing environmental demands (35). Past studies have shown how cultures with a history of ecological and human-made threats (invasions, natural disasters, and pathogens) are higher on tightness (i.e., have strictly enforced social norms and low tolerance for deviant members)

**Table 1. Results of regression with ARIMA errors**

Indicator	Threat				Threat and GDP per capita			
	( <i>p, d, q</i> ) ( <i>P, D, Q</i> )*	<i>B</i> ( <i>SE</i> )	<i>t</i>	<i>p</i>	( <i>p, d, q</i> ) ( <i>P, D, Q</i> )*	<i>B</i> ( <i>SE</i> )	<i>t</i>	<i>p</i>
Tightness	(1, 1, 1) (2, 0, 0)	0.08 (0.02)	5.28	<0.001	(1, 1, 1) (2, 0, 0)	0.09 (0.02)	5.46	<0.001
Collectivism	(2, 1, 2) (0, 0, 2)	0.54 (0.05)	9.97	<0.001	(0, 1, 2) (2, 0, 0)	0.56 (0.06)	10.14	<0.001
Anti-immigration	(0, 1, 0)	0.35 (0.12)	3.02	<0.01	(1, 0, 0)	0.34 (0.12)	2.90	<0.01
Presidential approval	(2, 1, 2) (2, 0, 0)	0.06 (0.02)	3.19	<0.01	(1, 0, 2) (2, 0, 0)	0.06 (0.02)	3.14	<0.01
Republican partisanship	(0, 1, 0)	0.24 (0.12)	2.05	0.04	(1, 0, 1)	0.20 (0.11)	1.70	0.09
S&P 500	(3, 2, 1)	-0.01 (0.003)	-4.07	<0.001	(0, 2, 4)	-0.01 (0.003)	-4.00	<0.001
DJIA	(3, 1, 0) (2, 1, 0)	-0.03 (0.01)	-4.85	<0.001	(0, 2, 1)	-0.02 (0.004)	-4.07	<0.001
NASDAQ	(5, 2, 4) (0, 0, 2)	-0.01 (0.004)	-2.24	0.03	(2, 2, 3) (0, 0, 2)	-0.01 (0.004)	-2.38	0.02
Patents	(2, 2, 2)	-0.10 (0.03)	-3.51	<0.001	(1, 1, 1)	-0.12 (0.03)	-3.96	<0.001

The set of results on the left features the parameters and coefficients from ARIMA models with only threat as a predictor. Estimates presented on the right side of the table belong to models with threat as a predictor, in addition to controlling for real GDP per capita.

\*The ARIMA model parameters are specified by its nonseasonal components (*p, d, q*) and seasonal components (*P, D, Q*). For example, we fit a regression with ARIMA (2, 2, 2) errors for the model regressing patent numbers on threat levels over time.

(36, 37). Likewise, ecological disasters have been correlated with cultural collectivism or greater group orientation that supersedes individualistic priorities (38, 39). We tested these relationships with a more expansive conceptualization of threats as they have taken place over time. We expected variance in threat levels in the United States to be associated with shifts in America’s tight and collectivistic leanings. To quantify year to year change in these cultural properties from our corpus of newspapers, we used previously validated linguistic measures of tightness–looseness (40) and collectivism (16, 41). Words associated with tight cultures include “constrain,” “comply,” and “dictate,” and sample words that represent loose cultures are “allow,” “leeway,” and “limitless.” For our collectivistic terms, second-person pronouns were used, such as “we” and “us.” We tallied the relative frequencies of these keywords in our corpus of newspapers to obtain time series on monthly change in these cultural properties from 1900 to 2020. Consistent with our predictions, the results from our ARIMA models showed that variations in threat were positively associated with America’s cultural tightness ( $B = 0.08$ , 95% CI [0.05, 0.11],  $P < 0.001$ ) and collectivistic leanings ( $B = 0.54$ , 95% CI [0.43, 0.65],  $P < 0.001$ ).


