Inattention and Differential Exposure: How Media Questioning of Election Fraud Misinformation Often Fails to Reach the Public

Mathieu Lavigne
Dartmouth College

Brian Fogarty
University of Notre Dame

John Carey
Dartmouth College

Brendan Nyhan
Dartmouth College

Jason Reifler
University of Southampton

Abstract

Why are high-profile misperceptions like the myth of widespread fraud in the 2020 U.S. presidential election so persistent and pervasive? Many observers blame partisan demand for congenial news and resistance to corrective information. Others blame media amplification of false claims. However, nationally representative survey and behavioral data from the U.S. during the periods around the 2020 and 2022 elections show that skewed online information diets are rare and that even supporters of Donald Trump change their views about fraud claims when randomly exposed to fact-checks. The problem is not the prevalence of accurate information, either; the fraud-related content that people saw online overwhelmingly questioned fraud claims. We instead conclude that a key factor is inattention (i.e., a lack of exposure) — most people encountered relatively little fraud-related content in their web browsing. Among those who did see such content, most supporters of both Joe Biden and Trump saw more content that questioned fraud claims. However, Trump supporters were differentially likely to be exposed to articles that did not question claims of widespread fraud, including after encountering more skeptical coverage — a pattern that has been shown to undermine correction effects. The persistence of fraud beliefs thus appears to be attributable to the combination of low levels of attention and differential exposure to congenial content undermining the effects of more accurate information.

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, 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, and seminars at the Harvard/Northeastern misinformation speaker series, American University, Max Planck Institute, and Boston University. We thank Yang Xu for assistance with data collection. All errors are our own.

False and 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, persistent, and consequential misperception in American politics is the 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 presidential 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 inspired, which appear to be durable and largely sincere (Graham and Yair 2024), cast lasting doubts on the legitimacy of Joe Biden's presidency and 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 in turn is necessary for the survival of democracy (Przeworski 2019). Given these stakes, it is critical to understand why false beliefs in widespread voter fraud persist.

We consider how both information consumption preferences and the prevalence of accurate information in online media diets may affect the persistence of misperceptions. We focus on three particular elements of news consumers' demand for information: levels of attention to news and information about politics; differences in exposure to congenial and uncongenial news content; and responsiveness to corrective information.

Prior scholarship on misperceptions emphasizes two demand-related explanations for why misperceptions endure. The first focuses on selective exposure to congenial content. Audience demand for news and information that reinforces one's existing preferences and beliefs is well-documented in political communication research (Iyengar and Hahn 2009; Stroud 2008; Taber and Lodge 2006) and can lead to greater exposure to false or misleading claims (Garrett et al. 2019, 2016; Meirick 2013; Meirick and Bessarabova 2016). Selective exposure is typically defined as a tendency to prefer or predominantly select information that supports one's views, which in turn suggests news diets that tilt toward congenial content (though measurement approaches and empirical findings of course vary; see, e.g., Garrett 2009; Hart et al. 2020; Stroud 2008, 2010). Critics often worry that such tendencies are exacerbated online, describing people as trapped in "echo chambers" or "filter bubbles" that reinforce viewpoint homogeneity (Pariser 2011; Sunstein 2001).

A second demand-driven explanation for the persistence of misperceptions emphasizes the role of directional motivations in how people reason about facts (Kunda 1990; Taber and Lodge 2006). Under this account, people tend to uncritically accept politically congenial claims that are false or unsupported and reject accurate information that is uncongenial (Flynn et al. 2017; Peterson and Iyengar 2021). In extreme cases, such resistance could potentially even cause people to embrace false claims more in response to exposure to corrective information (Nyhan and Reifler 2010; Nyhan et al. 2013).

However, these explanations are incomplete. Another important demand-related factor is inattention. In short, some people consume very little news and political content. Audience size ratings and behavioral data show that survey self-reports vastly overstate news consumption (Jerit et al. 2016; Prior 2009a,b). News consumption tends to be low for most people, who are typically uninformed about political matters (Lupia 2016). For some people, these patterns of inattention may reflect news avoidance(Robertson 2025; Skovsgaard and Andersen 2020; Toff and Kalogeropoulos 2020; Toff et al. 2023; Villi et al. 2022), but we cannot distinguish between news avoidance and other causes of inattention in the data we consider.

Beyond these demand-related propositions, the prevalence of accurate information available online could also contribute to the persistence of misperceptions. Journalists may feel pressure to create artificial balance in news coverage of factual or scientific disputes, even when an expert consensus exists (Boykoff 2008; Boykoff and Boykoff 2004) — "teaching the controversy" over taking sides. These types of situations help explain why journalistic norms at times may stand in the way of decisively debunking false or unsupported claims in standard reporting (Fahy 2017; Graves 2016). Some coverage has been shown to uncritically amplify false claims made by those in power (Bennett 1990; Bennett et al. 2008; Hayes and Guardino 2010), which may be particularly damaging during elections as exposure to news content reporting fraud claims by political elites promotes distrust in election results (Berlinski et al. 2023; Clayton et al. 2021; Justwan and Williamson 2022; Lyons and Workman 2022).

Given these concerns, it is worth considering whether enough accurate and reliable information is available to effectively inform people about the extent of election fraud and election legitimacy.

Prior scholarship testing the first two demand-driven explanations raises doubts that they can fully account for the persistence of fraud misperceptions. Contrary to fear of widespread echo chambers, most Americans have relatively balanced news diets on average (Guess et al. 2019, 2018). Correspondingly, exposure to untrustworthy websites and other dubious content 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). Similarly, initial concerns about motivated resistance (Nyhan and Reifler 2010; Nyhan et al. 2013) have not been supported by more recent work showing that fact-checking is generally effective at correcting misperceptions and produces patterns of parallel updating across groups rather than backfire effects among the most resistant ones (Carnahan et al. 2021; Coppock et al. 2023; Nyhan 2021; Walter et al. 2020; Walter and Murphy 2018; Wood and Porter 2019). In some cases, fact-checks may even be more effective among the people who are most likely to believe misinformation (Carey et al. 2022).

We advance this literature by expanding the set of demand-related factors that could drive consumption of online content related to election fraud and by estimating the prevalence of accurate information related to fraud. In so doing, we provide further evidence that questions whether echo chambers or resistance to corrective information can account for the endurance of misperceptions about fraud. First, we present findings from a nationally representative survey experiment indicating that corrective information successfully reduces inaccurate beliefs about election fraud in the 2020 presidential election, including among the most susceptible (or vulnerable) groups. Second, we estimate the share of content questioning the existence of widespread election fraud in the set of fraud-related articles Americans encountered in online media content (other than the text of social media posts) during the 2020 and 2022 election cycles. We employ advanced methods by using a large language model (LLM) to identify fraud-related articles based on their content rather than ratings of domain trustworthiness. More specifically, we use the LLM to classify articles on whether they push back against fraud claims. This approach allows us to measure individuals' consumption of fraud-promoting versus fraud-questioning content *at the article level*. The use of LLMs to successfully code political content has already shown some promise (Haroon et al. 2025). This granular classification provides a more reliable assessment of fraud-related information exposure, recognizing that trustworthy and even untrustworthy sources can contain a mix of questioning and non-questioning articles. It enables us to investigate patterns of selective exposure within source types and to evaluate the relative reach of fact-checking initiatives, corrective information, and counter-narratives about election fraud.

Our data suggest that the two most prominent demand-based explanations — selective exposure to congenial information and resistance to corrective information — cannot fully account for the persistence of fraud beliefs.

Contrary to worries about popular accounts of "echo chambers" or worries about media amplification of false claims, the fraud-related information Americans encountered from online sources was overwhelmingly skeptical toward fraud claims. In particular, we do not observe evidence that Trump supporters predominantly encountered content that flatters their own predispositions by uncritically endorsing fraud claims. Even Trump supporters saw substantially more content that questioned these claims than content that did not.

In addition, we find that exposure to corrective information does reduce false beliefs about widespread voter fraud and improve discernment between true and false claims about elections. These effects are observed across a variety of subgroups that we might expect to be especially resistant to content debunking fraud claims, including Trump approvers, Republicans, and people with conspiratorial predispositions.

Two related factors seem to better explain the persistence of fraud misperceptions. First, inattention is widespread. Even when false claims are widely questioned in news coverage and online information, most people encounter this content infrequently (Guess et al. 2020). Specifically, most Americans visited few URLs related to fraud during the 2020 and 2022 elections, limiting their exposure to corrective information, including exposure to fact-checks about fraud, which was particularly rare.

Second, we observe a pattern of differential exposure by candidate support that is more subtle than most accounts of echo chambers and filter bubbles (Pariser 2011; Sunstein 2001). As noted above, Trump voters who encountered fraud content were not insulated from skeptical coverage —

indeed, most were exposed to more questioning than unquestioning content. However, exposure to content questioning fraud claims was relatively lower among Trump voters than Biden voters, a gap that reflects both differences in outlet preferences and differences in article selection within outlets. Moreover, Trump supporters were disproportionately likely to consume content that failed to question the falsehood, including after exposure to skeptical content, a pattern that can undermine the effect of exposure to corrective information (Nyhan et al. 2022).

These results suggest that high-profile misperceptions persist not because corrections and skeptical media coverage are unavailable or ineffective, but because people consume so little political news and the effects of exposure to accurate information may be blunted by exposure to congenial content among vulnerable groups.

Results

We present results as follows. First, we establish the persistence of beliefs in widespread fraud in U.S. elections over time using nationally representative survey data. We then use a survey experiment conducted after the 2020 election to demonstrate how exposure to fact-checks reduces belief in fraud claims, suggesting that motivated resistance to corrective information does not seem to explain fraud belief endurance. Finally, using individual-level web trace data paired with our survey data we measure the content of fraud-related information online and patterns of exposure to it during and after the 2020 and 2022 U.S. elections.

Our online exposure data consist of all URLs viewed by participants, including links they clicked on social media (but not social media content itself, which we cannot access). Building on prior research analyzing information exposure at the domain level using online behavior data (Guess et al. 2020; Moore et al. 2023), we instead examine exposure to fraud-related content at the article level using coding by a large language model of online behavior data. Our results illustrate clear patterns of inattention and differential exposure. About half of Americans were not exposed to fraud-related content at all in their online media diet (excluding social media). Among those who saw any fraud-related content, most encountered more online content questioning fraud claims, but Trump supporters saw relatively more non-questioning content, including after exposure to questioning content.

The persistence of election fraud misperceptions

We first evaluate how beliefs in election fraud changed, and endured, over time, drawing on data from six survey waves, from August 2020 to January 2023. Figure 1 tracks over-time levels of belief that there were thousands of cases or more of six distinct types of election fraud among both Trump and Clinton/Biden supporters. 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 voters endorsing claims of widespread fraud had decreased slightly among Trump supporters but remained alarmingly high given the lack of evidence supporting the beliefs reported by participants. Absolute levels and over-time trends are similar when we instead disaggregate by Trump approval and partisan identification — see Figures S1 and S2 in the Appendix.

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 always reaching 85% or more. Stability is consistently lower when measuring the perceived prevalence of officials preventing absentee voters from voting (71–76% in post-2020 election waves). Additional visualizations of over-time stability in election fraud beliefs are provided in Figures S3 and S4.

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

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

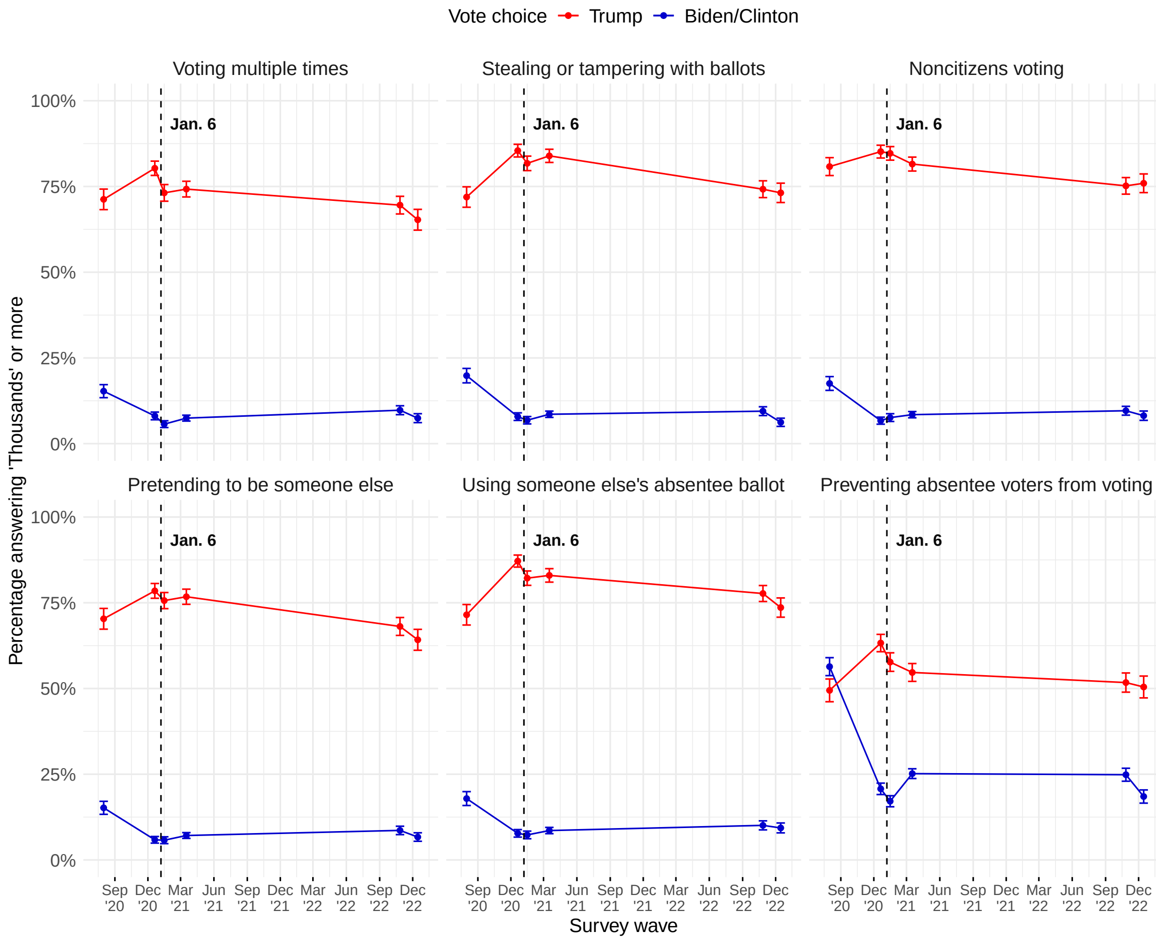

Percentage of respondents indicating that there are "Thousands" of cases or more of each type of election fraud. Categories included "Less than ten" (post-2020 election waves only), "Less than a hundred," "Hundreds,""Thousands," "Tens of thousands," "Hundreds of thousands," and "A million or more". Fraud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

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). Participants were then asked to evaluate the truthfulness of eight targeted statements about election fraud (four false and four true) on four-point scales. (For additional details, see the Materials and Methods section.)

Based on previous research, we expected that fact-checks would be effective at reducing belief in false claims, increasing belief in accurate information, and improving discernment between the two (Carnahan et al. 2021; Chan et al. 2017; Coppock et al. 2023; Walter et al. 2020; Walter and Murphy 2018; Wood and Porter 2019).

