

Outlet- and Topic-Level Selective Avoidance Limit the Reach of Content Challenging Election Fraud Misinformation

Mathieu Lavigne
McGill University

Brian Fogarty
University of Notre Dame

John Carey
Dartmouth College

Brendan Nyhan
Dartmouth College

Jason Reifler
University of Southampton

Abstract

If exposure to corrective information typically reduces false beliefs, why are high-profile misperceptions like the myth of widespread fraud in the 2020 U.S. election so persistent? One explanation is selective exposure to pro-attitudinal claims, but we show using nationally representative survey and online behavior data that most supporters of Donald Trump encountered little fraud-related content of any kind around the 2020 and 2022 elections (which we measure using a large language model). Instead, we observe a pattern of selective *avoidance* among Trump supporters, who were less likely than were Biden supporters to see articles that challenge fraud claims with facts and evidence. Counterfactual permutation tests indicate that this gap was driven by differences in news outlet choice (preferring sources that challenge fraud claims less, including untrustworthy ones) and topical attention (consuming less news about fraud). These findings suggest that the most difficult challenge in correcting misinformation is reaching vulnerable audiences.

We are grateful for funding support from an Evolving Election Administration Landscape Grant awarded by MIT Election Data + Science Lab (funds provided by Election Performance Project, LLC, a subsidiary of the Pew Charitable Trusts). This project received funding from the National Science Foundation (award number 2103262) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number 682758). Mathieu Lavigne acknowledges funding from the Social Sciences and Humanities Research Council of Canada in the form of a postdoctoral fellowship. We thank Shun Yamaya, Thomas Zeitzoff, Matt Levendusky, and Hwayong Shin for helpful comments as well as attendees at the 2024 Knowledge Resistance conference at Stockholm University, the 2025 Information and Misinformation in Elections Conference at Caltech, the 2024 Canadian Political Science Association conference, the 2025 American Political Science Association conference, and seminars at the Harvard/Northeastern misinformation speaker series, American University, Max Planck Institute, Boston University, and Northwestern University. We thank Yang Xu for assistance with data collection. All errors are our own.

False or unsupported beliefs about high-profile issues like climate change (Marlon et al. 2024), Barack Obama’s birthplace (Brown 2013), and vaccines (Reinhart 2020) can linger for years, confusing public debate and influencing policy. Strikingly, these misperceptions can persist even when information demonstrating they are false is widely available (Berinsky 2012; Nyhan 2020).

A particularly prominent and consequential misperception in American politics is the longstanding belief that election fraud is widespread (Mayer 2012; Minnite 2010; YouGov America 2017). Following the 2016 election, Donald Trump made repeated false claims about the prevalence of widespread voter fraud despite winning the election. Trump fanned these flames further during and after the 2020 election, using false fraud claims to try to overturn his defeat and ultimately inspiring a violent insurrection at the U.S. Capitol in January 2021. The false beliefs he promoted, which appear to be durable and largely sincere (Graham and Yair 2024), undermined confidence in the integrity of the U.S. electoral system (Bright Line Watch 2024). Confidence in elections is essential to citizens’ willingness to vote (Birch 2010; Franklin 2004) and to the willingness of election losers to accept defeat (Anderson et al. 2005; Nadeau and Blais 1993), which is necessary for democracy to survive (Przeworski 2019). Given the stakes, it is critical to understand why false beliefs in widespread election fraud persist.

Until recently, scholars typically attributed the durability of misperceptions about controversial issues like election fraud to directionally motivated reasoning. According to this account, people tend to uncritically accept politically congenial claims that are false or unsupported and to reject accurate information that is uncongenial (Flynn et al. 2017; Peterson and Iyengar 2021). However, numerous studies show that fact-checking is generally effective at correcting misperceptions (Carnahan et al. 2021; Coppock et al. 2023; Nyhan 2021; Walter et al. 2020; Walter and Murphy 2018; Wood and Porter 2019; see Nyhan 2021 for a review). In particular, so-called “backfire effects” appear to be exceedingly rare (Prike and Ecker 2023; Swire-Thompson et al. 2020); instead, fact-checks and other forms of corrective information typically produce patterns of parallel belief updating rather than backfire even among groups who might find them uncongenial (Carey et al. 2020, 2025).

Given the lack of evidence for directionally motivated reasoning, we instead consider a prior step of belief and attitude formation — information exposure — and ask to what extent people are being exposed to information that would challenge their misconceptions about election fraud.

The leading scholarly explanation for how people’s information diets might reinforce and/or fail to correct election fraud misperceptions is selective exposure, which includes both selective approach (a preference for consuming congenial content) and selective avoidance (a tendency to avoid uncongenial content; see, e.g., Arceneaux and Johnson 2019). Though theoretical frameworks and empirical findings vary, audience demand for pro-attitudinal news and information is well-documented (Iyengar and Hahn 2009; Nyhan et al. 2023; Stroud 2008; Taber and Lodge 2006) and has often been found to be stronger than preferences for avoiding counter-attitudinal information (Garrett 2009; Garrett et al. 2013; Hart et al. 2009, though see, e.g., Garrett and Stroud 2014). Critics fear that people’s tendencies toward selective exposure, which can lead to greater exposure to false or misleading claims (Garrett et al. 2019, 2016; Meirick 2013; Meirick and Bessarabova 2016), are exacerbated online, where “echo chambers” or “filter bubbles” of like-minded content may reinforce viewpoint homogeneity (Pariser 2011; Sunstein 2001). Contrary to these fears, though, most Americans have relatively balanced news diets on average (Barrie et al. 2025; Guess et al. 2019, 2018). Moreover, exposure to untrustworthy websites and other dubious content and sources of information online is rare and concentrated among small groups of people with extreme preferences (Chen et al. 2023; Eady et al. 2023; Grinberg et al. 2019; Guess et al. 2020; Moore et al. 2023; see Budak et al. 2024 for a review).

While research on selective exposure emphasizes how political preferences affect choices about *which* news to consume, fewer studies focus on inattention — whether to consume any news or information at all — and “news avoidance,” in part because these behaviors are difficult to study in surveys that condition on attention (though see, e.g., Eady et al. 2019; Garrett et al. 2019; Prior 2009a). Survey self-reports vastly overstate news consumption, which is why measuring information exposure using behavioral data of the sort used in this study is so important (Jerit et al. 2016; Prior 2009a,b). In general, most people consume very little news or political content and, as a result, are typically uninformed about political matters (Lupia 2016). Some of the observed lack of exposure may stem from intentional news avoidance (people who actively avoid news content) but much of it may be unintentional (e.g., people preferring other types of content) (Skovsgaard and Andersen 2020, 2022). However, the prevalence of inattention does not mean no one is being exposed to news. For instance, levels of attention can vary over time — when the news agenda is uncongenial, partisans can reduce

overall news consumption (Kim and Kim 2021) and/or selectively avoid coverage of specific topics (Skovsgaard and Andersen 2022).

These patterns of inattention and potential news avoidance likely shape public opinion. If people in the real world never see the types of corrective messages that have been shown to work in experiments (conditional on exposure), their beliefs and attitudes will necessarily remain unchanged. In particular, the prior research described above generally suggests that detailed debunkings reduce false beliefs more than simple refutation or false tags (Chan and Albarracin 2023; Chan et al. 2017; Ecker et al. 2020; Swire et al. 2017). It is therefore important to measure whether people are ever exposed to such messages (rather than, say, headlines), which may be more likely to provide novel information that causes people to update their mental models or understanding (Carey et al. 2025; Thorson 2024).

Accordingly, this study reports the results of a comprehensive descriptive analysis of behavioral data measuring the American public's exposure to information about election fraud, including content that challenges misperceptions about the issue with facts and evidence (the detailed debunkings we expect to have the greatest effects). Unlike prior studies, we examine patterns of exposure at the topic (rather than source) level using behavioral data (rather than survey self-reports) and article-specific content data (rather than source-level measures of the political slant or trustworthiness of the domain). We find a pattern of selective news avoidance in which Trump supporters — the group most likely to hold misperceptions about election fraud — see less voter fraud content than Biden supporters and this deficit is greater for content challenging fraud claims.

Our results proceed as follows. First, we establish a puzzle — the persistence of fraud misperceptions despite the widespread availability of corrective information. Findings from a nationally representative survey experiment provide no evidence that the reason for fraud misperception persistence is resistance to corrective information. As we show, exposure to corrective information successfully reduces inaccurate beliefs about election fraud in the 2020 presidential election, including among the groups who should be most motivated to maintain those beliefs.

We therefore instead consider the role of information exposure to fraud-related content during the 2020 and 2022 election study periods. Using a large language model (LLM) and online behavior data, we specifically classify the fraud-related content that Americans encountered in web browsing (ex-

cluding social media) based on the extent to which the articles challenge fraud claims. This granular article-level classification process (a promising application of LLMs; see, e.g., Haroon et al. 2025) allows us to estimate exposure to content challenging claims of widespread fraud coverage in people's information diets and to see how it varies by factors such as candidate support and source trustworthiness.

These comprehensive measures of article-level exposure reveal widespread inattention to election fraud coverage as well as large differences in exposure by candidate support. Fewer than half of Americans were exposed to *any* election fraud content in their web browsing around the 2020 and 2022 elections and even fewer saw content that meaningfully challenged those claims using facts and evidence. Moreover, in both 2020 and 2022, Biden supporters consumed more overall election content than did Trump supporters and the content that they encountered was more likely to challenge false claims of fraud. In the 2020 election period, for instance, Biden supporters saw 14.3 articles challenging fraud claims on average, compared to 3.5 among Trump supporters. Using counterfactual permutation analyses, we show that these differences in exposure to content challenging fraud claims are driven primarily by topical attention and news outlet choice, not selective exposure to congenial versus uncongenial election fraud articles within outlets. Although Trump and Biden supporters consumed similar amounts of news overall, Trump supporters paid less attention to election fraud as a topic and, when exposed to fraud-related coverage, drew disproportionately on news sources that challenged election fraud claims less frequently. In particular, what little information Trump supporters did see about fraud often comes from untrustworthy sources.

These results suggest that high-profile misperceptions persist not because corrective information is unavailable or ineffective, but because of a pattern of selective avoidance of challenging content about relevant topics or issues. People who are most likely to hold inaccurate beliefs consume even less content about election fraud than do other Americans and the content that they do encounter is disproportionately unlikely to challenge fraud claims with facts and evidence.

Results

Our analysis proceeds in four steps. Using nationally representative survey data, we first establish the puzzle our study seeks to explain — the persistence of beliefs in widespread fraud in U.S. elections over time despite the widespread availability of corrective information. Second, we use a survey experiment conducted after the 2020 election to demonstrate that exposure to fact-checks reduces belief in fraud claims, suggesting that motivated resistance to corrective information does not seem to explain the persistence of fraud beliefs. We then turn to our main results, which leverage paired individual-level survey and online behavior data to measure exposure to different types of fraud-related content online at the article/URL level during the periods around the 2020 and 2022 U.S. elections. Our third finding shows that exposure to fraud-related content is lower among Trump supporters, especially content that challenges false claims with facts and evidence. Our final result examines the mechanism for this pattern of selective avoidance, which we demonstrate is driven by topic avoidance and outlet choice using counterfactual permutation tests.

Participant descriptives and online exposure data

Our study uses survey data from three nationally representative panel studies by YouGov conducted between May 2020 and January 2023 (29,884 total observations), including a survey experiment conducted in January 2021 (n=3847). We also examine behavioral data from participants who voluntarily provided passively collected online browsing data from laptop/desktop computers and/or mobile devices for three or more months in the September 13, 2020–January 29, 2021 period (n=1,596) or three or more months in the August 31, 2022–January 31, 2023 period (n=1,518). After applying post-stratification weights provided by YouGov, these participants closely resemble our 2020 and 2022 general population survey samples in terms of age, gender, education, race, and partisanship, though the 2022 sample is somewhat younger and less likely to support Trump (see Materials and Methods and Tables [S1–S2](#) in the Supplementary Information for more information).

We leverage this behavioral data to collect the text content from all publicly accessible URLs viewed by participants at domains rated for trustworthiness by Lin et al. (2023) or NewsGuard. Unlike prior research, which typically measures information exposure based on domain level visits using

online behavior data (e.g., Dahlke and Hancock 2025; Guess et al. 2020; Moore et al. 2023), we use a LLM to code online content, enabling us to measure exposure to fraud-related content at the article level. Our results indicate that about half of Americans were not exposed to fraud-related content at all in their online media diet (excluding social media). Most of the fraud-related content that people saw did question fraud claims at least minimally, but exposure to content that meaningfully challenged these claims with facts and evidence was relatively rare, especially among Trump supporters. (See Materials and Methods and the Supplementary Information for more information.)

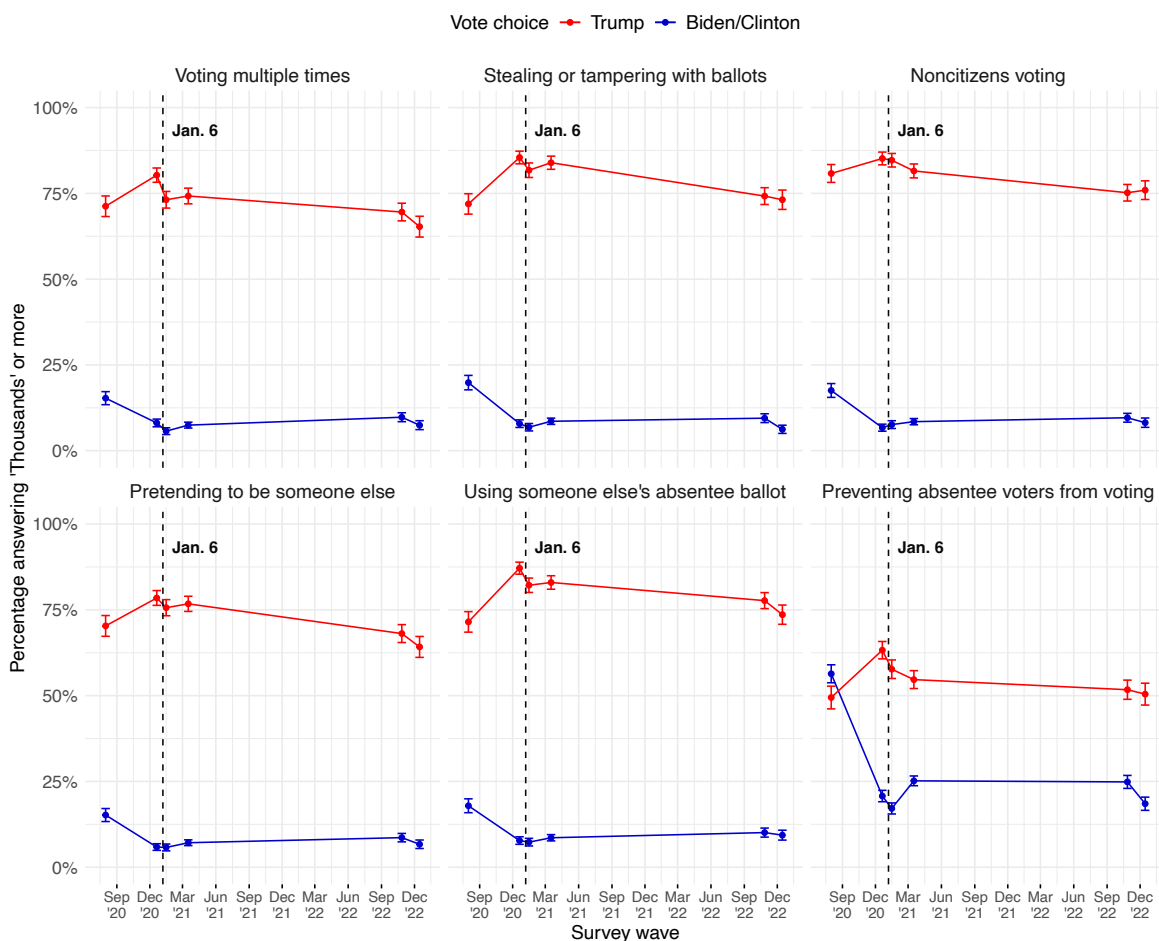
The persistence of election fraud misperceptions

We demonstrate the persistence of belief in election fraud over time, drawing on data from six survey waves conducted between August 2020 and January 2023. Figure 1 tracks changes over time in the belief that at least thousands of cases occurred across six distinct types of election fraud and Trump and Clinton/Biden supporters (categorized by their reported vote choice in the 2016 and 2020 elections). For all measures except one, Trump supporters were more likely than were Clinton supporters to believe in widespread fraud before the 2020 election. After the election, Biden supporters became less likely to say that fraud was widespread and Trump supporters became more likely to do so. This finding is consistent with prior research showing that election outcomes affect beliefs about election legitimacy (Anderson and Guillory 1997; Anderson and LoTempio 2002; Clarke and Acock 1989; Ginsberg and Weissberg 1978; Maldonado and Seligson 2014; Sinclair et al. 2018) as well as recent research showing how elite claims of fraud can decrease election confidence (Berlinski et al. 2023; Clayton et al. 2021; Justwan and Williamson 2022; Lyons and Workman 2022). By 2022, the share of Trump supporters endorsing claims of widespread fraud had decreased slightly but remained puzzlingly high given the widespread availability of corrective information. Both overall rates of belief and their changes over time follow a similar pattern when we instead disaggregate by Trump approval or partisan identification — see Figures S3 and S4 in the Supplementary Information (SI).

Consistent with Graham and Yair (2024), we also find high levels of within-respondent stability in perceived fraud prevalence. For five of the six measures, more than 75% of respondents remained in the same binary category (thousands of cases or more versus hundreds of cases or fewer) between the

pre- and post-election waves in 2020. Stability increases in later waves, with the percentage remaining in the same category reaching 85% or more. However, stability is consistently lower for the perceived prevalence of officials preventing absentee voters from voting (71–76% in post-2020 election waves) than for other types of fraud. Additional visualizations of over-time stability in election fraud beliefs are provided in Figures S5 and S6.

Figure 1: Perceived prevalence of election fraud over time by candidate support



Percentage of respondents indicating that there are “Thousands” of cases or more of each type of election fraud. Categories included “Less than ten” (post-2020 election waves only), “Less than a hundred,” “Hundreds,” “Thousands,” “Tens of thousands,” “Hundreds of thousands,” and “A million or more”. Fraud prevalence measured for U.S. elections in general (Aug. 2020) and for the 2020 U.S. presidential election specifically (subsequent waves). Group means computed using each wave’s post-stratification weights. We measure candidate support using 2016 vote choice as reported in YouGov surveys fielded before the 2020 election and 2020 vote choice as measured in subsequent waves.

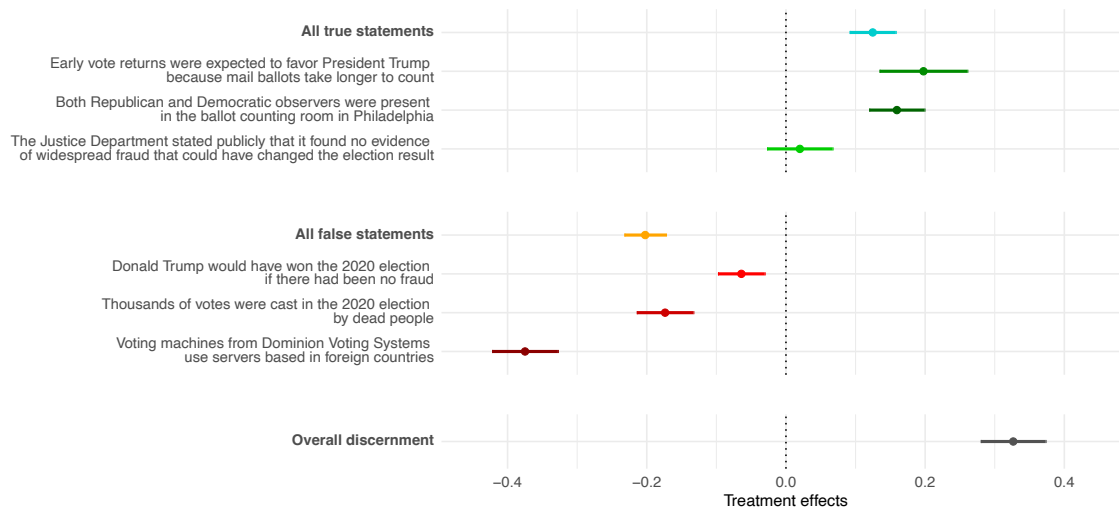
The effectiveness of fact-checking election fraud claims

We next investigate the potential role of motivated resistance to corrective information in explaining the persistence of fraud misperceptions using a preregistered survey experiment conducted in January 2021 with a nationally representative YouGov survey sample. Participants were assigned with equal probability to either a fact-check condition or a control condition. Those in the fact-check condition were exposed to a shortened version of a fact-check article from the Associated Press debunking five unfounded claims Trump made about fraud (Yen et al. 2020, see SI Section [S2.2.1](#) for the treatment article). Participants were then asked to evaluate the truthfulness of six targeted statements about election fraud (three false and three true) on four-point scales. (For additional details, see the Materials and Methods section.)

Based on previous research (e.g., Carnahan et al. 2021; Coppock et al. 2023; Nyhan 2021; Walter et al. 2020; Walter and Murphy 2018; Wood and Porter 2019), we expected that fact-checks would reduce belief in false claims, increase belief in accurate information, and improve discernment between them. As shown in [Figure 2](#), fact-check exposure significantly decreases mean beliefs in false claims about election fraud ($\beta = -0.20$, 95% CI=-0.23, -0.17; $d = -0.18$) and increases both mean beliefs in true claims ($\beta = 0.12$, 95% CI=0.09, 0.16; $d = 0.16$) and overall truth discernment ($\beta = 0.33$, 95% CI=0.28, 0.37; $d = 0.19$), confirming our expectations. Effects are consistent using a simple differences of means and among the subsample of participants whose online behavior we analyze below (see SI [Figure S7](#) and [Section S3](#)). When we disaggregate these results by item, we find that exposure to fact-checking has the expected effect on the perceived accuracy of each targeted statement, although the effect fails to reach statistical significance for one of the six statements.

We find no measurable evidence of heterogeneous treatment effects. SI [Figure S8](#) reports average treatment effects among preregistered subgroups of interest. Consistent with previous studies (Carey et al. 2022), we find that fact-check exposure increases truth discernment across levels of partisanship, Trump approval, feelings towards Trump, trust in authoritative sources of information, conspiratorial thinking, and truth discernment in the previous survey wave. Exploratory analyses also suggest that exposure to fact-checking of false election fraud claims increases truth discernment irrespective of beliefs that Biden is the rightful winner of the 2020 election or prior exposure to fraud-challenging

Figure 2: Treatment effects of exposure to fact-check article



Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of targeted statements. OLS regression coefficients with 95% confidence intervals; estimated using covariates measured in a previous survey wave selected by lasso. Outcomes measured on a four-point scale ranging from “1-not at all accurate” to “4-very accurate”. Overall discernment represents the mean difference in perceived accuracy of true and false claims. Complete regression tables including lasso-selected covariates are provided in SI Tables S4 and S5.

content and fact-checks (Figure S34).

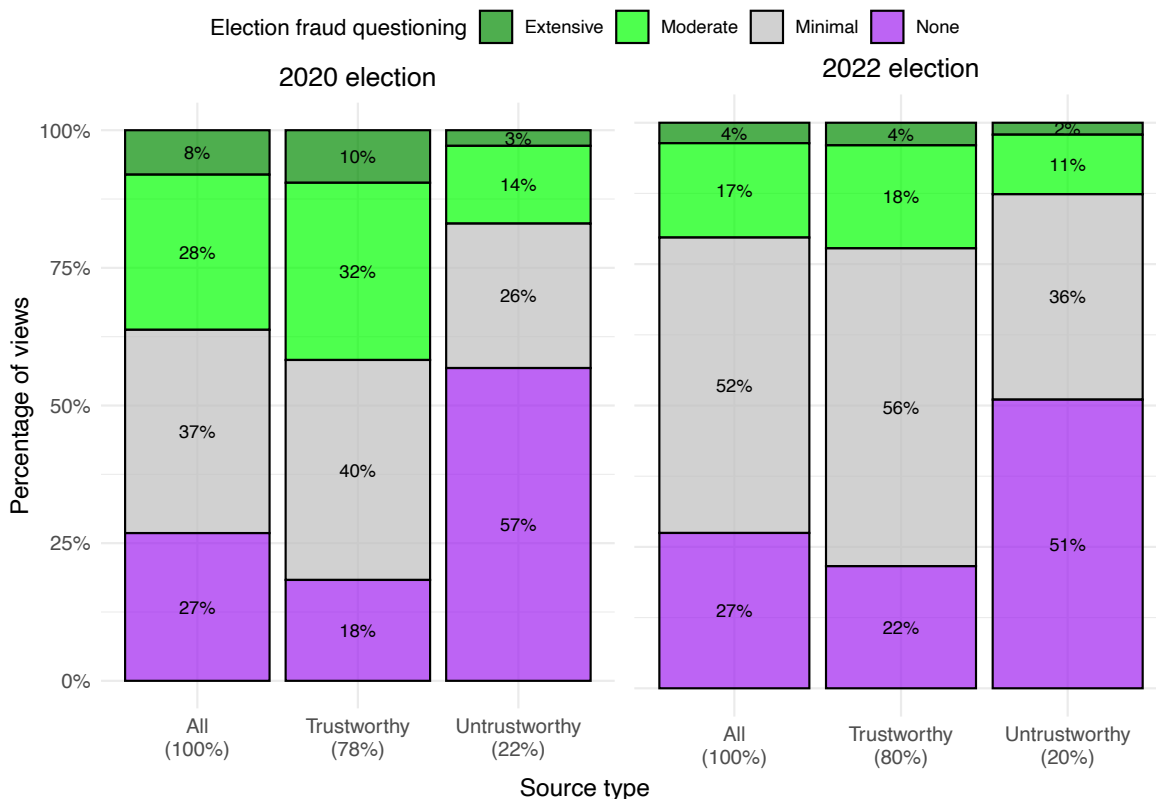
Total exposure to content challenging fraud claims: 2020–2022

The findings above demonstrate that misperceptions about election fraud persist over time and provide no evidence that motivated resistance to corrective information is the cause of that persistence. Instead, we find that exposure to corrective information can reduce fraud misperceptions. We thus instead consider the prevalence of exposure to different types of fraud-related information online, using a LLM to analyze the online content encountered by participants at the article/URL level. The model classifies each article seen by a participant for whether it refers to election fraud and, if so, whether it meaningfully challenges claims that election fraud is widespread or could change the outcome of one or more elections. Table 2 in Materials and Methods summarizes coding rules and examples; SI Section S4.7 provides additional coding details.

We first provide aggregate measures of the fraud-related content that people saw around the 2020

and 2022 elections. Figure 3 shows the share of views of fraud-related articles by levels of questioning and source trustworthiness. Contrary to fears of pervasive exposure to content that uncritically accepted false claims of widespread fraud, we find that most of what people saw online about fraud questioned these claims at least minimally. During the 2020 election period (the left panel of Figure 3), 73% of the fraud-related content people saw online questioned false claims at least minimally. However, only about one-third of these views (36%) consisted of content that challenged false claims moderately or extensively.

Figure 3: Percentage of views of election fraud by coverage type, source type, and year



Data disaggregated by source trustworthiness based on ratings from NewsGuard and Lin et al. (2023). Weighted percentages based on post-stratification weights.

During the 2022 election period (the right panel of Figure 3), the share of viewed fraud-related content that questioned false claims at least minimally was similar to 2020 levels (73%). However, only 20% of viewed content moderately or extensively questioned such claims (which we refer to as

“challenging”).

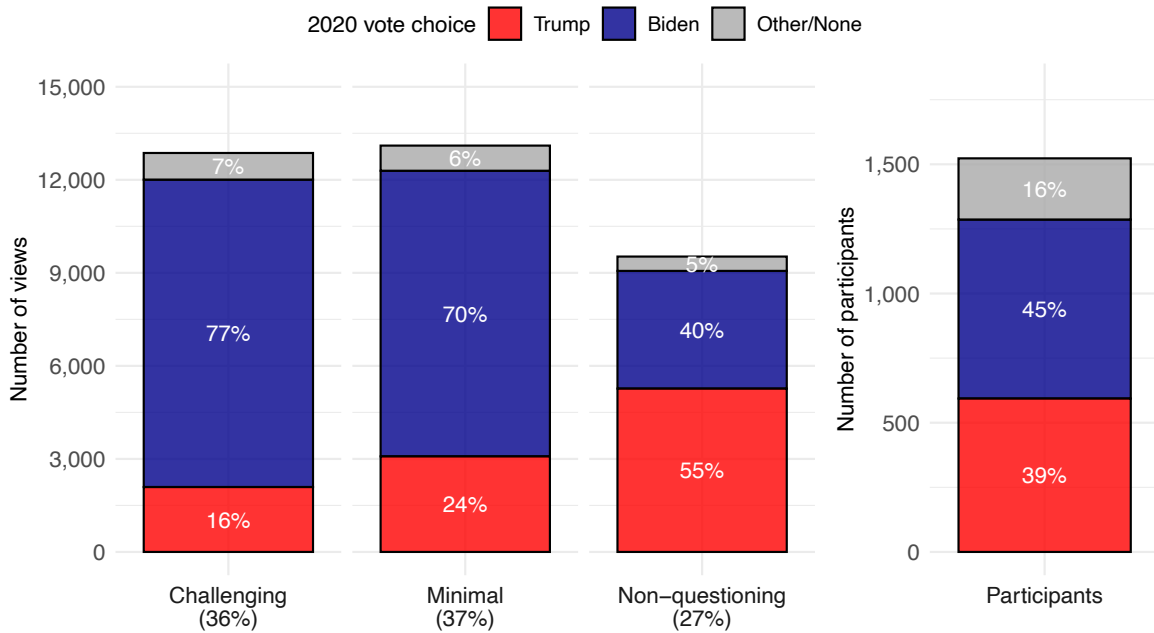
Notably, these patterns vary systematically when we distinguish between trustworthy and untrustworthy sources, which we classify using NewsGuard ratings or corresponding Lin et al. (2023) ratings (see Materials and Methods). Untrustworthy sources were far less likely to question or challenge fraud claims. For example, among untrustworthy domains in 2020 (which were responsible for 22% of total fraud-related views in that election period), only 43% of viewed content questioned false claims even minimally. In contrast, 82% of fraud-related content viewed from trustworthy domains (78% of total fraud-related views) questioned false claims at least minimally. (Results are similar in the 2022 election period.)