**Political Shifts.** In political science and other behavioral science disciplines, times of great threat have been linked with higher leadership popularity and political conservatism (42–44). The uncertainty–threat model, an integration of these theories, holds that under dangerous conditions, humans wish to maintain the status quo in order to mitigate their feelings of uncertainty and fear (45). This manifests in political preferences and attitudes that are more conservative, favorable toward current institutional authorities, and ethnocentric. Studies on the threat–conservatism relationship have primarily used experimental or case study methods (45–47). Our method enables us to examine these patterns at the population level over a 100-year period. Additionally, past empirical work demonstrating how crises lead to more popular support of national leaders, known as the “rally around the flag” effect (48), mainly focused on specific foreign conflicts and wars (49, 50). Here, we offer a broader conceptualization of societal threats and their political ramifications over longer periods of US history.

The original source for data on the political indicators of this study came from Gallup, the longest polling service that has tracked attitudinal trends among Americans. For our analysis on leadership support, we examined US presidential approval numbers from the popular Gallup survey item: “Do you approve or disapprove of the way [enter President name] is handling his job as President?” The frequent administration of this survey item at daily to weekly intervals from 1945 made it

amenable to aggregate as a monthly average (51). As a measure of ethnocentrism, we examined the annual percentage of Americans who reported preferring fewer immigrants in the United States, which has been available since 1965 (monthly data were not available). Lastly, Republican party identification is based on Gallup’s polling question on which political party Americans most identify with—captured by taking the annual percentage of Republican identifiers out of the sum of Republican and Democratic identifiers (52, 53). The results of our findings showed that increases in threat levels coincided with Americans’ approval of their sitting president from Harry Truman all the way through to Donald Trump ( $B = 0.06$ , 95% CI [0.02, 0.10],  $P < 0.01$ ); increases in ethnocentric attitudes (i.e., percentage of Americans who wanted fewer immigrants in the country;  $B = 0.35$ , 95% CI [0.12, 0.58],  $P < 0.01$ ); and greater Republican identification ( $B = 0.24$ , 95% CI [0.01, 0.46],  $P < 0.05$ ), suggesting conservative gains during high-threat periods.

To further probe this interplay between threat sensitivity and political conservatism, we tested whether public remarks made by US presidents showed differences along party lines in their references to threats. Political analysts have highlighted the strategic communication of threats by Republican leaders from Richard Nixon’s fear appeals on rising crime rates to George W. Bush’s rhetoric on the threat of terrorism after the September 11th attacks (54–57). Turning to the Miller Center’s dataset (58) of famous speeches made by US presidents during their time in office, we found that as compared with Democratic presidents ( $M = 1.37$ ,  $SD = 0.77$ ), Republican presidents ( $M = 1.59$ ,  $SD = 1.01$ ) indeed alluded to threats at higher levels:  $t(356) = -2.39$ ,  $P = 0.02$ . Fig. 3 provides a comparison of US presidents from the past 70 years and their average use of threat words in speeches. For this comparison, we analyzed data from 1948 to 2020, with 1948 marking the year in which the two major parties ideologically realigned and the Democratic Party officially adopted a more liberal political agenda (59, 60).

**Macroeconomic Shifts.** Finally, we examined the macroeconomic response to societal threats. Climatoeconomic theory describes how communities that deal with extreme hot and cold climates are less risk-taking and more prone to holding on to their monetary assets (61). Past researchers have shown how different categories of national security threats—armed conflicts (62), disease outbreaks (63), natural disasters (64), and terrorism (65)—adversely impact a country’s economic development and innovation capabilities. Our threat measure enables an assessment of the economic consequence of reported threats across the last century based on a variety of threats. As a barometer of the country’s financial health, we compiled daily closing stock



**Fig. 3.** Percentage of threat dictionary terms found in US presidential speeches. The bars represent the average percentages of threat dictionary words found across all speeches and public remarks made by US presidents from 1948 to 2020, with red and blue shading corresponding to each president's political party affiliation (Republican and Democrat, respectively).