As shown in Figure 2, fact-check exposure significantly *decreases* mean beliefs in false claims about election fraud (β = -0.20, 95% CI=-0.23,-0.17) and *increases* both mean beliefs in true claims (β = 0.12, 95% CI=0.09,0.16) and overall truth discernment (β = 0.33, 95% CI=0.28,0.37), confirming our expectations. These effects correspond to -0.18, 0.16, and 0.19 standard deviations, respectively. When we disaggregate these results by item, 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. (Note: Effects are consistent when computing simple differences of means — see Figure S5.)

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

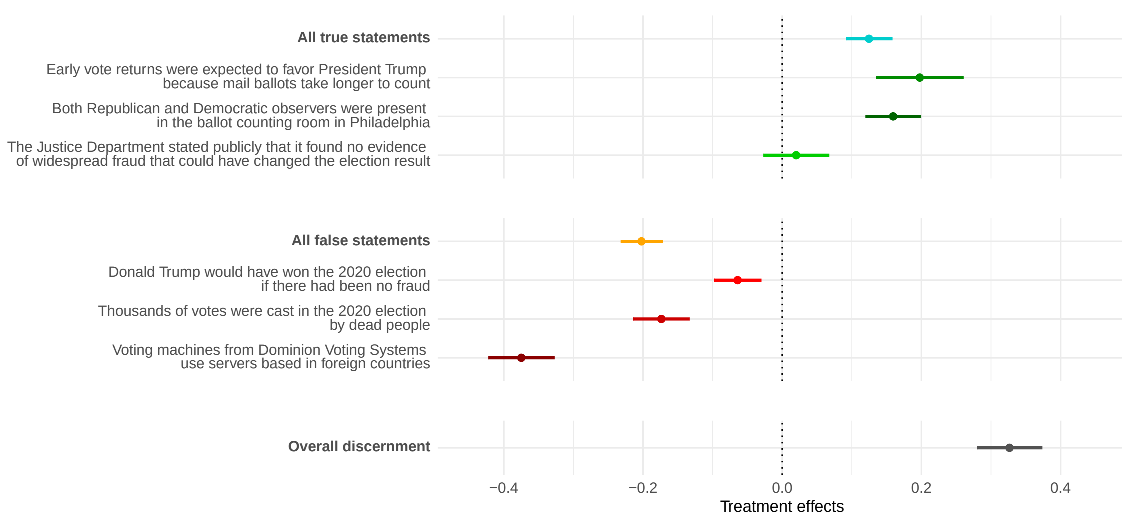

Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of targeted statements. OLS regression coefficients with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Outcomes measured on a four-point scale ranging from "1-not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

There is no indication that these effects vary by whether the content of the fact-check is congenial or not. Figure S6 in the Appendix reports average treatment effects among preregistered subgroups of interest. Consistent with previous studies (Carey et al. 2022), the positive effects of fact-checks on truth discernment are observed 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 questioning content and fact-checks about election fraud (Figure S21).

Prevalence of content questioning fraud claims: 2020–2022

Given the lack of evidence that people resist fact-checking of fraud claims, we next examine the content of fraud-related information online. What we find runs counter to fears of uncritical amplification.

Specifically, we use a large language model to code all fraud-related online media content encountered by our study participants as to whether it questions claims of widespread fraud. To qualify as questioning fraud claims, an article needed to include a statement questioning or contradicting claims that election fraud is widespread or could change the outcome of one or more elections under current or past practices (see Materials and Methods; Appendix Section S4.7 provides coding details). We also consider how this content varies by source trustworthiness.

The fraud-related content our study participants encountered online overwhelmingly questioned claims of widespread fraud. These results are summarized in Figure 3, which plots both the share of unique fraud-related articles seen that question fraud claims and the share of views of fraud-related articles that question fraud claims.

Overall, more than 80% of the unique fraud-related articles our participants encountered questioned fraud claims in both elections. Correspondingly, more than 80% of views of fraud-related articles were of articles that questioned fraud claims. These values exceeded 90% in each case for articles from sources we classify as trustworthy, which we define as sources with NewsGuard ratings above 60 or corresponding Lin et al. (2023) ratings (see Materials and Methods). By contrast, the percentage of articles and views questioning fraud claims from sources we classify as untrustworthy was lower but still well above 50% (57.6% and 60.4% in the 2020 election data and 66.6% and 68.5%, respectively, in the 2022 data).

Figure 3: Percentage of articles and views of election fraud content that question fraud claims by source type and year

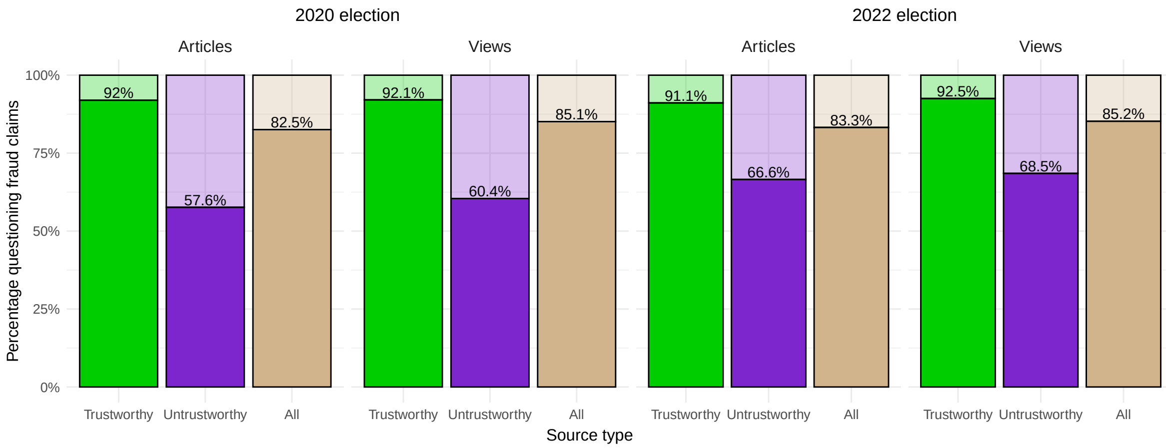

The first three columns for the 2020 and 2022 elections show the percentage of unique election fraud articles seen by respondents that include a statement questioning claims of widespread fraud, while the last three columns for 2020 and 2022 show the percentage of views of election fraud articles questioning fraud claims. Data disaggregated by source trustworthiness based on NewsGuard and Lin et al.'s (2023) ratings: green represents trustworthy sources, purple represents untrustworthy sources, and tan represents all sources combined. Darker shades indicate articles and views that include a statement questioning claims of widespread fraud, while lighter shades represent the remaining non-questioning articles and views.

Exposure to fraud-related content: 2020–2022

Incidence of exposure versus non-exposure to different types of fraud-related content

We next turn to individual-level patterns of exposure of fraud-related content, examining whether or not participants encountered different types of fraud-related content in their web browsing, including content that questioned narratives of widespread fraud, and measuring how exposure levels vary between Trump and Biden supporters.

The first pattern we observe 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. As Figure 4 shows, only about half of our participants (55% of Biden supporters and 44% of Trump supporters) encountered any fraud-related articles from online sources during the period we monitored around the 2020 election. We find distinct patterns of selective exposure to (and/or avoidance of) election fraud content by candidate support: Trump supporters were significantly less likely to encounter fraud-related articles than were Democrats despite having similar overall levels of news consump-

Figure 4: Exposure to specific types of election fraud content

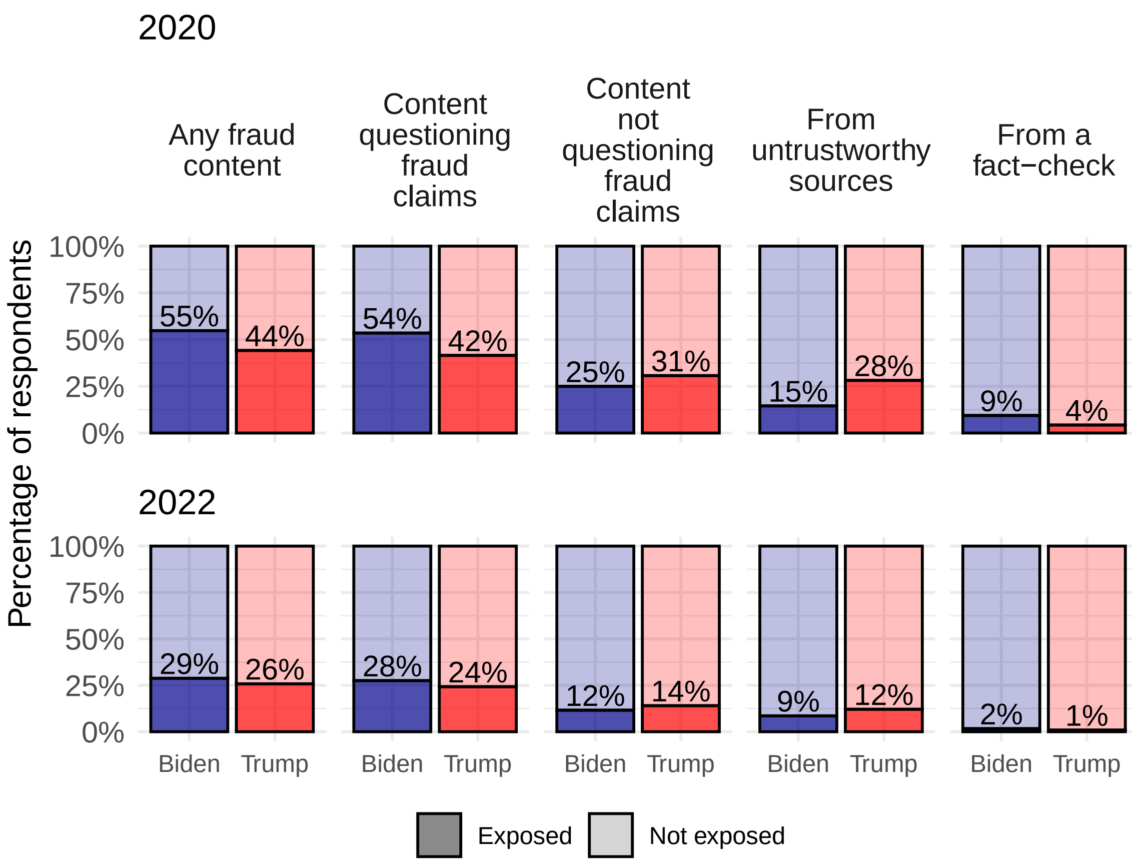

Weighted percentages based on post-stratification weights.

tion (Figure S14). During the 2022 midterms (when interest was lower and fraud claims were less prevalent), fraud-related content exposure rates dropped to 29% and 26%, respectively.

To test for selective exposure, we next consider whether or not people encountered different types of fraud-related content by candidate support. (Figure S12 shows the same statistics conditional on exposure to election fraud content.) Moving from left to right in Figure 4, we see first that Biden supporters were more likely than Trump supporters to have seen any content that questioned fraud claims (54% to 42% in the 2020 election data, 28% to 24% in the 2022 data). Correspondingly, Trump supporters were more likely to encounter content that failed to question fraud claims (31% to 25% in the 2020 election data, and 14% to 12% in the 2022 data). In the 2020 data, Trump supporters were almost twice as likely as Biden supporters to have seen fraud-related content from untrustworthy sources (28% to 15%), whereas Biden supporters were more than twice as likely to view at least one fact-check article about fraud claims (9% to 4%). Both of those gaps narrowed substantially in 2022: 12% of Trump supporters encountered any fraud-related content from untrustworthy sources versus 9% for Biden supporters; the corresponding figures were 1.7% versus 0.9%, respectively, for fraud-related fact-checks.

These findings suggest that the people who were most vulnerable to false beliefs about voter fraud (Trump supporters) were less likely to have encountered any information online questioning the claims of their preferred candidate.

The challenges posed by low exposure to corrective information may be reinforced by the timing and sequence of exposure to fraud-related content. For instance, prior research shows that science news coverage can increase belief accuracy about climate change, but subsequent exposure to skeptical opinion content eliminates its effects (Nyhan et al. 2022). We find that Trump supporters who saw content that questioned fraud-related claims were less likely to have those doubts reinforced and more likely to have them undermined by the content they saw immediately afterward.

As Figure 5 illustrates, Trump supporters saw systematically different types of content in the week after exposure to an article questioning fraud claims. During that post-exposure period, Trump supporters were more likely than Biden supporters to encounter non-questioning content (in both 2020 and 2022) and less likely to encounter questioning content (only in 2022). Specifically, the mean participant-level rate of exposure to non-questioning content in the week after skeptical content exposure was higher for Trump supporters than for Biden supporters (34% versus 17% in the 2020 election data; 24% versus 8% in the 2022 data). In the 2020 data, moreover, the rate of exposure to questioning content in the week afterward was higher for Biden supporters (68% versus 51% for Trump supporters; rates were similar in 2022 at 52.0% and 51.3%, respectively). Relatedly, Figure S17 in the Appendix shows that the probability that a participant encountered content questioning fraud claims decreases based on the order of views of fraud-related content among Trump supporters, while it remains stable over time among Biden supporters. The observed patterns of exposure suggest a failure among Trump supporters to lock in the accuracy effects of content that pushed back against fraud claims.

Overall, the patterns of exposure we observe suggest that content questioning fraud claims fre-

Figure 5: Differences in subsequent fraud content exposure between Biden and Trump supporters

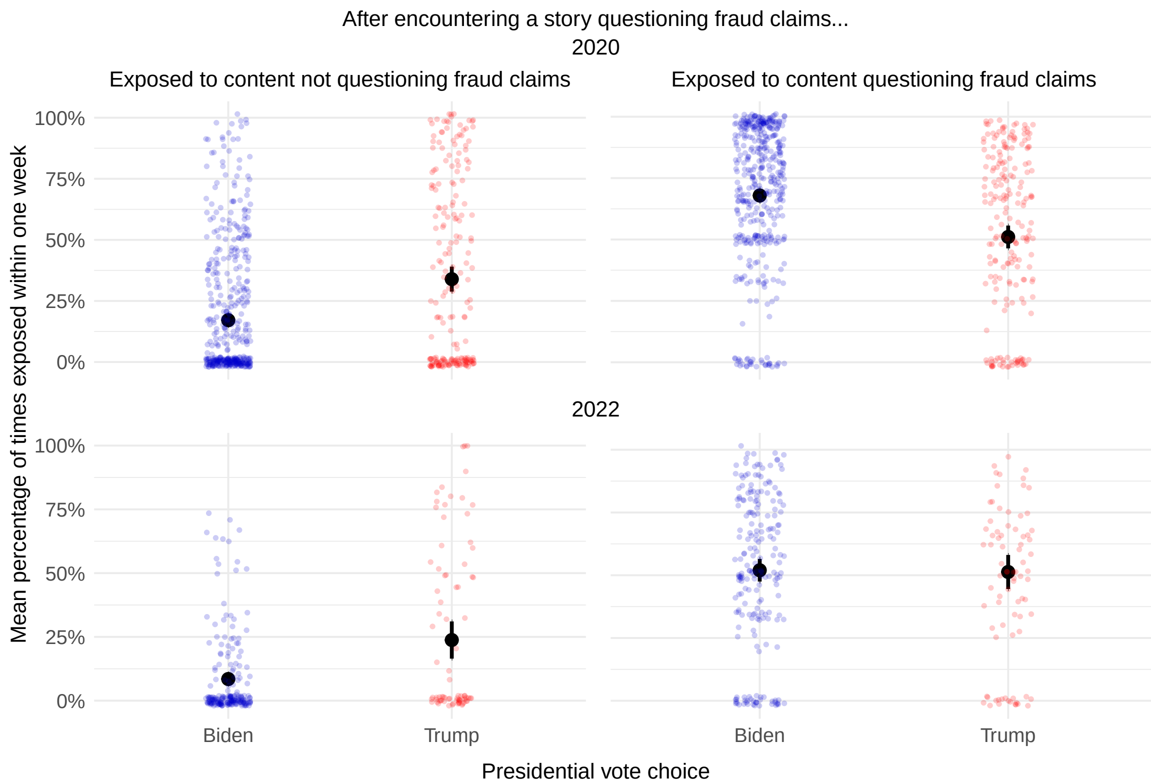

Mean percentage of times exposed to content non-questioning and questioning fraud claims within one week of exposure to content questioning fraud claims. Calculated from participant-level means after exposure to content questioning fraud claims; see Materials and Methods for more.

quently failed to reach people who would be most vulnerable to the false narrative of widespread fraud — a potential explanation for the belief persistence we observe in longitudinal survey data. We confirm this by separating participants directly by their beliefs about whether Biden was the rightful winner of the 2020 election. As Figure 6 shows, people who believe Biden was not the rightful winner in 2020 were less likely to have encountered any stories questioning fraud claims (40.6% versus 48% among those who affirmed the 2020 result). Similarly, they were only about half as likely to have viewed any fact-checks as people who endorsed the legitimacy of Biden's victory (3.8% versus 7.9%). (Patterns of exposure are similar between those who believe that Biden is "definitely not" versus "probably not" the rightful winner —see Figure S13.)