To investigate how these patterns vary based on people’s political views, we disaggregate aggregate exposure to different types of fraud content vary by candidate preference in the 2020 election in Figure 4. The total height of the left three bars represents the total number of views of that type of content; the percentages indicate the share of those views coming from supporters of Biden (45% of our sample), Trump (39%), or neither (16%). Moving from left to right, the first column shows that only 16% of views of content that challenges fraud claims with facts and evidence in 2020 (which made up 36% of all fraud-related exposure during that period) came from Trump supporters compared with 77% for Biden supporters, echoing findings in prior studies that fact-checks often fail to reach vulnerable audiences (Guess et al. 2020; Wack et al. 2024). Similarly, only 24% of views of minimally questioning content (which made up 37% of fraud-related content exposure in the period) come from Trump supporters versus 70% for Biden supporters. By contrast, the majority of views of non-questioning content (accounting for 27% of all fraud-related exposure) came from Trump supporters (55%) compared with 40% for Biden supporters. (The fourth column shows the relative prevalence of each group in the sample for reference.) The corresponding results for the 2022 election period, when exposure levels were lower, are reported in SI Figure S12.

Individual-level exposure to fraud-related content

We now move from examining aggregate exposure to instead consider the distribution of exposure to different types of fraud-related content at the individual level. We specifically examine how often

Figure 4: Differences in fraud-related information exposure by candidate support (2020)



Weighted percentages based on post-stratification weights. Analysis restricted to the 2020 election period. The corresponding results for the 2022 election period are reported in SI Figure S12.


participants encountered content that challenged narratives of widespread fraud and how those patterns varied by 2020 vote choice.

Binary measures of individual exposure to fraud-related content

Figure 5 presents binary measures of exposure to any fraud content and to different types of it (challenging content, minimally questioning content, and non-questioning content) for supporters of both Biden and Trump. We also disaggregate our binary measure of exposure to non-questioning fraud content between trustworthy and untrustworthy sources. (The same statistics are presented conditional on exposure to election fraud content in SI Figure S16; SI Tables S21–S23 summarize how participant demographics vary for binary measures of exposure to different types of fraud-related content.)

The first pattern we observe in Figure 5 is inattention. A substantial number of both Trump and Biden supporters encountered no fraud-related content at all in their web browsing in either election cycle despite being exposed to a substantial amount of news content (see SI Figure S15 and Figure 8

Figure 5: Exposure to specific types of election fraud content



Percentage of respondents who were exposed to each type of content at least once by candidate support in the 2020 and 2022 elections; all values calculated based on post-stratification weights.

below). As shown in the left column of the figure, only about half of Americans encountered any fraud-related online articles during the 2020 election period. We also find clear differences by vote choice in the frequency of exposure to election fraud content. Despite having similar overall levels of news consumption (see Figure 8 and SI Figure S13), Trump supporters (48%) were less likely to encounter fraud-related articles than were Biden supporters (55%). During the 2022 midterm election period (when political interest was lower and fraud claims were less prevalent), fraud-related content exposure rates dropped to 36% among both Biden and Trump supporters.

We next consider whether Biden and Trump supporters encountered different types of fraud-related content. Moving from left to right in Figure 5, we see that Biden supporters were more likely than were Trump supporters to have encountered every type of fraud-related content. Notably, Biden supporters were more likely to have seen content that challenged fraud claims (45% vs. 36% in the 2020 study period, 24% vs. 19% in 2022). However, though levels of overall exposure to non-questioning content were higher among Biden supporters (43% versus 38%) Trump supporters were more than twice as likely as Biden supporters to see non-questioning content from untrustworthy sources in 2020 (26% vs.

11%). Levels of exposure were lower for every group and type of content in 2022, but the differences are similar. (Because participants rarely encountered formal fact-checks about election fraud, we do not discuss them further in this section — see SI Figure S15.)


Relative measures of individual exposure to fraud-related content

We next examine individual-level variation in *how much* fraud-related content of different types participants encountered in their online browsing and how those exposure levels vary by candidate support. In these analyses, we move from binary measures of exposure to instead focus on relative levels of exposure, which are likely to be important given that the net balance of information received can have an important effect on people’s opinions (e.g., Zaller 1992). We specifically focus on relative exposure to the two types of content where our theoretical expectations about message effects are clearest: content that challenges fraud claims with facts and evidence versus content that fails to question fraud claims at all. (Theoretical expectations about the effects of exposure to content that only minimally questions of fraud claims are less clear, but see SI Figures S17–S19, which replicate the analyses below but distinguish instead between views of content with any questioning and content with none.)

Our theoretical expectation is that it matters not just whether people were exposed to challenging content but how much they encountered relative to content that fails to question fraud claims. Figure 6 therefore groups participants into one of four categories: those who were not exposed to fraud-related content, those who saw more content that failed to question fraud claims than challenging content, those who saw equal quantities of the two, and those who saw more content challenging fraud claims. In total, 37% of Biden supporters consumed more online stories that challenged fraud claims than stories that failed to question them compared to just 15% of Trump supporters during the 2020 study period. Additionally, 24% of Trump supporters in the 2020 period saw more stories that failed to question fraud claims than those that challenged them compared to just 9% of Biden supporters. Findings were similar in 2022, but with lower levels of exposure to any fraud-related content, especially for challenging content (16% of Biden supporters saw more challenging than non-questioning content compared to 9% of Trump supporters).

We next consider exactly how much content of each type people saw rather than just classifying

Figure 6: Individual relative exposure to different types of election fraud content

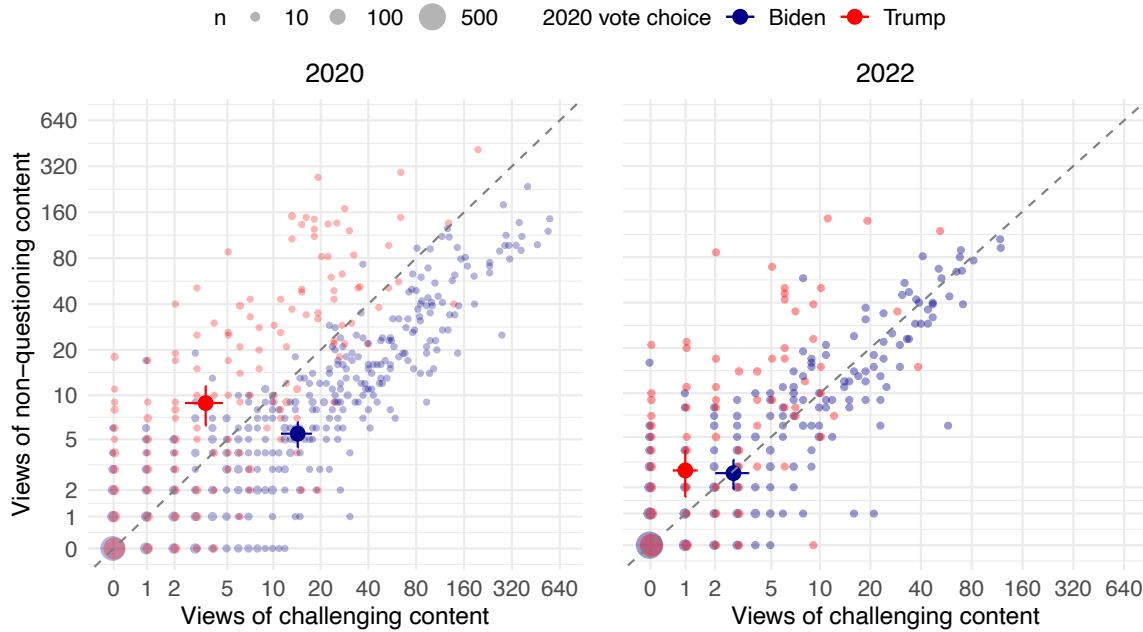


Weighted percentages based on post-stratification weights. Respondents who saw content minimally questioning fraud claims but no fraud-challenging or non-questioning content are included in the “Equally exposed” category.

people based on their relative balance of content consumption. Figure 7 plots the full distributions of participant-level views of fraud-related articles that challenged fraud claims (horizontal axis) versus views of articles that did not question those claims (vertical axis) among Biden and Trump supporters (see Figure S18 for a similar visualization of how patterns of exposure vary by beliefs about the legitimacy of the 2020 election). The axes use a log scale to better illustrate the distributions as most participants are clustered at low numbers of articles but a small subset of highly engaged participants viewed dozens or even hundreds of fraud-related articles.

On average, Biden supporters were exposed to more election fraud content overall than were Trump supporters. Their information diets tilted toward content challenging fraud claims in 2020 (a mean of 14.3 views of challenging content vs. 5.5 views of non-questioning content). In 2022, however, exposure to both types of content among Biden supporters had dropped sharply and was more balanced (averages of 2.8 challenging views and 2.6 non-questioning views). Trump supporters, in contrast, saw more non-questioning content than content challenging fraud claims on average in both study periods

Figure 7: Views of content challenging and not questioning election fraud claims by election year and candidate support



Highlighted markers represent mean values, with 95% confidence intervals, for Trump and Biden supporters. Axes use a log scale (with zero views transformed to an infinitesimally small number). The dashed line represents equal exposure to both types of content.

(8.9 vs. 3.5 in 2020; 2.8 vs. 1.0 in 2022). However, we note that these differences in average exposure are inflated by skew in the exposure distributions (further details are provided below; see also SI Figures S20–S21).

Roles of topic, outlet, and article choice in exposure differences to challenging content

Why do we observe differences in exposure to information challenging claims of widespread fraud between Biden and Trump supporters? One common explanation emphasizes differences in exposure to untrustworthy websites by candidate support (Guess et al. 2020; Moore et al. 2023). Source-level partisan variation in news consumption is well-documented (Eady et al. 2019; Guess 2021; Peterson and Iyengar 2021; Robertson et al. 2023). However, recent evidence suggests that within-outlet selective exposure to congenial articles also contributes to differences in information diets (Braghieri et al. 2024; Green et al. 2025).

We build on recent scholarship in this area by using counterfactual permutation tests to measure the relative roles of topic attention (fraud-related content versus not), outlet choice, and article choice within outlets in producing differences between Biden and Trump supporters in exposure to content challenging claims of widespread fraud.


Before reporting these tests, we first provide descriptive evidence on each potential mechanism: attention to fraud content, outlet selection, and within-outlet article choice. As described above, we observe differences in fraud-related exposure by candidate preference in 2020. One possible explanation is differences in total news exposure. However, as Figure 8 illustrates, Biden and Trump supporters had comparable levels of overall news exposure, both on average (650 vs. 668) and at the median (158 vs. 165). Despite these similar levels of news consumption, Biden supporters encountered substantially more fraud-related content, suggesting topic-level selective avoidance among Trump supporters. On average, Biden supporters saw nearly twice as many fraud-related articles as did Trump supporters (33 vs. 18), a gap that widens at higher levels of exposure (2 vs. 0 at the median, 16 vs. 7 at the 75th percentile, and 91 vs. 38 at the 90th percentile; see Table S39 for additional exposure statistics). These results suggest that Trump supporters generally paid less attention to the topic of election fraud than did Biden supporters in the period around the 2020 election.

We also observe aggregate-level differences in the sources and types of fraud-related content seen by Biden and Trump supporters. As Figure 9 shows, Trump supporters not only saw less fraud-related content than did Biden supporters but were far more likely to consume fraud-related content from untrustworthy sources when they did encounter it at all.

In total, 91% of the 33 election fraud views we observe from Biden supporters on average were from trustworthy sources. Of these, 84% of the views were to content that challenged (44%) or minimally questioned (40%) fraud claims. By contrast, Trump supporters only had 18 election fraud views on average and just 47% were to trustworthy sources, where just 70% of views were of challenging (32%) or minimally questioning (38%) articles. Among the 53% of their views of fraud content that came from untrustworthy sources, 69% of articles did not question fraud claims.

Figure 10 illustrates these findings at the outlet level using the 50 most popular sources of fraud-related content by candidate support in the 2020 data, which represent 88% of all views for Biden

Figure 8: Views of news and election fraud content by candidate support (2020)



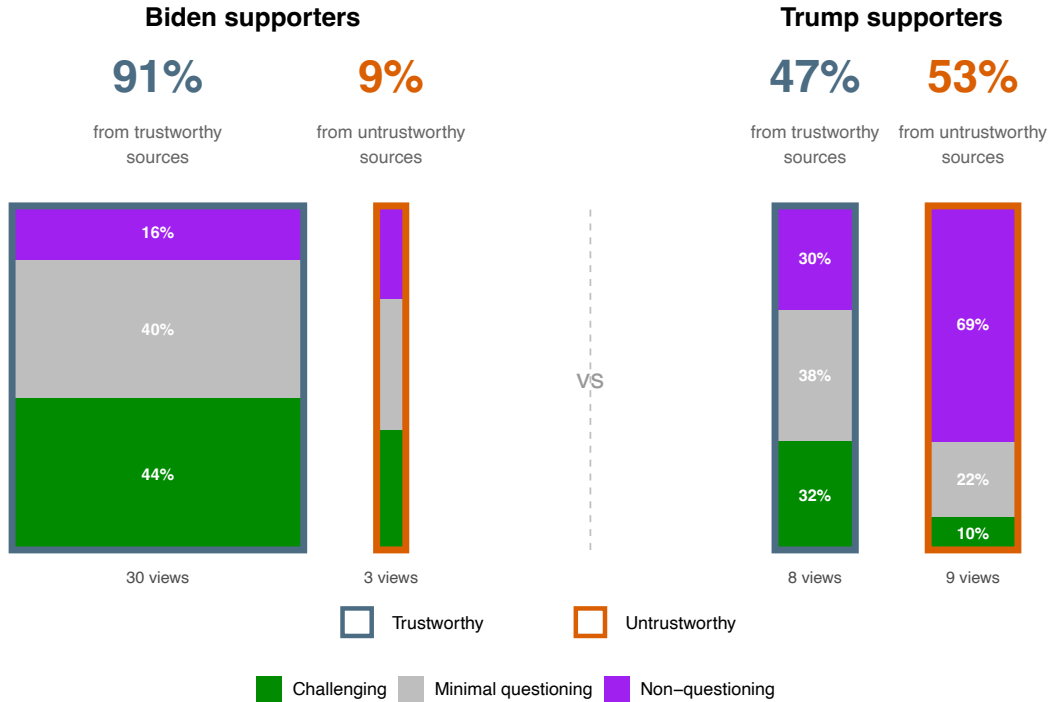
Weighted means and medians calculated using post-stratification weights, with 95% intervals around the means. Analysis restricted to 2020 election data (see SI Figure S22 for analysis based on 2022 election data).

supporters and 73% for Trump supporters. As the figure indicates, outlet selection plays an important role in the differences we observe in exposure to challenging fraud content — the sources seen by Biden supporters were more likely to feature content that challenged fraud claims. (SI Figure S24 provides equivalent data for 2022.)

However, when we examine aggregate-level views of content *within* source for the top 50 outlets in the 2020 study period, we find relatively similar patterns of exposure between Biden and Trump supporters. As Figure 11 shows, Biden and Trump supporters have similar shares of exposure to challenging content within source as a proportion of their fraud-related views (Pearson’s correlation weighted by total fraud views per outlet is 0.59; $p < 0.001$). These findings suggest that the mix of fraud-related content seen within outlet does not vary widely by candidate preference.

We therefore conduct independent permutation tests to better understand the respective contributions of topic selection, outlet selection, and within-topic article selection to the gap in fraud-challenging views between Biden (mean of 14.3 views) and Trump supporters (mean of 3.5 views)

Figure 9: Source trustworthiness of election fraud views by candidate support (2020)

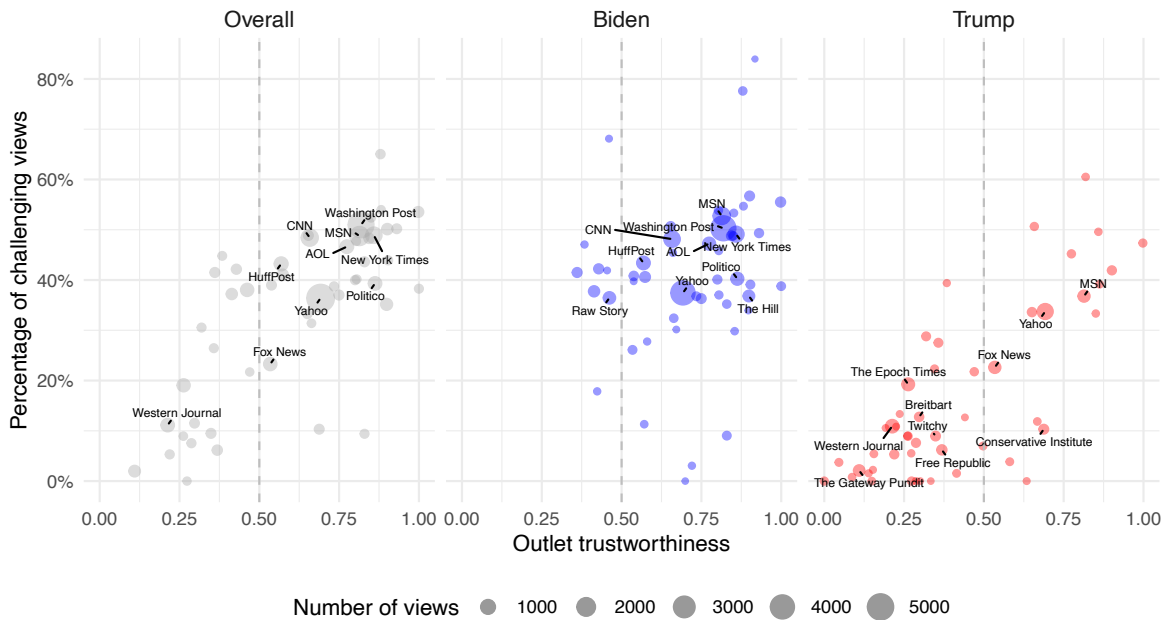


Weighted means calculated using post-stratification weights. Analysis restricted to 2020 election data (see SI Figure S23 for analysis based on 2022 election data).

during the 2020 study period. These tests ask how much of the observed gap in fraud-challenging content exposure would remain if we equalized one factor holding all else constant. We use permutation tests to compare the magnitude of the effects of each factor in this way rather than conducting a regression-based decomposition because many outlets are consumed by only one candidate’s supporters, which creates a separation issue in regression analysis. In addition, our three mechanisms (topical avoidance, within-outlet selection, and overall outlet choice) are defined on different subsets of the data (news content versus election fraud content), making it difficult to decompose their contributions within a single regression framework. Finally, conducting these permutation tests independently allows us to avoid the complexities of order effects and interactions between factors (Fortin et al. 2011).

In the first test, we randomly permute content types (fraud-related versus non-fraud-related) within each news outlet holding participants’ overall news consumption and outlet composition constant.


Figure 10: Share of challenging views by outlet trustworthiness and candidate support (2020)



Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of fraud-related views within each group. Labels identify the top 10 outlets per group. Point size indicates the total number of fraud-related views by outlet. The x-axis measures outlet trustworthiness based on ratings from NewsGuard and Lin et al. (2023). The dashed vertical line represents the threshold between sources classified as trustworthy and those classified as untrustworthy. Analysis restricted to 2020 election data (see SI Figure S24 for analysis based on 2022 election data and Figures S26–S28 for same analysis based on the percentage of views including any questioning of fraud claims).

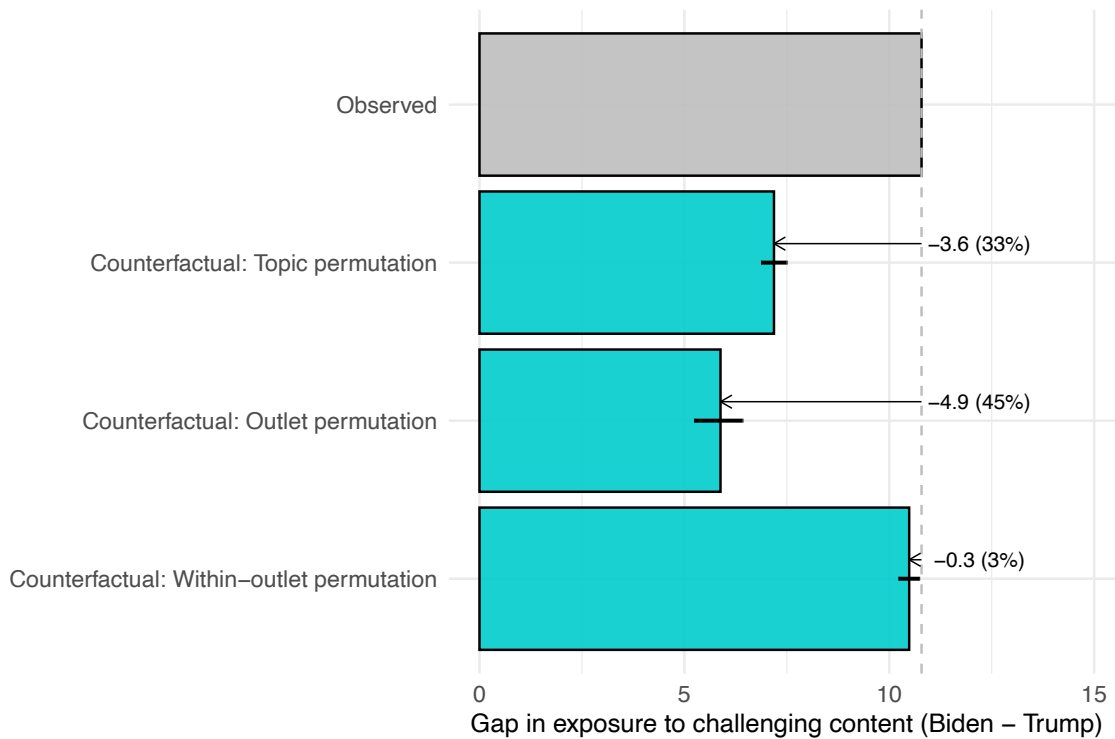
This test isolates the extent to which the gap reflects partisan differences in the likelihood of selecting fraud-related content versus other types of content. In the second test, we permute which outlets users encountered for fraud-related content while maintaining each outlet’s overall share of fraud-related views and balance of challenging versus non-challenging articles. Specifically, each participant’s outlet blocks of fraud-related views (e.g., 10 from The Washington Post, 5 from CNN) were randomly replaced by equivalently-sized blocks from other outlets (e.g., 10 from Fox News, 5 from The New York Times). This test measures how much of the gap comes from differences in outlet choices across partisan groups. In the third test, we randomly reassign challenging and non-challenging views within outlets among fraud-related articles to assess how much of the gap stems from selective exposure to (or avoidance of) congenial versus uncongenial fraud-related coverage. We ran 5,000 permutations in each case to generate a null distribution for the remaining gap after each test and estimated the

Figure 11: Comparing challenging content exposure rates within outlet by candidate support (2020)



Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of election fraud-related views within each group. Labels identify the top 10 outlets per group (red for Trump supporters, blue for Biden supporters, black for both). Points below the 45-degree line indicate that, for a given outlet, Biden supporters encountered a larger share of content challenging fraud claims than Trump supporters. Points above the line indicate the opposite. Point size indicates the total number of election fraud content views by outlet. The blue-red gradient of the points indicates the proportion of fraud-related views from that outlet coming from Trump supporters. Analysis restricted to 2020 election data (see SI Figure S25 for analysis based on 2022 election data and Figures S27–S29 for same analysis based on the percentage of views including any questioning of fraud claims).

Figure 12: Permutation tests of the Biden/Trump supporter gap in challenging content exposure (2020)



Weighted estimates based on post-stratification weights. The figure reports independent permutation tests in which we randomly permuted content types (fraud-related versus non-fraud-related) within outlet holding overall news and outlet composition constant; outlet selection holding overall outlet content fixed; and challenging vs. non-challenging views within outlets holding outlet selection fixed (see Materials and Methods for further details). Analysis restricted to 2020 election data (see SI Figure S30 for equivalent results for the 2022 study period).

effects of independently varying each factor by comparing the permuted and observed gaps. This approach allows us to evaluate their respective contribution to the observed difference in challenging content exposure between Biden and Trump supporters. (Further details are provided in Materials and Methods.)

The results reported in Figure 12 indicate that the pattern of selective avoidance we observe is primarily driven by outlet and topic selection, which account for most of the gap in challenging fraud content exposure. Specifically, the Biden-Trump supporter gap of 10.8 views of challenging fraud content is reduced by 4.9 and 3.6 views, respectively, when outlet and topic selection are randomly permuted. Put another way, 45% of the challenging fraud content exposure gap would be eliminated if Trump and Biden supporters consumed fraud news from the same distribution of outlets (eliminating

differences in untrustworthy status and the frequency with which outlets challenged fraud claims). Likewise, 33% of the gap would be eliminated if Trump and Biden supporters consumed fraud news equally often. By contrast, permuting exposure to congenial versus uncongenial fraud coverage within outlets only reduces the gap by 0.3 points on average (just 3% of the total difference), suggesting that within-outlet selection of congenial articles plays a negligible role in differences in attention to challenging fraud content. (Equivalent results for the 2022 study period are reported in SI Figure S30.)

Discussion

The persistence of misperceptions about widespread election fraud raises an important puzzle. Why are people not updating their beliefs about this issue when so much corrective information is available? One possibility is directionally motivated resistance to accurate information. However, we find in a survey experiment that exposure to fact-checks of fraud claims made by Trump and his allies increased belief accuracy even among Trump supporters, suggesting that directional motivations do not prevent the processing of corrective information.

We therefore instead consider the role of information exposure. One key factor is inattention. Despite extensive media coverage and rebuttals of election fraud allegations, roughly half of Americans never encountered any news content mentioning election fraud in their web browsing. Even fewer came across coverage that challenged fraud claims with facts or evidence. These patterns were especially pronounced among Trump supporters, who we show engage in a pattern of what we call selective avoidance. Rather than seeking out high levels of pro-attitudinal content about fraud, Trump supporters consumed much less fraud content than did Biden supporters. In particular, they were especially unlikely to encounter content that challenges fraud claims with facts and evidence — the type of detailed debunking that previous research and our survey experiment show can change people’s minds (Chan and Albarracin 2023; Chan et al. 2017; Ecker et al. 2020; Swire et al. 2017). Finally, we demonstrate that Trump supporters’ lower exposure to fraud-challenging content is driven by differences in topic attention and outlet choice (including much greater exposure to fraud content from untrustworthy sources) rather than selective exposure to congenial articles within outlets.

These findings represent a profound challenge for democracy. The claim of widespread elec-

tion fraud comes from the most high-profile figure in American politics (Donald Trump), concerns an incredibly salient issue (whether the 2020 presidential election was rightfully decided), and was repeatedly shown to be false by election experts, courts, and the media. This evidence was widely covered in the press and easily available online. Nonetheless, people's exposure to coverage challenging fraud claims was often minimal and the people who were most vulnerable to false claims (Trump supporters) were differentially less likely to encounter it. This pattern of differential exposure to uncongenial topical content is more subtle than most accounts of echo chambers and filter bubbles would suggest (e.g., Pariser 2011; Sunstein 2001). These theories typically consider overall tendencies toward congenial information consumption rather than avoidance and do not distinguish between issues or topics. Future research should seek to determine the conditions under which news media coverage and related efforts like fact-checking overcome these obstacles to reach broader audiences, including people who might be especially likely to be inattentive, and have lasting effects on their beliefs and attitudes (Bowles et al. 2025; Nyhan et al. 2022).