prices of the three major market indices listed on the US stock exchange: the Standard and Poor (S&P) 500, the Dow Jones Industrial Average (DJIA), and the National Association of Securities Dealers Automated Quotations (NASDAQ) Composite. Daily price returns were averaged at monthly intervals from the start of each stock's established indices up to December 2020. For example, data on the S&P 500 were collected from 1957 when the index first became a 500-stock composite to the end of 2020 (66). DJIA performance was measured from 1928 to 2020, and we collected the closing prices of the NASDAQ from 1971 to 2020. Lastly, we indexed the number of utility patent applications from 1900 to 2019. This annual count is provided by the US Patent and Trademark Office (USPTO) for utility patents, which include new inventions and improvements to existing products from 1900 to 2019. Patent counts have been commonly used in past studies as a national measure of inventive activities and potential for creativity (40, 67). We found that greater threat levels in newspapers were significantly negatively associated with stock market returns for the S&P 500 ( $B = -0.01$ , 95% CI  $[-0.04, -0.02]$ ,  $P < 0.001$ ), DJIA ( $B = -0.03$ , 95% CI  $[-0.04, -0.02]$ ,  $P < 0.001$ ), and NASDAQ ( $B = -0.01$ , 95% CI  $[-0.02, -0.001]$ ,  $P < 0.05$ ). Increases in threat were also negatively related to the number of USPTO-reported patent applications ( $B = -0.10$ , 95% CI  $[-0.15, -0.04]$ ,  $P < 0.001$ ).

Table 1 summarizes the results of these ARIMA models and shows how the relationship between threat and these indicators remained significant even after controlling for real gross domestic product (GDP) per capita, except for Republican Party partisanship. We also provide correlational results in *SI Appendix*, Table S8 to demonstrate alternative methods for examining these joint processes.

To understand the directionality of these cross-temporal relationships, Granger tests of predictive causality were next conducted to study whether threat levels precede these cultural, political, and economic changes. Given two time series, Granger causality models test whether past values of a predictor variable can explain the changing outcome of another variable over and beyond the ability of this outcome variable's prior observations predicting its own future values (68). Across two sets of Granger analyses, two potential directional possibilities were tested with 1) threat modeled as a predictor and 2) threat modeled as an outcome. In the first model, if threat significantly predicts a societal outcome, while the reverse direction is not significant in the second model, these two sets of results

provide strong evidence of a specific temporal ordering wherein changes in a societal outcome likely follow threats. If the second model is significant, but the first model is not, this would indicate that changes in threat levels likely follow a societal trend. When both models are significant, this can indicate bidirectionality or possible influence of an exogenous variable causing changes to both series.

To conduct Granger tests, we removed the time dependencies from each individual series using the previously described ARIMA procedure and then, extracted the residuals from each series. Table 2 summarizes the findings of our Granger tests at a select lag of up to 5 years, as with the lag specifications of previous research (see refs. 38 and 40). Following a model comparison of all lags possible for up to a period of 5 years, the reported models correspond to the lag lengths representing the best model fit according to AIC estimates.

The results indicated that, under the optimal lag of our threat measure, threat levels significantly predicted cultural tightness, collectivism, the S&P 500, and the DJIA—over and beyond lagged values of the criterion predicting its own current values. The reverse directionality was also found to be significant for these indicators. For example, just as stock market performance suffers from news of threats, stock market downturns themselves would also instate a national financial threat. Meanwhile, lags of threat significantly predicted NASDAQ outcomes, whereas the reverse direction was not significant. While significantly correlative, our Granger models showed no significant lagged causal links, modeled in both directions, between threat and the remaining indicators. This could suggest a periodicity issue wherein directionality is sensitive to high-frequency monthly data. For example, data on changes in the percentage of Republican Party affiliates, anti-immigration views, and patents numbers were assessed at annual levels based on available data. It is possible that these societal indicators are fast changing or quick to normalize, thereby explaining why directionality would not be captured by low-frequency yearly lags.

## Discussion

A primary goal of the present study was to show how language can be harnessed to estimate when diverse threats occur throughout history and their effects on societal culture, politics, and macroeconomic activity. Our dictionary development process introduced the application of multimodel computational techniques to identify threat-relevant words and clustering methods to inform the final threat dictionary. We also validated our threat dictionary based on its convergence with documented threats that occurred in real life, demonstrating the measure's sensitivity to major categories of threats, including wars, pathogen stress, and natural disasters. The convergence between these resulting threat indices and actual high-threat events over time supports the stability and generalizability of our threat measure.