In short, while false claims about election fraud were widely questioned by online information

Figure 6: Exposure to fraud fact-checks and skeptical coverage by 2020 election beliefs

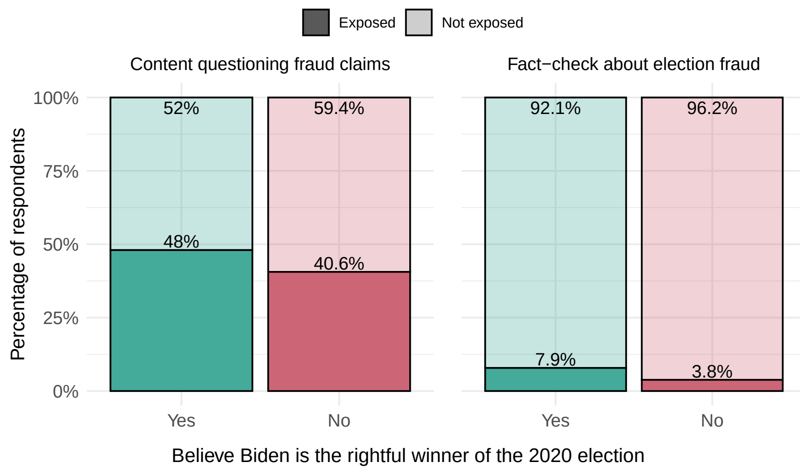

Weighted percentages based on post-stratification weights.

sources, many Americans never encountered skeptical coverage or fact-checks of these claims. People who supported Trump and those who deny the legitimacy of Biden's victory were especially unlikely to encounter questioning content. When Trump supporters did encounter such content, they were more likely to be subsequently exposed to non-questioning content over the next week. These patterns of inattention and offsetting exposure seemingly contribute to the persistence of misperceptions that we observe in longitudinal survey data.

Balance of exposure to questioning versus non-questioning content

Unlike the previous analyses, which examined the prevalence of exposure to *any* fraud-related online media content, we now examine variation in *how much* questioning versus non-questioning fraud-related content participants encountered in their web browsing and how those exposure levels vary by candidate support.

Biden and Trump supporters encountered notably different information, conditional on exposure. As Figure 7 shows, 97% of Biden supporters and 70% of Trump supporters who saw fraud-related content during the 2020 election cycle consumed more online stories that questioned fraud claims than stories that did not. By contrast, 23% of Trump supporters and only 2% of Biden supporters read more stories that presented fraud claims without challenge. In 2022, when fraud content exposure rates were lower, we again see that Biden supporters encountered more content challenging fraud claims than not (89% versus 75% of Trump supporters) and Trump supporters reached more content that does not question fraud claims than content that does (13% versus 4% of Biden supporters).

Figure 7: Individual relative exposure to different types of election fraud content

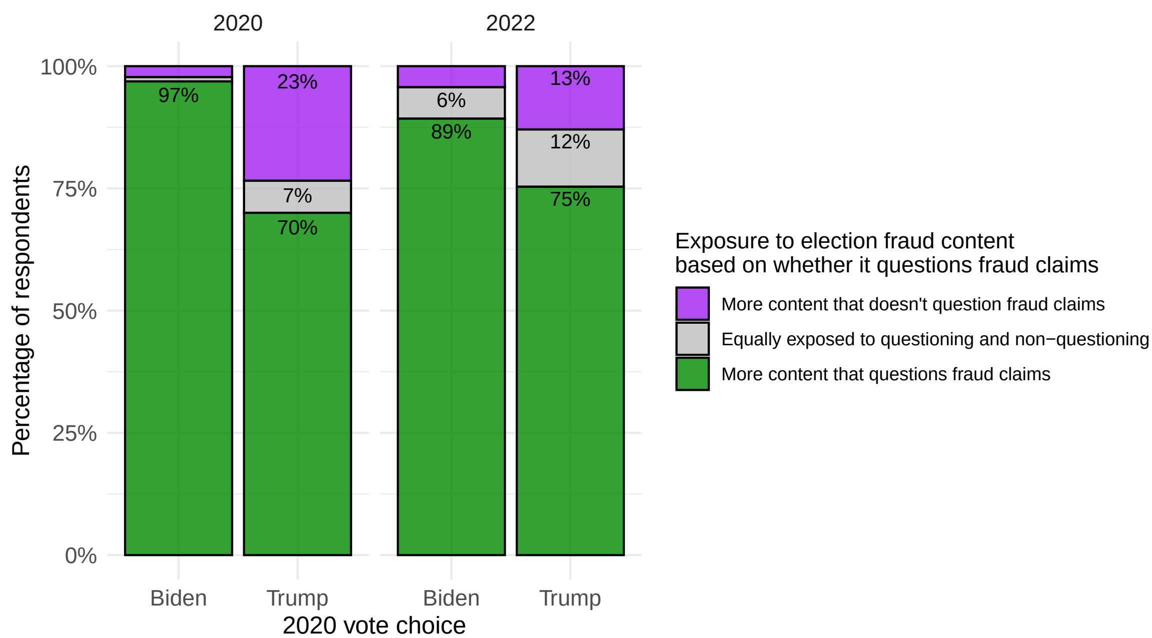

Weighted percentages based on post-stratification weights, restricted to users exposed to election fraud content.

We next consider exactly how much content of each type people saw rather than just classifying people based on their relative balance of content consumption. Figure 8 plots the full distributions of participant-level views of fraud-related articles that questioned fraud claims (horizontal axis) versus views of articles that did not question those claims (vertical axis) among Biden and Trump supporters (see Figure S15 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 (with zero views transformed to an infinitely small number) to better illustrate the distributions, where most participants are clustered at low numbers of page views, while a small subset of highly engaged participants viewed dozens or even hundreds of fraud-related articles.

On average, both Biden and Trump supporters were exposed to more questioning than non-questioning

Figure 8: Views of questioning and non-questioning fraud content by election and candidate preference

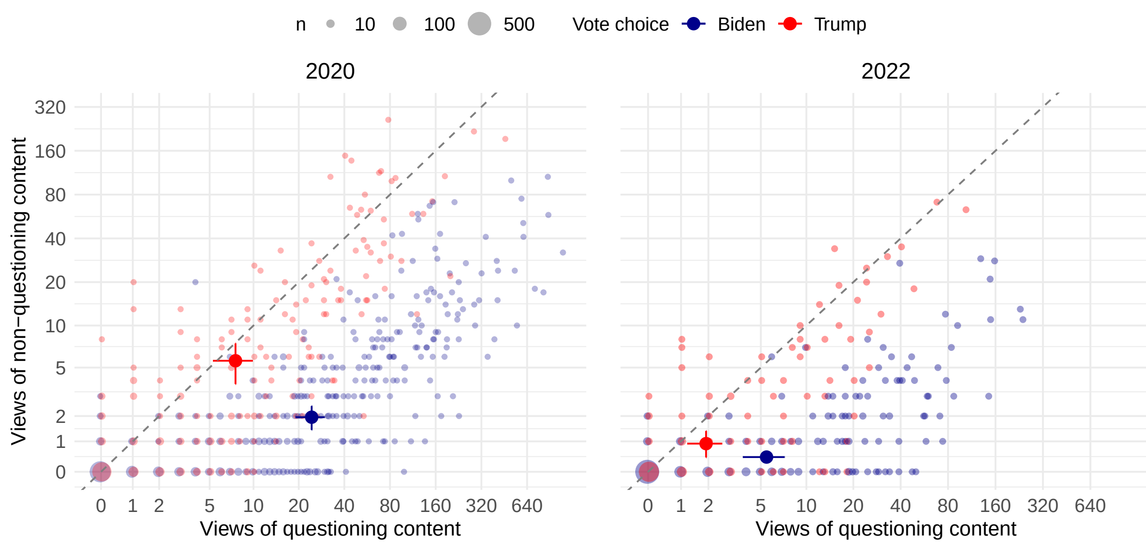

Highlighted markers represent mean values, with 95% confidence intervals, for Trump and Biden supporters. The dashed line represents the 45-degree line and aims to show how many participants were more exposed to one type of content or the other.

content. Trump supporters viewed a mean of 7.6 questioning articles versus 5.6 non-questioning articles in the 2020 election data and 1.9 versus 0.9 in the 2022 election data. The scarcity of points above the 45-degree line reinforces the point from Figure 7 that few Trump supporters – and almost no Biden supporters – in our data were exposed to more non-questioning content than questioning content. By contrast, Biden supporters encountered more election fraud content overall than did Trump supporters and a greater amount of questioning content in both the 2020 and 2022 elections – a mean of 24.4 questioning articles versus 1.9 non-questioning articles in 2020 and 5.5 questioning versus 0.5 non-questioning in 2022. However, the partisan differences observed in Figure 8 are partially exaggerated by skewed exposure distributions.

Figure 9 shows how highly concentrated overall election fraud exposure and types of exposure were in the 2020 and 2022 elections. In 2020, 1% of participants were responsible for 29% of election-fraud views, 31% of questioning views, 39% of non-questioning 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 accounting for 47% of election-fraud views, 48% of questioning views, 54% of non-questioning views, and all 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 half (54%) of fact-checking views compared to 77% among Trump supporters, while in 2022 these percentages increase to 81% and 100%, respectively.

These results demonstrate that the vast majority of both candidates' supporters were exposed to more election-fraud content that questioned claims than did not question claims during the 2020 and 2022 elections, but the relative balance of the exposure diverged along partisan lines.

Figure 9: Empirical cumulative distribution functions of exposure to election fraud content by election and candidate preference

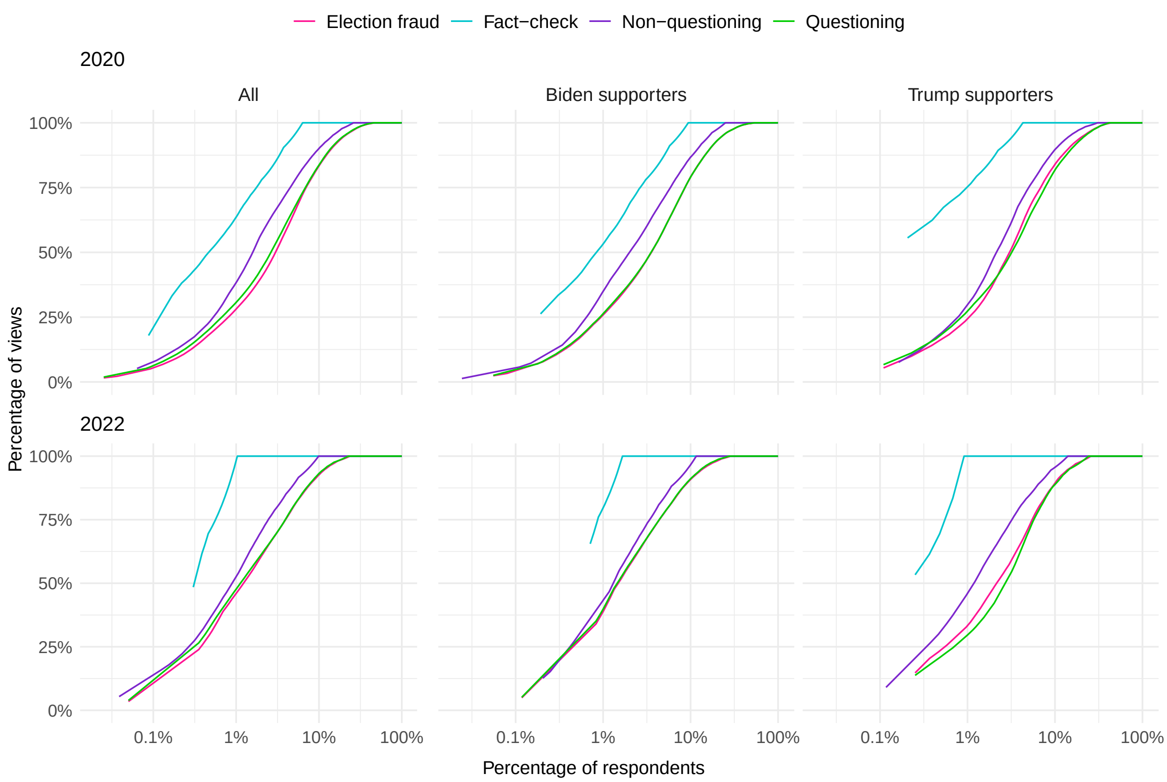

Weighted cumulative distribution functions based on post-stratification weights. The x-axis represents weighted percentage of participants responsible for a given percentage (y-axis) of all exposures. The lines have different starting points because the participant(s) with the most views do not have 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 S16.

How outlet and article choices affect differences in exposure to questioning content

Why do we observe differences in aggregate exposure to fraud-questioning online media content between Biden and Trump supporters? One explanation emphasizes differences in exposure to untrustworthy websites by candidate support (Guess et al. 2020; Moore et al. 2023). Source-level partisan variations in news consumption like these are 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 add to this scholarship by using our article-level data to measure the relative roles of differences in exposure between types of outlets (trustworthy versus untrustworthy) and variations in article selection within each outlet type.

Table 1 illustrates the differences between Biden and Trump supporters in the proportion of fraud-related content they saw during the 2020 election that questions claims of widespread fraud using frequency weights to account for varying levels of exposure (See Materials and Methods for more detailed explanation. Table S11 reproduces the results with probability weights). As previously discussed, Biden supporters have a significantly greater proportion of their views of election fraud content that question fraud claims than Trump supporters (92% vs. 59%). On average, when accounting for exposure levels, Biden supporters received 86% of their views of election fraud content from trustworthy sources compared to 43% among Trump supporters. Conversely, Trump supporters received a significantly larger proportion of their views from untrustworthy sources (57%) than Biden supporters (14%).

We also observe significant differences in exposure to election fraud articles within each source type. When consuming content from trustworthy sources, Biden supporters saw a greater proportion of content questioning fraud claims (94%) than did Trump supporters (75%). However, the largest difference in exposure to content questioning fraud claims occurs among untrustworthy sources. Only half of Trump supporters' views from these sources questioned fraud claims, compared to 84% of Biden supporters' views. Therefore, in addition to receiving a small proportion (15%) of fraud-related content during the 2020 election from untrustworthy sources, the content from such sources encountered by Biden supporters disproportionally questioned fraud claims.