Of course, this study has limitations. First, our study relies on a combination of descriptive evidence and a survey experiment (which validates evidence from numerous prior studies). Researchers should develop externally valid ways to directly estimate the causal effects of behavioral exposure to topic-specific information in real-world settings, including content that challenges misinformation about election fraud and other topics (an extremely difficult task, to be sure). Second, we cannot directly measure the causes of selective avoidance of content or the precise mechanisms by which it takes place. For instance, some people are likely to actively engage in directionally motivated news avoidance behaviors (Robertson 2025; Skovsgaard and Andersen 2020; Toff and Kalogeropoulos 2020; Toff et al. 2023; Villi et al. 2022) while others may have weaker motivations but just rely on passive defaults (e.g., existing media diets). Our design does not allow us to differentiate between these motivations or the mechanisms by which avoidance behavior takes place. Future research should seek to measure the causes and mechanisms of selective avoidance more effectively in behavioral data. Third, we are only able to measure the content that YouGov Pulse participants viewed in desktop and mobile web browsers, but people of course encounter information from other sources such as mobile apps, social media platforms, television, radio, etc. In the future, scholars should measure exposure to specific

information such as fraud claims across as many types of media as possible with as representative of a sample as possible (Allen et al. 2020; Aridor et al. 2025; Broockman and Kalla 2025). Finally, though we followed current best practices, it would be desirable to further improve the accuracy and reliability of LLM text coding. Recent research highlights how coding results can change with more recent or more powerful LLMs and how different LLMs vary in their correspondence with human coding (Yang et al. 2025). Our process of developing a codebook that achieves high reliability among human coders and also between the consensus human judgment and the LLM may be a useful framework for future research.

Despite these limitations, this paper makes a number of important contributions. Most notably, we help reconcile the observed persistence of misperceptions with the apparent effectiveness of strategies intended to address them. By combining a fact-check survey experiment with a content-specific examination of data on individual-level media diets, we challenge the notion that persistence is driven by resistance to corrective information or heavy or disproportionate consumption of pro-attitudinal content. Instead, we find that vulnerable populations are differentially *less* likely to see content that challenges claims of widespread fraud with facts and evidence, in part due to their reliance on untrustworthy sources.

Finally, we move beyond prior source-level analyses of exposure to untrustworthy online outlets and instead measure whether and how misinformation is presented online across source types. To accomplish this, we demonstrate that an LLM can accurately estimate online exposure to specific topics or claims using digital behavior data. This approach and the findings it generates have great potential to deepen our understanding of why people believe false information for so long after it has been widely debunked — a critical issue for democracy in the U.S. and countries around the world.

Materials and Methods

This study analyzes panel surveys conducted by YouGov on national samples of Americans during the COVID-19 pandemic (4,399 participants), the 2020 United States elections (4,312 participants), and the 2022 United States elections (3,772 participants) as well as passively collected online browsing data from a subset of these survey participants. This subset consisted of members of YouGov’s Pulse

panel who, after providing informed consent, allowed their desktop and mobile device web browsing activity to be tracked. This study was approved by the Institutional Review Boards of Dartmouth College, the University of Exeter, and the University of Notre Dame. The datasets used in this study are summarized in Table 1; see Table S1 for demographics.

Table 1: Summary of data

	COVID-19	2020 election	2022 election
Survey sample	n=4,399	n=4,312	n=3,772
Democrat	2,279	2,227	1,969
Republican	1,442	1,381	1,188
Wave 1	May 20–June 3, 2020 n=4,399	Dec. 17, 2020– Jan. 5, 2021 n=4,312	Oct. 18–Nov. 7, 2022 n=3,772
Wave 2	June 25–July 12, 2020 n=3,680	Jan. 13–19, 2021 n=3,847	Dec. 7–20, 2022 n=2,896
Wave 3	July 28–Aug. 19, 2020 n=2,983		Jan. 21–30, 2023 n=2,100
Wave 4	March 9–23, 2021 n=5,575 (2,464 recontacts)		
Behavioral data		n=1,596	n=1,518
Democrat		833	809
Republican		535	456
		Sept. 13, 2020– Jan. 29, 2021	Aug. 31, 2022– Jan. 31, 2023

We draw participants from three studies in total:

- **A prior study of attitudes toward COVID-19** consisting of a YouGov sample built from three sampling frames: 1,096 drawn from their general population panel, 2,238 drawn from their Pulse panel (participants who consented to share their web data), and 1,065 drawn from areas with high COVID-19 prevalence. Post-stratification weights were constructed based on 2016 presidential vote and a four-way stratification of gender, age, race, and education. Election fraud items were first asked in Wave 3 of this study, but we use data from Wave 1 in one analysis reported in the SI.

- **A study conducted around the 2020 election** consisting of a YouGov sample built from four sampling frames: whether participants are recontacts from the COVID study or not and are Pulse participants or not (1,026–1,168 participants per frame). Post-stratification weights were constructed based on 2016 and 2020 presidential vote and a four-way stratification of gender, age, race, and education.
- **A study conducted around the 2022 election** consisting of a YouGov sample drawn from their general population panel (2,643 of the 3,772 participants are recontacts from our 2020 election study). Post-stratification weights were constructed based on 2016 and 2020 presidential vote and a four-way stratification of gender, age, race, and education.

Survey measures

We use survey responses from our participants to measure presidential candidate support and election fraud beliefs. For analyses after the 2020 U.S. election, we measure Trump support using self-reported 2020 presidential vote choice. For analyses that precede the 2020 election (e.g., August 2020 data points in Figure 1), we use self-reported 2016 vote choice. We measure fraud beliefs in two different ways. The first measure is the perceived prevalence of six types of fraud (e.g., voting more than once in an election) on a scale ranging from “Less than a hundred” to “A million or more.” Our main descriptive outcome is the percentage of respondents indicating that there are “Thousands” of cases or more of each type of fraud. The second measure assesses perceptions of whether Joe Biden was the rightful winner of the 2020 election. Respondents answering “Definitely” or “Probably” were coded as 1 (Yes), while those who answered “Definitely not” or “Probably not” were coded as 0 (No). Demographics for survey participants and a list of the dates the surveys were administered are provided in Table S1. Questionnaires are available online for each survey at <https://osf.io/h89wa/>.

Fact-checking survey experiment

In the survey experiment we conducted, participants in the treatment condition were randomly assigned to read a shortened version of an Associated Press article fact-checking five claims about supposed election fraud during the 2020 election (e.g., that the election was stolen, the vote counting process was

corrupt, and votes were being counted in foreign countries). After reading the fact-check, participants evaluated the accuracy of three true and three false election fraud claims from the article on a four-point scale from “not at all accurate” to “very accurate”. (A full questionnaire including the treatment article and question wording is available at <https://osf.io/h89wa/>.)

We estimate the effects of exposure to the fact-check article on the perceived accuracy of true and false fraud claims and truth discernment (the difference between the two; i.e., “additive” truth discernment) using Ordinary Least Squares (OLS) regression with robust (HC2) standard errors. For each model, pre-treatment control variables were selected via lasso to improve precision (Bloniarz et al. 2016) from a preregistered list: pre-treatment outcome measures (Clifford et al. 2021), partisan identification, ideology, political interest, political knowledge, trust in authoritative sources of information, education, and race (pre-treatment outcome measures and partisan identification were selected most often). All experimental results are consistent when restricting the analysis to participants for whom online behavior data is available (see SI Section S3).

We test for heterogeneous treatment effects among groups that may be differentially vulnerable to misinformation (a preregistered research question) using Bayesian Causal Forest models. More details about how each variable is measured are provided in SI Section S2.2. We also report other preregistered analyses in the SI.

The survey experiment preregistration is available at <https://osf.io/2scah>.

Online behavior data

Respondents were invited to install software tracking web traffic on all their internet browsers, which they could disable or uninstall at any time. Identifying information, passwords, and financial transactions were not recorded.

Our web browsing data consist of 1,716 participants from September 13, 2020–January 29, 2021 and 1,756 participants from August 31, 2022–January 31, 2023. We focus on participants who were active online for the majority of months (at least three of five) in each election study period (1,596 participants in the 2020 election data; 1,518 participants in the 2022 data). SI Section S4.1 shows that most participants stayed active throughout the relevant election study period and that there were

no differences between Democrats and Republicans with regards to deactivation. These participants are broadly representative of the general population in terms of age, gender, education, race, and partisanship (see SI Section S1).

Among active participants, we have laptop/desktop data for 1,209 participants (76%) and mobile (smartphone, tablet, other) data for 502 participants (31%) in 2020, with 118 contributing data from both device types (device type is missing for three respondents). In 2022, we have laptop/desktop data for 965 participants (64%) and mobile (smartphone, tablet, other) data for 621 participants (41%), with 71 participants providing data from both device types (device type is missing for three respondents). As can be expected, participants who provided laptop/desktop data tend to be older and more educated compared to those who provided mobile data (Table S3). Exposure to news content, including election fraud content, is higher on laptop/desktop (Table S41), but the pattern of results we observe is similar across device types (Figure S31). Although we cannot determine the specific content that mobile users viewed within news apps, we show in SI Section S1.1 that mobile users are more frequently exposed to news content through web browsing than through news apps (Aridor et al. 2025 reach the same conclusion). Moreover, using news apps is positively correlated with visiting news websites, which suggests that the two forms of news consumption are complements rather than substitutes.

Online exposure measures

News-related content We identified all news-related URLs that participants were exposed to around the 2020 and 2022 elections using domains rated by Lin et al. (2023) or NewsGuard (more details below) after using [Shallalist](#) to remove domains from those lists that are search engines, social media, shopping, or sports domains or otherwise topically irrelevant (e.g., cars, fortune telling, etc.).

Election fraud content We scraped the content of all news-related URLs visited by participants ($N = 626,205$) and cleaned the output to remove as much content irrelevant to the body of the article as possible (e.g., invitations to subscribe or sign in, information about the use of cookies, ads, etc.). We applied an extensive dictionary (see SI Section S4.6) to identify possible fraud-related information in the content that people viewed. We then used GPT-5 (2025-08-07 version) to classify these articles (37,564 articles in the 2020 election data; 14,515 articles in the 2022 data) on three dimensions: 1)

whether the article mentions election fraud; 2) if yes, whether it includes a statement questioning or contradicting claims that election fraud is widespread or could change the outcome of one or more elections; and 3) if yes, the extent of questioning or debunking election fraud claims. These distinctions are motivated by research suggesting that effectively debunking false or unsupported claims can increase the accuracy of people’s factual beliefs (Chan and Albarracin 2023; Chan et al. 2017; Ecker et al. 2020; Swire et al. 2017). Table 2 presents examples for each level of questioning.

Table 2: Coding levels for extent of questioning

Extent of questioning	Examples (excerpts)
<p>0 = None — No questioning/debunking</p>	<p>‘There Was in Fact Fraud That Took Place’: FEC Chairman Trey Trainor “The massive amounts of affidavits that we see in these cases show that there was in fact fraud that took place,” he said during an interview with “Just the News AM” on Friday.</p> <p>Accountant testifies Trump claimed decade of huge tax losses Sen. Lindsey Graham, a Trump ally, testified before a Georgia grand jury probing alleged 2020 election interference.</p>
<p>1 = Minimal — Questioning/debunking appears only in the body (not in the headline or lede); limited space (≤ 2 sentences or $\sim \leq 10\%$ of the article); no explanation or at most a single brief reason with no elaboration; not a primary focus.</p>	<p>Hassan holds 13-point lead in New Hampshire reelection bid: survey Bolduc had supported Trump’s unfounded claims of mass electoral fraud during the 2020 presidential election, but quickly reversed his position... [single mention]</p>
<p>2 = Moderate — Questioning/debunking appears in the headline, lede, or body; receives some space ($\sim 10\text{--}33\%$ or a brief piece overall); includes some explanation or evidence (at least one concrete reason or source) that goes beyond a passing phrase; may be a primary focus if the article is short or provides little detail, but otherwise is not the main focus.</p>	<p>‘Don’t be weak fools!’ Trump tries to enlist more GOP lawmakers to prevent Biden presidency Twitter slapped a disclaimer onto the post, as the social media service has been doing since the election, which all 50 states have certified for Biden, who also won the Electoral College vote on Monday. Trump’s legal team has lost at least 60 cases in their election challenge, which the U.S. Supreme Court has also refused to hear after 126 Republican lawmakers signed on to a lawsuit filed by Texas attorney general Ken Paxton. [segment included at the end of a longer article]</p>
<p>3 = Extensive — Questioning/debunking is a primary focus: it appears in the headline or lede, occupies a substantial share of the article ($\geq \sim 34\%$, typically most), and provides detailed explanations/evidence (multiple reasons/sources or an in-depth analysis of one).</p>	<p>False claims persist over shredded election trash in Georgia AP’S ASSESSMENT: False. Cobb County election officials hired a shredding company to destroy non-relevant leftover election materials. The images show inner privacy envelopes voters used to seal their ballots. The actual ballots and the outside envelopes with voter signatures have been preserved, along with other relevant materials.</p> <p>Among the images that spread on social media was a photo of trash can with a paper labeled “ABSENTEE BALLOT” inside. But Eveler said that was an inner privacy envelope used by voters to seal absentee ballots, and had “no evidentiary value.” County officials will hold onto the actual absentee ballots, as well as the outer envelopes signed by voters, for two years. [segment of a fact-check article on this topic]</p>

We validated the LLM output by comparing it to the consensus coding from four human coders for a random sample of 50 articles, ensuring it met minimal reliability thresholds. For identifying whether an article mentioned election fraud, the LLM achieved 100% agreement with human coders and a Krippendorff's alpha of 1. For determining whether an article included a statement questioning widespread fraud claims, agreement was 96% with an alpha of 0.92. Finally, the Krippendorff's alpha was 0.94 for the full ordinal measure of the extent of questioning or debunking, 0.95 for the grouped ordinal measure corresponding to the categories we use in this paper ("Moderate" or "Extensive" questioning = 2, "Minimal" questioning = 1, "Absent" = 0), and 1 for a dichotomization contrasting fraud-challenging content with content containing no questioning at all ("Moderate" or "Extensive" questioning = 1, "Absent" questioning = 0, "Minimal" questioning = Excluded). More information about the codebook, how it was developed, and LLM coding accuracy are provided in Appendices [S4.7–S4.9](#).

In SI Tables [S11–S14](#), we list the most frequently visited domains for election fraud content in the period around the 2020 and 2022 elections by trustworthiness ratings and by total views of content challenging and non-questioning fraud claims. Tables [S15](#) and [S16](#) show the percentage of fraud-related articles that challenged election fraud claims across the 25 outlets with the highest number of fraud-related articles in the period around the 2020 and 2022 elections. Tables [S17–S20](#) present weighted and unweighted versions of the same statistics at the view level.

As a robustness check for whether our approach accurately captures participants' exposure to online fraud-related content, we draw on a dataset of Twitter posts about voter fraud surrounding the 2020 U.S. election (Abilov et al. 2021). This dataset includes all external URLs that appeared in the posts. Using this list, we identify the URLs viewed by our participants that were excluded by our dictionary prior to LLM classification. In total, 531 URLs (accounting for 1,105 views) excluded by the dictionary were classified by the LLM as containing fraud-related information (these URLs would represent 3% of election fraud views if included), and 819 URLs (accounting for 2,397 views) were classified as not containing fraud-related information.

Trustworthiness We rely on two sources of data for measuring the trustworthiness of the online sources to which participants were exposed: ratings from NewsGuard (February 2021 version) and Lin

et al. (2023). NewsGuard provides a binary label (trustworthy and not trustworthy) for 7,109 domains. Lin et al. (2023) combine expert ratings from five sources (Ad Fontes Media, Iffy index of unreliable sources, Media Bias/Fact-Check, Lasser et al. 2022, and professional fact-checkers from Pennycook and Rand 2019) into a single principal component score. Their dataset contains ratings for 11,520 domains. To increase the number of rated domains in our dataset, we combined the two sets of ratings by converting Lin et al.'s principal component score into a binary measure of trustworthiness. The conversion was performed using the R package `cutpointnr` (Thiele and Hirschfeld 2021). These two sets of ratings are highly consistent with one another: among the domains with ratings from both sources, 93% have the same rating in our 2020 Pulse data, while 94% have the same rating in our 2022 Pulse data. When both sources rated the same domain, we used Lin et al.'s rating based on the assumption that the combination of five expert ratings should be given more weight than a single one.

More details about the creation of each behavioral variable are provided in SI Section [S4](#). Descriptive statistics about exposure to all types of content, overall and by candidate support, are included in SI Section [S4.11](#).

Permutation tests

We conduct independent permutation tests to better understand the roles of topic selection, outlet selection, and within-topic article selection in explaining the gap in the number of fraud-challenging views between Biden and Trump supporters. In the first test, we randomly permuted content types (fraud-related versus non-fraud-related) within each news outlet, holding participants' overall news consumption and outlet composition constant. This test isolates the extent to which the gap reflects partisan differences in the likelihood of selecting fraud-related content versus other types of content. In the second test, we permuted outlet selection for fraud-related content while maintaining each outlet's overall share of fraud-related views and balance of challenging and non-challenging articles. Specifically, each participant's outlet blocks of fraud-related views (e.g., 10 from *The Washington Post*, 5 views from CNN) were randomly replaced by equivalently-sized blocks from other outlets (e.g., 10 from Fox News, 5 from *The New York Times*). This test measures how much of the gap results from differences in outlet choices across partisan groups. In the third test, we randomly reassigned chal-

lenging and non-challenging views within outlets among fraud-related articles to assess how much of the gap stems from selective exposure to (or avoidance of) congenial versus uncongenial fraud-related coverage.

We ran 5,000 permutations in each case to generate a null distribution for the remaining gap after each test and estimated the effects of independently varying each factor by comparing the permuted and observed gaps. As noted in the main text, each test is conducted independently to avoid the complexities of order effects and interactions between factors (Fortin et al. 2011). We instead separately compare the magnitude of the effects of each factor using the counterfactuals created by the permutation tests (which therefore do not add up to 100%).

Weighting

We do not apply post-stratification weights in our experimental analyses because they can introduce bias if untestable assumptions about sample selection and treatment effect heterogeneity are violated, cause covariate imbalance, and lead to significant loss in statistical power (Franco et al. 2017). All other survey and online-based analyses use probability weights constructed by YouGov to more accurately represent the U.S. adult population.

References

- Abilov, Anton, Yiqing Hua, Hana Matatov, Ofra Amir, and Mor Naaman. 2021, May. “Voter-Fraud2020: A Multi-modal Dataset of Election Fraud Claims on Twitter”. *Proceedings of the International AAAI Conference on Web and Social Media* 15 : 901–912.
- Allen, Jennifer, Baird Howland, Markus Mobius, David Rothschild, and Duncan J Watts. 2020. “Evaluating the fake news problem at the scale of the information ecosystem”. *Science Advances* 6 (14): eaay3539.
- Anderson, Christopher J., André Blais, Shaun Bowler, Todd Donovan, and Ola Listhaug. 2005, January. *Losers’ Consent: Elections and Democratic Legitimacy*. Oxford University Press.

- Anderson, Christopher J and Christine A Guillory. 1997. "Political institutions and satisfaction with democracy: A cross-national analysis of consensus and majoritarian systems". *American Political Science Review* 91 (1): 66–81.
- Anderson, Christopher J and Andrew J LoTempio. 2002. "Winning, losing and political trust in America". *British Journal of Political Science* 32 (2): 335–351.
- Arceneaux, Kevin and Martin Johnson. 2019. "Selective avoidance and exposure". In *Oxford Research Encyclopedia of Communication*. Oxford University Press.
- Aridor, Guy, Tevel Dekel, Rafael Jiménez Durán, Ro'ee Levy, and Lena Song. 2025. "Digital news consumption: Evidence from smartphone content in the 2024 US elections". Downloaded December 16, 2025 from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5364054.
- Barrie, Christopher, Aybüke Atalay, and Alia ElKattan. 2025. "Information diets are more diverse in attention than in engagement". OSF. Downloaded September 5, 2025 from https://osf.io/preprints/osf/kvbzd_v1.
- Berinsky, Adam. 2012. "The birthers are back". YouGov, February 3, 2012. Downloaded December 18, 2024 from <https://today.yougov.com/politics/articles/2862-birthers-are-back>.
- Berlinski, Nicolas, Margaret Doyle, Andrew M Guess, Gabrielle Levy, Benjamin Lyons, Jacob M Montgomery, Brendan Nyhan, and Jason Reifler. 2023. "The effects of unsubstantiated claims of voter fraud on confidence in elections". *Journal of Experimental Political Science* 10 (1): 34–49.
- Birch, Sarah. 2010. "Perceptions of electoral fairness and voter turnout". *Comparative Political Studies* 43 (12): 1601–1622.
- Bloniarz, Adam, Hanzhong Liu, Cun-Hui Zhang, Jasjeet S Sekhon, and Bin Yu. 2016. "Lasso adjustments of treatment effect estimates in randomized experiments". *Proceedings of the National Academy of Sciences* 113 (27): 7383–7390.

- Bowles, Jeremy, Kevin Croke, Horacio Larreguy, Shelley Liu, and John Marshall. 2025. "Sustaining Exposure to Fact-Checks: Misinformation Discernment, Media Consumption, and Its Political Implications". *American Political Science Review* 119 (4): 1864–1887.
- Braghieri, Luca, Sarah Eichmeyer, Ro'ee Levy, Markus M. Mobius, Jacob Steinhardt, and Ruiqi Zhong. 2024. "Article-Level Slant and Polarization of News Consumption on Social Media".
- Bright Line Watch. 2024. "America looks ahead to a second trump term". Downloaded December 18, 2024 from <https://brightlinewatch.org/america-looks-ahead-to-a-second-trump-term/>.
- Broockman, David E and Joshua L Kalla. 2025. "Selective exposure and echo chambers in partisan television consumption: Evidence from linked viewership, administrative, and survey data". *American Journal of Political Science* 69 (3): 847–865.
- Brown, Katie. 2013. "Trayvon Martin, Obama, and the Birthers: How race and suspicion continue over 35 years". Institute for Social Research, Center for Political Studies, University of Michigan. July 29, 2013. Downloaded July 19, 2024 from <https://cpsblog.isr.umich.edu/?p=54>.
- Budak, Ceren, Brendan Nyhan, David M. Rothschild, Emily Thorson, and Duncan J. Watts. 2024. "Misunderstanding the harms of online misinformation". *Nature* 630 (8015): 45–53.
- Carey, John M., Victoria Chi, D.J. Flynn, Brendan Nyhan, and Thomas Zeitzoff. 2020. "The effects of corrective information about disease epidemics and outbreaks: Evidence from Zika and yellow fever in Brazil". *Science Advances* 6 (5): eaaw7449.
- Carey, John M, Elizabeth Chun, Alice Cook, Brian J Fogarty, Leyla Jacoby, Brendan Nyhan, Jason Reifler, and Lilian Sweeney. 2025. "The narrow reach of targeted corrections: No impact on broader beliefs about election integrity". *Political Behavior* 47 (2): 737–750.
- Carey, John M, Brian Fogarty, Marília Gehrke, Brendan Nyhan, and Jason Reifler. 2025. "Prebunking and credible source corrections increase election credibility: Evidence from the US and Brazil". *Science Advances* 11 (35): eadv3758.

- Carey, John M, Andrew M Guess, Peter J Loewen, Eric Merkley, Brendan Nyhan, Joseph B Phillips, and Jason Reifler. 2022. “The ephemeral effects of fact-checks on COVID-19 misperceptions in the United States, Great Britain and Canada”. *Nature Human Behaviour* 6 (2): 236–243.
- Carnahan, Dustin, Daniel E. Bergan, and Sangwon Lee. 2021. “Do Corrective Effects Last? Results from a Longitudinal Experiment on Beliefs Toward Immigration in the U.S.”. *Political Behaviour* 43 (3): 1227–1246.
- Chan, Man-pui Sally and Dolores Albarracín. 2023. “A meta-analysis of correction effects in science-relevant misinformation”. *Nature Human Behaviour* 7 (9): 1514–1525.
- Chan, Man-pui Sally, Christopher R. Jones, Kathleen Hall Jamieson, and Dolores Albarracín. 2017. “Debunking: A meta-analysis of the psychological efficacy of messages countering misinformation”. *Psychological Science* 28 (11): 1531–1546.
- Chen, Annie Y, Brendan Nyhan, Jason Reifler, Ronald E Robertson, and Christo Wilson. 2023. “Subscriptions and external links help drive resentful users to alternative and extremist YouTube channels”. *Science Advances* 9 (35): eadd8080.
- Clarke, Harold D and Alan C Acock. 1989. “National elections and political attitudes: The case of political efficacy”. *British Journal of Political Science* 19 (4): 551–562.
- Clayton, Katherine, Nicholas T. Davis, Brendan Nyhan, Ethan Porter, Timothy J. Ryan, and Thomas J. Wood. 2021. “Elite rhetoric can undermine democratic norms”. *Proceedings of the National Academy of Sciences* 118 (23): e2024125118.
- Clifford, Scott, Geoffrey Sheagley, and Spencer Piston. 2021. “Increasing precision without altering treatment effects: Repeated measures designs in survey experiments”. *American Political Science Review* 115 (3): 1048–1065.
- Coppock, Alexander, Kimberly Gross, Ethan Porter, Emily Thorson, and Thomas J. Wood. 2023. “Conceptual Replication of Four Key Findings about Factual Corrections and Misinformation during the 2020 US Election: Evidence from Panel-Survey Experiments”. *British Journal of Political Science* 53 (4): 1328–1341.

- Dahlke, Ross and Jeffrey Hancock. 2025, September. "Untrustworthy Website Exposure and Election Beliefs: Selective Exposure and Ideological Asymmetry". *Journal of Online Trust and Safety* 3 (1): 1–37.
- Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky, and Joshua A Tucker. 2019. "How many people live in political bubbles on social media? Evidence from linked survey and twitter data". *Sage Open* 9 (1): 2158244019832705.
- Eady, Gregory, Tom Paskhalis, Jan Zilinsky, Richard Bonneau, Jonathan Nagler, and Joshua A Tucker. 2023. "Exposure to the Russian Internet Research Agency foreign influence campaign on twitter in the 2016 US election and its relationship to attitudes and voting behavior". *Nature Communications* 14 (1): 62.
- Ecker, Ullrich KH, Ziggy O'Reilly, Jesse S Reid, and Ee Pin Chang. 2020. "The effectiveness of short-format refutational fact-checks". *British Journal of Psychology* 111 (1): 36–54.
- Flynn, D.J., Brendan Nyhan, and Jason Reifler. 2017. "The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics". *Political Psychology* 38 (S1): 127–150.
- Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. "Decomposition methods in economics". In *Handbook of Labor Economics*, Volume 4, pp. 1–102. Elsevier.
- Franco, Annie, Neil Malhotra, Gabor Simonovits, and L. J. Zigerell. 2017. "Developing Standards for Post-Hoc Weighting in Population-Based Survey Experiments". *Journal of Experimental Political Science* 4 (2): 161–172.
- Franklin, Mark N. 2004. *Voter turnout and the dynamics of electoral competition in established democracies since 1945*. Cambridge University Press.
- Garrett, R Kelly. 2009. "Politically motivated reinforcement seeking: Reframing the selective exposure debate". *Journal of Communication* 59 (4): 676–699.

- Garrett, R Kelly, Dustin Carnahan, and Emily K Lynch. 2013. "A turn toward avoidance? selective exposure to online political information, 2004–2008". *Political Behavior* 35 (1): 113–134.
- Garrett, R Kelly, Jacob A Long, and Min Seon Jeong. 2019. "From partisan media to misperception: Affective polarization as mediator". *Journal of Communication* 69 (5): 490–512.
- Garrett, R Kelly and Natalie Jomini Stroud. 2014. "Partisan paths to exposure diversity: Differences in pro-and counterattitudinal news consumption". *Journal of Communication* 64 (4): 680–701.
- Garrett, R Kelly, Brian E Weeks, and Rachel L Neo. 2016. "Driving a wedge between evidence and beliefs: How online ideological news exposure promotes political misperceptions". *Journal of Computer-Mediated Communication* 21 (5): 331–348.
- Ginsberg, Benjamin and Robert Weissberg. 1978. "Elections and the mobilization of popular support". *American Journal of Political Science* 22 (1): 31–55.
- Graham, Matthew H and Omer Yair. 2024. "Expressive responding and belief in 2020 election fraud". *Political Behavior* 46 (2): 1349–1374.
- Green, Jon, Stefan McCabe, Sarah Shugars, Hanyu Chwe, Luke Horgan, Shuyang Cao, and David Lazer. 2025. "Curation Bubbles". *American Political Science Review* 119 (4): 1704–1722.
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer. 2019. "Fake news on Twitter during the 2016 US presidential election". *Science* 363 (6425): 374–378.
- Guess, Andrew, Jonathan Nagler, and Joshua Tucker. 2019. "Less than you think: Prevalence and predictors of fake news dissemination on Facebook". *Science Advances* 5 (1): eaau4586.
- Guess, Andrew, Brendan Nyhan, Benjamin Lyons, and Jason Reifler. 2018. "Avoiding the echo chamber about echo chambers". *Knight Foundation* 2 (1): 1–25.
- Guess, Andrew M. 2021. "(Almost) everything in moderation: New evidence on Americans' online media diets". *American Journal of Political Science* 65 (4): 1007–1022.
- Guess, Andrew M., Brendan Nyhan, and Jason Reifler. 2020. "Exposure to untrustworthy websites in the 2016 US election". *Nature Human Behaviour* 4 (5): 472–480.