Our linguistic tool offers predictive insights into societal responses to mass communicated threats. Using several time series analytical methods, we demonstrated how historical patterns in threat levels coincided with stock market trends, conservative political attitudes, presidential approval numbers, and changing cultural norms. This is an attempt to empirically examine different cultural, political, and economic cross-temporal effects at once using one comprehensive measure of threat over many time points. While the current study's scope is limited to the English language, especially as it relates to dynamics found in the United States, this work's methods and findings can be extended to additional languages and national contexts. In addition, future research can examine how the dictionary relates to other phenomena beyond those we studied (e.g., religiosity, foreign policy, and economic investment

**Table 2.** F statistic for Granger causality tests at maximum lags of  $t - 5$  y

Indicator	Threat precedes indicator	Indicator precedes threat	Lag order ( $t$ ) <sup>†</sup>
Tightness	$F(2,482, 2,542) = 1.35^*$	$F(2,482, 2,542) = 1.63^{**}$	$t - 60$ mo
Collectivism	$F(2,503, 2,560) = 2.06^{***}$	$F(2,503, 2,560) = 1.78^{***}$	$t - 57$ mo
Anti-immigration	$F(101, 102) = 0.12$	$F(101, 102) = 3.00$	$t - 1$ y
Presidential approval	$F(1,805, 1,806) = 0.65$	$F(1,805, 1,806) = 3.27$	$t - 1$ mo
Republican identification	$F(123, 124) = 1.46$	$F(123, 124) = 0.08$	$t - 1$ y
S&P 500	$F(1,215, 1,260) = 2.87^{***}$	$F(1,215, 1,260) = 1.85^{***}$	$t - 45$ mo
DJIA	$F(1,897, 1,942) = 0.27^{***}$	$F(1,897, 1,942) = 2.24^*$	$t - 45$ mo
NASDAQ	$F(986, 1,016) = 4.52^{***}$	$F(986, 1,016) = 0.24$	$t - 30$ mo
Patent applications	$F(231, 232) = 0.03$	$F(231, 232) = 0.06$	$t - 1$ y

\* $P < 0.05$ ; \*\* $P < 0.01$ ; \*\*\* $P < 0.001$ .

<sup>†</sup>The  $t$  is the number of time points lagged. The AIC was used as an objective benchmark for the optimal order selection of lags to report.

among others). In addition, future research that seeks to measure specific types of threat linguistically can apply our measurement development process to create more specialized dictionaries—for instance, on mass shootings or cybersecurity threats.

The threat dictionary may be useful for understanding a number of critical societal issues. History is replete with examples of organizations and political leaders who have been culpable of inflating threats and peddling fear to obtain popular support or undermine democratic principles (3, 69). The threat dictionary enables comparisons and evaluations of historical and contemporary leaders' talking points to assess these exaggerations of threats or strategic fabrications of collective dangers and their consequences. This linguistic tool can also be used to examine how threat—whether real or manipulated—propagates across social media and its negative effects online and ultimately, offline. Indeed, on social media platforms, millions of users and news outlets are jockeying to be heard and seen. We have shown how the prevalence of threatening language within tweets increases their widespread dissemination capabilities. While the current paper only tested its contagious properties with Twitter behavior during the COVID-19 pandemic, future directions include generalizing these findings to other collective threats. Excessive use of threat words may also explain the disconcerting popularity of counterproductive online content. For instance, an analysis of online groups that capitalize on threat-fueled chatter can reveal which networks are prone to using threatening language to motivate followers and how threat travels through different online spaces (i.e., “highways of hate”) (70, 71). The use of threat words is a salient and powerful rhetorical device that may have substantive

implications for future research aimed at curbing the misinformation, prejudice, and mass panic that frequently unravels within these online ecologies.

As more rich textual resources become accessible and more researchers integrate NLP techniques into their methodological tool kits, reliable instruments for analyzing written works such as dictionaries are in high demand. Currently, existing dictionaries tap into many useful sociopsychological constructs, such as positive affect or whether people's moral beliefs are grounded in fairness or purity (11, 33). We build upon these important efforts to linguistically examine the multifaceted construct of threat, which has long been of critical theoretical and practical significance to both researchers and policy makers. The threat dictionary provides opportunities for collecting data on changing threat levels with high temporal resolution, across media platforms, and across different levels of analysis. Future research will now be able to formally test extant theoretical work that previously lacked an adequate long-term measure to empirically verify hypotheses on societal responses to collective threats. Much like a searchlight, the threat dictionary's ability to scan texts for the prevalence of threat-related words is informative not only of people's daily exposure to this terminology but also, of the linguistic footprints that make interesting societal patterns across history visible.