Table 1: Differential exposure to fraud content between Biden and Trump supporters in the 2020 election cycle

Variables	**Biden**	**Trump**	**Diff**
Percentage of views questioning fraud claims	92.02	58.89	33.13
Source-level			
Percentage of views from trustworthy sources	86.39	42.98	43.41
Percentage of views from untrustworthy sources	13.61	57.02	-43.41
Article-level			
Percentage of trustworthy views questioning fraud claims	93.95	75.31	18.64
Percentage of trustworthy views not questioning fraud claims	6.05	24.69	-18.64
Percentage of untrustworthy views questioning fraud claims	84.04	50.15	33.89
Percentage of untrustworthy views not questioning fraud claims	15.96	49.85	-33.89

Note: Means calculated with frequency weights (number of views of election fraud articles).

We use Blinder-Oaxaca decomposition to evaluate how much of the 33 percentage point difference in the proportion of views questioning fraud claims between Biden supporters (92%) and Trump supporters (59%) reflects differences in rates of source exposure (trustworthy sources versus not) versus the choice of articles within sources (articles that question fraud claims versus not).

The results are presented in Figure 10. We again use frequency weights so these estimates reflect individual-level differences in total fraud-related exposure (results are consistent when we instead use probability weights — see Figure S18). The endowment term estimates the effect of Trump supporters being more likely to view content from untrustworthy sources. The results suggest that if Trump supporters received the same proportion of their election fraud views from untrustworthy sources as did Biden supporters and the effect of exposure remained the same, the gap in the percentage of fraud-questioning views would be reduced by nearly two-thirds — 20 percentage points (out of 33). However, that still leaves a large gap in the amount of views of questioning content between Biden and Trump supporters through article-level choices.

The coefficient term measures how differences in the rate of exposure to questioning content *within* source type (i.e., trustworthy versus untrustworthy) contribute to the overall Biden/Trump supporter gap in exposure to questioning content. The results suggest that if Trump supporters had the same pattern of article exposure to questioning content within source types as Biden supporters but main-

Figure 10: Decomposing differential exposure to skeptical fraud content between Biden and Trump supporters

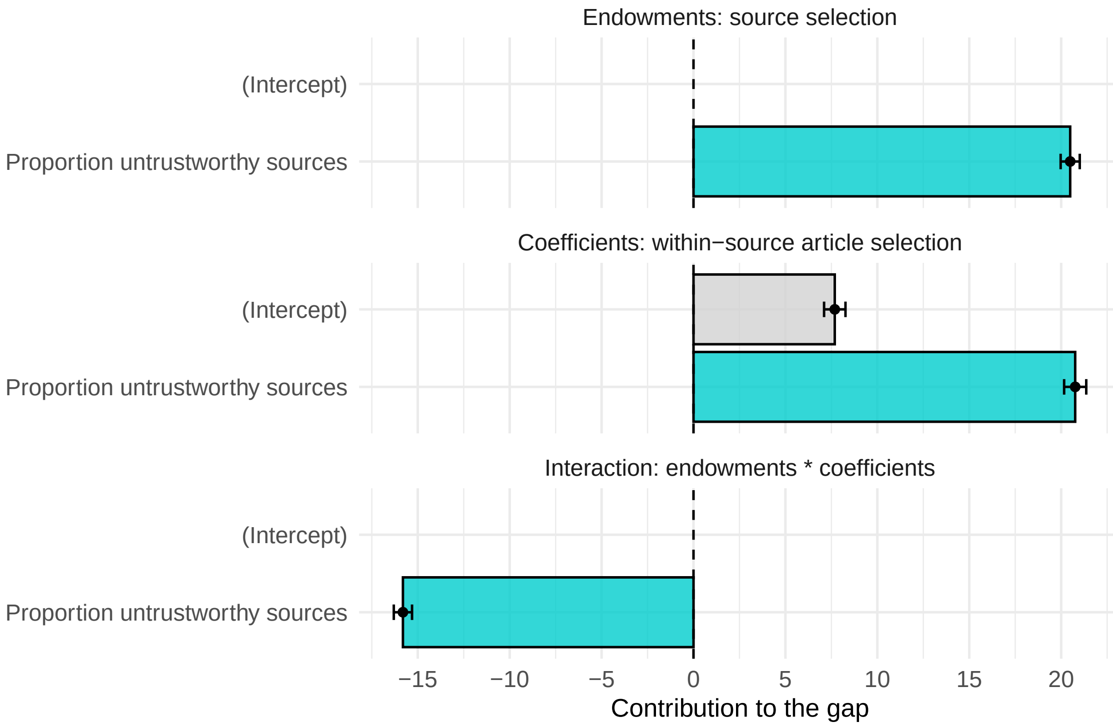

Blinder-Oaxaca decomposition of exposure to skeptical fraud-related content based on the proportion of fraud-related views from untrustworthy sources and Trump support. Estimated with frequency weights representing the number of views of fraud-related content.

tained their share of views coming from untrustworthy sources, the gap in exposure to skeptical content would be reduced by 21 percentage points. Therefore, the negative relationship between exposure to untrustworthy sources and exposure to fraud-questioning content is stronger among Trump supporters than it is among Biden supporters. There is also an estimated intercept of eight percentage points that reflects the expected difference between Biden and Trump supporters in exposure to skeptical fraud content conditional on rates of exposure to trustworthy versus untrustworthy sources and difference in exposure to questioning content conditional on type.

Finally, the negative interaction term of -16 percentage points illustrates that the combined effects of source exposure (endowments) and article selection (coefficients) partially cancel each other out. This finding likely reflects a ceiling effect: since most content viewed from trustworthy sources questions fraud claims, Biden supporters' higher exposure to these sources has a limited effect on increasing their proportion of fraud-questioning views.

Overall, we find evidence that the gap in exposure to fraud-questioning content between Trump and Biden supporters is about equally explained by differences in their rates of exposure to trustworthy versus untrustworthy sources and in their article selection within source types.

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? Prior research identifies two demand-related factors that could explain this persistence: selective exposure and directionally motivated resistance to accurate information. Other observers worry that the prevalence of accurate information is limited due to media practices that amplify or fail to challenge false claims. However, we find that the online media content consumed by both Biden and Trump supporters in their web browsing around the 2020 and 2022 elections overwhelmingly questioned unfounded fraud claims, making it unlikely that access to accurate information was a primary cause. We also found that exposure to fact-checks of fraud claims increased belief accuracy among both Biden and Trump supporters, suggesting that directional motivations not preventing the processing of corrective information.

A more compelling explanation is a lack of demand for political news in general. Our data show that most people saw little to no online media content about election fraud even in the immediate aftermath of the 2020 election and January 6 insurrection. We additionally find support for a more nuanced understanding of how selective exposure may affect the persistence of fraud beliefs. We demonstrate that while Trump supporters were exposed to questioning content, they were differentially exposed to non-questioning content about election fraud around the 2020 and 2022 U.S. elections, including after engaging with more skeptical coverage. This form of relative selective exposure may have reduced or eliminated the effects of exposure to fraud-questioning content, suggesting that even subtle amounts of selective exposure can have significant consequences.

These findings represent a profound challenge for democracy. The claim of widespread election fraud comes from the most high-profile figure in American politics (Donald Trump) and concerns an incredibly salient issue (whether the 2020 presidential election was rightfully decided). News coverage and online content overwhelmingly questioned and debunked this claim. Nonetheless, people's exposure to that coverage was often minimal and the people who were most vulnerable (Trump supporters) were differentially likely to encounter non-questioning content after exposure to questioning content, potentially undermining its effects. Future research should seek to determine the conditions under which news media coverage and related efforts like fact-checking reach broad audiences, including people who mostly or entirely avoid political news, and have lasting effects on their beliefs and attitudes (Bowles et al. 2025; Nyhan et al. 2022).

Of course, this paper has limitations. Although our analysis of the effect of fact-checking on fraud beliefs is causally identified, we cannot establish a causal relationship between broader patterns of exposure to fraud-related content and belief in widespread fraud. We instead assess which potential explanations of the persistence of fraud beliefs are most consistent with the available data. Future researchers should develop externally valid ways to directly estimate the causal effects of behavioral exposure to the factors we identify, including skeptical content (an extremely difficult task, to be sure). Second, we are only able to measure the URLs participants visited in web browsers, but people of course encounter information from other sources such as mobile browsing, social media platforms (both article links/headlines and organic content), television, radio, newspapers, and word of mouth. In the future, scholars should measure exposure to specific information such as fraud claims across as many types of media as possible (Allen et al. 2020) and test other mechanisms of exposure. Third, we are limited by the composition of the YouGov Pulse panel and the online behavior data collected from participants; researchers must improve both the representativeness of online behavior data and measurement of both the content seen and the attention it receives, especially on mobile devices where our data are more limited. Finally, though we followed current best practices, it would be desirable to further improve the accuracy and reliability of LLM text coding and to incorporate the uncertainty of these estimates into our results. Model updates and availability continue to pose a challenge to the replicability of most LLM-based research.

Despite these limitations, this paper makes a number of significant 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 a pattern of "echo chamber"-style information consumption. Instead, our data suggest that the lack of change in beliefs we observe is driven by a combination of widespread inattention and what we call relative selective exposure among vulnerable populations, who still see mostly skeptical content but may be exposed to enough non-questioning content to undo or offset the effects of critical coverage.

Finally, we move beyond prior source-level analyses of exposure to untrustworthy online outlets and instead measure whether and how misinformation is presented in news articles across source types. To accomplish this, we demonstrate that large language models can be applied to 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 (2,983 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 2; see Table S1 for demographics. We draw participants from three studies in total:

- **A prior study of attitudes toward COVID-19** consisting of a YouGov sample built from three

Table 2: 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

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.

- **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 Trump support and to measure 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. For measuring fraud beliefs, we rely on two main measures. 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." 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 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, ranging 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 Section S3 in the Appendix).

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 Appendix S2.2.

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

Digital trace 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 consists 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 period (1,596 participants in the 2020 election data; 1,518 participants in the 2022 data). Appendix S4.1 shows that most participants stayed active throughout the entire election 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 Appendix 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. In 2022, we have laptop/desktop data for 965 participants (64%) and mobile (smartphone, tablet, other) data for 621 participants (41%) (71 participants provided data from both types). 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 S26), but the proportion of fraud-questioning views does not differ by device type (Figure S10).

Behavioral 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 ($N = 612,806$) 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 Appendix S4.6) to identify possible fraud-related information. We then used GPT-4o mini (2024-07-18 version) to classify these articles (32,963 articles in the 2020 election data; 10,938 articles in the 2022 data) on two dimensions: 1) whether the article discusses election fraud, and 2) whether it includes a statement questioning or contradicting claims that election fraud is widespread or could change the outcome of one or more elections. We use a broad definition of questioning that includes, among other things, qualifying fraud claims as "baseless," "unfounded," "false," "unsupported," or "absurd"; describing those making these claims as election deniers; questioning the legality or legitimacy of efforts to change or interfere with the certification of election results; or describing courts rejecting or denying election fraud lawsuits. Mentions of efforts to overturn or refuse to concede the 2020 election and descriptions of the certification process or January 6 insurrection are not classified as questioning content if they do not challenge fraud allegations. We validated the LLM output by comparing it to the consensus coding from four human coders for a random sample of 100 articles (50 from 2020 and 50 from 2022), ensuring it met minimal reliability thresholds. For identifying whether an article mentioned election fraud, the LLM achieved 87% agreement with human coders and a Krippendorff's alpha of 0.72. For determining whether an article included a statement questioning widespread fraud claims, agreement was 89% with an alpha of 0.78. When weighted by the total number of fraud-related articles or views of fraud-related articles from each election in our Pulse data, the percent agreement and Krippendorff's alpha was at least 91% and

0.79.

In light of concerns about the replicability of LLM coding (Barrie et al. 2024), we took several steps to increase confidence in our approach. In line with best practices, we set the temperature at 0 (ensuring the most consistent output over time) and specify exactly which model version we use in this manuscript. We also show that our coding is highly consistent across two different models (see Materials and Methods and Appendix S4). More information about the codebook, how it was developed, and GPT coding accuracy are provided in Appendices S4.7–S4.9.

Trustworthiness We rely on two sources of data for measuring the trustworthiness of the online sources to which participants were exposed: ratings from Lin et al. (2023) and NewsGuard (February 2021 version). 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. NewsGuard provides a binary label (trustworthy and not trustworthy) for 7,109 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 variable. The conversion was performed using the R package `cutpointr` (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.

Fact-checks 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.

More details about the creation of each behavioral variable, including the full list of keywords for fact-check article URLs, are provided in Appendix S4. Descriptive statistics about exposure to all types of content, overall and by candidate support, are included in Appendix S4.11.

Participant-level aggregation

Analyses of the behavioral data described above are conducted at the participant-level except when we consider the prevalence of questioning content, which is instead analyzed at the article and view level.

When we analyze *subsequent* exposure at the participant level, we first identify all views of fraud-related content for each participant. For each view, we determine if the participant encountered non-questioning and questioning content over the next seven days (coded as 1) or not (coded as 0). We then separately calculate the mean of these values by participant, which represents the proportion of times they were subsequently exposed to each type of content after an initial exposure to fraud-related content.

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 (except for the decomposition analysis described below) use probability weights constructed by YouGov to more accurately represent the U.S. adult population. When evaluating within-respondent stability in fraud beliefs, we generally use the weights from the more recent wave for each between-wave comparison, given that their sample size 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. We rely on frequency weights in the decomposition analysis to account for varying levels of total exposure to fraud-related content between participants (both the outcome and explanatory variables in the Blinder-Oaxaca model are proportions that do not take total exposure into account). Using frequency weights ensures that our results reflect actual differences between Biden and Trump supporters in aggregate exposure to questioning and non-questioning fraud-related content.

References

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. OUP Oxford.

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.

Barrie, Christopher, Alexis Palmer, and Arthur Spirling. 2024, December. "Replication for Language Models: Problems, Principles, and Best Practice for Political Science".

Bennett, W. Lance. 1990. "Toward a theory of press-state relations in the United States". Journal of Communication 40 (2): 103–127.

Bennett, W Lance, Regina G Lawrence, and Steven Livingston. 2008. When the press fails: Political power and the news media from Iraq to Katrina. University of Chicago Press.

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, February. "Sustaining Exposure to Fact-Checks: Misinformation Discernment, Media Consumption, and Its Political Implications". American Political Science Review: 1–24.

Boykoff, Maxwell T. 2008. "Lost in translation? United States television news coverage of anthropogenic climate change, 1995–2004". Climatic Change 86 (1): 1–11.

Boykoff, Maxwell T. and Jules M. Boykoff. 2004. "Balance as bias: global warming and the US prestige press". Global Environmetal Change 14 : 125–136.

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/.

Brown, Katie. 2013. "Trayvon Martin, Obama, and the Birthers: How race and suspicion continue over 35 years". Institute for Social Research, Center for Political Studes, 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., Elizabeth Chun, Alice Cook, Brian J. Fogarty, Leyla Jacoby, Brendan Nyhan, Jason Reifler, and Lilian Sweeney. 2024. "The Narrow Reach of Targeted Corrections: No Impact on Broader Beliefs About Election Integrity". Political Behavior.

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 Behavior 43 (3): 1227–1246.

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". The 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.

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.

Fahy, Declan. 2017. "Objectivity, false balance, and advocacy in news coverage of climate change". In Oxford Research Encyclopedia of Climate Science. Oxford University Press.

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.

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, 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, 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.