- Haroon, Muhammad, Magdalena Wojcieszak, and Anshuman Chhabra. 2025. "'Whose side are you on?' Estimating ideology of political and news content using large language models and few-shot demonstration selection". March 31, 2025. <https://doi.org/10.48550/arXiv.2503.20797>.
- Hart, William, Dolores Albarracín, Alice H Eagly, Inge Brechan, Matthew J Lindberg, and Lisa Merrill. 2009. "Feeling validated versus being correct: A meta-analysis of selective exposure to information.". *Psychological Bulletin* 135 (4): 555–588.
- Iyengar, Shanto and Kyu S. Hahn. 2009. "Red media, blue media: Evidence of ideological selectivity in media use". *Journal of Communication* 59 (1): 19–39.
- Jerit, Jennifer, Jason Barabas, William Pollock, Susan Banducci, Daniel Stevens, and Martijn Schoonvelde. 2016. "Manipulated vs. measured: Using an experimental benchmark to investigate the performance of self-reported media exposure". *Communication Methods and Measures* 10 (2-3): 99–114.
- Justwan, Florian and Ryan D. Williamson. 2022. "Trump and Trust: Examining the Relationship between Claims of Fraud and Citizen Attitudes". *PS: Political Science & Politics* 55 (3): 462–469.
- Kennedy, Ian, Morgan Wack, Andrew Beers, Joseph S. Schafer, Isabella Garcia-Camargo, Emma S. Spiro, and Kate Starbird. 2022. "Repeat Spreaders and Election Delegitimization: A Comprehensive Dataset of Misinformation Tweets from the 2020 U.S. Election". *Journal of Quantitative Description: Digital Media* 2.
- Kim, Jae Wook and Eunji Kim. 2021. "Temporal selective exposure: How partisans choose when to follow politics". *Political Behavior* 43 : 1663–1683.
- Lasser, Jana, Segun Taofeek Aroyehun, Almog Simchon, Fabio Carrella, David Garcia, and Stephan Lewandowsky. 2022. "Social media sharing by political elites: An asymmetric American exceptionalism". arXiv:2207.06313 [cs].
- Lin, Hause, Jana Lasser, Stephan Lewandowsky, Rocky Cole, Andrew Gully, David G Rand, and

- Gordon Pennycook. 2023. “High level of correspondence across different news domain quality rating sets”. *PNAS Nexus* 2 (9): pgad286.
- Lupia, Arthur. 2016. *Uninformed: Why people know so little about politics and what we can do about it*. Oxford University Press.
- Lyons, Benjamin A. and Kaitlyn S. Workman. 2022. “Research note: Explicit voter fraud conspiracy cues increase belief among co-partisans but have broader spillover effects on confidence in elections”. *Harvard Kennedy School Misinformation Review*.
- Maldonado, Arturo and Mitchell A Seligson. 2014. “Electoral trust in Latin America”. In P. Norris, R. W. Frank, and F. Martinez (Eds.), *Advancing electoral integrity*, pp. 229–245. Oxford University Press.
- Marlon, Jennifer, Matthew Ballew, Seth Rosenthal, Edward Maibach, and Anthony Leiserowitz. 2024. “Explore climate change in the American mind”. Yale Program on Climate Change Communication, May 20, 2024. Downloaded July 19, 2024 from <https://climatecommunication.yale.edu/visualizations-data/americans-climate-views/>.
- Mayer, Jane. 2012. “The voter-fraud myth”. *The New Yorker*, October 12, 2012.
- Meirick, Patrick C. 2013. “Motivated misperception? Party, education, partisan news, and belief in “death panels””. *Journalism & Mass Communication Quarterly* 90 (1): 39–57.
- Meirick, Patrick C and Elena Bessarabova. 2016. “Epistemic factors in selective exposure and political misperceptions on the right and left”. *Analyses of Social Issues and Public Policy* 16 (1): 36–68.
- Minnite, Lorraine C. 2010. *The myth of voter fraud*. Cornell University Press.
- Moore, Ryan C., Ross Dahlke, and Jeffrey T. Hancock. 2023. “Exposure to untrustworthy websites in the 2020 US election”. *Nature Human Behaviour* 7 (7): 1096–1105.
- Nadeau, Richard and André Blais. 1993. “Accepting the election outcome: the effect of participation on losers’ consent”. *British Journal of Political Science* 23 (4): 553–563.

- NewsGuard. 2022, October. “More Than 2,100 Websites Improve their Journalism Trust Practices Through NewsGuard’s Rating Process”.
- Nyhan, Brendan. 2020. “Facts and myths about misperceptions”. *Journal of Economic Perspectives* 34 (3): 220–236.
- Nyhan, Brendan. 2021. “Why the backfire effect does not explain the durability of political misperceptions”. *Proceedings of the National Academy of Sciences* 118 (15): e1912440117.
- Nyhan, Brendan, Ethan Porter, and Thomas J Wood. 2022. “Time and skeptical opinion content erode the effects of science coverage on climate beliefs and attitudes”. *Proceedings of the National Academy of Sciences* 119 (26): e2122069119.
- Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Andrew M. Guess, Edward Kennedy, Young Mie Kim, David Lazer, Neil Malhotra, Devra Moehler, Jennifer Pan, Daniel Robert Thomas, Rebekah Tromble, Carlos Velasco Rivera, Arjun Wilkins, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker. 2023. “Like-minded sources on Facebook are prevalent but not polarizing”. *Nature* 620 (7972): 137–144.
- Painter, David Lynn and Juliana Fernandes. 2024. ““The Big Lie”: How Fact Checking Influences Support for Insurrection”. *American Behavioral Scientist* 68 (7): 892–912.
- Pariser, Eli. 2011. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.
- Pennycook, Gordon and David G. Rand. 2019. “Fighting misinformation on social media using crowdsourced judgments of news source quality”. *Proceedings of the National Academy of Sciences* 116 (7): 2521–2526.
- Peterson, Erik and Shanto Iyengar. 2021. “Partisan gaps in political information and information-seeking behavior: Motivated reasoning or cheerleading?”. *American Journal of Political Science* 65 (1): 133–147.

- Prike, Toby and Ullrich KH Ecker. 2023. "Effective correction of misinformation". *Current Opinion in Psychology* 54 : 101712.
- Prior, Markus. 2009a. "The immensely inflated news audience: Assessing bias in self-reported news exposure". *Public Opinion Quarterly* 73 (1): 130–143.
- Prior, Markus. 2009b. "Improving media effects research through better measurement of news exposure". *The Journal of Politics* 71 (3): 893–908.
- Przeworski, Adam. 2019. *Crises of democracy*. Cambridge University Press.
- Reinhart, R.J. 2020. "Fewer in u.s. continue to see vaccines as important". Gallup, January 14, 2020. Downloaded July 19, 2024 from <https://web.archive.org/web/20240105212603/https://news.gallup.com/poll/276929/fewer-continue-vaccines-important.aspx>.
- Robertson, Craig T. 2025. "People are turning away from the news. Here's why it may be happening". Reuters Institute for the Study of Journalism, February 21, 2025. <https://reutersinstitute.politics.ox.ac.uk/news/people-are-turning-away-news-heres-why-it-may-be-happening>.
- Robertson, Ronald E, Jon Green, Damian J Ruck, Katherine Ognyanova, Christo Wilson, and David Lazer. 2023. "Users choose to engage with more partisan news than they are exposed to on Google Search". *Nature* 618 (7964): 342–348.
- Sinclair, Betsy, Steven S Smith, and Patrick D Tucker. 2018. "'It's largely a rigged system': Voter confidence and the winner effect in 2016". *Political Research Quarterly* 71 (4): 854–868.
- Skovsgaard, Morten and Kim Andersen. 2020. "Conceptualizing news avoidance: Towards a shared understanding of different causes and potential solutions". *Journalism Studies* 21 (4): 459–476.
- Skovsgaard, Morten and Kim Andersen. 2022. "News avoidance". *The SAGE Encyclopedia of Journalism* 1 : 1099–1103.

- Stroud, Natalie Jomini. 2008. "Media use and political predispositions: Revisiting the concept of selective exposure". *Political Behavior* 30 (3): 341–366.
- Sunstein, Cass R. 2001. *Republic.com*. Princeton University Press.
- Swire, Briony, Ullrich KH Ecker, and Stephan Lewandowsky. 2017. "The role of familiarity in correcting inaccurate information.". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 43 (12): 1948.
- Swire-Thompson, Briony, Joseph DeGutis, and David Lazer. 2020. "Searching for the backfire effect: Measurement and design considerations". *Journal of Applied Research in Memory and Cognition* 9 (3): 286–299.
- Taber, Charles S. and Milton Lodge. 2006. "Motivated skepticism in the evaluation of political beliefs". *American Journal of Political Science* 50 (3): 755–769.
- Thiele, Christian and Gerrit Hirschfeld. 2021, June. "Cutpointr: Improved Estimation and Validation of Optimal Cutpoints in R". *Journal of Statistical Software* 98 : 1–27.
- Thorson, Emily. 2024. *The Invented State: Policy Misperceptions in the American Public*. Oxford University Press.
- Toff, Benjamin and Antonis Kalogeropoulos. 2020. "All the news that's fit to ignore: How the information environment does and does not shape news avoidance". *Public Opinion Quarterly* 84 (S1): 366–390.
- Toff, Benjamin, Ruth Palmer, and Rasmus Kleis Nielsen. 2023. *Avoiding the News: Reluctant Audiences for Journalism*. Columbia University Press.
- Törnberg, Petter. 2024. "Best Practices for Text Annotation with Large Language Models". *Sociologica* 18 (2): 67–85.
- Uscinski, Joseph E., Casey Klofstad, and Matthew D. Atkinson. 2016. "What Drives Conspiratorial Beliefs? The Role of Informational Cues and Predispositions". *Political Research Quarterly* 69 (1): 57–71.

- Villi, Mikko, Tali Aharoni, Keren Tenenboim-Weinblatt, Pablo J Boczkowski, Kaori Hayashi, Eugenia Mitchelstein, Akira Tanaka, and Neta Kligler-Vilenchik. 2022. "Taking a break from news: A five-nation study of news avoidance in the digital era". *Digital Journalism* 10 (1): 148–164.
- Wack, Morgan, Kayla Duskin, and Damian Hodel. 2024. "Political fact-checking efforts are constrained by deficiencies in coverage, speed, and reach". arXiv preprint arXiv:2412.13280.
- Walter, Nathan, Jonathan Cohen, R. Lance Holbert, and Yasmin Morag. 2020, May. "Fact-Checking: A Meta-Analysis of What Works and for Whom". *Political Communication* 37 (3): 350–375.
- Walter, Nathan and Sheila T. Murphy. 2018, July. "How to unring the bell: A meta-analytic approach to correction of misinformation". *Communication Monographs* 85 (3): 423–441.
- Wood, Thomas and Ethan Porter. 2019. "The elusive backfire effect: Mass attitudes' steadfast factual adherence". *Political Behavior* 41 : 135–163.
- Yang, Eddie, Zoey Wang, Carl Zhou, and Yaosheng Xu. 2025. "Data annotation with large language models: Lessons from a large empirical evaluation". October 20, 2025. Downloaded October 21, 2025 from https://www.eddieyang.net/research/llm_annotation.pdf.
- Yen, Hope, Ali Swenson, and Amanda Seitz. 2020. "AP FACT CHECK: Trump's claims of vote rigging are all wrong". Associated Press, December 3, 2020.
- YouGov America. 2017. "America's friends and enemies". February 2, 2017. Downloaded May 23, 2022 from <https://today.yougov.com/topics/politics/articles-reports/2017/02/02/americas-friends-and-enemies>.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. Cambridge University Press.

Supplementary Information

Contents

S1	Sample characteristics	S2
S1.1	Evaluating online news consumption and news app usage among mobile users	S5
S2	Additional results	S6
S2.1	Persistence of misperceptions	S6
S2.2	Preregistered experimental analyses	S11
S2.2.1	Experimental treatment	S11
S2.2.2	Experimental results	S12
S2.3	Exposure to election fraud content	S25
S3	Robustness checks	S59
S3.1	Experimental results among active Pulse participants	S59
S4	Online behavior data	S69
S4.1	Pulse data	S69
S4.2	Domain selection process	S69
S4.3	Ratings of domains	S70
S4.4	Identifying outliers with unusual browsing behavior	S71
S4.5	Scraping process	S72
S4.6	Text parsing and wrangling	S72
S4.7	Developing the LLM prompt	S73
S4.8	Validation	S78
S4.9	Coding articles using LLM	S79
S4.10	Fact-checks	S79
S4.11	Descriptive statistics	S80

S1 Sample characteristics

Table S1 provides information about the demographic composition of each of our survey samples. (Note: Wave 2 of the prior COVID study and Wave 3 of the 2022 election study were not used in this research.) Table S2 presents the same information among Pulse participants who were active for the majority of months during the 2020 and 2022 study periods (our eligibility criterion for behavioral data). Finally, Table S3 shows the demographic composition of the sample based on whether they provided laptop/desktop or mobile web browsing data in 2020 and 2022 (note: some participants provided both).

Table S1: Demographics of survey participants

Sample	Prior COVID study			2020 election study		2022 election study		
	Wave 1 (2020)	Wave 3 (2020)	Wave 4 (2021)	Wave 1 (2020–21)	Wave 2 (2021)	Wave 1 (2022)	Wave 2 (2022)	Wave 3 (2023)
Dates	May 20– June 3	July 28– Aug. 19	March 9– March 23	Dec. 17– Jan. 5	Jan. 13–19	Oct. 18– Nov. 9	Dec. 7–20	Jan. 21–30
Gender								
Male	47.6	48.2	48.4	48.7	48.5	48.0	48.2	48.6
Female	52.4	51.8	51.6	51.3	51.5	52.0	51.8	51.4
Age								
18–29 years	13.9	12.7	19.6	18.8	18.8	20.4	19.9	20.7
30–44 years	29.8	32.7	25.6	25.5	25.9	25.6	25.7	25.3
45–59 years	23.4	23.9	23.5	22.4	22.7	21.8	21.3	21.3
60+ years	32.8	30.8	31.2	33.4	32.6	32.2	33.1	32.6
Education								
4-year college	30.1	29.6	31.4	30.2	30.0	31.3	30.6	30.4
Race								
White	64.8	63.8	64.7	64.2	64.2	63.8	63.7	63.5
Black	12.0	12.0	12.3	11.9	11.9	12.1	12.3	12.4
Hispanic	14.5	15.6	14.1	15.5	15.3	15.2	15.1	15.2
Asian	4.2	4.5	3.9	3.4	3.4	3.0	2.6	2.6
Party ID/Vote								
Democrat	45.5	44.7	46.9	45.2	46.1	45.4	47.2	46.7
Republican	36.5	36.9	35.3	37.1	35.4	36.6	35.6	36.2
Clinton/Biden supporter	35.5	35.6	44.5	45.8	45.7	42.1	42.7	43.0
Trump supporter	33.5	33.8	40.5	39.6	39.2	39.4	39.1	39.3
N	4399	2983	5575	4312	3847	3772	2896	2100

Participants are YouGov panel members. All listed statistics other than the sample sizes are percentages. Participants in Waves 2 and 3 of each study are recontacts from Wave 1. COVID Wave 4 includes both recontacts and new participants. We use 2016 vote choice for waves conducted before the 2020 election and 2020 vote choice for waves conducted after the election. Percentages calculated using post-stratification weights.

Table S2: Demographics of Pulse participants

Variables	2020 Pulse sample	2022 Pulse sample
Gender		
Male	48.0	48.7
Female	52.0	51.3
Age		
18–29 years	15.5	28.8
30–44 years	24.9	27.3
45–59 years	25.2	19.4
60+ years	34.4	24.5
Education		
4-year college	27.9	26.9
Race		
White	68.0	61.2
Black	10.2	13.5
Hispanic	13.9	16.6
Asian	2.5	2.8
Party ID/candidate support		
Democrat PID	45.7	47.8
Republican PID	38.4	32.7
Biden supporter (2020)	46.6	42.7
Trump supporter (2020)	40.0	33.0
N	1596	1518

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Table S3: Demographics of Pulse participants by device type

Variables	2020 Laptop/Desktop	2020 Mobile	2022 Laptop/Desktop	2022 Mobile
Gender				
Male	49.9	47.4	49.6	48.4
Female	50.1	52.6	50.4	51.6
Age				
18–29 years	15.9	17.7	28.0	29.5
30–44 years	21.6	29.3	19.4	37.1
45–59 years	25.5	25.3	19.8	18.9
60+ years	36.9	27.6	32.7	14.5
Education				
4-year college	29.3	24.5	33.6	18.6
Race				
White	72.0	60.5	69.7	51.5
Black	8.1	13.0	8.1	19.8
Hispanic	11.0	20.6	15.6	17.5
Asian	2.8	2.1	3.0	2.9
Party ID/candidate support				
Democrat PID	45.5	43.8	48.7	45.8
Republican PID	39.1	38.7	35.1	30.5
Biden supporter (2020)	46.2	44.7	44.8	39.5
Trump supporter (2020)	41.5	38.8	38.6	27.3
N	1209	502	965	621

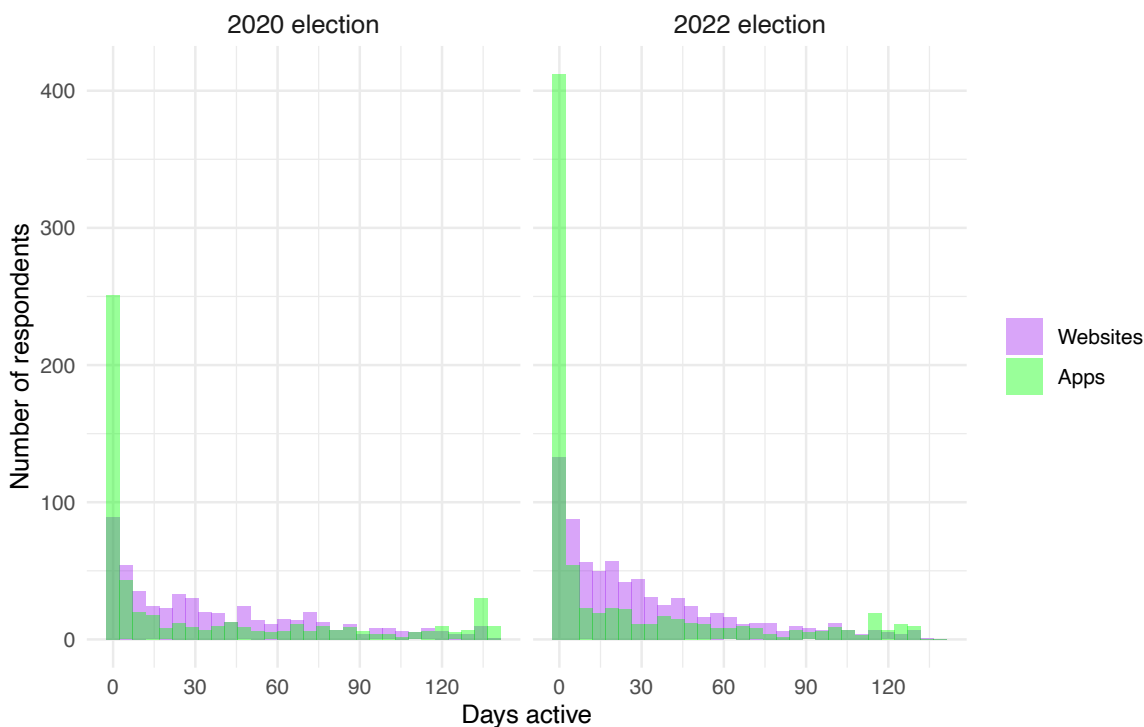
Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

S1.1 Evaluating online news consumption and news app usage among mobile users

We cannot directly observe the specific articles that mobile device users access via apps rather than web browsers. However, Pulse data still enable us to compare the relative usage of apps versus online browsers for news. These data show that mobile device users were more likely to access news through websites than through apps. Among participants for whom we have mobile data around the 2020 election, only 9% did not visit any news website, whereas 38% did not use any news app. In 2022, these figures were 8% and 46%, respectively.

When looking at the number of days of usage during our study periods, we find that mobile participants accessed news websites more frequently than news apps, with an average usage of 39 days in 2020 and 36 days in 2022 for websites, compared to 33 days in 2020 and 22 days in 2022 for apps. This difference becomes even more pronounced when examining the medians: 28 days (2020) and 23 days (2022) for websites versus just 4 days (2020) and 1 day (2022) for apps. These results suggest only a small subset of mobile participants used news apps on a regular basis. Figure S1 shows the full distribution of days of news website and app usage.

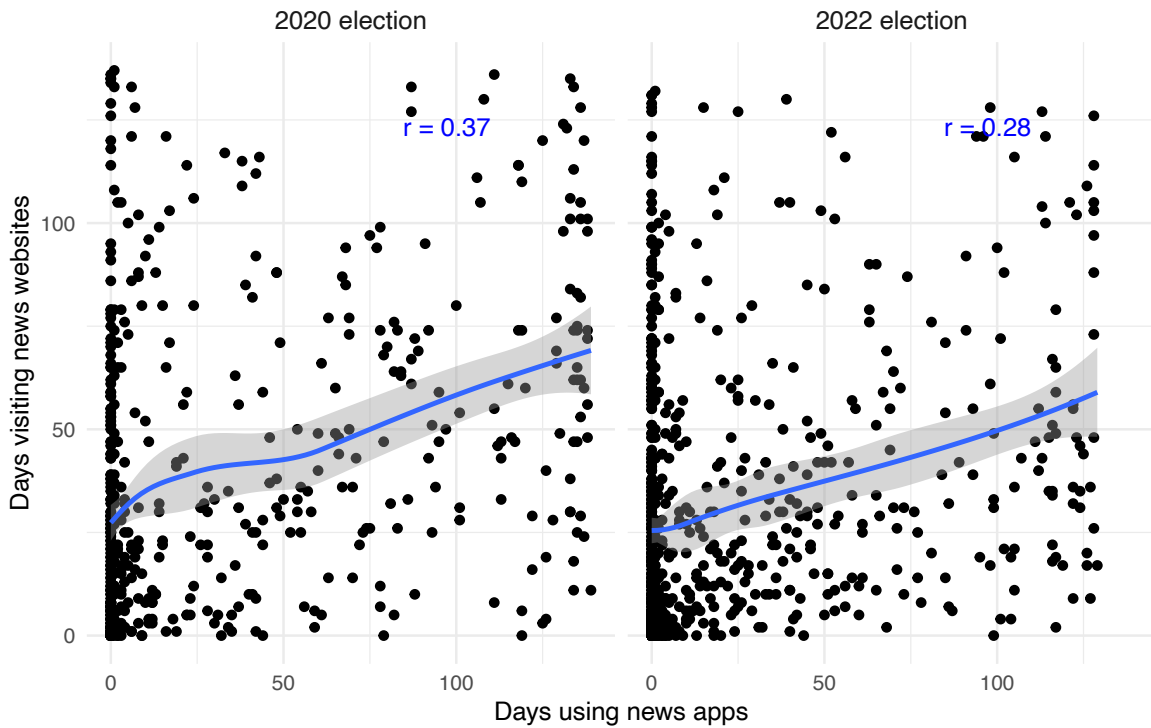
Figure S1: Distribution of the number of days participants visited news websites and used news apps during the 2020 and 2022 study periods



Based on data from participants who provided mobile data. The 2020 election period includes web browsing data from September 13, 2020–January 29, 2021 and app data from September 14, 2020–January 30, 2021 (138 days). The 2022 election period includes web browsing data from August 31, 2022–January 31, 2023 (153 days) and app data from September 1, 2022–January 11, 2023 (133 days). To ensure comparability in this figure, the 2022 web browsing data was truncated to align with the time period covered by the 2022 app data.

We also examined whether participants accessed news exclusively through websites or apps, or used a combination of both. As shown in Figure S2, there is a positive correlation between the number of days that participants visited news websites and the number of days that they used news apps, which suggests that people who are more interested in the news tend to use both (put another way, people do not appear to use news apps as a substitute for visiting news websites).

Figure S2: Number of days participants visited news websites and used news apps during the 2020 and 2022 study periods



Based on data from participants who provided mobile data. To ensure comparability in this figure, the 2022 web browsing data was truncated to align with the time period covered by the 2022 app data.


S2 Additional results

S2.1 Persistence of misperceptions

This subsection shows that variation in the perceived prevalence of election fraud over time is consistent when disaggregating the data by Trump approval (Figure S3) and partisan identification (Figure S4). Trump approval was measured using the following question: “Do you approve or disapprove of the way Donald Trump is handling his job as President?” We split the data into approvers (“Strongly approve” or “Somewhat approve”) and disapprovers (“Strongly disapprove” or “Somewhat disapprove”). Partisan identification was originally measured on a seven-point scale, which we col-

lapse by combining all Democratic and Republican identifiers, including leaners, into nominal categories.

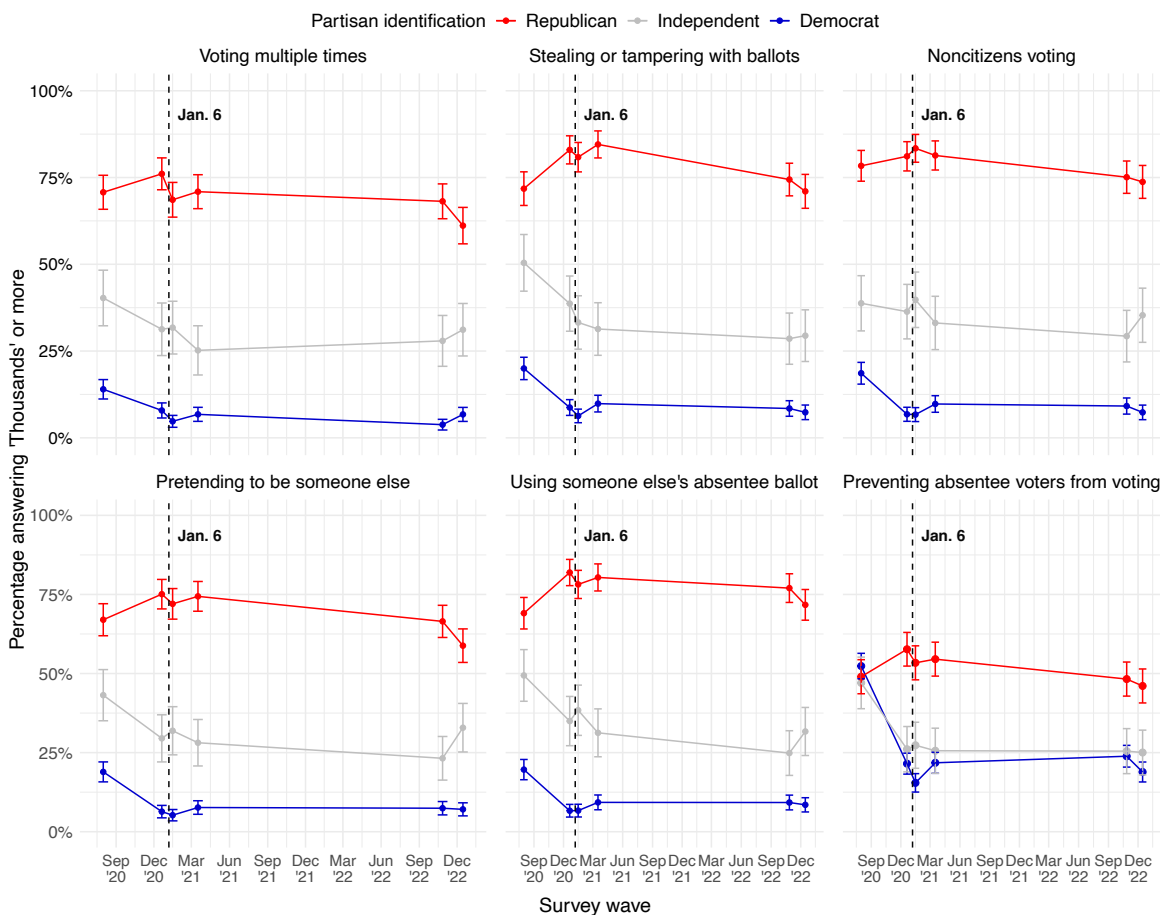
Figure S3: Perceived prevalence of election fraud over time by presidential approval



Percentage of respondents indicating that there are “Thousands” of cases or more of each type of election fraud. Categories included “Less than ten” (post-2020 election waves only), “Less than a hundred,” “Hundreds,” “Thousands,” “Tens of thousands,” “Hundreds of thousands,” and “A million or more”. Fraud prevalence measured for U.S. elections in general (Aug. 2020) and in the 2020 U.S. presidential election (subsequent waves). Group means computed using each wave’s post-stratification weights. Results are disaggregated by approval of Donald Trump as measured in May 2020 among participants who participated in that wave (N=1,074).


Figures S5 and S6 use alluvial plots to visualize within-respondent stability in perceptions of various measures related to election fraud prevalence and election confidence, showing how individuals moved between categories across survey waves. These analyses use post-stratification weights from the more recent wave for each between-wave comparison given that the sample size for that wave is closest to that of the merged dataset. The only exception is the comparison between the March 2021 and October 2022 waves, where the earlier wave provides a better sample size match. In all cases, using weights from the other wave does not change the results.