**Data Availability.** Code and data for these results are available at the Open Science Framework (OSF; <https://osf.io/eydaqb/>). All other data are included in the article and/or *SI Appendix*.

**ACKNOWLEDGMENTS.** We thank both Dylan Pieper and Ioanna Galani for their research assistance. The present research was funded in part by Office of Naval Research grant N000141912407 (M.J.G.). The information in this article does not imply or constitute an endorsement of the views therein by the Office of Naval Research, US Navy, or Department of Defense.

- D. L. Altheide, The news media, the problem frame, and the production of fear. *Sociol. Q.* **38**, 647–668 (1997).
- D. L. Altheide, R. S. Michalowski, Fear in the news: A discourse of control. *Sociol. Q.* **40**, 475–503 (1999).
- B. Glassner, *The Culture of Fear: Why Americans Are Afraid of the Wrong Things: Crime, Drugs, Minorities, Teen Moms, Killer Kids, Muta* (Hachette UK, 2010).
- T. Rose-Stockwell, This is how your fear and outrage are being sold for profit. *Medium*, 14 July 2017. <https://tobiasrose.medium.com/the-enemy-in-our-feeds-e86511488de>. Accessed 26 July 2021.
- A. Olsson, E. A. Phelps, Social learning of fear. *Nat. Neurosci.* **10**, 1095–1102 (2007).
- P. Vuilleumier, J. Armony, R. Dolan, Reciprocal links between emotion and attention. *Human Brain Function* **2**, 419–444 (2003).
- M. Appel, C. Marker, T. Gnams, Are social media ruining our lives? A review of meta-analytic evidence. *Rev. Gen. Psychol.* **24**, 60–74 (2020).
- C. Klemm, T. Hartmann, E. Das, Fear-mongering or fact-driven? Illuminating the interplay of objective risk and emotion-evoking form in the response to epidemic news. *Health Commun.* **34**, 74–83 (2019).
- M. B. Tannenbaum et al., Appealing to fear: A meta-analysis of fear appeal effectiveness and theories. *Psychol. Bull.* **141**, 1178–1204 (2015).
- J. P. Walsh, Social media and moral panics: Assessing the effects of technological change on societal reaction. *Int. J. Cult. Stud.* **23**, 840–859 (2020).
- 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).
- J. W. Pennebaker, R. L. Boyd, K. Jordan, K. Blackburn, The development and psychometric properties of LIWC2015 (2015). [https://repositories.lib.utexas.edu/bitstream/handle/2152/31333/LIWC2015\\_LanguageManual.pdf](https://repositories.lib.utexas.edu/bitstream/handle/2152/31333/LIWC2015_LanguageManual.pdf). Accessed 23 June 2020.
- T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space. arXiv [Preprint] (2013). <https://arxiv.org/abs/1301.3781> (Accessed 23 June 2020).
- J. Pennington, R. Socher, C. D. Manning, “GloVe: Global Vectors for Word Representation” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Stroudsburg, PA, 2014), pp. 1532–1543.
- L. E. Jehl, “Machine translation for Twitter,” Dissertation, University of Edinburgh, Edinburgh, UK (2010).
- J. M. Twenge, W. K. Campbell, B. Gentile, Generational increases in agentic self-evaluations among American college students, 1966–2009. *Self. Ident.* **11**, 409–427 (2012).
- S. Park, H. M. Kim, *Improving the Accuracy and Diversity of Feature Extraction from Online Reviews Using Keyword Embedding and Two Clustering Methods* (American Society of Mechanical Engineers, 2020).