Graves, Lucas. 2016. Deciding what's true: The rise of political fact-checking in American journalism. Columbia University Press.

Green, Jon, Stefan Mccabe, Sarah Shugars, Hanyu Chwe, Luke Horgan, Shuyang Cao, and David Lazer. 2025. "Curation Bubbles". American Political Science Review.

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, Kyle Richardson, Gregory K Tortoriello, and Allison Earl. 2020. "'You are what you read:' Is selective exposure a way people tell us who they are?". British Journal of Psychology 111 (3): 417–442.

Hayes, Danny and Matt Guardino. 2010. "Whose views made the news? Media coverage and the march to war in Iraq". Political Communication 27 (1): 59–87.

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.

Kunda, Ziva. 1990. "The case for motivated reasoning". Psychological bulletin 108 (3): 480–498.

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 and Jason Reifler. 2010. "When corrections fail: The persistence of political misperceptions". Political Behavior 32 (2): 303–330.

Nyhan, Brendan, Jason Reifler, and Peter A. Ubel. 2013. "The hazards of correcting myths about health care reform". Medical Care 51 (2): 127–132.

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.

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.

Stroud, Natalie Jomini. 2008. "Media use and political predispositions: Revisiting the concept of selective exposure". Political Behavior 30 (3): 341–366.

Stroud, Natalie Jomini. 2010. "Polarization and partisan selective exposure". Journal of communication 60 (3): 556–576.

Sunstein, Cass R. 2001. Republic.com. Princeton University Press.

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.

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, December. Avoiding the News: Reluctant Audiences for Journalism. Columbia University Press. Pages: 288 Pages.

Törnberg, Petter. 2024, October. "Best Practices for Text Annotation with Large Language Models". Sociologica 18 (2): 67–85. Number: 2.

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.

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.

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.

Supplementary Information

Contents

S1 Sample characteristics

Table S1 provides information about the demographic composition of each of our survey sample. (Note: Wave 2 of the Prior COVID study was 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.

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 voter	35.5	35.6	44.5	45.8	45.7	42.1	42.7	43.0
Trump voter	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 participants of Waves 2 and 3 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/Vote		
Democrat PID	45.7	47.8
Republican PID	38.4	32.7
Biden voter (2020)	46.6	42.7
Trump voter (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/Vote				
Democrat PID	45.5	43.8	48.7	45.8
Republican PID	39.1	38.7	35.1	30.5
Biden voter (2020)	46.2	44.7	44.8	39.5
Trump voter (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.

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 S1) and partisan identification (Figure S2). 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 measured on a seven-point scale. The "Democrat" and "Republican" categories combine "Lean," "Not very strong," and "Strong" Democrats/Republicans.

Figure S1: Perceived prevalence of election fraud over time by presidential approval

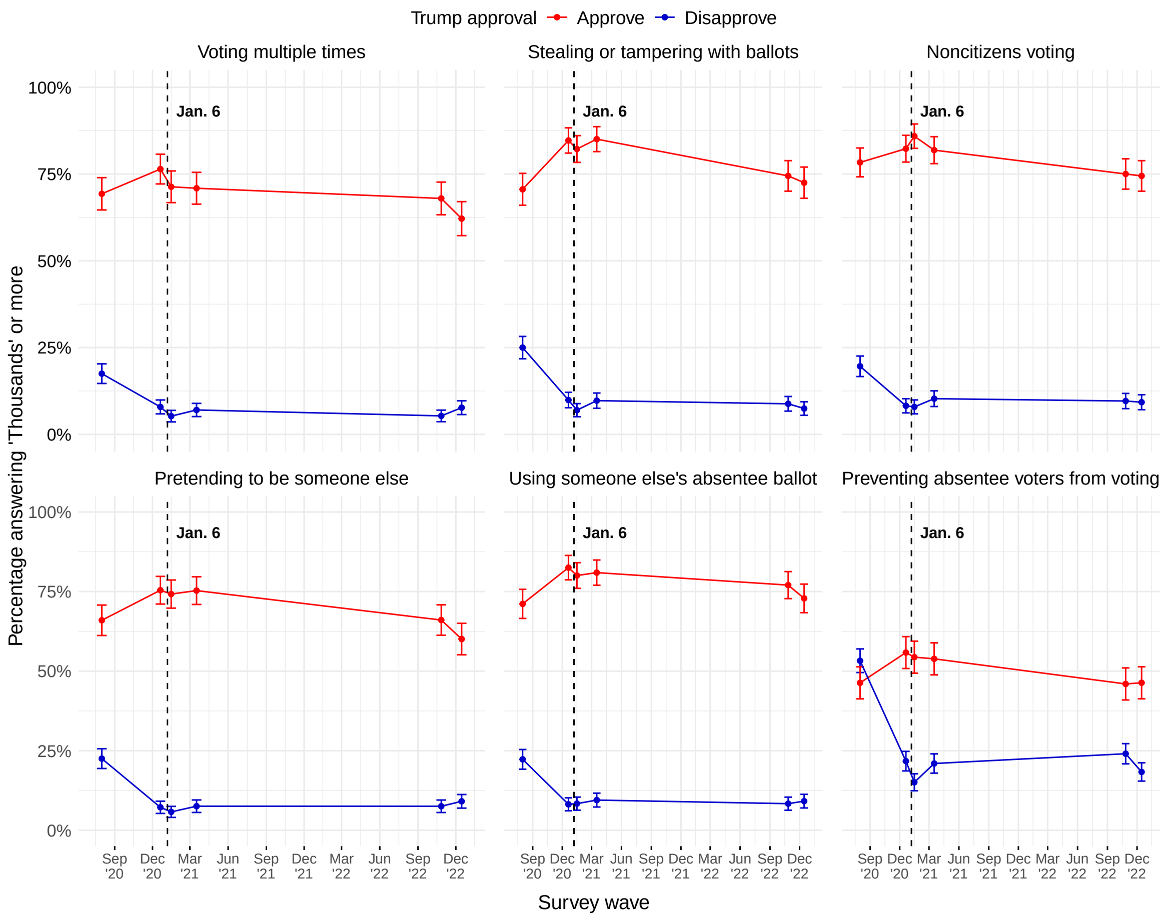

Percentage of respondents indicating that there are “Thousands” of cases or more of each type of election fraud. Categories included “Less than ten” (post-2020 election waves only), “Less than a hundred,” “Hundreds,”“Thousands,” “Tens of thousands,” “Hundreds of thousands,” and “A million or more”. Fraud 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 survey wave (N=1,074).

Figure S2: Perceived prevalence of election fraud by party

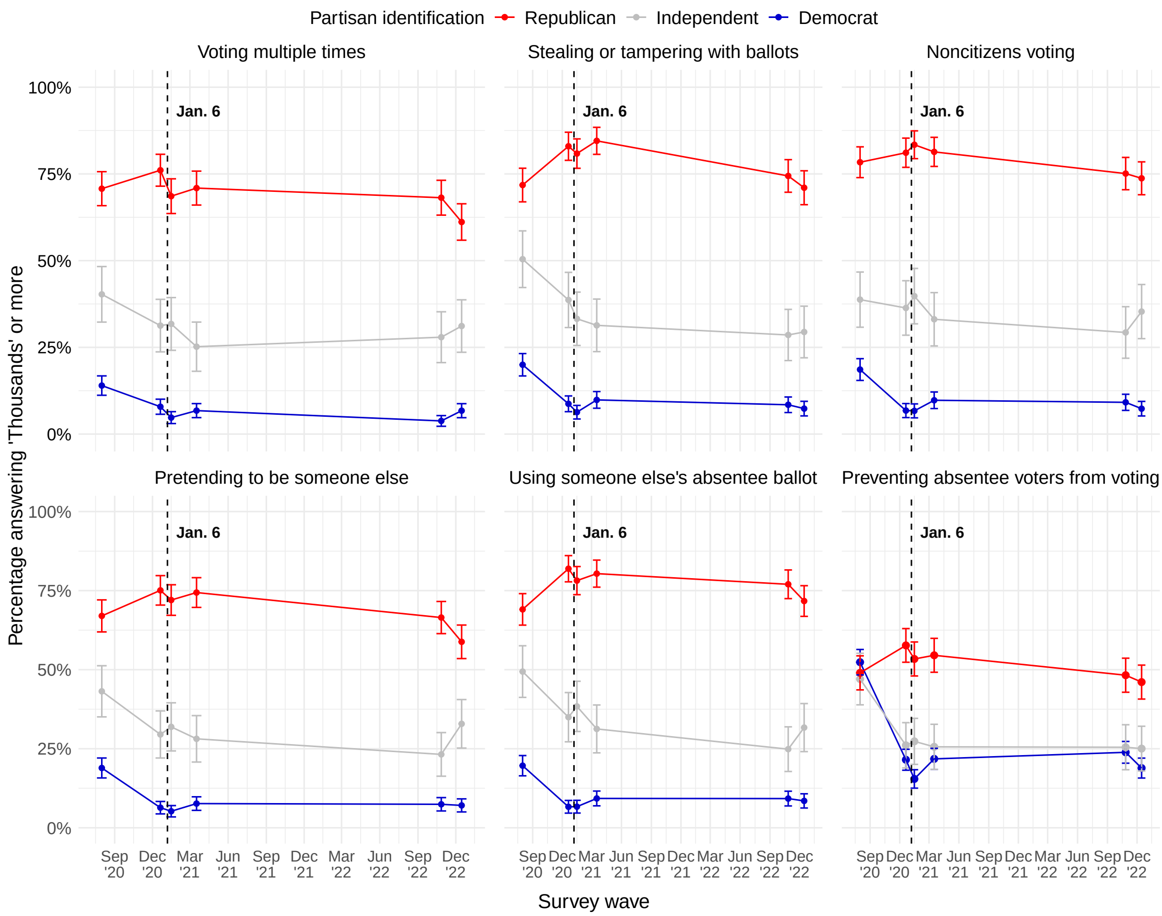

Percentage of respondents indicating that there are "Thousands" of cases or more of each type of election fraud. Categories included "Less than ten" (post-2020 election waves only), "Less than a hundred," "Hundreds,""Thousands," "Tens of thousands," "Hundreds of thousands," and "A million or more". Fraud 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 survey wave (N=1,074).

Figure S3: Within-respondent durability in the perceived prevalence of election fraud

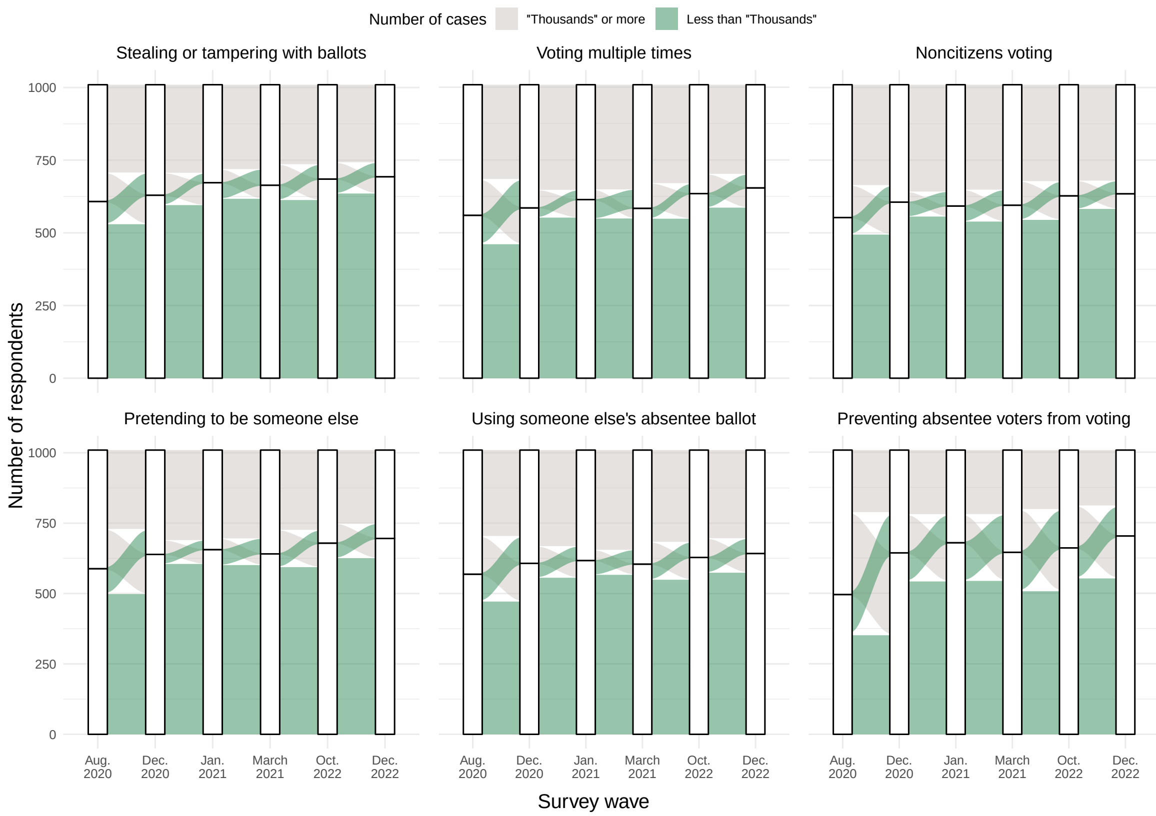

Weighted statistics based on post-stratification weights. Analysis restricted to participants who participated in all survey waves (N=1,074).

Figure S4: Within-respondent durability in perceived election legitimacy

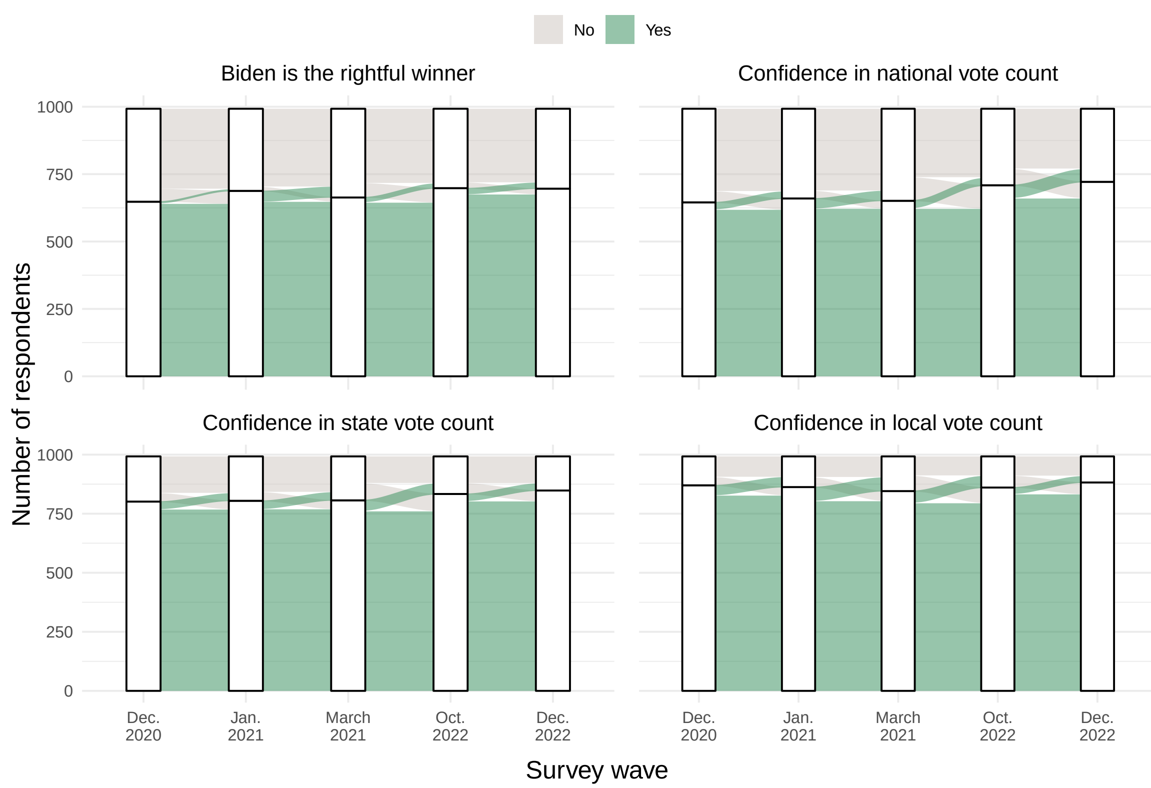

Weighted statistics based on post-stratification weights. Analysis restricted to participants who participated in all survey waves (N=1,074).