Figure S4: Perceived prevalence of election fraud by party



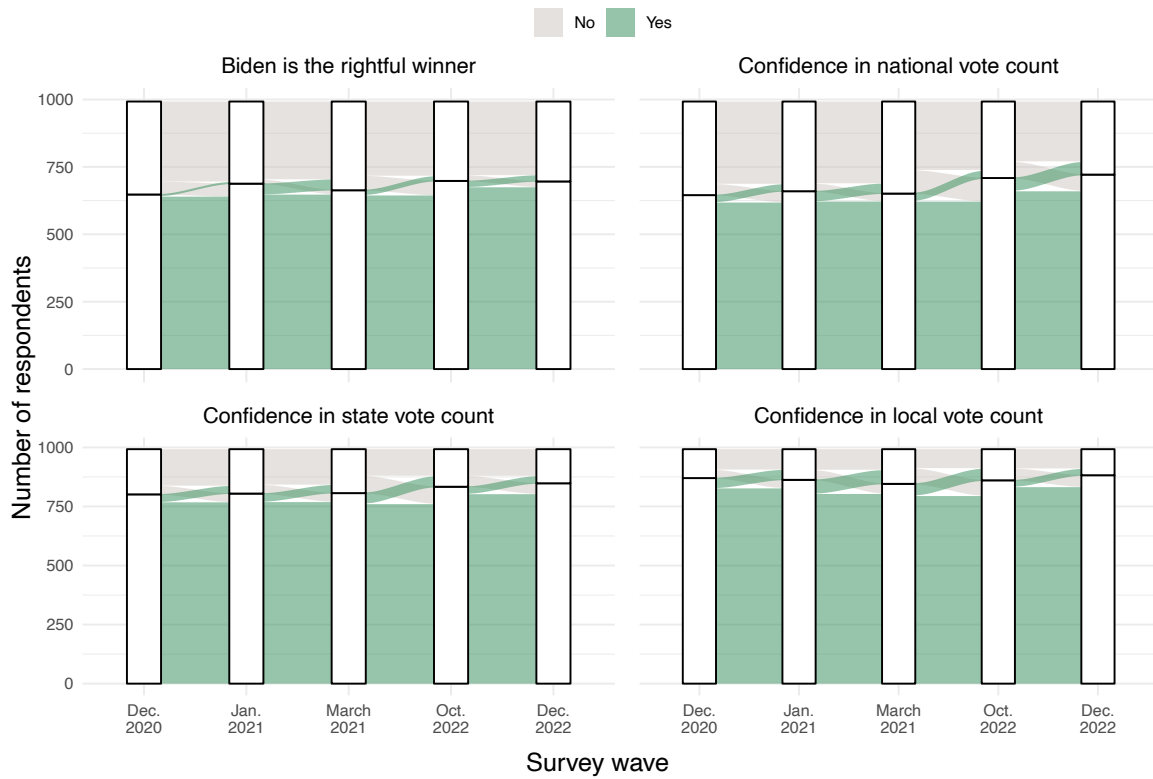
Percentage of respondents indicating that there are “Thousands” of cases or more of each type of election fraud. Categories included “Less than ten” (post-2020 election waves only), “Less than a hundred,” “Hundreds,” “Thousands,” “Tens of thousands,” “Hundreds of thousands,” and “A million or more”. Fraud prevalence measured for U.S. elections in general (Aug. 2020) and in the 2020 U.S. presidential election (subsequent waves). Group means computed using each wave’s post-stratification weights. Results are disaggregated by partisan identification as measured in August 2020 among participants in that wave (N=1,074).

Figure S5: Within-respondent durability in the perceived prevalence of election fraud



Weighted statistics based on post-stratification weights. Analysis restricted to participants who participated in all survey waves (N=1,074). Each survey wave is identified by the vertical bar, with the number of respondents in each wave answering *less than "thousands"* indicated by the horizontal line.

Figure S6: Within-respondent durability in perceived election legitimacy



Weighted statistics based on post-stratification weights. Analysis restricted to participants who participated in all survey waves (N=1,074). Each survey wave is identified by the vertical bar, with the number of respondents in each wave answering that Biden was “probably” or “definitely” the rightful winner or who are “somewhat” or “very” confident that votes were counted as intended indicated by the horizontal line.

S2.2 Preregistered experimental analyses

S2.2.1 Experimental treatment

Please read the following article carefully and take your time while reading. This page includes the first part of the article. The next page includes the second part of the article. We will ask you some questions about it afterward.

FACT CHECK: Trump’s claims of vote rigging are wrong

By HOPE YEN, ALI SWENSON and AMANDA SEITZ

Associated Press

Clinging to notions of widespread vote rigging that his own attorney general has disputed, President Donald Trump repeated a litany of baseless assertions Wednesday of political corruption, machine tampering and mysterious votes appearing out of nowhere that allowed Joe Biden to steal the election.

“This election is about great voter fraud, fraud that has never been seen like this before,” Trump said in a 46-minute address posted on social media.

“It’s about poll watchers who were not allowed to watch. So illegal. It’s about ballots that poured in and nobody but a few knew where they came from.... It’s about machinery that was defective, machinery that was stopped.”

None of those assertions are true.

A look at the claims and reality:

VOTER FRAUD

TRUMP: “You can’t let another person steal that election from you.”

THE FACTS: To be clear, no election was stolen from Trump. Others in his administration have already said the election was secure. Attorney General William Barr said the Justice Department had seen no evidence of widespread fraud to overturn Biden’s margin of victory. </p> <p>Trump’s allegations of massive voting fraud have been refuted by a variety of judges, state election officials and an arm of his own administration’s Homeland Security Department. Many of his campaign’s lawsuits across the country have been thrown out of court. And his administration has already agreed to allow the formal transition of power to Biden to begin.

BALLOT ‘DUMPING’

TRUMP: “I go from winning by a lot to losing a tight race. It’s corrupt.”

THE FACTS: No mass corruption happened. Trump is actually describing a legitimate vote counting process. Indeed, news organizations and officials had warned before the election that in-person votes, which tend to be counted more quickly, would likely favor the president, who had spent months warning his supporters to avoid mail-in voting and to vote in person either early or on Election Day. In addition mail-in ballots, which take longer to count, would favor Biden. That pattern was exacerbated by the fact that many states prohibited early counting of mail-in votes that arrived before Election Day.

(Article continues on next page. The forward button will appear momentarily.)

Please read the rest of this article carefully and take your time while reading. This page includes the second part of the article that you read on the previous page. We will ask you some questions about it afterward.

FACT CHECK: Trump’s claims of vote rigging are wrong (continued)

By HOPE YEN, ALI SWENSON and AMANDA SEITZ

Associated Press

DOMINION

TRUMP, on Dominion Voting Systems electoral software used in many states: “When you look at who’s running the company, who’s in charge, who owns it — which we don’t know — where are the votes counted — which we think are counted in foreign countries.”

THE FACTS: Servers that run Dominion software are in local election offices, not in foreign countries. Claims that the company has foreign servers or ties to Germany or Venezuela are false. According to a joint statement released by the federal Cybersecurity and Infrastructure Security Agency, “there is no evidence that any voting system deleted or lost votes, changed votes, or was in any way compromised” in the 2020 election.

‘DEAD’ VOTERS

TRUMP: “Dead people — and we have many examples — filled out ballots, made applications and then voted.”

THE FACTS: He’s repeating a false claim of dead people voting, particularly in Pennsylvania and Michigan. But there’s no evidence that this occurred, and officials in both states say the claims are unfounded. Experts told the AP that it is common for state voter rolls to include voters with birthdates that make them appear impossibly old, but these are usually explained by human error, software quirks or voter confidentiality issues.

POLL WATCHERS

TRUMP: “In Pennsylvania, large amounts of mail-in and absentee ballots were processed illegally and in secret in Philadelphia and Allegheny counties without our observers present. They were not allowed to be present....They were thrown out of the building.”

THE FACTS: He’s incorrect. Trump is wholly misrepresenting a court case in the state and what happened at voting places. No one tried to ban poll watchers representing each side in the election. Democrats did not try to stop Republican representatives from being able to observe the process. The main issue in the case was how close observers representing the parties could get to election workers who were processing mail-in ballots in Philadelphia. Trump’s lawyers even admitted in court that their campaign had observers in the room.

(The forward button will appear momentarily.)

S2.2.2 Experimental results

Regression tables showing the main effects of exposure to the fact-checking treatment on the average perceived truthfulness of targeted true and false statements and overall discernment, as well the perceived truthfulness of each individual statement, are included in Tables [S4](#) and [S5](#), respectively. Figure [S7](#) displays preregistered differences of means between the treatment conditions.

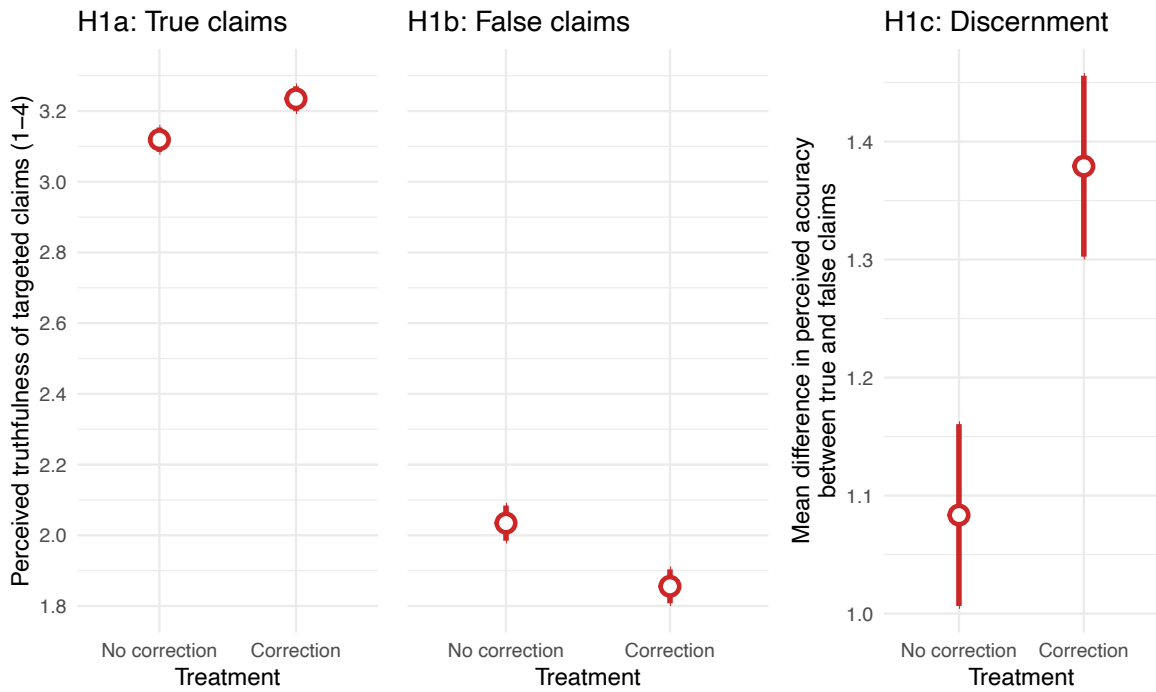
The regression tables report results of OLS models with robust standard errors, with pre-treatment control variables selected via lasso to improve precision (Bloniarz et al. 2016) from a preregistered list: pre-treatment outcome measures (Clifford et al. 2021), partisan identification, ideology, political interest, political knowledge, information trust, education, and race. Partisan identification was measured on a seven-point scale. The Democrat and Republican categories include partisan leaners. Ideology is measured using a seven-point scale from very liberal (1) to very conservative (7). Political interest is measured using a five-point scale from not at all interested (1) to extremely interested (5). Political knowledge represents the number of correct responses to 5 items (0-5): 1) how many years there are in one full term of office for a U.S. Senator, 2) the number of times an individual can be elected President of the United States under current laws, 3) the number of Senators from each state, 4) the name of the Prime Minister of the United Kingdom, and 5) how many years there are in one full term of office for a U.S. House member. Information trust is based on the average level of trust in national news organizations, local news organizations, political leaders in the federal government and political leaders in their state (each measured on four-point scales). Education is a dichotomous variable coded as 1 when respondents have a four-year degree or postgraduate education and 0 otherwise. Finally, race is coded as 0 if respondents identify as white and 1 if they identify Black, Hispanic, Asian, Native American, Mixed, or Other.

Table S4: Main treatment effects of exposure to fact-check article

	Targeted true	Targeted false	Discernment
Fact-check	0.125*** (0.017)	-0.202*** (0.015)	0.327*** (0.024)
Lasso controls	✓	✓	✓
Num.Obs.	3624	3784	3772

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, ideology, information trust, education (Targeted true model), and political knowledge (Targeted true model).

Figure S7: Combined treatment effects of exposure to fact-check article



Group means with 95% confidence intervals. t-test results: true claims: $t = 4.60$, $p < 0.001$; false claims: $t = -5.08$, $p < 0.001$; discernment: $t = 5.33$, $p < 0.001$.

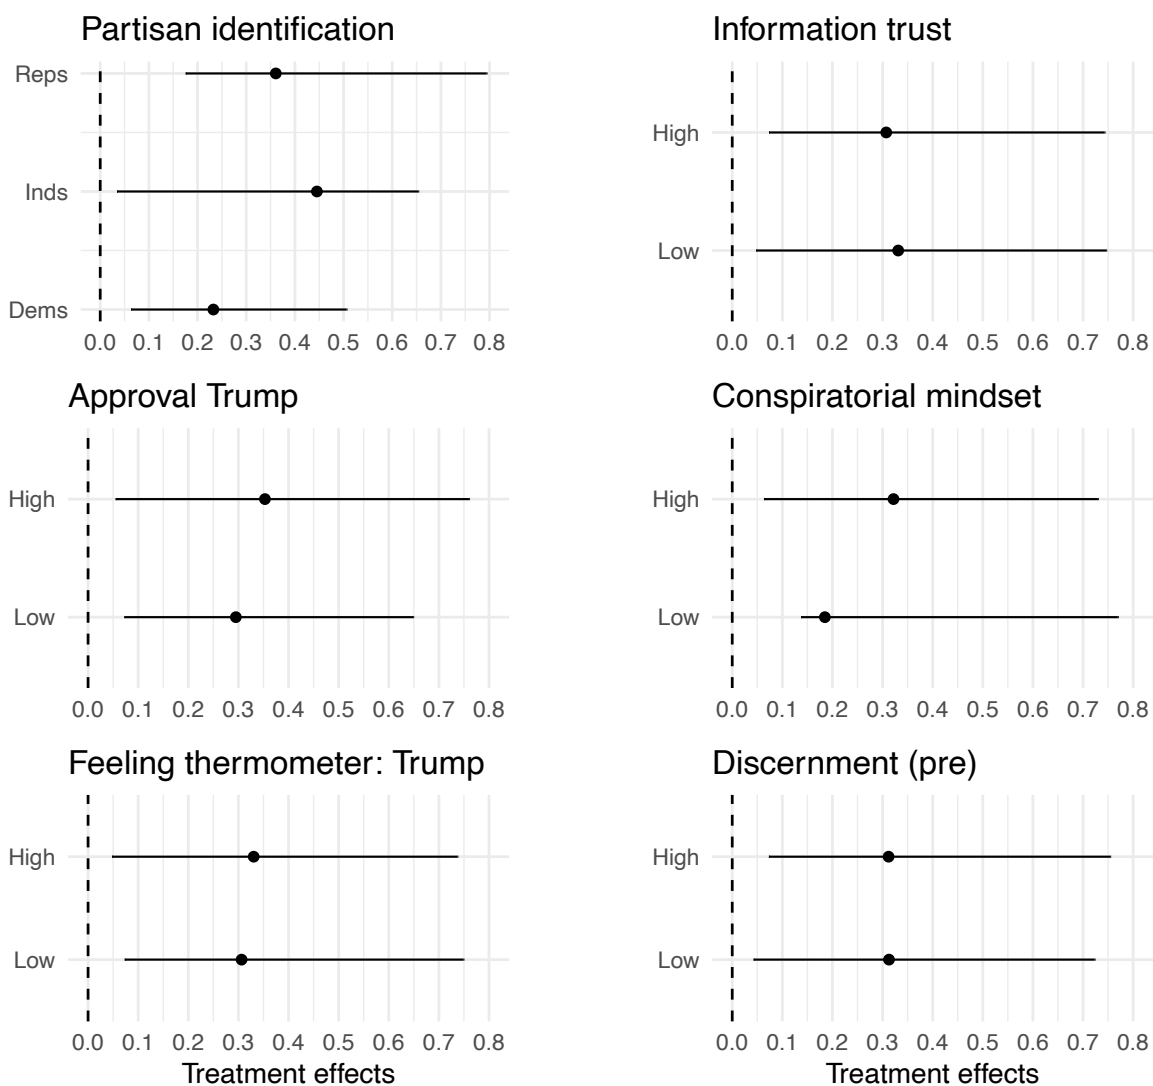
Table S5: Treatment effects of exposure to fact-check article on the perceived accuracy of targeted statements

	Early votes	Observers	Justice Dept.	Trump won	Dead people	Dominion
Fact-check	0.198*** (0.032)	0.159*** (0.020)	0.020 (0.024)	-0.064*** (0.017)	-0.174*** (0.021)	-0.375*** (0.024)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	3678	3786	3633	3789	3789	3634

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (all except Early votes), ideology (all except Justice Dept.), information trust, political knowledge (Early votes, Justice Dept., Dominion), and non-white (Early votes).

Figure S8 uses Bayesian Causal Forest models (`bcf` R package) to perform preregistered tests for heterogeneous treatment effects by partisanship, Trump approval, feelings towards Trump, trust in authoritative sources of information, conspiratorial thinking, and truth discernment in the previous survey wave. Subsection S2.1 details how partisan identification and Trump approval are measured. Feelings towards Donald Trump are measured using a feeling thermometer (0–100), with ratings greater or equal to 50 coded as “High” and ratings lower than 50 coded as “Low”. Conspiratorial predispositions are based on participants’ level of agreement on a five-point scale with four statements such as “Much of our lives are being controlled by plots hatched in secret places” (Uscinski et al. 2016). Responses were combined into an additive index and divided into “Low” and “High” categories using a median split. Finally, our pre-treatment measure of discernment relies on the same six targeted statements used for the outcome variable (truth discernment) as measured in the previous wave (Dec. 17, 2020–Jan. 5, 2021). Respondents with scores greater or equal to 0 on the -3 to 3 scale are coded as having high discernment, while those with a score below 0 are coded as having low discernment.


Figure S8: Heterogeneous treatment effects of exposure to fact-check on truth discernment



Heterogeneous treatment effects based on Bayesian Causal Forest models. Median effect by group with 95% credible intervals (2.5th and 97.5th percentiles), which can be asymmetric around the median.

We also conducted exploratory analyses examining whether fact-check exposure influences beliefs in non-targeted claims. The results show that fact-checks increase truth discernment for non-targeted claims about election fraud, although the effects are substantively smaller than for targeted claims and only significant for three claims out of six tested (Figure S9 and Tables S6 and S7).

Figure S9: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted claims



Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of non-targeted statements. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Outcomes measured on a four-point scale ranging from “not at all accurate” (1) to “very accurate” (4).

Table S6: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted claims

	Non-targeted true	Non-targeted false	Non-targeted discernment
Fact-check	0.046** (0.016)	-0.025 (0.014)	0.067** (0.021)
Lasso controls	✓	✓	✓
Num.Obs.	3619	3620	3606

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, ideology, information trust, and political knowledge.

Table S7: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted statements

	Mail voting	Obama margin	Legal challenges	Online voting	Biden won Florida	Illegal ballots
Fact-check	0.065** (0.023)	0.060* (0.027)	0.017 (0.022)	-0.007 (0.022)	-0.012 (0.027)	-0.051* (0.022)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	3629	3630	3685	3684	3688	3782

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (Mail voting, Obama margin, Illegal ballots), ideology (Mail voting, Obama margin, Illegal ballots), information trust (Mail voting, Obama margin, Illegal ballots), political knowledge (all except Illegal ballots), and political interest (Mail voting, Biden won Florida).

Finally, we assessed the impact of exposure to fact-checking about election fraud on general attitudes about the 2020 election, the vote counting process, and the January 6 insurrection.


2020 election attitudes were measured using the following six questions:

1. The result of the 2020 election (“Donald Trump won” coded as 1, “Joe Biden won,” “Winner not yet known,” and “Don’t know” coded as 0)
2. Who will be inaugurated as President (“Donald Trump” coded as 1, “Joe Biden” or “Someone else” coded as 0)
3. Whether Joe Biden is the rightful winner of the 2020 election (4-point scale from “Definitely not” to “Definitely”)
4. Whether Joe Biden’s victory is legitimate (4-point scale from “Not legitimate at all” to “Entirely legitimate”)
5. Whether Joe Biden’s victory was fair (4-point scale from “Not at all fair” to “Very fair”)
6. The likelihood that election fraud was involved in the outcome of the election (4-point scale from “Very unlikely” to “Very likely”).

Confidence in vote counting was measured using four questions asking participants how confident they are, on a four-point scale ranging from not at all confident to very confident, that a) their vote, b) votes in their local area, c) votes in their state, and d) votes nationwide were counted as intended in the 2020 election. Regarding January 6 attitudes, participants were first asked about what the most accurate description of the people who stormed the U.S. Capitol is (“People trying to make Trump look bad” coded as 1, “Trump supporters” and “Don’t know” coded as 0). They were then questioned about their approval of the actions of the people who stormed the U.S. Capitol (four-point scale), what they think of the actions of the people who stormed the U.S. Capitol (three-point scale including “They were mostly wrong”, “They went too far, but they had a point”, and “They were mostly right”), and whether they approve of lawmakers’ continued efforts to block the certification of Biden’s victory (four-point scale).

Results reported in Figure S10 and Tables S8, S9, and S10 suggest that exposure to fact-checks can contribute to making citizens more confident in election legitimacy and vote counting at the local, state, and national level. However, fact-checks did not affect perceptions that Trump would be inaugurated, support for the Capitol insurrection, or support for lawmakers’ attempts to block Biden’s inauguration. The results are somewhat different from those of Carey et al. (2025), as fact-checks in the current study did affect broader beliefs about election integrity. This difference could potentially be explained by the fact that Carey et al.’s study was conducted in 2022, when people’s beliefs about the 2020 election had become relatively stable (Graham and Yair 2024). However, in line with previous studies (Painter and Fernandes 2024), the findings demonstrate the limited impact of fact-checks on opinions about the January 6 insurrection.

Figure S10: Treatment effects of exposure to fact-check article on broader attitudes about the 2020 election



Sample average treatment effects of exposure to a fact-checking article on broader attitudes about the legitimacy of the 2020 election and the January 6 insurrection. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso.

Table S8: Treatment effects of exposure to fact-check article on the perceived legitimacy of the election

	Trump won	Trump inaugurated	Rightful	Legitimate	Fair	Fraud influence
Fact-check	-0.024** (0.008)	0.000 (0.006)	0.064*** (0.014)	0.063*** (0.014)	0.061*** (0.014)	-0.059** (0.018)
Lasso controls	✓	✓	✓	✓	✓	✓
N	3792	3846	3793	3796	3793	3792

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (all except Trump inaugurated), ideology (Trump won, Fraud influence), and information trust (Trump won, Rightful, Fair, Fraud influence).

Table S9: Treatment effects of exposure to fact-check article on confidence in vote counts

	Individual	Local	State	National
Fact-check	0.025 (0.019)	0.039* (0.018)	0.060*** (0.017)	0.114*** (0.017)
Lasso controls	✓	✓	✓	✓
N	3393	3794	3791	3792

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, ideology (State, National), and information trust.

Table S10: Treatment effects of exposure to fact-check article on attitudes about the January 6 insurrection

	Make Trump look bad	Approve insurrection	Insurrection right	Block certification
Fact-check	-0.014 (0.011)	0.007 (0.018)	0.011 (0.013)	0.029 (0.027)
Lasso controls	✓	✓	✓	✓
N	3792	3794	3793	3792


*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: partisan identification, ideology (Make Trump look bad, Block certification), information trust, political interest (Make Trump look bad, Block certification), and education (Make Trump look bad).

S2.3 Exposure to election fraud content

This subsection presents additional analyses about exposure to election fraud content online during the 2020 and 2022 study periods.

Figure S11 replicates Figure 3 using articles (i.e., unique URLs after deduplication) instead of views as the unit of analysis.

Figure S11: Percentage of articles mentioning election fraud by coverage type, source type, and year



Data disaggregated by source trustworthiness based on ratings from NewsGuard and Lin et al. (2023).

Tables S11 and S12 list the ten most-viewed domains for election fraud content during the 2020 election period, ranked by trustworthiness rating and whether the content challenges fraud claims, respectively. Tables S13 and S14 report the same statistics for the period surrounding the 2022 election. Tables S15 and S16 show the percentage of fraud-related articles that challenged election fraud claims across the 25 outlets with the highest number of fraud-related articles in the period around the 2020 and 2022 elections. Tables S17–S20 report the unweighted and weighted percentage of views of election fraud content that challenged fraud claims across the 25 outlets with the highest number of fraud-related views around the 2020 and 2022 elections.

Table S11: Most-viewed domains for election fraud content by trustworthiness rating in the period around the 2020 election (unweighted)

	Trustworthy	Views	%	Untrustworthy	Views	%
1	yahoo.com	9566	22.83	rawstory.com	1683	15.17
2	washingtonpost.com	5003	11.94	dailykos.com	930	8.38
3	msn.com	3508	8.37	theepochtimes.com	850	7.66
4	cnn.com	2705	6.46	alternet.org	759	6.84
5	nytimes.com	2528	6.03	westernjournal.com	722	6.51
6	huffpost.com	1814	4.33	thegatewaypundit.com	420	3.79
7	aol.com	1130	2.70	palmerreport.com	302	2.72
8	politico.com	1025	2.45	freerepublic.com	264	2.38
9	foxnews.com	978	2.33	dailymail.co.uk	237	2.14
10	thehill.com	739	1.76	breitbart.com	197	1.78
	Total (Top 10)	28,996	69.21	Total (Top 10)	6364	57.37

Table S12: Most-viewed domains for content challenging and not questioning election fraud claims in the period around the 2020 election (unweighted)

	Challenging	Views	%	Non-questioning	Views	%
1	yahoo.com	3577	17.94	yahoo.com	1870	14.71
2	washingtonpost.com	2504	12.56	washingtonpost.com	706	5.56
3	msn.com	1722	8.64	msn.com	508	4.00
4	cnn.com	1348	6.76	theepochtimes.com	433	3.41
5	nytimes.com	1223	6.13	rawstory.com	382	3.01
6	huffpost.com	792	3.97	westernjournal.com	378	2.97
7	rawstory.com	600	3.01	thegatewaypundit.com	372	2.93
8	aol.com	529	2.65	foxnews.com	364	2.86
9	politico.com	401	2.01	cnn.com	305	2.40
10	dailykos.com	344	1.73	nytimes.com	250	1.97
	Total (Top 10)	13,040	65.39	Total (Top 10)	5568	43.81

Table S13: Most-viewed domains for election fraud content by trustworthiness rating in the period around the 2022 election (unweighted)

	Trustworthy	Views	%	Untrustworthy	Views	%
1	yahoo.com	3265	26.57	rawstory.com	881	27.80
2	msn.com	1530	12.45	westernjournal.com	293	9.25
3	nytimes.com	1362	11.08	dailykos.com	265	8.36
4	cnn.com	742	6.04	alternet.org	153	4.83
5	washingtonpost.com	587	4.78	theepochtimes.com	129	4.07
6	aol.com	437	3.56	redstate.com	128	4.04
7	huffpost.com	397	3.23	welovetrump.com	110	3.47
8	apnews.com	228	1.86	thegatewaypundit.com	100	3.16
9	nationalreview.com	215	1.75	dailymail.co.uk	95	3.00
10	nbcnews.com	159	1.29	townhall.com	88	2.78
	Total (Top 10)	8922	72.60	Total (Top 10)	2242	70.75

Table S14: Most-viewed domains for content challenging and not questioning election fraud claims in the period around the 2022 election (unweighted)

	Challenging	Views	%	Non-questioning	Views	%
1	yahoo.com	838	25.78	yahoo.com	749	18.06
2	msn.com	341	10.49	msn.com	408	9.84
3	cnn.com	252	7.75	rawstory.com	274	6.61
4	rawstory.com	221	6.80	westernjournal.com	171	4.12
5	washingtonpost.com	164	5.04	cnn.com	143	3.45
6	aol.com	128	3.94	nytimes.com	103	2.48
7	nytimes.com	121	3.72	aol.com	95	2.29
8	huffpost.com	91	2.80	washingtonpost.com	93	2.24
9	apnews.com	84	2.58	welovetrump.com	87	2.10
10	dailykos.com	60	1.85	thegatewaypundit.com	86	2.07
	Total (Top 10)	2300	70.75	Total (Top 10)	2209	53.27

Table S15: Percentage of articles challenging fraud claims by outlet in the period around the 2020 election (top 25 outlets with largest number of articles about election fraud)

Domain	Articles	% fraud-challenging articles
yahoo.com	3066	39.1
msn.com	2773	48.3
washingtonpost.com	1412	46.2
cnn.com	1103	46.7
rawstory.com	928	37.7
nytimes.com	921	47.7
huffpost.com	654	42.7
politico.com	494	36.4
theepochtimes.com	489	15.3
thehill.com	482	34.0
foxnews.com	466	22.5
dailykos.com	465	38.7
aol.com	435	41.8
westernjournal.com	402	10.7
theguardian.com	358	36.3
alternet.org	357	42.9
usatoday.com	316	47.8
thegatewaypundit.com	309	2.6
apnews.com	269	52.8
thedailybeast.com	269	37.5
nationalreview.com	261	36.4
freerepublic.com	242	6.6
nbcnews.com	226	42.5
palmerreport.com	202	41.6
conservativeinstitute.org	190	10.5

Table S16: Percentage of articles challenging fraud claims by outlet in the period around the 2022 election (top 25 outlets with largest number of articles about election fraud)