18. A. Y. Ng, M. I. Jordan, Y. Weiss, "On spectral clustering: Analysis and an algorithm" in *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic* (MIT Press, Vancouver, British Columbia, Canada, 2001), pp. 849–856.
19. Newspapers.com, Experience the headlines like they did. <https://www.newspapers.com/>. Accessed 3 January 2021.
20. S. Pinker, *The Better Angels of Our Nature: Why Violence Has Declined* (Penguin Group USA, 2012).
21. J. S. Goldstein, *Winning the War on War: The Decline of Armed Conflict Worldwide* (Plume Books, 2012).
22. T. R. Gurr, Historical trends in violent crime: A critical review of the evidence. *Crime Justice* **3**, 295–353 (1981).
23. J. Mueller, War has almost ceased to exist: An assessment. *Polit. Sci. Q.* **124**, 297–321 (2009).
24. R. J. Hyndman, Y. Khandakar, Automatic time series forecasting: The forecast package for R. *J. Stat. Softw.* **27**, 1–22 (2008).
25. R. Hyndman *et al.*, Forecasting functions for time series and linear models. R Package version 8 (2019). <http://pkg.robjhyndman.com/forecast>. Accessed 4 January 2021.
26. A. T. Jebb, L. Tay, Introduction to time series analysis for organizational research: Methods for longitudinal analyses. *Organ. Res. Methods* **20**, 61–94 (2017).
27. P. J. Brockwell, P. J. Brockwell, R. A. Davis, R. A. Davis, *Introduction to Time Series and Forecasting* (Springer, 2016).
28. M. J. Gelfand, *Rule Makers, Rule Breakers: How Culture Wires Our Minds, Shapes Our Nations and Drives Our Differences* (Robinson, 2018).
29. J. Henrich, *The Secret of Our Success* (Princeton University Press, 2015).
30. P. Turchin, *Ultrasociety: How 10,000 Years of War Made Humans the Greatest Cooperators on Earth* (Beresta Books, Chaplin, CT, 2016).
31. P. D. Bliese, J. W. Lang, Understanding relative and absolute change in discontinuous growth models: Coding alternatives and implications for hypothesis testing. *Organ. Res. Methods* **19**, 562–592 (2016).
32. W. J. Brady, J. A. Wills, J. T. Jost, J. A. Tucker, J. J. Van Bavel, Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 7313–7318 (2017).
33. J. Graham, J. Haidt, B. A. Nosek, Liberals and conservatives rely on different sets of moral foundations. *J. Pers. Soc. Psychol.* **96**, 1029–1046 (2009).
34. A. Goldenberg, J. J. Gross, Digital emotion contagion. *Trends Cogn. Sci.* **24**, 316–328 (2020).
35. R. Inglehart, W. E. Baker, Modernization, cultural change, and the persistence of traditional values. *Am. Sociol. Rev.* **65**, 19–51 (2000).
36. M. J. Gelfand *et al.*, Differences between tight and loose cultures: A 33-nation study. *Science* **332**, 1100–1104 (2011).
37. J. C. Jackson, M. Gelfand, C. R. Ember, A global analysis of cultural tightness in non-industrial societies. *Proc. Biol. Sci.* **287**, 20201036 (2020).
38. I. Grossmann, M. E. Varnum, Social structure, infectious diseases, disasters, secularism, and cultural change in America. *Psychol. Sci.* **26**, 311–324 (2015).
39. H. C. Triandis, Individualism-collectivism and personality. *J. Pers.* **69**, 907–924 (2001).
40. J. C. Jackson, M. Gelfand, S. De, A. Fox, The loosening of American culture over 200 years is associated with a creativity-order trade-off. *Nat. Hum. Behav.* **3**, 244–250 (2019).
41. L. D. Stone, J. W. Pennebaker, Trauma in real time: Talking and avoiding online conversations about the death of Princess Diana. *Basic Appl. Soc. Psychol.* **24**, 173–183 (2002).
42. J. T. Jost *et al.*, Are needs to manage uncertainty and threat associated with political conservatism or ideological extremity? *Pers. Soc. Psychol. Bull.* **33**, 989–1007 (2007).
43. H. Thórisdóttir, J. T. Jost, Motivated closed-mindedness mediates the effect of threat on political conservatism. *Polit. Psychol.* **32**, 785–811 (2011).
44. G. D. Wilson, *The Psychology of Conservatism* (Academic Press, 1973).
45. J. T. Jost, J. Glaser, A. W. Kruglanski, F. J. Sulloway, Political conservatism as motivated social cognition. *Psychol. Bull.* **129**, 339–375 (2003).