S2.2 Preregistered experimental analyses

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 respectively included in Tables S4 and S5. Figure S5 displays preregistered differences of means between the treatment conditions.

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

	Targeted true	Targeted false	Discernment
Fact-check	0.125***	−0.202***	0.327***
	(0.017)	(0.015)	(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 S5: Combined treatment effects of exposure to fact-check article

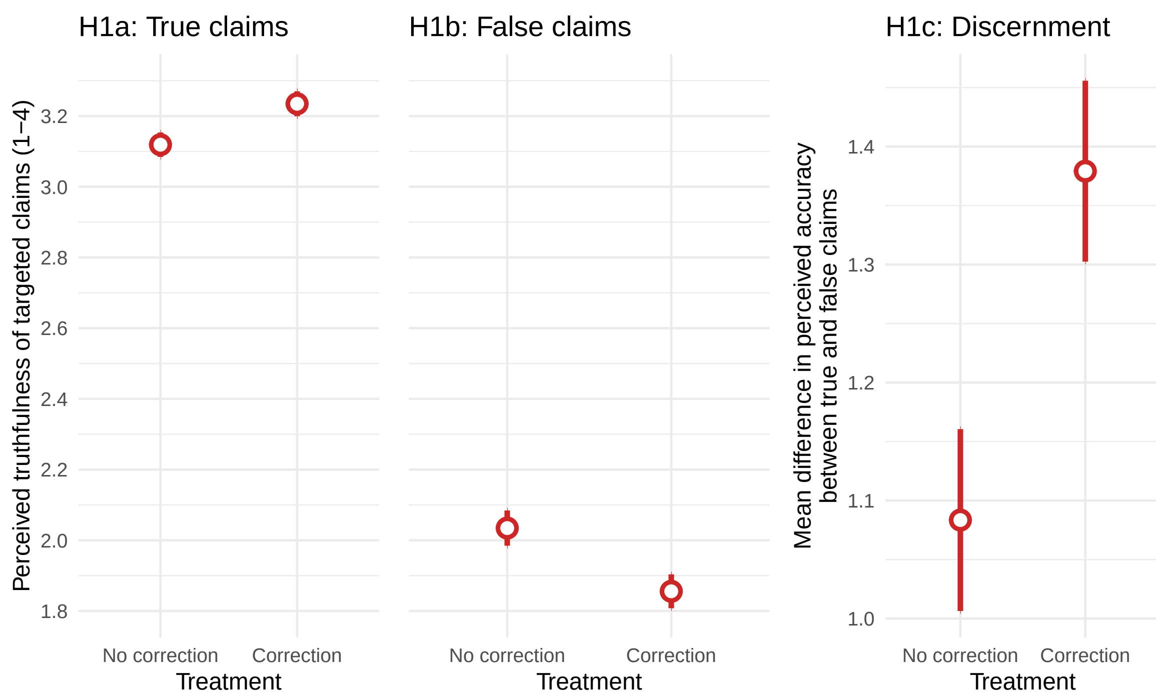

Group means shown with 95% confidence intervals.

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***	0.159***	0.020	−0.064***	−0.174***	−0.375***
	(0.032)	(0.020)	(0.024)	(0.017)	(0.021)	(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 S6 uses Bayesian Causal Forest models to test for preregistered heterogeneous treatment effects (i.e., 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". 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. Conspiratorial predispositions are based on participants' level of agreement (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 to 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 S6: Heterogeneous treatment effects of exposure to fact-check article on truth discernment

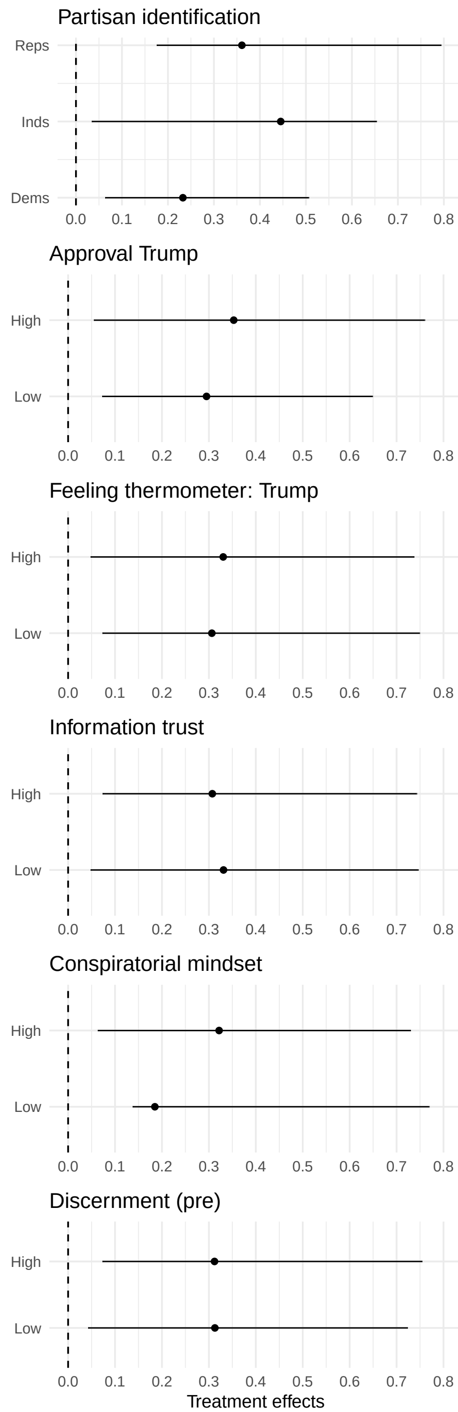

Heterogeneous treatment effects based on Bayesian Causal Forest models. Median effect by group shown with 95% confidence intervals (2.5th and 97.5th percentiles).

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 S7 and Tables S6 and S7).

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

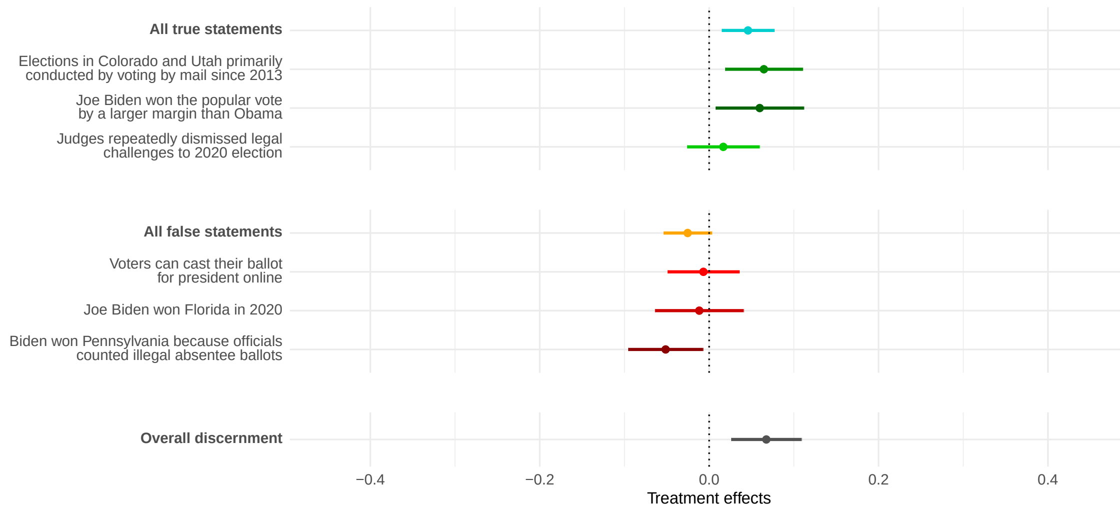

Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of non-targeted statements. OLS regression coefficients shown with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Outcomes measured on a four-point scale ranging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

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^{**}	-0.025	0.067^{**}
	(0.016)	(0.014)	(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**	0.060*	0.017	−0.007	−0.012	−0.051*
	(0.023)	(0.027)	(0.022)	(0.022)	(0.027)	(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 include six questions measuring beliefs about:

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 is based on 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 (4-point scale), what they think of the actions of the people who stormed the U.S. Capitol (3-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 (4-point scale).

Results reported in Figure S8 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. (2024), 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 were more ingrained compared to the period immediately after the election. 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 S8: Treatment effects of exposure to fact-check article on broader attitudes about the 2020 election

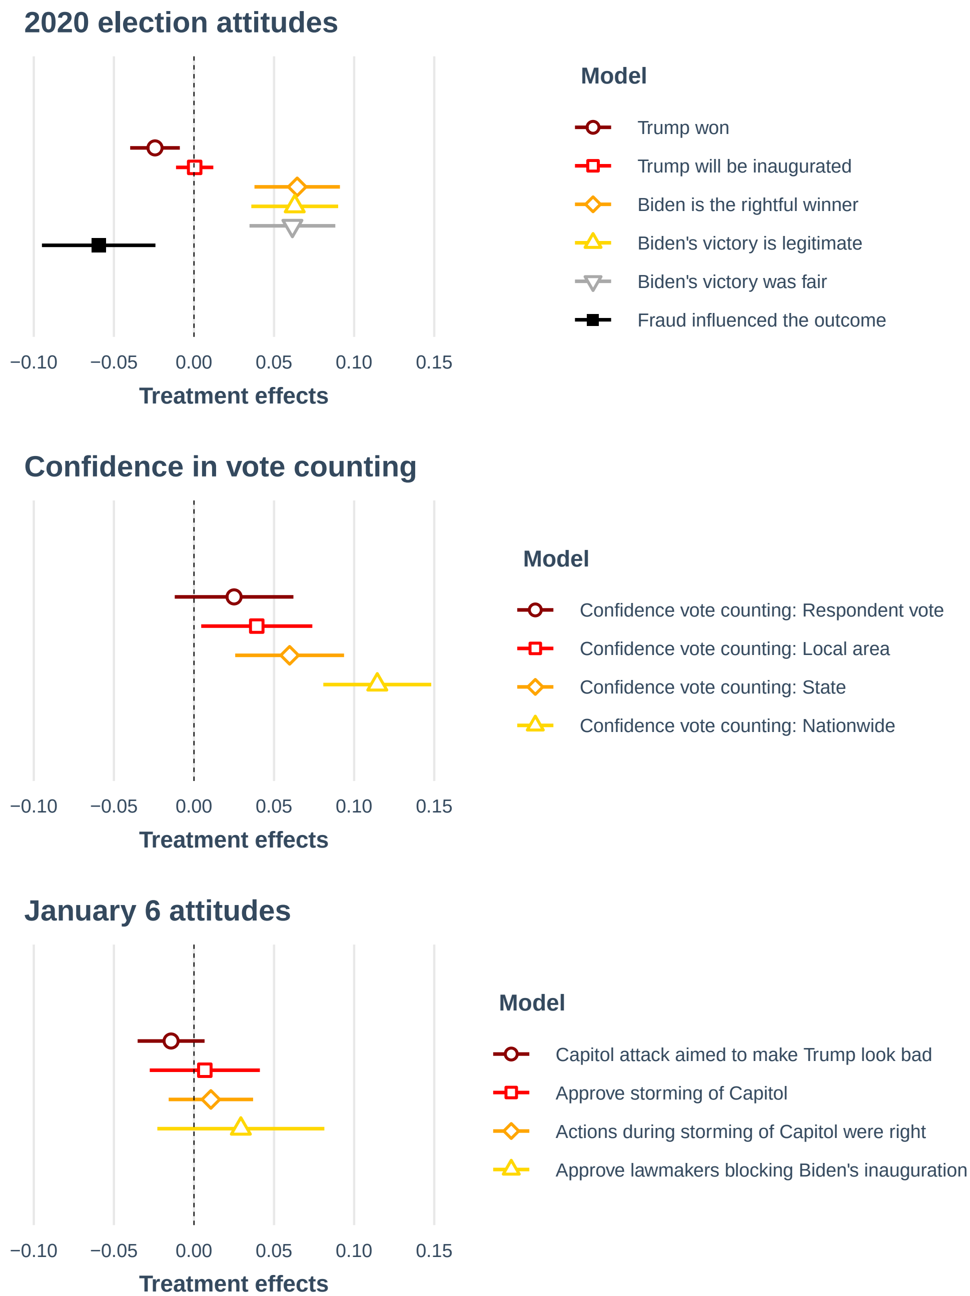

Sample average treatment effects of exposure to a fact-checking article on broader attitudes about the legitimacy of the 2020 election and the January 6 insurrection. OLS regression coefficients shown with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso.

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**	0.000	0.064***	0.063***	0.061***	−0.059**
	(0.008)	(0.006)	(0.014)	(0.014)	(0.014)	(0.018)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	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	0.039*	0.060***	0.114***
	(0.019)	(0.018)	(0.017)	(0.017)
Lasso controls	✓	✓	✓	✓
Num.Obs.	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	0.007	0.011	0.029
	(0.011)	(0.018)	(0.013)	(0.027)
Lasso controls	✓	✓	✓	✓
Num.Obs.	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 (Trump supporters, Block certification), information trust, political interest (Trump supporters, Block certification), and education (Trump supporters).

S2.3 Exposure to election fraud content

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

Figure S9 shows the distribution of all fraud-related news content views by the trustworthiness of the source and by our coding for whether the article itself did or did not question claims that fraud shaped the election outcome. Of all views of fraud-related content, 85% were to content that questioned fraud claims in some way. This skeptical coverage overwhelmingly came from trustworthy sources (71% of all fraud-related views versus 14% for content from untrustworthy sources). The remaining 15% of page views of fraud-related content advanced fraud claims without questioning them. Much of this non-questioning content exposure came from untrustworthy sources (9% of total fraud-related views), though trustworthy sources also contributed (6% of total fraud-related views).