Domain	Articles	% fraud-challenging articles
yahoo.com	1603	25.6
msn.com	1005	23.3
rawstory.com	548	25.2
nytimes.com	463	15.1
cnn.com	337	27.9
washingtonpost.com	316	28.8
aol.com	264	28.4
huffpost.com	243	24.7
westernjournal.com	178	9.0
apnews.com	157	34.4
dailykos.com	135	20.7
nationalreview.com	125	21.6
alternet.org	123	23.6
redstate.com	112	5.4
politico.com	96	26.0
foxnews.com	94	5.3
nbcnews.com	86	19.8
thehill.com	85	20.0
welovetrump.com	85	4.7
npr.org	83	39.8
dailymail.co.uk	82	25.6
theepochtimes.com	82	17.1
msnbc.com	80	35.0
thegatewaypundit.com	74	2.7
townhall.com	73	8.2

Table S17: Percentage of respondents exposed to election fraud content and percentage of views of fraud-challenging content by outlet in the period around the 2020 election (top 25 outlets with largest number of views of election fraud content, unweighted)

Domain	Views	% exposed	% fraud-challenging views
yahoo.com	9566	20.2	37.4
washingtonpost.com	5003	16.3	50.0
msn.com	3508	9.8	49.1
cnn.com	2705	17.7	49.9
nytimes.com	2528	16.1	48.4
huffpost.com	1814	7.6	43.7
rawstory.com	1683	3.5	35.7
aol.com	1130	3.2	46.8
politico.com	1025	11.2	39.2
foxnews.com	978	9.6	22.4
dailykos.com	930	4.1	37.3
theepochtimes.com	850	3.9	17.5
alternet.org	759	1.4	39.4
thehill.com	739	8.3	32.7
westernjournal.com	725	2.8	10.8
theguardian.com	556	6.3	37.3
nationalreview.com	541	1.6	40.0
usatoday.com	505	9.1	50.1
thedailybeast.com	460	6.6	37.6
forbes.com	423	5.5	47.9
thegatewaypundit.com	420	2.9	2.4
apnews.com	390	8.2	54.6
nbcnews.com	377	8.5	42.2
vox.com	345	4.3	47.2
palmerreport.com	302	0.7	42.1

Table S18: Percentage of respondents exposed to election fraud content and percentage of views of fraud-challenging content by outlet in the period around the 2020 election (top 25 outlets with largest number of views of election fraud content, weighted)

Domain	Views	% exposed	% fraud-challenging views
yahoo.com	5659	14.6	36.4
washingtonpost.com	4753	13.1	51.0
msn.com	2119	7.4	48.7
cnn.com	1530	15.5	48.4
nytimes.com	1343	12.6	48.9
huffpost.com	804	4.3	43.2
aol.com	733	2.4	46.6
westernjournal.com	726	2.6	11.1
politico.com	717	9.1	39.3
foxnews.com	687	7.9	23.3
rawstory.com	671	2.1	38.0
theepochtimes.com	660	3.6	19.1
thehill.com	480	6.6	35.2
thegatewaypundit.com	427	3.0	2.0
usatoday.com	379	7.5	50.1
dailykos.com	373	2.4	37.2
nationalreview.com	341	1.2	41.1
apnews.com	331	6.5	53.5
nbcnews.com	264	6.6	48.3
freerepublic.com	256	0.3	6.2
alternet.org	236	0.7	42.1
forbes.com	225	3.9	43.6
palmerreport.com	222	0.4	41.5
theguardian.com	222	3.8	37.0
twitchy.com	214	0.5	9.5


Table S19: Percentage of respondents exposed to election fraud content and percentage of views of fraud-challenging content by outlet in the period around the 2022 election (top 25 outlets with largest number of views of election fraud content, unweighted)

Domain	Views	% exposed	% fraud-challenging views
yahoo.com	3265	9.4	25.7
msn.com	1530	6.9	22.3
nytimes.com	1362	8.7	8.9
rawstory.com	881	2.0	25.1
cnn.com	742	8.8	34.1
washingtonpost.com	587	5.0	27.9
aol.com	437	1.6	29.3
huffpost.com	397	3.4	22.9
westernjournal.com	293	1.4	9.6
dailykos.com	265	1.6	25.5
apnews.com	228	3.1	36.8
nationalreview.com	215	0.5	22.0
nbcnews.com	159	4.5	12.6
alternet.org	153	0.8	24.8
foxnews.com	148	2.4	4.2
npr.org	137	3.6	37.2
politico.com	136	3.1	25.0
theepochtimes.com	129	1.5	15.5
redstate.com	128	0.7	5.5
welovetrump.com	110	0.5	4.5
thehill.com	108	2.4	20.4
politicalwire.com	107	0.3	3.3
thegatewaypundit.com	100	1.1	2.0
msnbc.com	96	2.0	29.2
dailymail.co.uk	95	1.1	28.4

Table S20: Percentage of respondents exposed to election fraud content and percentage of views of fraud-challenging content by outlet in the period around the 2022 election (top 25 outlets with largest number of views of election fraud content, weighted)

Domain	Views	% exposed	% fraud-challenging views
yahoo.com	2480	8.2	23.2
aol.com	1435	1.2	28.0
nytimes.com	1209	7.8	7.8
msn.com	793	5.1	21.1
cnn.com	611	7.2	39.3
rawstory.com	592	1.4	22.2
washingtonpost.com	388	4.3	26.1
westernjournal.com	253	1.3	10.4
huffpost.com	238	1.9	23.1
apnews.com	210	2.5	37.2
redstate.com	203	0.5	4.5
theepochtimes.com	189	2.0	13.9
nbcnews.com	177	4.7	10.7
dailykos.com	164	0.8	25.9
nationalreview.com	135	0.2	22.5
foxnews.com	132	1.8	4.2
npr.org	129	3.5	33.3
townhall.com	127	0.7	7.7
welovetrump.com	118	0.5	5.2
politico.com	116	2.8	22.6
twitchy.com	96	0.3	3.8
fivethirtyeight.com	95	2.3	7.5
thegatewaypundit.com	88	0.9	1.2
thehill.com	83	2.1	23.5
msnbc.com	72	1.7	25.5

Figure S12: Differences in fraud-related information exposure by candidate support in the period around the 2022 election



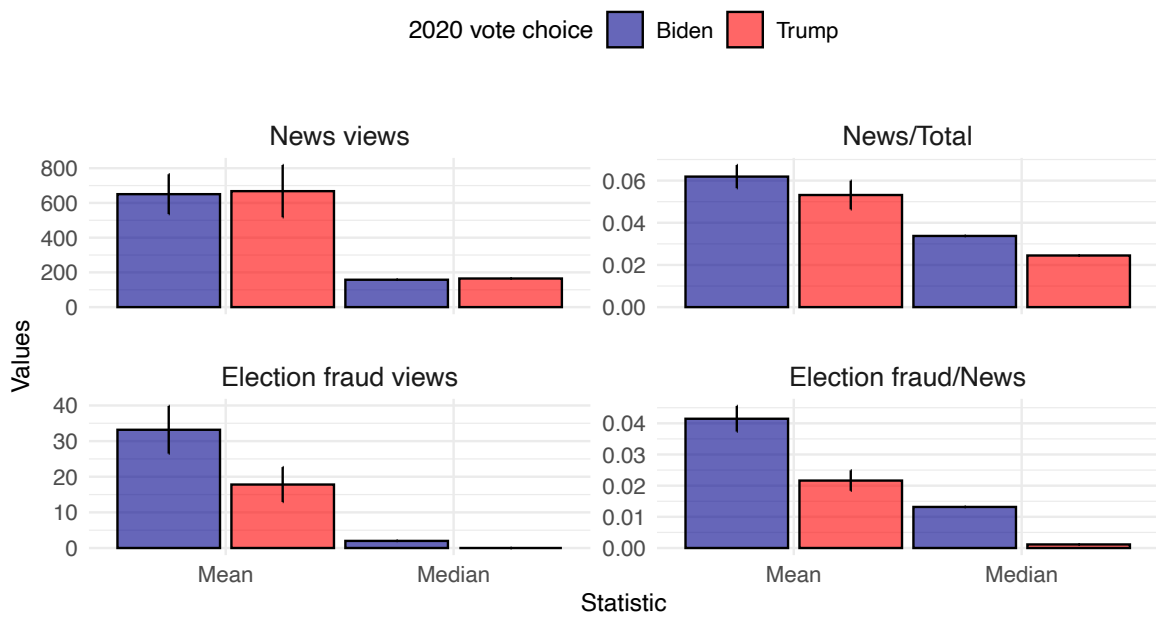
Weighted statistics based on post-stratification weights. The corresponding results for the 2020 election period are reported in Figure 4. Totals may not sum to 100% due to rounding.

Figure S12 replicates Figure 4 from the main text using data from the 2022 election period.

Figures S13 and S14 show the mean and median number and proportion of views of news and election fraud content in the period surrounding the 2020 and 2022 elections, respectively.


Figure S15 replicates Figure 5 in the main text adding exposure to fact-checking, while Figure S16 instead replicates Figure 5 conditional on having been exposed to election fraud content. To identify fraud-related fact-check articles, we relied on lists of fact-checkers compiled by Poynter’s International Fact-Checking Network (IFCN) and the Duke Reporters Lab. We included all U.S. fact-checking sites. When the website is only or primarily doing fact-checking (e.g., FactCheck.org, Politifact.com, Snopes.com), we coded all URLs from these websites as fact-check articles. For broader media organizations (e.g., Reuters, CNN, ABC), we identified sections of their website devoted to fact-checking (subdomains or subdirectories) and labeled all content from these sections as fact-check articles. Finally, given that fact-check articles are not always identified as such on news websites, we used the description of each fact-checking initiative on Duke Reporters’ Lab website and manual searches on each news website to build a comprehensive list of the keywords used in fact-check article URLs (e.g., reality-check, trust-index, fact-brief). We then labeled all content coming from listed fact-checking initiatives where the URL contained at least one of the fact-checking keywords as a fact-checking article. The full list of keywords and descriptive statistics are provided in SI Section S4.10.

Figure S13: Mean and median exposure to news content and election fraud content in the period around the 2020 election by candidate support



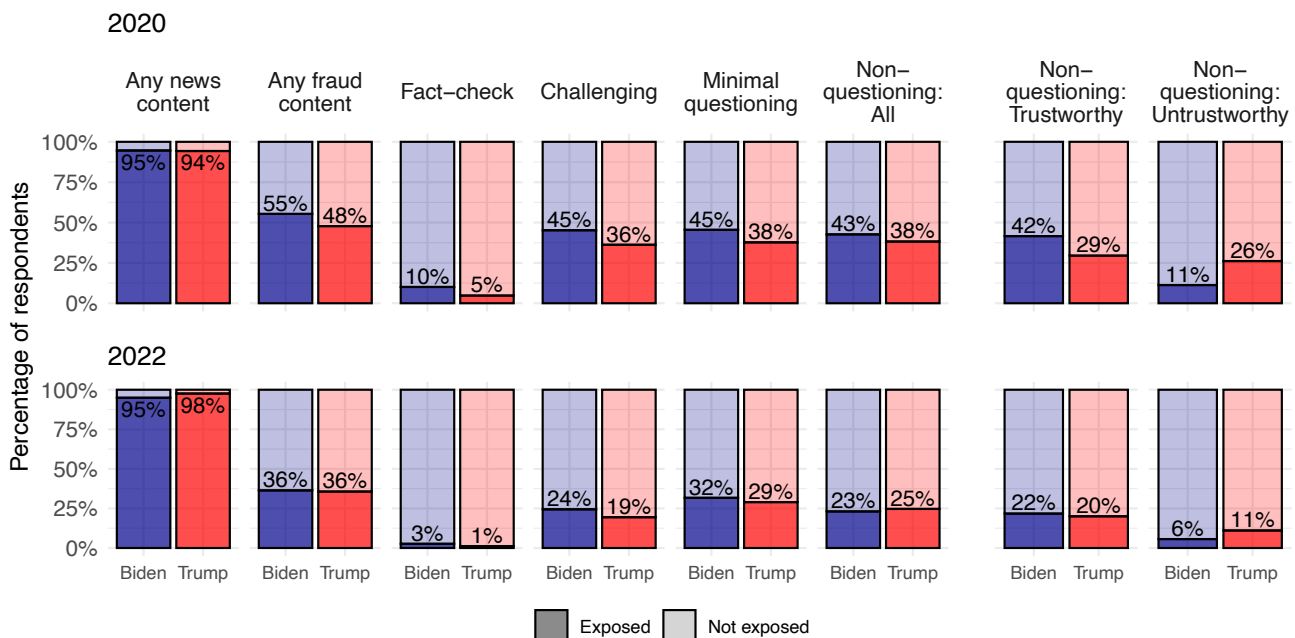
Weighted percentages based on post-stratification weights.

Figure S14: Mean and median exposure to news content and election fraud content in the period around the 2022 election by candidate support



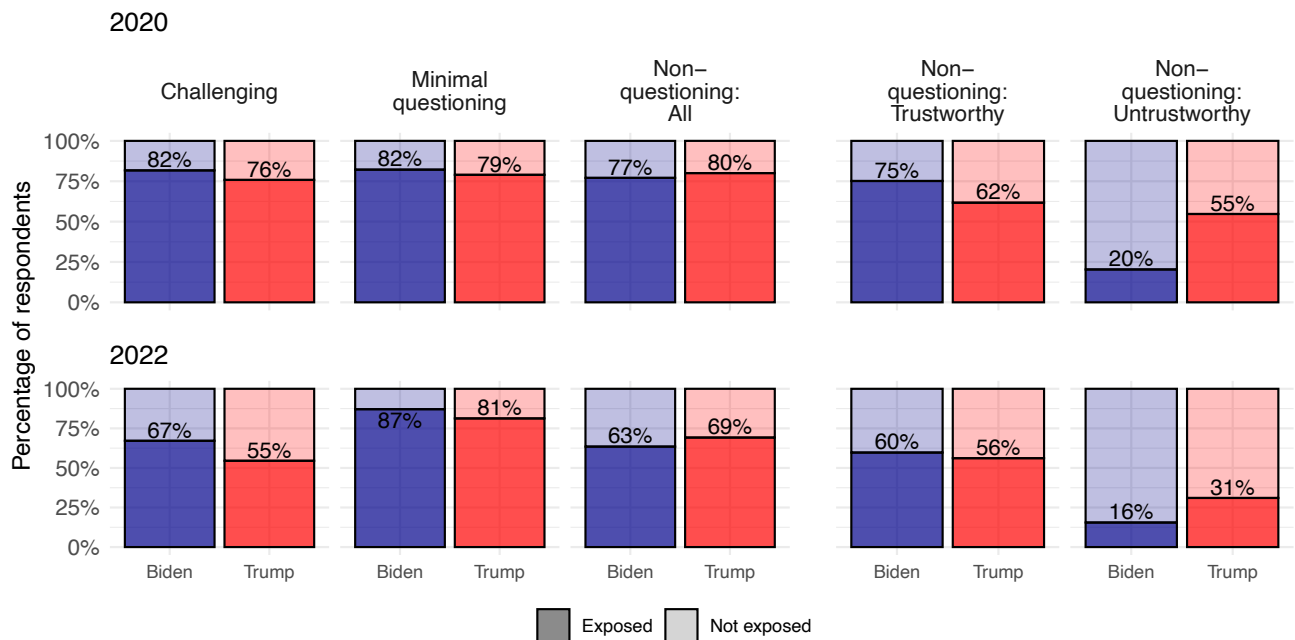
Weighted percentages based on post-stratification weights.

Figure S15: Individual exposure to different types of content



Percentage of respondents who were exposed to each type of content at least once by candidate support in the 2020 and 2022 elections; all values calculated based on post-stratification weights.

Figure S16: Individual exposure to different types of content conditional on exposure to election fraud content



Percentage of respondents who were exposed to each type of content at least once by candidate support in the 2020 and 2022 elections; all values calculated based on post-stratification weights.

Table S21: Characteristics of respondents exposed and not exposed to election fraud content in the period around the 2020 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	20.2	25.5	16.1	6.6
	35-44 years	23.3	20.6	12.9	14.4
	45-54 years	11.5	12.9	11.6	15.8
	55-64 years	25.3	17.6	28.5	29.5
	65+ years	19.7	23.3	30.9	33.7
Gender	Female	59.0	51.2	43.3	49.0
	Male	41.0	48.8	56.7	51.0
Education	College	29.1	35.4	22.5	28.5
	No college	70.9	64.6	77.5	71.5
Race	Non-white	49.4	35.6	14.1	16.0
	White	50.3	64.4	85.9	83.3
Region	Midwest	19.2	26.6	25.8	25.0
	Northeast	24.4	12.1	19.3	14.7
	South	39.6	35.3	35.1	42.0
	West	16.8	25.9	19.8	18.3
Ideology	Left/Moderate	89.3	95.0	24.9	21.0
	Right	10.7	5.0	75.1	79.0
Trump approval	Approve	6.5	0.7	95.3	94.0
	Disapprove	93.5	99.3	4.7	6.0
Feelings toward Trump	Negative	91.5	95.7	6.2	6.1
	Positive	6.8	1.5	93.2	93.8
Political interest	High	28.8	35.2	30.8	35.1
	Low	71.2	64.8	69.2	64.9
Information trust	High	86.0	85.9	35.5	35.0
	Low	14.0	14.1	64.2	65.0
Conspiracism	High	81.0	71.3	96.0	95.3
	Low	19.0	28.7	4.0	4.6
Truth discernment	High	95.6	99.7	29.5	23.6
	Low	3.9	0.3	70.2	76.4
Biden rightful winner	No	0.4	0.2	80.6	83.3
	Yes	99.6	99.8	19.4	16.7

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Table S22: Characteristics of respondents exposed and not exposed to challenging content in the period around the 2020 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	22.3	24.1	15.4	4.7
	35-44 years	22.0	21.6	12.0	16.6
	45-54 years	11.0	13.9	11.5	17.4
	55-64 years	23.6	17.9	30.8	25.8
	65+ years	21.0	22.5	30.3	35.6
Gender	Female	58.5	50.1	44.4	48.8
	Male	41.5	49.9	55.6	51.2
Education	College	30.2	35.4	23.2	29.1
	No college	69.8	64.6	76.8	70.9
Race	Non-white	49.1	32.9	15.5	14.2
	White	50.6	67.1	84.5	84.9
Region	Midwest	19.3	28.2	24.5	27.0
	Northeast	22.1	12.1	18.3	14.9
	South	39.6	34.4	37.5	40.0
	West	19.0	25.4	19.7	18.1
Ideology	Left/Moderate	91.1	94.1	26.4	17.1
	Right	8.9	5.9	73.6	82.9
Trump approval	Approve	5.3	0.9	94.2	95.6
	Disapprove	94.7	99.1	5.8	4.4
Feelings toward Trump	Negative	92.9	95.0	7.1	4.6
	Positive	5.6	1.9	92.5	95.4
Political interest	High	31.0	34.0	31.2	35.8
	Low	69.0	66.0	68.8	64.2
Information trust	High	86.3	85.5	35.1	35.6
	Low	13.7	14.5	64.7	64.4
Conspiracism	High	79.4	71.1	95.9	95.3
	Low	20.6	28.9	4.0	4.7
Truth discernment	High	96.2	99.8	29.7	21.3
	Low	3.3	0.2	70.0	78.7
Biden rightful winner	No	0.4	0.2	80.1	85.0
	Yes	99.6	99.8	19.9	15.0

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Table S23: Characteristics of respondents exposed and not exposed to non-questioning content in the period around the 2020 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	11.5	9.5	6.1	4.4
	35-44 years	15.1	14.8	10.7	8.0
	45-54 years	14.5	14.8	14.6	13.3
	55-64 years	30.5	23.1	28.8	33.8
	65+ years	28.4	37.8	39.8	40.4
Gender	Female	58.4	50.6	45.0	46.7
	Male	41.6	49.4	55.0	53.3
Education	College	43.3	47.9	26.5	37.3
	No college	56.7	52.1	73.5	62.7
Race	Non-white	31.1	20.1	12.0	13.3
	White	68.7	79.9	87.7	86.7
Region	Midwest	22.0	27.1	24.3	28.0
	Northeast	20.0	13.5	18.4	16.4
	South	32.8	34.1	36.9	35.1
	West	25.2	25.3	20.4	20.4
Ideology	Left/Moderate	95.1	96.5	23.0	19.6
	Right	4.9	3.5	77.0	80.4
Trump approval	Approve	2.1	1.0	95.8	94.2
	Disapprove	97.9	99.0	4.2	5.8
Feelings toward Trump	Negative	96.2	96.7	6.8	5.3
	Positive	2.6	1.8	92.6	94.2
Political interest	High	33.7	44.9	28.5	39.1
	Low	66.3	55.1	71.5	60.9
Information trust	High	85.3	89.0	36.9	34.2
	Low	14.7	11.0	62.8	65.8
Conspiracism	High	72.5	69.4	94.5	93.8
	Low	27.5	30.6	5.5	5.8
Truth discernment	High	98.3	99.5	28.5	19.6
	Low	1.3	0.5	71.2	80.4
Biden rightful winner	No	0.2	0.3	78.3	85.8
	Yes	99.8	99.7	21.7	14.2

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Table S24: Characteristics of respondents exposed and not exposed to election fraud content in the period around the 2022 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	34.7	23.2	19.5	25.3
	35-44 years	23.5	20.4	24.8	13.4
	45-54 years	10.6	16.0	14.2	11.0
	55-64 years	17.5	21.6	21.8	22.7
	65+ years	13.7	18.8	19.7	27.6
Gender	Female	58.9	46.1	48.0	39.0
	Male	41.1	53.9	52.0	61.0
Education	College	29.2	48.3	25.0	29.2
	No college	70.8	51.7	75.0	70.8
Race	Non-white	51.8	38.1	19.3	15.6
	White	48.2	61.9	80.7	84.4
Region	Midwest	18.2	21.5	19.6	20.9
	Northeast	12.6	15.9	14.0	14.6
	South	45.9	40.5	45.3	41.3
	West	23.3	22.1	21.2	23.2
Ideology	Left/Moderate	94.4	98.7	27.1	16.7
	Right	5.6	1.3	72.9	83.3
Biden approval	Approve	86.7	94.1	7.5	4.7
	Disapprove	13.3	5.9	92.5	95.3
Feelings toward Trump	Negative	92.7	96.7	21.5	22.8
	Positive	5.1	2.6	77.7	77.0
Political interest	High	23.5	46.2	27.9	37.9
	Low	76.5	53.8	72.1	62.1
Information trust	High	79.8	83.4	44.5	38.5
	Low	20.2	16.5	55.5	61.5
Conspiracism	High	78.0	66.8	93.2	86.2
	Low	21.8	33.2	6.7	13.8
Biden rightful winner	No	1.1	1.8	63.0	61.8
	Yes	98.9	98.2	37.0	38.2

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Table S25: Characteristics of respondents exposed and not exposed to challenging content in the period around the 2022 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	33.2	22.1	19.4	30.4
	35-44 years	22.9	20.7	22.4	14.1
	45-54 years	11.6	15.5	14.6	6.8
	55-64 years	18.3	21.2	23.0	18.6
	65+ years	14.0	20.4	20.7	30.2
Gender	Female	58.2	41.9	47.4	34.2
	Male	41.8	58.1	52.6	65.8
Education	College	32.5	47.4	26.0	28.7
	No college	67.5	52.6	74.0	71.3
Race	Non-white	49.9	37.3	18.1	17.7
	White	50.1	62.7	81.9	82.3
Region	Midwest	18.1	23.5	19.1	23.7
	Northeast	13.7	14.0	14.7	12.2
	South	44.9	41.0	46.9	31.0
	West	23.3	21.5	19.2	33.1
Ideology	Left/Moderate	94.9	99.1	25.8	13.4
	Right	5.1	0.9	74.2	86.6
Biden approval	Approve	88.2	93.1	6.0	8.7
	Disapprove	11.8	6.9	94.0	91.3
Feelings toward Trump	Negative	93.0	97.4	22.0	21.6
	Positive	4.8	2.3	77.2	78.4
Political interest	High	26.5	48.0	28.2	44.9
	Low	73.5	52.0	71.8	55.1
Information trust	High	80.9	81.6	43.8	36.5
	Low	19.0	18.4	56.2	63.5
Conspiracism	High	76.4	66.3	91.7	86.5
	Low	23.5	33.7	8.2	13.5
Biden rightful winner	No	1.4	1.1	62.4	63.0
	Yes	98.6	98.9	37.6	37.0

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

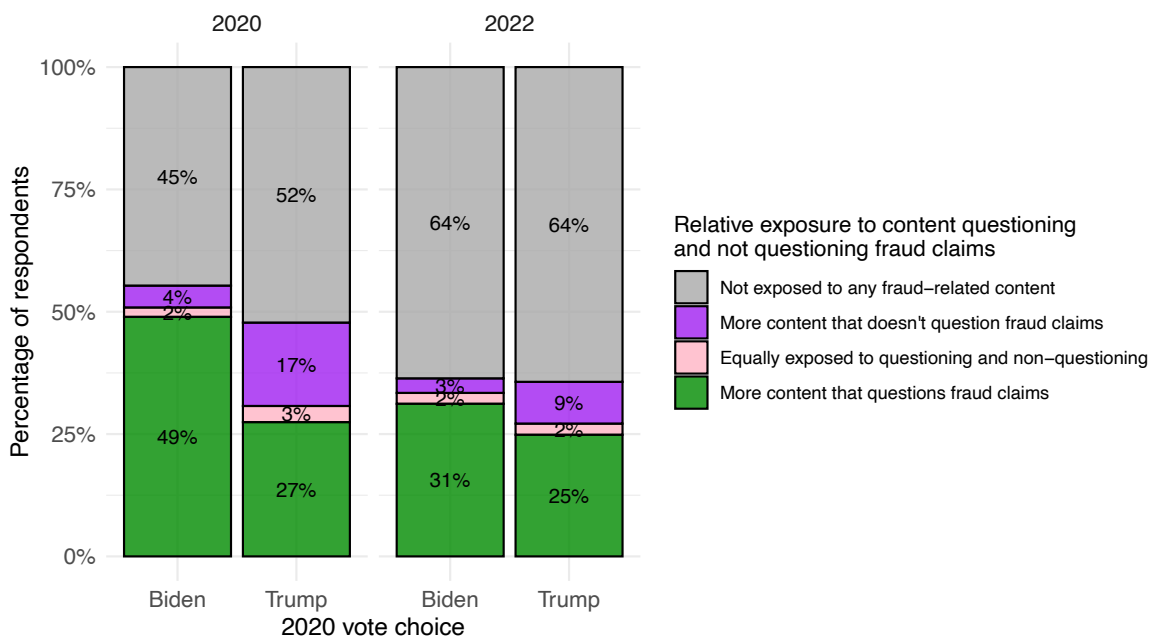
Table S26: Characteristics of respondents exposed and not exposed to non-questioning content in the period around the 2022 election

Variable	Value	Biden not exposed	Biden exposed	Trump not exposed	Trump exposed
Age	18-34 years	34.9	16.0	21.3	22.4
	35-44 years	23.3	19.1	23.0	13.8
	45-54 years	11.8	15.0	13.3	12.4
	55-64 years	16.8	26.3	22.5	20.9
	65+ years	13.2	23.5	19.9	30.5
Gender	Female	56.3	47.3	45.8	41.9
	Male	43.7	52.7	54.2	58.1
Education	College	31.6	51.3	25.2	30.5
	No college	68.4	48.7	74.8	69.5
Race	Non-white	51.8	30.3	18.7	15.8
	White	48.2	69.7	81.3	84.2
Region	Midwest	17.9	24.4	20.2	19.7
	Northeast	13.4	15.1	14.4	13.6
	South	45.9	37.3	45.2	39.8
	West	22.8	23.2	20.3	26.8
Ideology	Left/Moderate	95.1	98.8	25.9	15.7
	Right	4.9	1.2	74.1	84.3
Biden approval	Approve	88.0	94.1	6.7	6.0
	Disapprove	12.0	5.9	93.3	94.0
Feelings toward Trump	Negative	93.2	97.3	23.6	16.9
	Positive	4.8	2.1	75.7	82.7
Political interest	High	26.8	48.3	26.7	45.9
	Low	73.2	51.7	73.3	54.1
Information trust	High	79.2	87.3	45.6	32.5
	Low	20.7	12.7	54.4	67.5
Conspiracism	High	78.2	59.7	91.3	88.7
	Low	21.7	40.3	8.5	11.3
Biden rightful winner	No	1.5	0.7	59.6	71.6
	Yes	98.5	99.3	40.4	28.4

Participants are YouGov Pulse panel members active for the majority of months during the study period. Percentages calculated using post-stratification weights.

Figure S17 replicates Figure 6 in the main text but instead measures exposure to any questioning of election fraud claims by combining views of content that minimally questioned election fraud claims with views of content that meaningfully challenged these claims with facts and evidence.