46. J. T. Jost, D. M. Amodio, Political ideology as motivated social cognition: Behavioral and neuroscientific evidence. *Motiv. Emot.* **36**, 55–64 (2012).
47. J. T. Jost, C. Stern, N. O. Rule, J. Sterling, The politics of fear: Is there an ideological asymmetry in existential motivation? *Soc. Cogn.* **35**, 324–353 (2017).
48. J. E. Mueller, Presidential popularity from Truman to Johnson. *Am. Polit. Sci. Rev.* **64**, 18–34 (1970).
49. J. R. Oneal, A. L. Bryan, The rally 'round the flag effect in US foreign policy crises, 1950–1985. *Polit. Behav.* **17**, 379–401 (1995).
50. C. W. Ostrom, D. M. Simon, Promise and performance: A dynamic model of presidential popularity. *Am. Polit. Sci. Rev.* **79**, 334–358 (1985).
51. G. Peters, Presidential job approval. American Presidency Project. <https://www.presidency.ucsb.edu/statistics/data/presidential-job-approval>. Accessed 3 January 2021.
52. R. S. Erikson, M. B. MacKuen, J. A. Stimson, *The Macro Polity* (Cambridge University Press, 2002).
53. J. Stimson, James Stimson's Site » Data » Macropartisanship. <https://stimson.web.unc.edu/data>. Accessed 4 January 2021.
54. D. H. Bennett, *The Party of Fear: The American Far Right from Nativism to the Militia Movement* (Vintage, New York, NY, 1995).
55. D. E. Foster, Pulling at our heartstrings: The Republican Party's use of pathos in their presidential campaign rhetoric as an explanation for their success in recent presidential elections. *Ky. J. Commun.* **29**, 8–12 (2010).
56. C. Robin, *Fear: The History of a Political Idea* (Oxford University Press, 2004).
57. S. M. Lipset, E. Raab, The message of proposition 13. *Commentary* **66**, 42 (1978).
58. Miller Center, Famous presidential speeches (2016). <https://millercenter.org/the-presidency/presidential-speeches>. Accessed 4 January 2021.
59. J. Bell, The changing dynamics of American liberalism: Paul Douglas and the elections of 1948. *J. Ill. State Hist. Soc.* **96**, 368–393 (2003).
60. R. R. Preuhs, Racial realignment: The transformation of American liberalism, 1932–1965. By Eric Schickler. Princeton, NJ: Princeton University Press, 2016. 384p. \$35.00 cloth. *Perspect. Polit.* **16**, 527–529 (2018).
61. E. Van de Vliert, Climato-economic habitats support patterns of human needs, stresses, and freedoms. *Behav. Brain Sci.* **36**, 465–480 (2013).
62. G. Schneider, V. E. Troeger, War and the world economy: Stock market reactions to international conflicts. *J. Conflict Resolut.* **50**, 623–645 (2006).
63. R. J. Barro, J. F. Ursúa, J. Weng, *The Coronavirus and the Great Influenza Pandemic: Lessons from the 'Spanish Flu' for the Coronavirus's Potential Effects on Mortality and Economic Activity* (National Bureau of Economic Research, 2020).
64. E. Cavallo, S. Galiani, I. Noy, J. Pantano, Catastrophic natural disasters and economic growth. *Rev. Econ. Stat.* **95**, 1549–1561 (2013).
65. G. A. Karolyi, R. Martell, Terrorism and the stock market. *Int. Rev. Appl. Financial Issues Econom.* **2**, 285–314 (2010).
66. C. Valetkevitch, Key dates and milestones in the S&P 500's history. *Reuters*, 6 May 2013. <https://www.reuters.com/article/us-usa-stocks-sp-timeline-idUSBRE9450WL20130506>. Accessed 2 January 2021.
67. J. Schmookler, The level of inventive activity. *Rev. Econ. Stat.* **36**, 183–190 (1954).
68. C. A. Sims, Money, income, and causality. *Am. Econ. Rev.* **62**, 540–552 (1972).
69. M. J. Gelfand, R. Lorente, "Threat, tightness, and the evolutionary appeal of populist leaders" in *The Psychology of Populism: The Tribal Challenge to Liberal Democracy*, J. P. Forgas, W. D. Crano, K. Fiedler, Eds. (Routledge, New York, NY, 2021), pp. 276–294.
70. P. Ball, A. Maxmen, The epic battle against coronavirus misinformation and conspiracy theories. *Nature* **581**, 371–374 (2020).
71. N. Johnson *et al.*, Social media cluster dynamics create resilient global hate highways. arXiv [Preprint] (2018). <https://arxiv.org/abs/1811.03590> (Accessed 4 January 2021).