Figure S9: Percentage of views of fraud content that questions fraud claims by source type

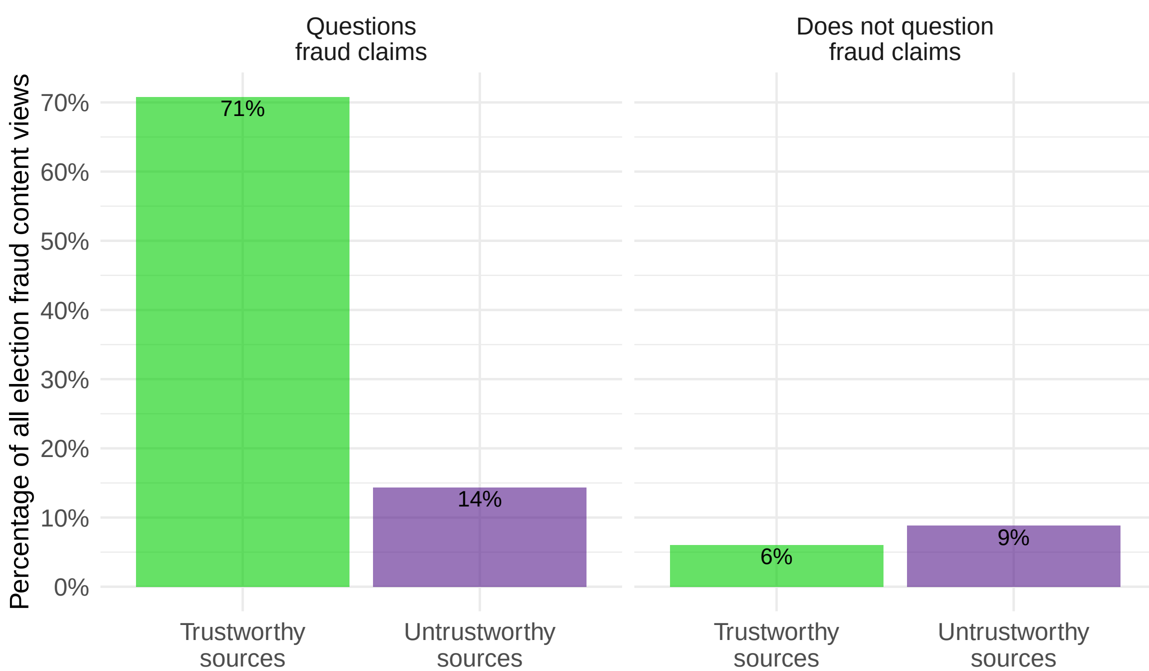

Weighted percentages based on post-stratification weights.

Figure S11 breaks out individual participant-level exposure to election-related content across pre- and post-election periods and by whether each participant has higher- or lower-than-median interest in politics. Figure S12 presents similar measures as in Figure 4, but estimated among the subset of participants with exposure to election fraud content. Figures S13 and S15 disaggregate participants' exposure to election fraud content by the strength of their beliefs that Biden was the rightful winner of the 2020 election. Figure S14 evaluates differences in overall news exposure and exposure to fraud-related content by 2020 candidate preference. Figure S16 replicates the empirical cumulative distribution functions included in Figure 9 without post-stratification weights. Figure S17 further

Figure S10: Percentage of views of fraud content that questions fraud claims by device type

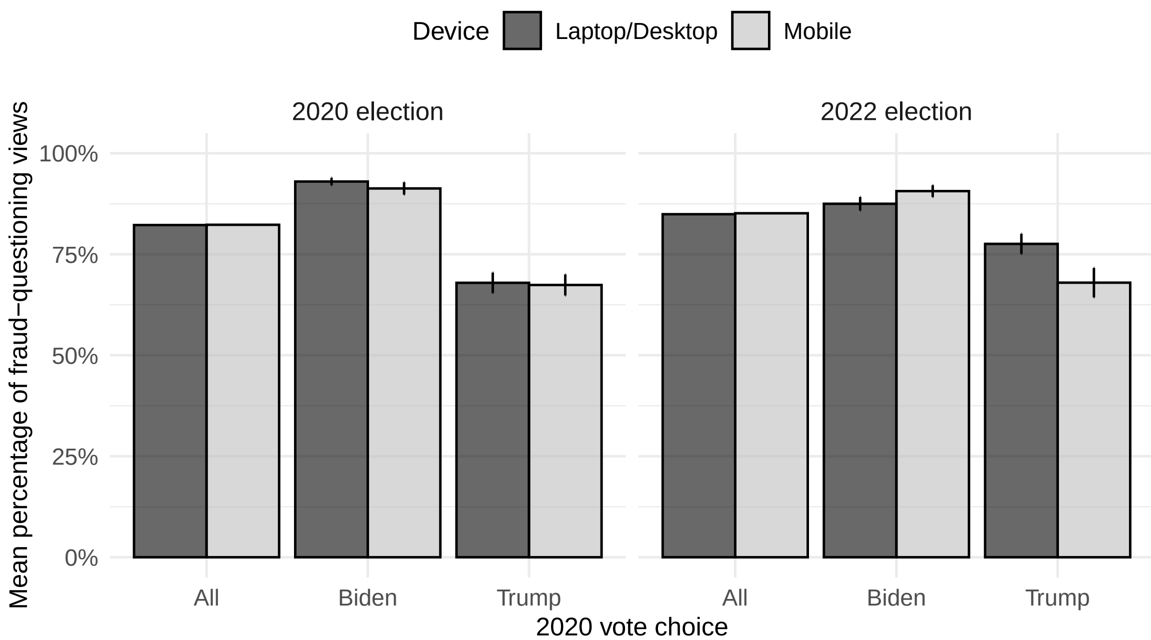

Weighted mean of individual-level percentages. Calculated using post-stratification weights.

evaluates patterns of engagement between the heaviest news consumers (top 10%) and the rest of the sample. In the figure, the horizontal axis indicates the sequence of fraud-related stories encountered up to 15 (if observed for a given participant). The vertical axis indicates the share of stories at that order in the sequence that pushed back against fraud claims. Finally, Table S11 and Figure S18 replicate the Blinder-Oaxaca Decomposition analysis using probability rather than frequency weights.

Figure S11: Exposure to election fraud content based on vote choice during the 2020 U.S. presidential election, political interest, and period

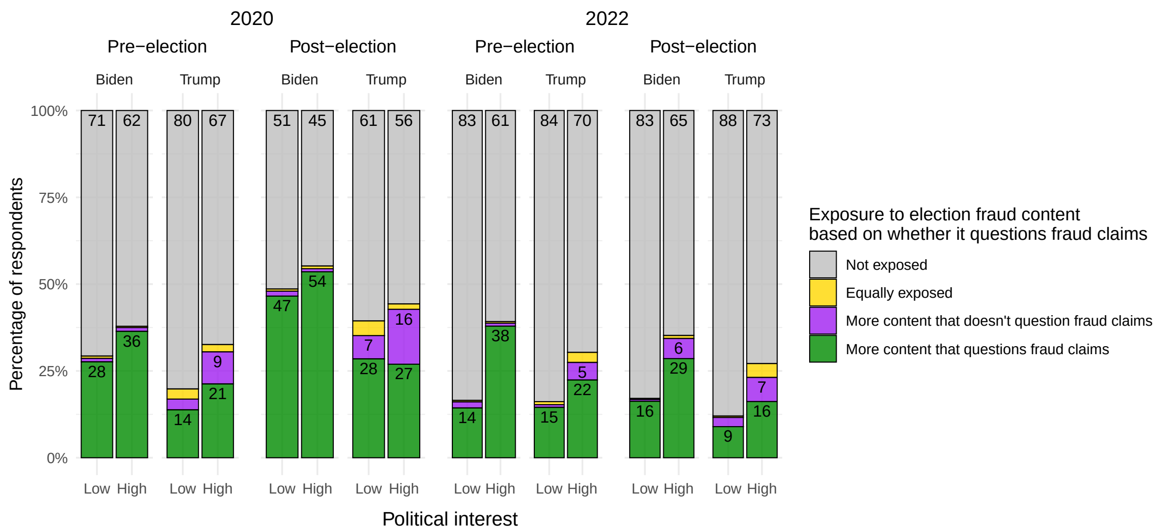

Weighted percentages based on post-stratification weights. Political interest is measured on a five-point scale ranging from “Not at all interested” to “Extremely interested.” We use a median split to categorize the sample into those with low and high interest.

Figure S12: Individual exposure to different types of content conditional on exposure to election fraud content

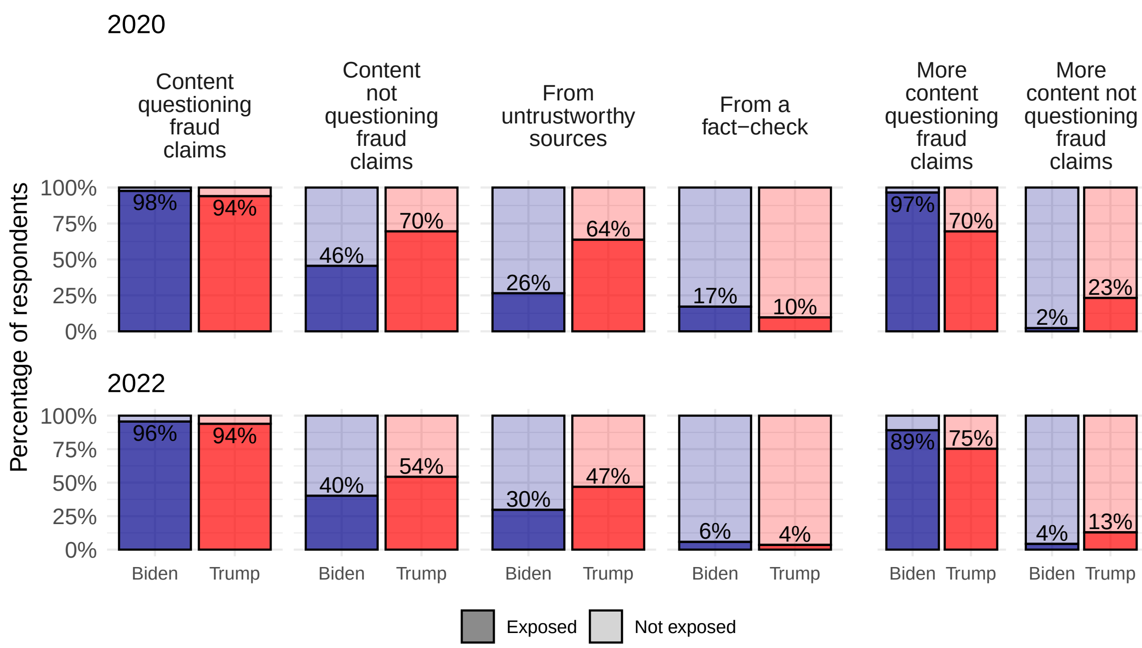

Weighted percentages based on post-stratification weights.

Figure S13: Exposure to specific types of election fraud content by 2020 election beliefs

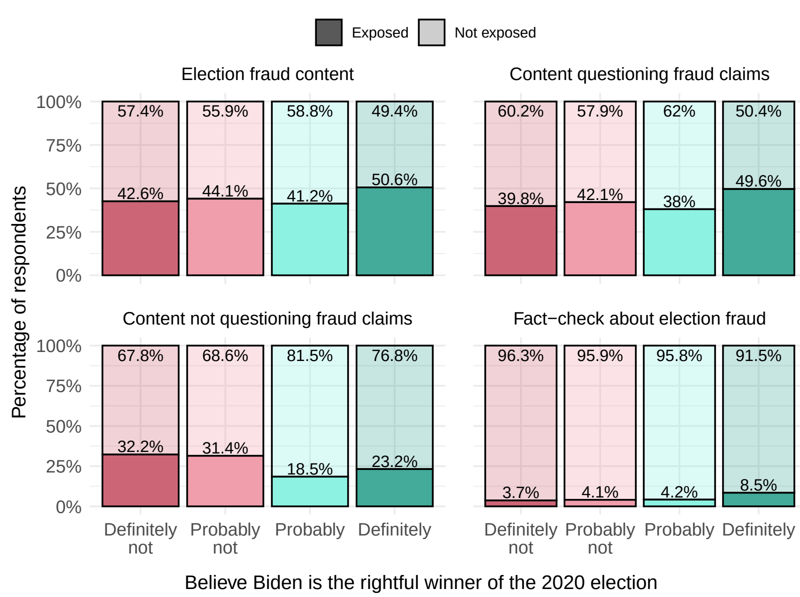

Weighted percentages based on post-stratification weights.

Figure S14: Mean and median exposure to news content and election fraud content during the 2020 election by vote choice

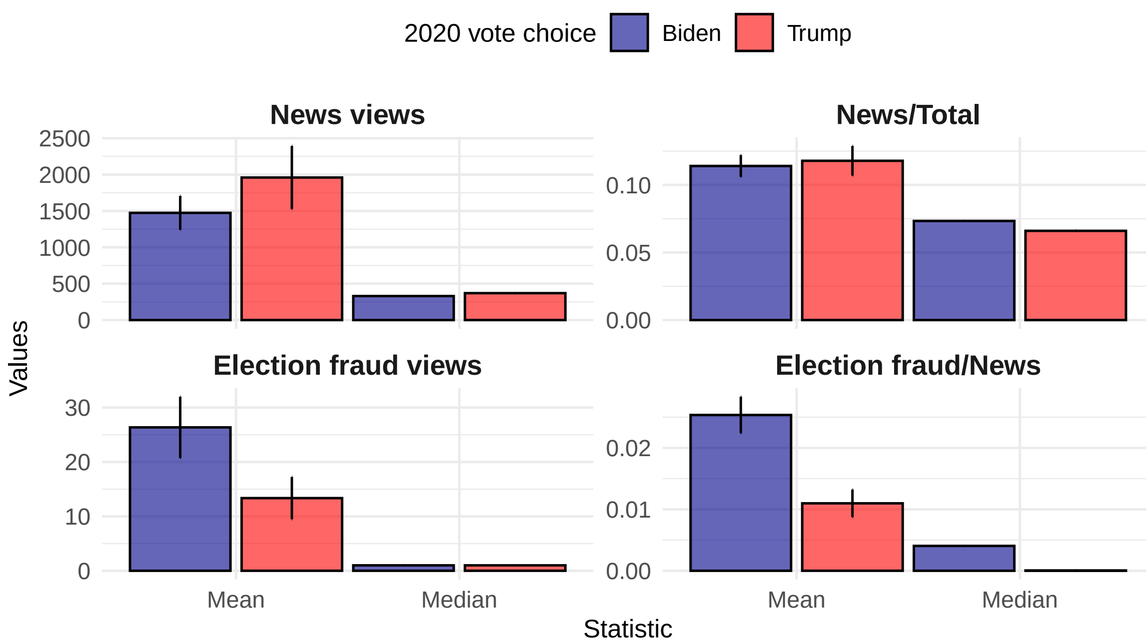

Weighted percentages based on post-stratification weights.

Figure S15: Views of questioning and non-questioning fraud content by election and 2020 election beliefs

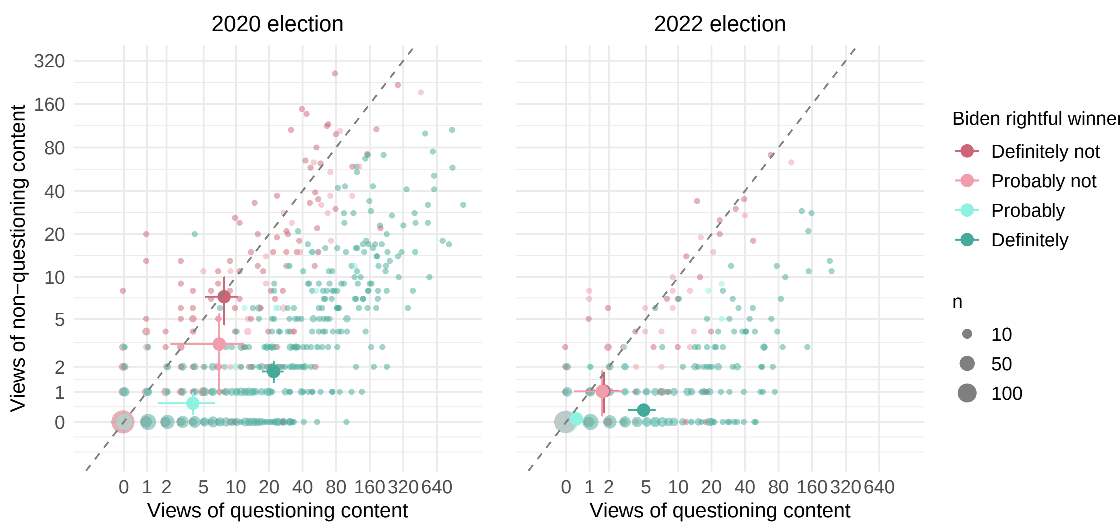

Highlighted markers represent mean values, with 95% confidence intervals, separated based on participants' views on whether Biden was the rightful winner of the 2020 election. The dashed line represents the 45-degree line and aims to show how many participants were more exposed to one type of content or the other.