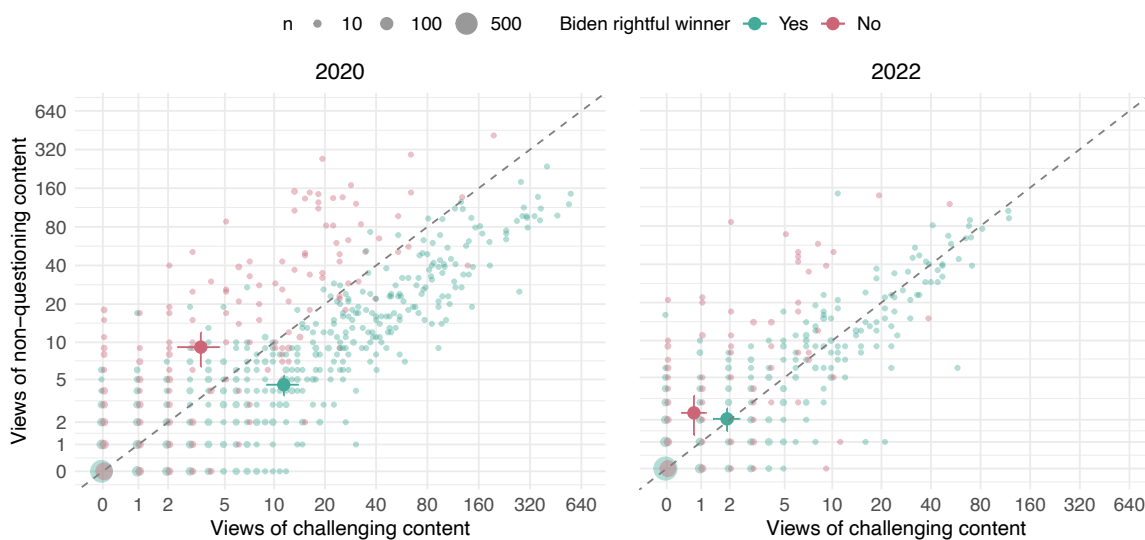
Figure S17: Individual relative exposure to different types of election fraud content



Weighted percentages based on post-stratification weights.

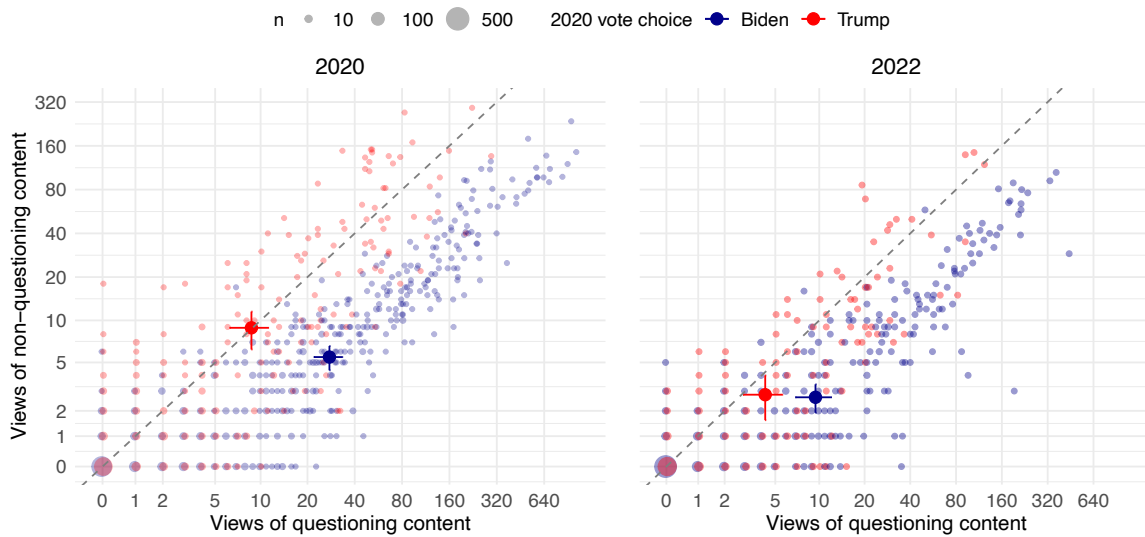
Figure S18 reproduces Figure 7 in the main text but groups respondents based on their belief about whether Joe Biden is the rightful winner of the 2020 U.S. presidential election rather than by candidate preference. Figure S19 replicates Figure 7 in the main text but groups views of content that minimally question election fraud claims with views of content that meaningfully challenged these claims with facts and evidence.

Figure S18: Views of questioning and non-questioning fraud content by election year and belief that Joe Biden was the rightful winner of the 2020 U.S. presidential election



Highlighted markers represent mean values, with 95% confidence intervals, among those who believe and do not believe that Biden was the rightful winner of the 2020 election. Axes use a log scale (with zero views transformed to an infinitesimally small number). The dashed line represents equal exposure to both types of content.

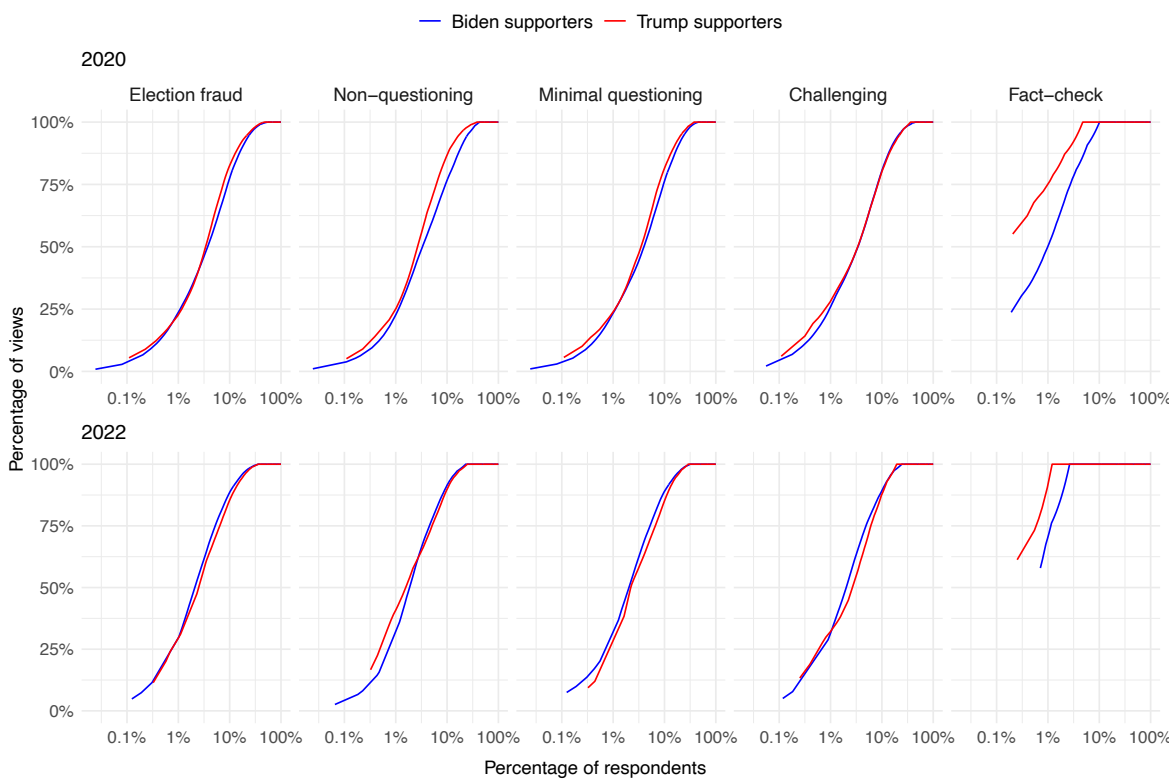
Figure S19: Views of questioning and non-questioning fraud content by election year and candidate support



Highlighted markers represent mean values, with 95% confidence intervals, for Trump and Biden supporters. Axes use a log scale (with zero views transformed to an infinitesimally small number). The dashed line represents equal exposure to both types of content.

To illustrate the extent of the skew we observe in fraud-related content exposure, Figure S20 shows how concentrated overall election fraud exposure and specific types of exposure were in the 2020 and 2022 elections. In 2020, 1% of participants were responsible for 27% of election-fraud views, 28% of non-questioning views, 27% of minimally-questioning views, 32% of challenging views, and 64% of views of fact-checking content. Additionally, more than 80% of views across all types of content originated from less than 10% of participants. The distribution was even more concentrated in 2022, where the top percentile accounted for 38% of election-fraud views, 41% of non-questioning views, 38% of minimally-questioning views, 42% of challenging views, and 85% of views of fact-checking content. Views were similarly concentrated among Biden and Trump supporters except for exposure to fact-checking content. In 2020, 1% of Biden supporters were responsible for about a third (35%) of fact-checking views compared to two thirds (68%) among Trump supporters. Figure S21 replicates the empirical cumulative distribution functions included in Figure S20 without post-stratification weights.

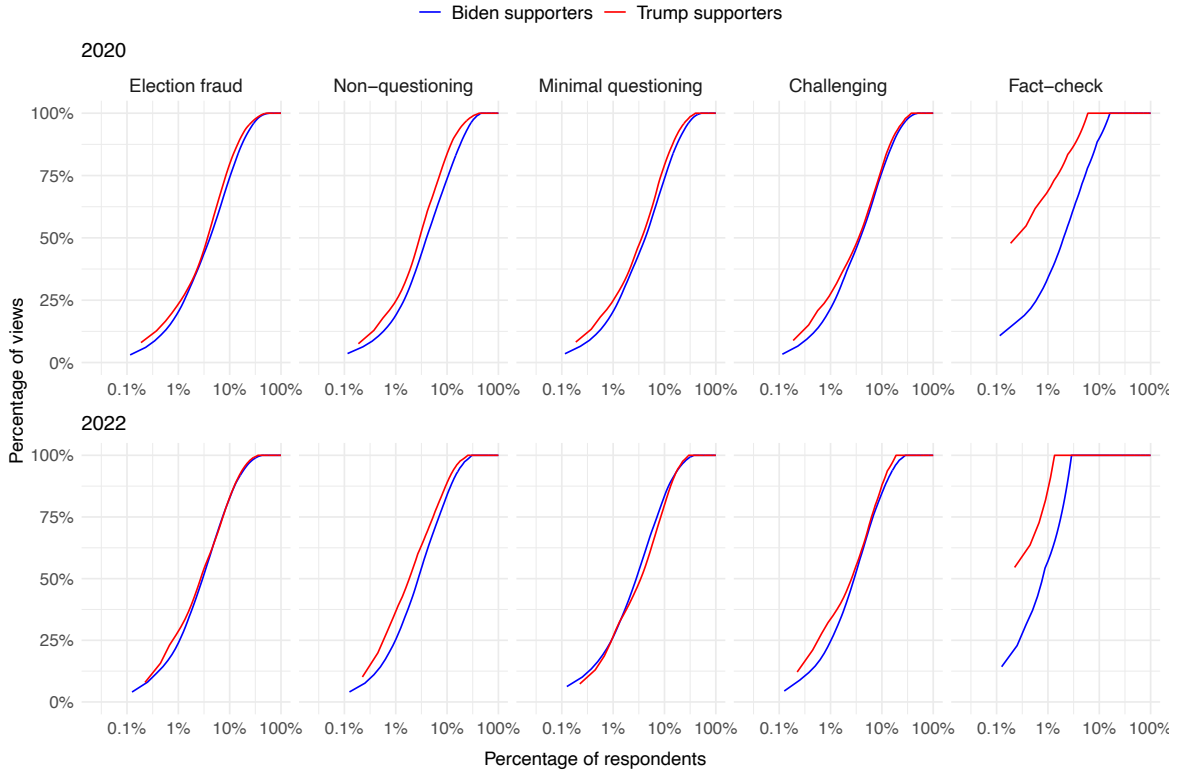
Figure S20: Empirical cumulative distribution functions of exposure to election fraud content by election and candidate support



Weighted cumulative distribution functions based on post-stratification weights. The x-axis represents the weighted percentage of participants responsible for a given percentage (y-axis) of all exposures. The lines have different starting points because the participant(s) with the most views do not have the same post-stratification weights and, consequently, do not account for the same weighted top percentile (additionally, more than one participant share the highest number of views of fact-checking in 2022). Unweighted CDFs are reported in Figure S21.

Figures S22, S24, and S25 replicate Figures 8, 10, and 11 using data from the 2022 rather than

Figure S21: Empirical cumulative distribution functions of exposure to election fraud content by election and candidate support




The x-axis represents percentage of participants responsible for a given percentage (y-axis) of all exposures. The lines have different starting points because the sample does not contain the same number of Trump and Biden supporters.

2020 election period. Figures S26–S29 reproduce Figures 10, 11, S24, and S25, looking at the percentage of views that include any questioning of election fraud claims (rather than views of content that meaningfully challenge fraud claims with facts and evidence).


Figure S30 replicates the permutation analysis in Figure 12 from the main text using the 2022 election study period data.

Figure S22: Views of news and election fraud content by candidate support (2022)




Weighted means and medians calculated using post-stratification weights, with 95% intervals around the means. Analysis restricted to 2022 election data (see Figure 8 for analysis based on 2020 election data).

Figure S23: Source trustworthiness of election fraud views by candidate support (2022)




Weighted means calculated using post-stratification weights. Analysis restricted to 2022 election data (see Figure 9 for analysis based on 2020 election data).

Figure S24: Share of challenging views by outlet trustworthiness and candidate support (2022)




Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of fraud-related views within each group. Labels identify the top 10 outlets per group. Point size indicates the total number of fraud-related views by outlet. The x-axis measures outlet trustworthiness based on ratings from NewsGuard and Lin et al. (2023). The dashed vertical line represents the threshold between sources classified as trustworthy and those classified as untrustworthy. Analysis restricted to 2022 election data (see Figure 10 for analysis based on 2020 election data).

Figure S25: Comparing challenging content exposure rates within outlet by candidate support (2022)




Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of election-fraud-related views within each group. Labels identify the top 10 outlets per group (red for Trump supporters, blue for Biden supporters, black for both). Points below the 45-degree line indicate that, for a given outlet, Biden supporters encountered a larger share of content challenging fraud claims than Trump supporters. Points above the line indicate the opposite. Point size indicates the total number of election fraud content views by outlet. The blue-red gradient of the points indicates the proportion of fraud-related views from that outlet coming from Trump supporters. Analysis restricted to 2022 election data (see Figure 11 for analysis based on 2020 election data).

Figure S26: Share of questioning views by outlet trustworthiness and candidate support (2020)




Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of election-fraud-related views within each group. Labels identify the top 10 outlets per group. Point size indicates the total number of fraud-related views by outlet. The x-axis measures outlet trustworthiness based on ratings from NewsGuard and Lin et al. (2023). The dashed vertical line represents the threshold between sources classified as trustworthy and those classified as untrustworthy. Analysis restricted to 2020 election data (see Figure S28 for analysis based on 2022 election data).

Figure S27: Comparing questioning content exposure rates within outlet by candidate support (2020)




Weighted percentages based on post-stratification weights. Restricted to the 50 outlets with the highest number of election-fraud-related views within each group. Labels identify the top 10 outlets per group (red for Trump supporters, blue for Biden supporters, black for both). Points below the 45-degree line indicate that, for a given outlet, Biden supporters encountered a larger share of content questioning fraud claims than Trump supporters. Points above the line indicate the opposite. Point size indicates the total number of election fraud content views by outlet. The blue-red gradient of the points indicates the proportion of fraud-related views from that outlet coming from Trump supporters. Analysis restricted to 2020 election data (see Figure S29 for analysis based on 2022 election data).

Figure S28: Share of questioning views by outlet trustworthiness and candidate support (2022)




Weighted percentages based on post-stratification weights. Restricted to the 50 outlets with the highest number of election-fraud-related views within each group. Labels identify the top 10 outlets per group. Weighted percentages based on post-stratification weights. Figure restricted to the 50 outlets with the highest number of fraud-related views within each group. Labels identify the top 10 outlets per group. Point size indicates the total number of fraud-related views by outlet. The x-axis measures outlet trustworthiness based on ratings from NewsGuard and Lin et al. (2023). The dashed vertical line represents the threshold between sources classified as trustworthy and those classified as untrustworthy. Analysis restricted to 2022 election data (see Figure S26 for analysis based on 2020 election data).

Figure S29: Comparing questioning content exposure rates within outlet by candidate support (2022)



Weighted percentages based on post-stratification weights. Restricted to the 50 outlets with the highest number of election-fraud-related views within each group. Labels identify the top 10 outlets per group (red for Trump supporters, blue for Biden supporters, black for both). Points below the 45-degree line indicate that, for a given outlet, Biden supporters encountered a larger share of content questioning fraud claims than Trump supporters. Points above the line indicate the opposite. Point size indicates the total number of election fraud content views by outlet. The blue-red gradient of the points indicates the proportion of fraud-related views from that outlet coming from Trump supporters. Analysis restricted to 2022 election data (see Figure S27 for analysis based on 2020 election data).

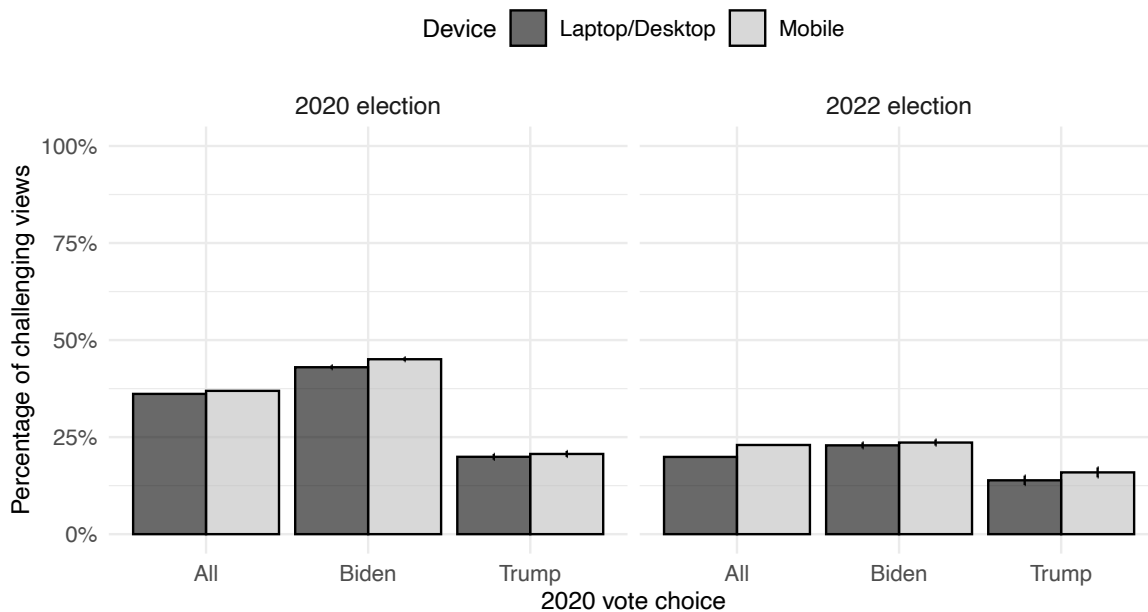
Figure S30: Permutation tests of the Biden/Trump supporter gap in challenging content exposure (2022)



Weighted estimates based on post-stratification weights. The figure reports independent permutation tests in which we randomly permuted content types (fraud-related versus non-fraud-related) within outlet holding overall news and outlet composition constant; outlet selection holding overall outlet content fixed; and challenging vs. non-challenging views within outlets holding outlet selection fixed (see Materials and Methods for further details). Analysis restricted to 2022 election data (see main text Figure 12 for equivalent results for the 2020 study period).

Figure S31 shows that the proportion of fraud-related views that question fraud claims is similar between participants who provided laptop/desktop web browsing data and participants who provided mobile data.

Figure S31: Percentage of views of fraud content that challenge fraud claims by device type




Weighted mean of individual-level percentages. Calculated using post-stratification weights.

S3 Robustness checks

S3.1 Experimental results among active Pulse participants

This subsection replicates all the preregistered experimental analyses among Pulse participants who were active online for the majority of months during the 2020 study period.

Figure S32: Treatment effects of exposure to fact-check article among active Pulse participants



Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of targeted statements. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted to Pulse participants active for the majority of months during the study period. Outcomes measured on a four-point scale ranging from “not at all accurate” to “very accurate”.

Table S27: Main treatment effects of exposure to fact-check article among active Pulse participants)

	Targeted true	Targeted false	Discernment
Fact-check	0.182*** (0.027)	-0.214*** (0.024)	0.397*** (0.038)
Lasso controls	✓	✓	✓
Num.Obs.	1374	1442	1439


*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, ideology (Targeted true model), information trust (Targeted false and Targeted true models), and political knowledge (Targeted true model).

Table S28: Treatment effects of exposure to fact-check article on the perceived truthfulness of targeted statements among active Pulse participants)

	Early votes	Observers	Justice Dept.	Trump won	Dead people	Dominion
Fact-check	0.275*** (0.053)	0.222*** (0.031)	0.027 (0.037)	-0.055* (0.025)	-0.190*** (0.033)	-0.411*** (0.038)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	1383	1444	1378	1443	1445	1376

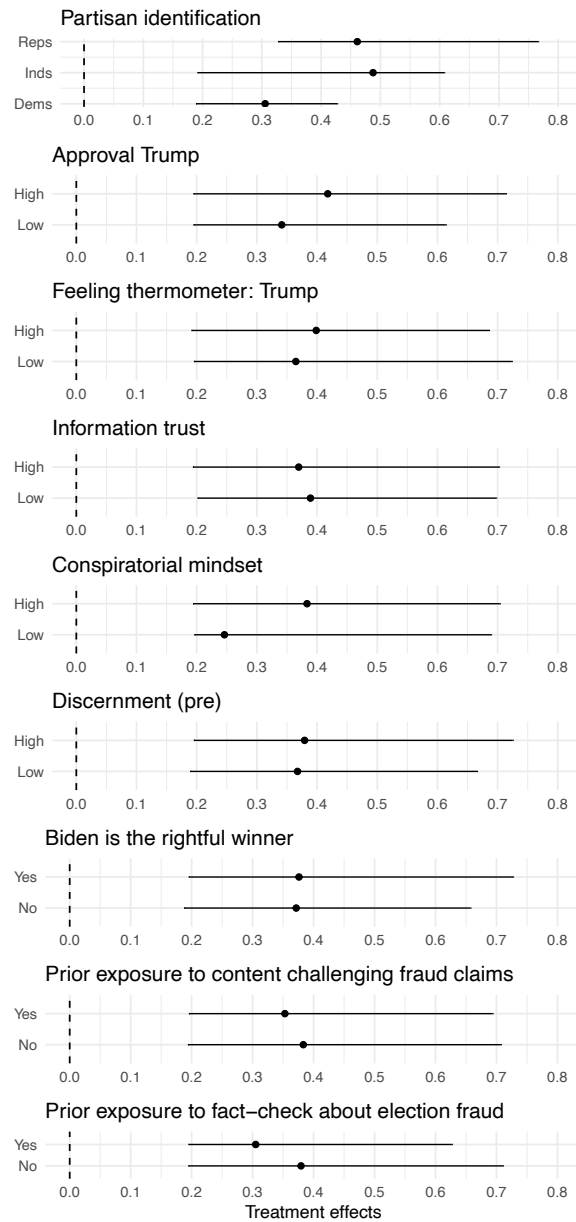
*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (all except Early votes), ideology (all except Justice Dept. and Trump won), information trust (Early votes, Dead people, Dominion), political knowledge (Early votes, Justice Dept., Dominion), and political interest (Early votes).

Figure S33: Combined treatment effects of exposure to fact-check article among active Pulse participants




Group means with 95% confidence intervals. Restricted to Pulse participants active for the majority of months during the study period.

Figure S34: Heterogeneous treatment effects of exposure to fact-check article on truth discernment among active Pulse participants



Heterogeneous treatment effects based on Bayesian Causal Forest models. Median effect by group with 95% credible intervals (2.5th and 97.5th percentiles), which can be asymmetric around the median. Restricted to Pulse participants active for the majority of months during the study period. Includes the following exploratory moderators (i.e., not preregistered): prior perceptions that Biden is the rightful winner of the 2020 election, prior exposure to content challenging fraud claims, and prior exposure to fact-checking about election fraud.

Figure S35: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted claims among active Pulse participants



Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of non-targeted statements. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted to Pulse participants active for the majority of months during the study period.

Table S29: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted claims among active Pulse participants

	Non-targeted true	Non-targeted false	Non-targeted discernment
Fact-check	0.055* (0.025)	-0.021 (0.022)	0.076* (0.033)
Lasso controls	✓	✓	✓
Num.Obs.	1372	1372	1367


*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, ideology (Non-targeted false), information trust (Non-targeted true, Non-targeted discernment), and political knowledge.

Table S30: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted statements among active Pulse participants

	Mail voting	Obama margin	Legal challenges	Online voting	Biden won Florida	Illegal ballots
Fact-check	0.046 (0.037)	0.078 (0.043)	0.040 (0.034)	-0.014 (0.035)	-0.005 (0.042)	-0.035 (0.033)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	1376	1375	1390	1388	1391	1442

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (Mail voting, Obama margin, Illegal ballots), ideology (Illegal ballots), information trust (Mail voting, Obama margin, Illegal ballots), political knowledge (Mail voting, Legal challenges, Biden won Florida), and political interest (Mail voting).

Figure S36: Treatment effects of exposure to fact-check article on broader attitudes about the 2020 election among active Pulse participants



Sample average treatment effects of exposure to a fact-checking article on broader attitudes about the legitimacy of the 2020 election and the January 6 insurrection. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted to Pulse participants active for the majority of months during the study period.

Table S31: Treatment effects of exposure to fact-check article on the perceived legitimacy of the election among active Pulse participants

	Trump won	Trump inaugurated	Rightful	Legitimate	Fair	Fraud influence
Fact-check	-0.033* (0.013)	0.009 (0.009)	0.110*** (0.020)	0.088*** (0.022)	0.087*** (0.021)	-0.083** (0.028)
Lasso controls	✓	✓	✓	✓	✓	✓
N	1445	1457	1446	1448	1447	1445

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification (all except Trump inaugurated and Fair), ideology (Trump won, Fraud influence), and information trust (Trump won).

Table S32: Treatment effects of exposure to fact-check article on confidence in vote counts among active Pulse participants

	Individual	Local	State	National
Fact-check	0.011 (0.030)	0.013 (0.027)	0.050 (0.027)	0.127*** (0.027)
Lasso controls	✓	✓	✓	✓
N	1309	1447	1444	1445

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: pre-treatment outcome, partisan identification, and information trust.

Table S33: Treatment effects of exposure to fact-check article on attitudes about the January 6 insurrection among active Pulse participants

	Make Trump look bad	Approve insurrection	Insurrection right	Block certification
Fact-check	-0.003 (0.017)	-0.007 (0.026)	0.004 (0.020)	0.013 (0.042)
Lasso controls	✓	✓	✓	✓
N	1445	1447	1446	1445

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$. OLS regression with robust standard errors. Control variables selected via lasso: partisan identification (Make Trump look bad, Insurrection right, Block certification), ideology (Make Trump look bad, Block certification), information trust (Make Trump look bad, Block certification), political interest (Block certification), and education (Make Trump look bad, Block certification).


S4 Online behavior data

S4.1 Pulse data

The data include daily web activity for 1,716 participants from September 13, 2020 to January 29, 2021 and for 1,756 participants from August 31, 2022 to January 31, 2023 of the YouGov Pulse panel.

Figure S37 shows the number of months of online activity among participants in the Pulse sample, with five months corresponding to having visited at least one website in September, October, November, December, and January of either 2020–2021 or 2022–2023. The results show that about 75% of the Pulse sample was active throughout the entire period during the 2020 and 2022 study periods, with no statistically significant differences between Trump and Biden supporters (using *t*-tests). In all analyses using the online behavior data, we kept participants who were active for the majority of months (at least three out of five months) in each study period (1,596 participants in the 2020 election period data and 1,518 participants in the 2022 election period data).

Figure S37: Months active in Pulse data by 2020 presidential vote choice



Months active measured from September 2020–January 2021 and September 2022 (starting on August 31)–January 2023. Participants are considered active when they visit at least one website during a calendar month.

S4.2 Domain selection process

To identify fraud content, we first subsetted the Pulse data to only keep news-related domains and subdomains (henceforth referred to as domains) that were rated by NewsGuard (February 2021 list) and/or Lin et al. (2023). NewsGuard rates 6,200 domains and Lin et al. (2023) rates 11,520 domains.

- 5,956 domains are rated by both NewsGuard and Lin et al.

- 244 domains are only rated by NewsGuard.
- 5,564 domains are only rated by Lin et al.
- 11,764 unique domains appear in NewsGuard and/or Lin et al. in total.

Next, we removed 83 domains that were classified as *platforms* (e.g., Reddit) or *satire* (e.g., the Onion) by NewsGuard, leaving us with 11,681 domains.

We then conducted an additional domain exclusion process using classifications from Shallalist to exclude any domain that appeared in 51 categories that we deemed topically irrelevant (e.g., autos, shopping, sports, etc.). We did not exclude any domains that appeared in the news, radio/television, politics, government, forums, hospital, and recreation/wellness categories, or domains that were not classified by Shallalist. Domains in the hospital and recreation/wellness categories were retained due to the salience of COVID-19 in the 2020 elections. This process only removed one domain, leaving 11,680 domains.

We also removed domains (and their subdomains) that received a high volume of participant visits, but were irrelevant or were impossible to scrape (e.g., social media sites, information aggregators whose URLs re-directed to another domain). These domains include: `bing.com`, `twitter.com`, `google.com`, `youtube.com`, `target.com`, `facebook.com`, `walmart.com`, `instagram.com`, `paypal.com`, `espn.com`, `zoom.com`, `instacart.com`, `groger.com`, `apple.com`, `gettr.com`, `blogspot.com`, and `wordpress.com`. This removed 64 domains, leaving 11,616 domains.

We removed any domains with “sport” (e.g., `nbcsports.com`) in the domain name that were not already excluded by Shallalist. Afterward, we manually identified and removed sports domains (e.g., `theathletic.com`) not caught by these previous steps. This process removed 98 domains, leaving 11,518 domains.

S4.3 Ratings of domains

In their February 2021 list, NewsGuard rated domains as trustworthy and not trustworthy. For each domain, NewsGuard calculates a trust score using nine criteria, where the score ranges from 0 to 100. Domains with a score of 60 or above were classified as trustworthy and domains with a score below 60 were classified as not trustworthy (NewsGuard 2022). NewsGuard changed its terminology in February 2023 (i.e., after the end of our study period). Sources with a score of 100 are now classified as “high credibility,” scores between 75 and 99 as “generally credible,” scores between 60 and 74 as “credible with exceptions,” scores between 40 and 59 as “proceed with caution,” and scores between 0 and 39 as “proceed with maximum caution.” We kept the binary ratings to make the analysis and visualization more straightforward.

Lin et al. (2023) use five different domain ratings^{S1} to calculate a trust/reliability score using principal component analysis. The score ranges from 0 to 1. However, they do not classify domains as trustworthy or not trustworthy using this score. To compare Lin et al.’s ratings with NewsGuard ratings, we used the `cutpointnr` package in R (Thiele and Hirschfeld 2021) to find the optimal cutpoint by maximizing the Youden Index (= sensitivity + specificity - 1). We determined that the optimal cutpoint value for mapping Lin et al.’s ratings score onto NewsGuard’s binary ratings was 0.5

^{S1}Lin et al. (2023) compare NewsGuard ratings to these five domain ratings, finding high level of correspondence, but do not include NewsGuard ratings as part of their publicly available principal component scores to satisfy NewsGuard’s data publication requirements.

(Youden index = 0.666, AUC = 0.849, accuracy = 0.854). Domains with scores greater than 0.5 were coded as trustworthy, while domains with scores lower than 0.5 were coded as untrustworthy (none have a score of exactly 0.5).

We merged the Pulse data with domain ratings at either the domain level, subdomain level (e.g., `news.yahoo.com`), or subdirectory (e.g., `yahoo.com/news`) level depending on whether a domain's subdomains or subdirectories were included in NewsGuard's or Lin et al.'s (2023) ratings (as noted above, we refer to this combined list as "domains"). To de-duplicate the data, we removed URL fragments (e.g., tracking parameters) and excluded repeated visits to the same URL if they occurred within 10 seconds of each other. When measuring exposure to news content (including election fraud content), we included a maximum of three visits to a given URL. This process left 487,036 URLs from 4,254 rated domains for 2020 and 361,132 unique URLs from 3,747 rated domains for 2022.

There is a high level of agreement between NewsGuard and Lin et al. (2023). Among domains visited by participants that were rated by both sources, 2,353 out of 2,536 (93%) have the same rating in our 2020 election data and 2,133 out of 2,267 (94%) have the same rating in our 2022 election data. When ratings differed, we used Lin et al.'s ratings, as they are based on five expert sources rather than a single one and thus likely to be more reliable.

S4.4 Identifying outliers with unusual browsing behavior

To rule out "bot-like" or other forms of seemingly automated browsing behavior, we first identified outliers and then manually examined their browsing behavior.

We defined outliers as participants whose average daily unique URL visits were 3+ standard deviations above the mean. For the 2020 Pulse data, outliers were participants who had average daily unique URL visits greater than or equal to 583 unique URL visits (mean = 106, standard deviation = 159). For the 2022 Pulse data, outliers were participants who had average daily unique URL visits greater than or equal to 338 unique URL visits (mean = 70.5, standard deviation = 89.2).

Next, for each outlier, we examined domain-level counts of URL visits and flagged participants whose browsing consisted of 95% or more of visits to five or fewer domains. We then manually examined all of the raw Pulse data of the flagged participants to check for automated browsing behavior. We defined automated browsing behavior as repeated browsing patterns between a main incentivized referral domain (e.g., `inboxdollars.com`) and a different domain (e.g., `news.yahoo.com`) over a specific period of time (e.g., 1–2 hours), where the URL visit durations were similar (e.g., 25–30 seconds). We did not remove flagged participants whose browsing data was all or mostly incentivized URL visits but who did not appear to be using an automated process.

For the 2020 Pulse data, we identified thirteen participants as outliers and flagged six participants for manual review. We identified and removed three participants from the Pulse data who met our definition of unusual browsing behavior — in this case, the repeated pattern was 1 second at `inboxdollars.com` or `mypoints.com` and always right around 30 seconds or 1 minute at another domain, then 1 second at `inboxdollars.com` or `mypoints.com`, etc. Removing these outliers does not impact any of the substantive conclusions of this paper.

For the 2022 Pulse data, we identified twelve participants as outliers and flagged four participants for manual review. None of the flagged participants met our definition of unusual browsing behavior. Consequently, we did not remove any participant from the 2022 data.

S4.5 Scraping process

We collected the text content from the remaining URLs using both generic and custom scrapers built with the Python package `Scrapy`. For non-paywalled domains, the scraper was configured with a concurrency of 1.0 per domain and a limit of 128 concurrent requests. This was done to avoid being blocked for intensive scraping by domains. For paywalled domains, we scraped the text content via the Internet Archive. Here, the scraper was configured to retrieve the most recent archived image for each URL and used a limit of 2–5 concurrent requests.

We manually examined the unscraped URLs for domains with 100 or more that were unscraped, included an error message in the text, or contained less than 500 characters. We first removed URLs that were either unreachable (e.g., the URL does not exist anymore because the content was taken down) or uninformative (e.g., URLs were image-only slideshows). We next created a list of exclusion terms based on each domain’s URLs and then filtered out URLs with those terms for the corresponding domains. The exclusion terms were for domain sections that were definitely not related to elections. For example, we removed URLs from *The New York Times* that included “realestate|crosswords|puzzles|dining|get-started|holiday-gift-guide.” For the remaining URLs, we scraped the Internet Archive for non-paywalled domains and manually downloaded the content for paywalled domains.

In the end, we collected text content from 378,890 URLs across 4,563 rated domains for 2020, and from 276,458 URLs across 3,607 rated domains for 2022.

S4.6 Text parsing and wrangling

We first removed non-English websites, URLs leading to the home page or the section of a website (e.g., <https://www.cnn.com/politics>)^{S2}, and scraped articles containing less than 25 characters (about five words).

News pages from `yahoo.com` also often include recent, trending, or recommended articles after the main article. We removed these subsequent articles given that the ones that were scraped may be different from the ones participants encountered at the time of viewing. We also removed video pages from Yahoo when the video script was unavailable.

This process left 358,200 URLs for 2020 and 268,005 URLs for 2022.

We filtered the output to keep only content related to election fraud using an exhaustive dictionary based on the main examples of fraud discussed during the 2020 and 2022 study periods (e.g., see Kennedy et al. 2022). This filtering was necessary due to high sparsity in the text data. We used the following string queries to identify content related to election fraud.^{S3} A dot followed by brackets means that the word can include any number of characters included in the provided interval. The last term of a string can contain additional characters. For example, “(dump|stuff){0,4}(ballot|vot)” would include “dumped votes” as well as “dumping voting machines.”

- “(ballot|elect|vot){0,20}(fraud|rigged|rigging|illegal|illegitimate|ineligible|irregularities|suppression|charges|allegations|cases|integrity|stolen|manipulat|interfere|tamper|overturn|deni|bogus|corrupt|impersonation|miscount|wrongdoing)”

^{S2}Website home pages and sections tend to change over time. As a result, the scraped content would not be the same as what our participants saw.

^{S3}We used the `str_detect()` function from the `stringr` R package. Bullet points are used to make it easier for the reader to visualize the different strings included in our dictionary. These strings were combined in a single input vector.

- “(fraud|rigged|rigging|illegal|illegitimate|ineligible|irregularities|suppression|charges|allegations|cases|integrity|stolen|steal|manipulat|interfere|tamper|overturn|deni|bogus|corrupt|miscount|wrongdoing).{0,20} (ballot|elect|vot)”
- “(ballot|vote).{1,1}(dumping|stuffing)”
- “(dump|stuff).{0,4}(ballot|vot)”
- “vot.{1,4}(multiple times|more than once)”
- “(deceased|dead|dead people|noncitizens) vot”
- “stop the steal”

Finally, websites vary greatly in structure and HTML code, making it challenging to scrape only the main body of an article without capturing extraneous text. To reduce noise in the documents, we split the articles containing election fraud keywords into sections based on line breaks and tabs, then removed any sections containing three words or fewer, as these often corresponded to website navigation elements (e.g., “Politics,” “Business,” “Lifestyle,” “Contact Us”). We also removed all sections that were repeated in more than 20 scraped articles in one or more domains (independently calculated and applied for the 2020 and 2022 election periods), a proxy for text not specific to the article of interest.^{S4} These sections were often headlines of other trending articles (which change over time); other sections of the website; invitations to sign in, subscribe, create an account, follow the news outlet on social media, share the article, report a correction, etc.; or information about cookies, ads, and user privacy. After that cleaning process, we reran the election fraud dictionary and removed articles that no longer contained election fraud keywords.

After removing duplicates, we were left with 37,564 unique articles containing election fraud keywords for 2020 and 14,515 articles for 2022 that we coded using a LLM.

S4.7 Developing the LLM prompt

We developed the coding scheme and LLM prompt using a systematic coding procedure based on Törnberg (2024). In the first step, human coders coded a random set of articles on three dimensions: (1) whether the article mentions election fraud (“yes” or “no”); (2) if so, whether it questions or contradicts election fraud claims (“yes” or “no”); and (3) if so, the extent of questioning or debunking on a four-point scale (“0=None,” “1=Minimal,” “2=Moderate,” “3=Extensive”). Notes were added by the coders explaining their coding decisions. They then discussed and resolved any disagreements. We subsequently updated the coding instructions as needed. This process was repeated using new random sets of articles until reaching a sufficient intercoder agreement. These steps were particularly useful for identifying the scope conditions included in the LLM prompt.

The intercoder reliability (Krippendorff’s alpha) among the four human coders for the final set of 50 articles was 0.74 for mentions of election fraud, 0.77 for whether the article questioned or contradicted fraud claims, 0.83 for the full ordinal coding of the extent of questioning or debunking, and 0.83 for the ordinal coding corresponding to the three categories we use in this paper (“0=No questioning,” “1=Minimal questioning,” “2/3=Challenging”). The final version of the prompt is presented in Table S34.

^{S4}Complete sentences or paragraphs started appearing below that threshold, making it difficult to determine whether sections were related to the main article or not.

Table S34: Prompt

Task overview

As an expert annotator, you will be asked to annotate news articles about U.S. politics to answer three questions:

- Q1. Does the article contain any reference to or suggestion of election fraud (as defined in Question 1 codebook) in a U.S. election?
- Q2. Does the article contain a statement that questions, rejects, or disproves the idea that election fraud is widespread or could change the outcome of one or more elections?
- Q3. How extensive is the questioning or debunking of election fraud claims in the article?

Eligible content definition

- ✓ Included: Main article text only (including title/subtitles)
 - ✗ Excluded: Comments, ads, navigation captured during scraping
 - ✗ Excluded: Related/additional articles captured during scraping
 - ✗ Excluded: Satire (e.g., The Onion, Babylon Bee)
 - ✗ Excluded: Letters to the editor
-

Table S35: Prompt (continued)

Question 1 codebook

Q1. Does the article contain any reference to or suggestion of election fraud (as defined in Question 1 codebook) in a U.S. election? (Code as Yes / No)

Overall definition: Election fraud is any illegal action that alters, blocks, or fabricates the casting, counting, or certification of votes in a U.S. federal, state, or local election.

Claims of election fraud can be general or suggest interference with one of the following:

- **Casting votes** (e.g., impersonating a voter, destroying or stealing mail/absentee ballots, illegally preventing voters from receiving or casting ballots, voting illegally)
- **Counting votes** (e.g., altering vote records, stuffing ballot boxes with fraudulent votes, mishandling ballots, hacking voting machines or election systems)
- **Certifying votes** (e.g., forging, discarding, or falsifying official tally sheets or certificates, bribing or coercing officials to misreport results, seeking to block the certification of legitimate results or to force the certification of false results)

Code Yes if **any** of the following are included in an article regardless of its overall stance:

- **A direct claim of or reference to fraud in U.S. elections** (e.g., “ballot-box stuffing,” “votes changed,” “illegally discarded ballots,” “hacked voting machines”)
 - Includes **hypothetical or general** claims (e.g., “fraud-prone,” “dead voters”)
 - Includes references to **fraud in other U.S. elections** besides 2020.
- A reference to a **U.S. election outcome being disputed/not accepted** that makes some reference to fraud claims (e.g., “Vance said the election was stolen/rigged”)
- A reference to Donald **Trump or his allies disputing/not accepting the 2020 election** without explicit fraud claims (e.g., “Trump’s attempts to overturn his loss”)
- **Statements questioning or refuting fraud claims about U.S. elections** (e.g., “election denier,” “election denial,” “baseless fraud claims,” “the Big Lie”)
- **References to U.S. election vulnerabilities** (e.g., “absentee ballot fraud is easy”)

The following **do not** qualify unless they otherwise meet the definition above:

- **Other campaign/political activity (regardless of legality):** email, leaking documents, foreign interference that does not directly affect the casting, counting, and certifying of votes, legal investigations, campaign finance issues, disinformation.
- **Legal (or not-yet-adjudicated) election-administration rules:** voter ID laws, voter roll maintenance, or other measures labeled “voter suppression” unless ruled illegal.
- **Post-election protests and unrest:** e.g., Jan. 6 doesn’t count without fraud content.
- **Generic references:** terms like “election integrity” or “conspiracy theories” with no explicit fraud allegation.

Implementation tip for LLMs:

- IF article contains ≥ 1 statement that meets the conditions specified for Q1 above
- THEN Q1 = Yes
- ELSE Q1 = No

If Q1 = 0, then code Q2 = No and Q3 = 0 (i.e., skip further fraud-related questions).

Table S36: Prompt (continued)

Question 2 codebook

Q2. Does the article contain a statement that questions, rejects, or disproves the idea that election fraud (as defined in Question 1 codebook) is widespread or could change the outcome of one or more elections under current or past rules?

(Code as Yes / No; can be the same statement used in coding for Q1)

(Evaluate Q2 only if Q1 = Yes; otherwise Q2 = No.)

Code **Yes** if any of the following are included in an article regardless of its overall stance:

- **Explicit questioning/refutation of fraud claims**, including court rulings (e.g., “baseless / unfounded / false / unsupported fraud claims,” “no evidence of widespread fraud,” reporting that audits confirm results, judge rejects fraud claims)
- **Court rejection of lawsuit alleging election fraud**
- **Affirmation of legitimate results** (e.g., “Biden was legitimately elected,” “Biden’s clear victory,” “results were accurate,” “election was free and fair”)
- **Wording** describes challenge to 2020 election as **illegitimate**, **election denial**, or **refusing to accept results** or saying fraud claims are a **conspiracy theory** (e.g., condemns attempts to block certification, refers to “election deniers” or “election denial,” “the Big Lie,” Trump “refused to accept” results)
- **Fraud too small to matter** (e.g., saying not enough fraud for Trump to win)

None of the following qualify unless they otherwise meet the definition above:

- **Mere mention** of efforts to “overturn” the election or speculation on their success
- **Institutional actions** (e.g., legislative procedures, election result certification, media organizations projecting/calling the winner) **not explicitly framed as rejecting fraud claims or affirming the legitimacy or accuracy of the election** results.
- Coverage that merely **reports fraud claims or controversies** (e.g., “Republicans are divided over Trump’s fraud claims”)
- References to people **not acting on fraud-related claims if they are not described explicitly as false** (e.g., “Pence refused to block certification of the election”)
- Coverage that **references the election outcome without explicitly endorsing its legitimacy** (e.g., saying “Biden won,” referring to Biden as “President-elect”)
- **Hypothetical discussion of future rule changes** (e.g., “fraud would not rise if universal mail voting were adopted”)

Implementation tip for LLMs:

- IF article contains ≥ 1 statement that meets the definitions and conditions specified above
- THEN Q2 = Yes
- ELSE Q2 = No

If Q2 = No, then code Q3 = 0 (skip further fraud-related questions).

Table S37: Prompt (continued)

Question 3 codebook

Q3. How extensive is the questioning or debunking of election fraud claims in the article?

(Only complete Q3 if Q2 = Yes. If Q2 = No, then code Q3 = 0.)

Code each article in the most appropriate category:

- **0 = None** — No questioning/debunking.
- **1 = Minimal** — Questioning/debunking appears only in the body (not in the headline or lede); limited space (≤ 2 sentences or $\sim \leq 10\%$ of the article); no explanation or at most a single brief reason with no elaboration; not a primary focus.
- **2 = Moderate** — Questioning/debunking appears in the headline, lede, or body; receives some space ($\sim 10\text{--}33\%$ or a brief piece overall); includes some explanation or evidence (at least one concrete reason or source) that goes beyond a passing phrase; may be a primary focus if the article is short or provides little detail, but otherwise is not the main focus.
- **3 = Extensive** — Questioning/debunking is a primary focus: it appears in the headline or lede, occupies a substantial share of the article ($\geq \sim 34\%$, typically most), and provides detailed explanations/evidence (multiple reasons/sources or an in-depth analysis of one).

Clarification:


- The lede is defined as the first two sentences of the article, or the entire first paragraph if it contains two sentences or fewer. This definition excludes text or headlines unrelated to the main article.
 - If the article contains **only a single brief refutation, even if it cites evidence or an expert**, code Q3 = 1. Coding Q3 = 2 requires multiple sentences with some explanation or evidence.
-

S4.8 Validation

To validate the LLM results and compare the performance of different models, we took a random sample of 50 articles containing election fraud keywords from the pooled sample of 2020 and 2022 articles^{S5}; articles coded for developing the LLM prompt were excluded. Each article was coded by four human coders and with GPT-5. We resolved all disagreements between the human coders before comparing their results with the results from the GPT model.

Figure S38 shows the confusion matrix. Table S38 presents the Krippendorff's alpha across the four human coders, as well as the percent agreement, Krippendorff's alpha, precision, and recall between the consensus human coder judgments and GPT-5.

Figure S38: Confusion matrix



^{S5}The validation sample was drawn before we identified scraping issues (e.g., missing articles, short outputs, or error messages in the text) and subsequently rescraped the content. At the time of validation, 39,435/52,079 (76%) of the final articles with election fraud keywords were included. Of these, 434 (1%) were later rescraped because the initial version contained an error message or fewer than 500 characters. Two of these problematic URLs were included in the final validation set.

Table S38: Intercoder reliability among human coders and between human coders and GPT-5

Validation	Metric	Q1 Mention	Q2 Question	Q3 Depth	Q3 Depth (0 vs. 1 vs. 2-3)	Q3 Depth (0 vs. 2-3)
Humans	Krippendorff's alpha	0.74	0.77	0.83	0.83	0.91
Humans	Agreement	100%	96%	82%	94%	100%
vs. GPT-5	Krippendorff's alpha	1	0.92	0.94 (ordinal)	0.95 (ordinal)	1
	Precision	1	0.96	0.76 (nominal)	0.93 (nominal)	1
	Precision (weighted)	1	0.96	0.90 (nominal)	0.94 (nominal)	1
	Recall	1	0.96	0.70 (nominal)	0.93 (nominal)	1
	Recall (weighted)	1	0.96	0.82 (nominal)	0.94 (nominal)	1

S4.9 Coding articles using LLM

The results presented in this paper are based on 52,079 articles (combining the 2020 and 2022 data) containing election fraud keywords that were coded using GPT-5 via OpenAI's batch API (implemented via the `OpenAI` package in Python). OpenAI returned error messages for 19 articles, which are excluded from the above count and all analyses and results.

S4.10 Fact-checks

To be as exhaustive as possible in identifying fact-check articles, we included all U.S. fact-checking sites listed by Poynter's International Fact-Checking Network (IFCN) and Duke Reporters' Lab database of global fact-checking sites (a total of 100 unique domains).

We identified fact-checking articles in three steps:

1. For websites dedicated primarily or exclusively to fact-checking (e.g., FactCheck.org, Politifact.com, Snopes.com), all scraped content was categorized as fact-check articles.
2. When the fact-checking initiative is implemented by a broader media organization (e.g., Reuters, CNN, ABC), we identified sections of their website devoted to fact-checking (subdomains) and labeled all content from these sections as fact-check articles.
3. Since fact-check articles are not always clearly identified on news websites, we used the description of each fact-checking initiative on Duke Reporters' Lab website and manual searches on each news website to build a comprehensive list of the keywords they use to label fact-check articles. We labeled all content from the list of 100 domains where the URL contained at least one of the fact-checking keywords as a fact-checking article.

The list of keywords was:

"4-investigates", "4_investigates", "4investigates", "ad-watch",
 "ad_watch", "adwatch", "based-on-science", "based_on_science",
 "basedonscience", "cronicas-desinformacion",
 "cronicas_desinformacion", "cronicasdesinformacion", "debunk",

"detector-de-mentiras", "detector_de_mentiras",
"detectordementiras", "digging-deeper", "digging_deeper",
"diggingdeeper", "el-detector", "el_detector", "eldetector",
"fact-brief", "fact-check", "fact-finder", "fact-squad",
"fact_brief", "fact_check", "fact_finder", "fact_squad",
"factbrief", "factcheck", "factfinder", "factfinder",
"facts-first", "facts_first", "factsfirst", "factsquad",
"fake-news", "fake_news", "fakenews", "get-the-facts",
"get_the_facts", "getthefacts", "i9-fact-check", "i9_fact_check",
"i9factcheck", "not-real-news", "not_real_news", "notrealnews",
"pinocchio", "politifact", "reality-check", "reality_check",
"realitycheck", "science-vs", "science_vs", "sciencevs",
"spin-control", "spin_control", "spincontrol", "t-verifica",
"t_verifica", "tfcn", "trust-index", "trust_index", "trustindex",
"truth-be-told", "truth-in-numbers", "truth-squad", "truth-test",
"truth-tracker", "truth_be_told", "truth_in_numbers",
"truth_squad", "truth_test", "truth_tracker", "truthbetold",
"truthinnumbers", "truthsquad", "truthtest", "truthtest",
"truthtracker", "tverifica", "verificacion", "verify".

For the 2020 study period, we scraped 2,512 unique fact-checking articles, of which 401 were coded by GPT as mentioning election fraud. For 2022, we scraped 591 unique fact-checking articles, of which 38 were coded by GPT as related to election fraud.

S4.11 Descriptive statistics

Tables [S39](#) and [S40](#) include descriptive statistics about exposure to different types of election fraud content overall and by vote choice during the 2020 and 2022 study periods. Table [S41](#) compares levels of exposure by device type (laptop/desktop vs. mobile).

Table S39: Distribution of content exposure during the 2020 U.S. election (Pulse data)

Variable	Mean	SD	p10	p25	Median	p75	p90
Overall							
News content	613.09	1655.29	5	31	136	564	1504
Trustworthy news	499.63	1366.48	3	28	113	433	1193
Untrustworthy news	113.46	663.60	0	0	5	33	155
Election fraud content	23.44	78.81	0	0	0	9	51
" challenging fraud claims	8.45	33.54	0	0	0	2	18
" minimally questioning fraud claims	8.60	29.60	0	0	0	3	21
" not questioning fraud claims	6.25	22.85	0	0	0	2	11
" from trustworthy sources	18.20	67.80	0	0	0	7	39
" from untrustworthy sources	5.24	28.12	0	0	0	0	5
" from fact-check	0.30	2.62	0	0	0	0	0
Biden Voters							
News content	650.45	1723.51	7	41	158	586	1624
Trustworthy news	583.87	1489.25	4	36	152	552	1454
Untrustworthy news	66.57	645.73	0	0	4	16	58
Election fraud content	33.18	100.53	0	0	2	16	91
" challenging fraud claims	14.31	46.49	0	0	0	6	38
" minimally questioning fraud claims	13.30	39.53	0	0	0	7	37
" not questioning fraud claims	5.48	16.72	0	0	0	3	12
" from trustworthy sources	30.26	92.86	0	0	1	14	81
" from untrustworthy sources	2.92	21.75	0	0	0	0	2
" from fact-check	0.46	3.03	0	0	0	0	1
Trump Voters							
News content	667.85	1772.71	6	44	165	584	1650
Trustworthy news	469.41	1351.61	6	31	115	382	1156
Untrustworthy news	198.44	790.63	0	1	13	80	410
Election fraud content	17.81	57.68	0	0	0	7	38
" challenging fraud claims	3.53	12.57	0	0	0	2	9
" minimally questioning fraud claims	5.19	17.10	0	0	0	2	12
" not questioning fraud claims	8.87	31.32	0	0	0	2	18
" from trustworthy sources	8.34	30.78	0	0	0	4	19
" from untrustworthy sources	9.47	37.84	0	0	0	1	13
" from fact-check	0.21	2.56	0	0	0	0	0

Participants are YouGov Pulse panel members who were active for the majority of months throughout the study period. Estimates calculated with post-stratification weights.

Table S40: Distribution of content exposure during the 2022 U.S. election (Pulse data)

Variable	Mean	SD	p10	p25	Median	p75	p90
Overall							
News content	458.15	1407.84	6	30	90	327	1083
Trustworthy news	373.27	999.22	6	28	82	290	943
Untrustworthy news	84.87	896.31	0	0	3	11	52
Election fraud content	7.81	34.17	0	0	0	1	12
" challenging fraud claims	1.58	7.63	0	0	0	0	3
" minimally questioning fraud claims	4.07	18.94	0	0	0	1	6
" not questioning fraud claims	2.14	10.28	0	0	0	0	3
" from trustworthy sources	6.24	30.15	0	0	0	1	10
" from untrustworthy sources	1.57	13.54	0	0	0	0	0
" from fact-check	0.03	0.36	0	0	0	0	0
Biden Voters							
News content	533.08	1737.16	5	32	96	383	1356
Trustworthy news	456.20	1201.21	5	31	90	350	1232
Untrustworthy news	76.88	1175.87	0	0	2	10	36
Election fraud content	12.14	46.45	0	0	0	3	21
" challenging fraud claims	2.78	10.99	0	0	0	0	5
" minimally questioning fraud claims	6.66	26.84	0	0	0	1	10
" not questioning fraud claims	2.66	10.73	0	0	0	0	5
" from trustworthy sources	10.68	43.54	0	0	0	2	17
" from untrustworthy sources	1.46	13.69	0	0	0	0	0
" from fact-check	0.06	0.47	0	0	0	0	0
Trump Voters							
News content	576.32	1394.04	9	46	136	415	1656
Trustworthy news	429.57	1023.66	6	41	109	338	1028
Untrustworthy news	146.75	794.35	0	1	8	33	163
Election fraud content	7.05	25.90	0	0	0	2	15
" challenging fraud claims	0.99	3.93	0	0	0	0	2
" minimally questioning fraud claims	3.25	11.53	0	0	0	1	8
" not questioning fraud claims	2.80	12.87	0	0	0	0	5
" from trustworthy sources	4.29	15.37	0	0	0	1	9
" from untrustworthy sources	2.76	17.60	0	0	0	0	1
" from fact-check	0.02	0.32	0	0	0	0	0

Participants are YouGov Pulse panel members who were active for the majority of months throughout the study period. Estimates calculated with post-stratification weights.

Table S41: Distribution of content exposure during the 2020 and 2022 U.S. election (Pulse data) by device type

Variable	Mean	SD	p10	p25	Median	p75	p90
2020 Laptop/Desktop							
News content	764.14	1901.92	5	45	202	703	1935
Election fraud content	28.92	84.62	0	0	1	15	82
" challenging fraud claims	10.40	35.35	0	0	0	4	24
" minimally questioning fraud claims	10.62	32.04	0	0	0	6	30
" not questioning fraud claims	7.73	25.21	0	0	0	4	18
2020 Mobile							
News content	175.36	461.69	1	11	50	148	368
Election fraud content	7.27	50.69	0	0	0	3	9
" challenging fraud claims	2.66	23.80	0	0	0	1	3
" minimally questioning fraud claims	2.65	18.13	0	0	0	1	3
" not questioning fraud claims	1.91	12.68	0	0	0	0	2
2022 Laptop/Desktop							
News content	658.94	1781.52	6	37	145	591	1711
Election fraud content	11.99	42.80	0	0	0	4	25
" challenging fraud claims	2.38	9.61	0	0	0	1	5
" minimally questioning fraud claims	6.34	23.97	0	0	0	2	13
" not questioning fraud claims	3.23	12.85	0	0	0	1	7
2022 Mobile							
News content	168.18	458.16	3	19	61	141	322
Election fraud content	1.95	13.40	0	0	0	0	2
" challenging fraud claims	0.45	2.86	0	0	0	0	0
" minimally questioning fraud claims	0.90	6.53	0	0	0	0	1
" not questioning fraud claims	0.60	4.34	0	0	0	0	0

Participants are YouGov Pulse panel members who were active for the majority of months throughout the study period. Estimates calculated with post-stratification weights.