Figure S16: Empirical cumulative distribution functions of exposure to election fraud content by election and candidate preference

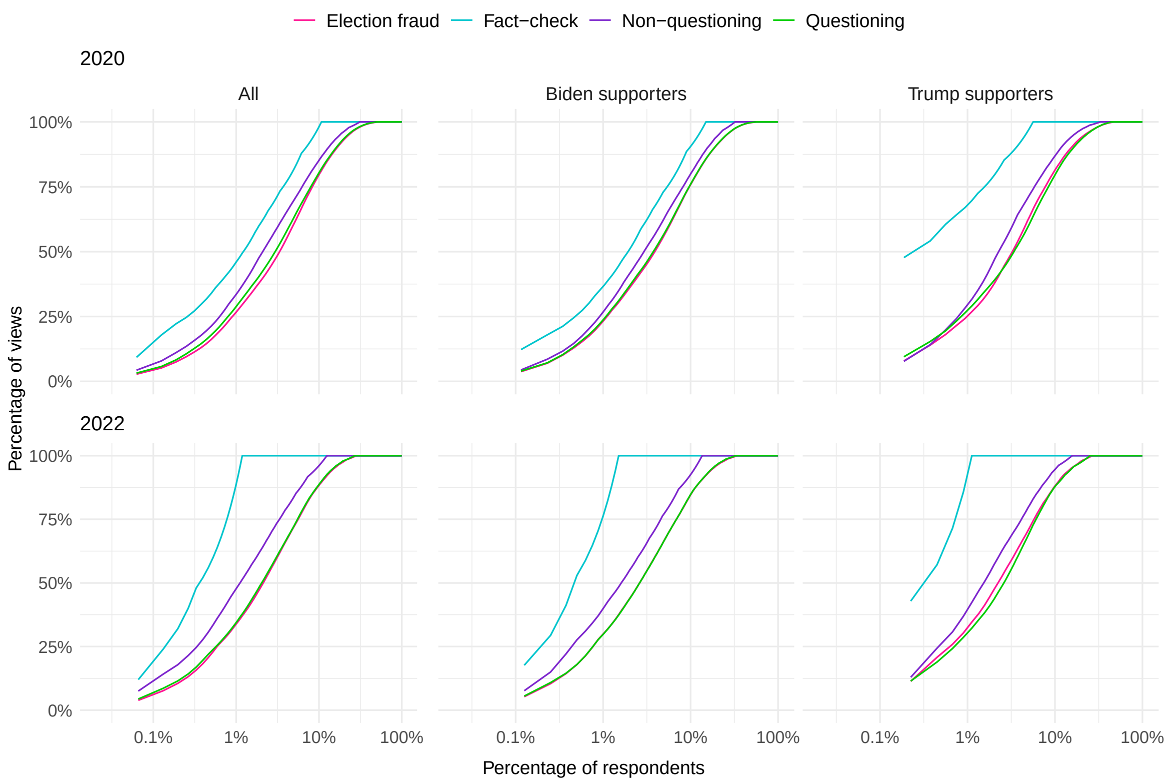

The x-axis represents percentage of participants responsible for a given percentage (y-axis) of all exposures.

Figure S17: How fraud content exposure varies by order of views, vote choice, and news consumption during the 2020 election

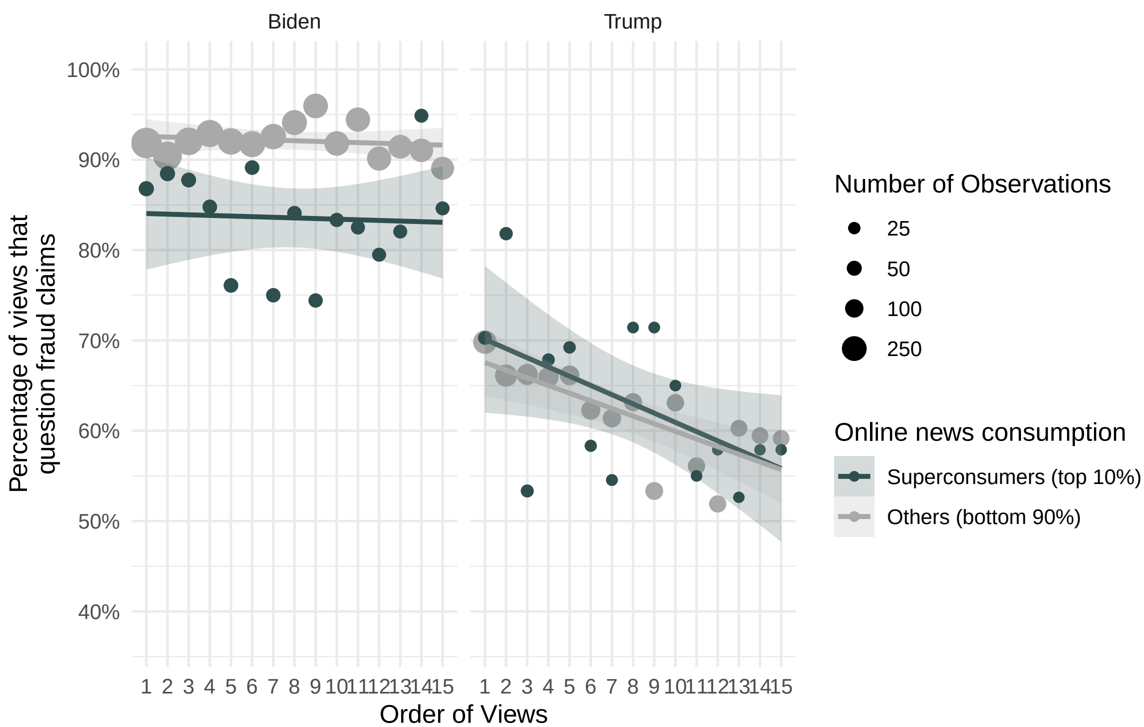

Weighted percentage of views questioning fraud claims, based on post-stratification weights. Quantile split based on participants' number of views of news-related URLs as identified by Lin et al. (2023) and NewsGuard.

Table S11: Differential exposure to fraud content between Biden and Trump supporters in the 2020 election cycle

Variables	**Biden**	**Trump**	**Diff**
Percentage of views questioning fraud claims	92.89	68.73	24.16
Source-level			
Percentage of views from trustworthy sources	94.13	62.75	31.38
Percentage of views from untrustworthy sources	5.87	37.25	-31.38
Article-level			
Percentage of trustworthy views questioning fraud claims	94.07	84.45	9.62
Percentage of trustworthy views not questioning fraud claims	5.93	15.55	-9.62
Percentage of untrustworthy views questioning fraud claims	93.44	62.67	30.77
Percentage of untrustworthy views not questioning fraud claims	6.56	37.33	-30.77

Note: Means calculated with probability weights.

Figure S18: Decomposing differential exposure to skeptical fraud content between Biden and Trump supporters

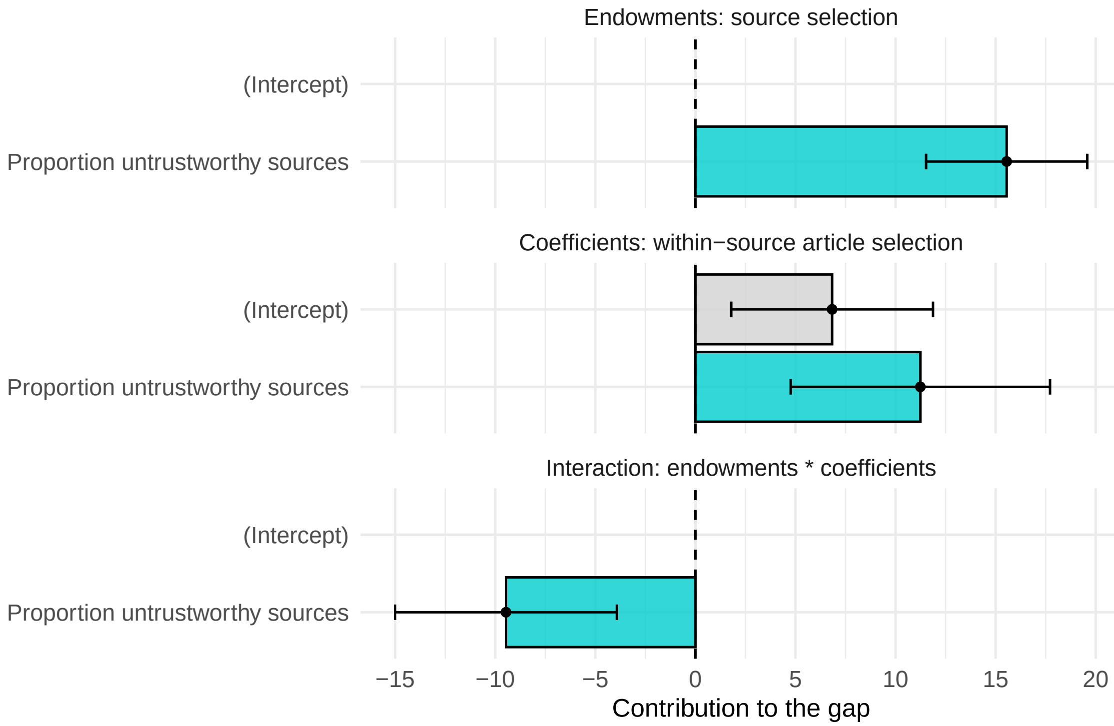

Blinder-Oaxaca decomposition of exposure to skeptical fraud content based on the proportion of fraud views from untrustworthy sources and Trump support. Estimated with probability weights.

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 election cycle.

Figure S19: Treatment effects of exposure to fact-check article among active Pulse participants

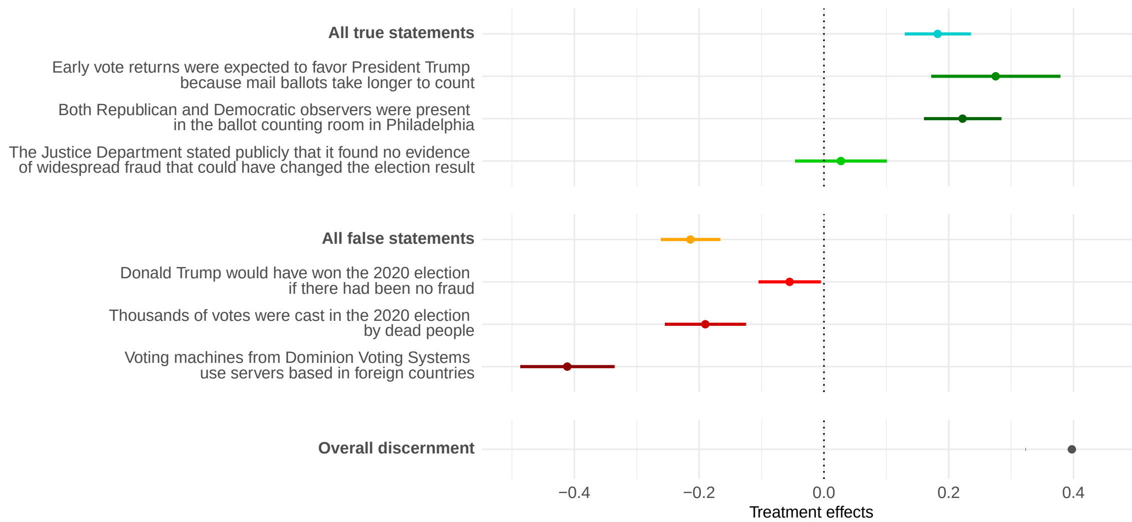

Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of targeted statements. OLS regression coefficients shown with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted to Pulse participants active for the majority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

Table S12: Main treatment effects of exposure to fact-check article among active Pulse participants)

	Targeted true	Targeted false	Discernment
Fact-check	0.182***	−0.214***	0.397***
	(0.027)	(0.024)	(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 S13: 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***	0.222***	0.027	−0.055*	−0.190***	−0.411***
	(0.053)	(0.031)	(0.037)	(0.025)	(0.033)	(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 S20: Combined treatment effects of exposure to fact-check article among active Pulse participants

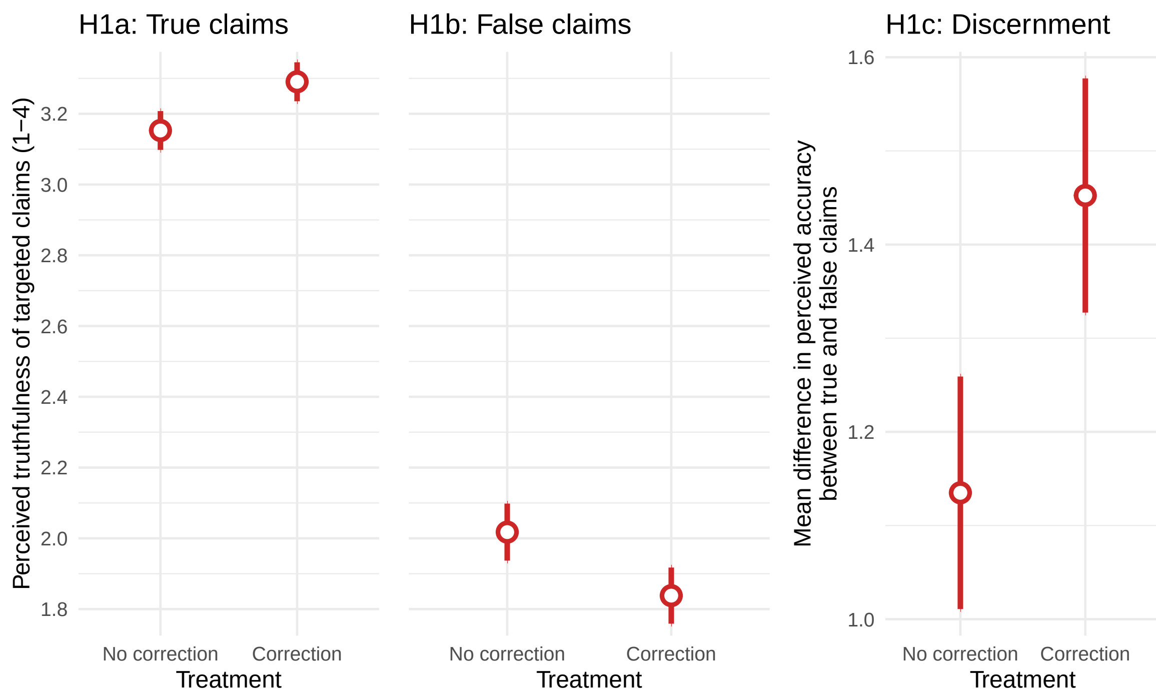

Group means shown with 95% confidence intervals. Restricted to Pulse participants active for the majority of months during the study period.

Figure S21: Heterogeneous treatment effects of exposure to fact-check article on truth discernment among active Pulse participants

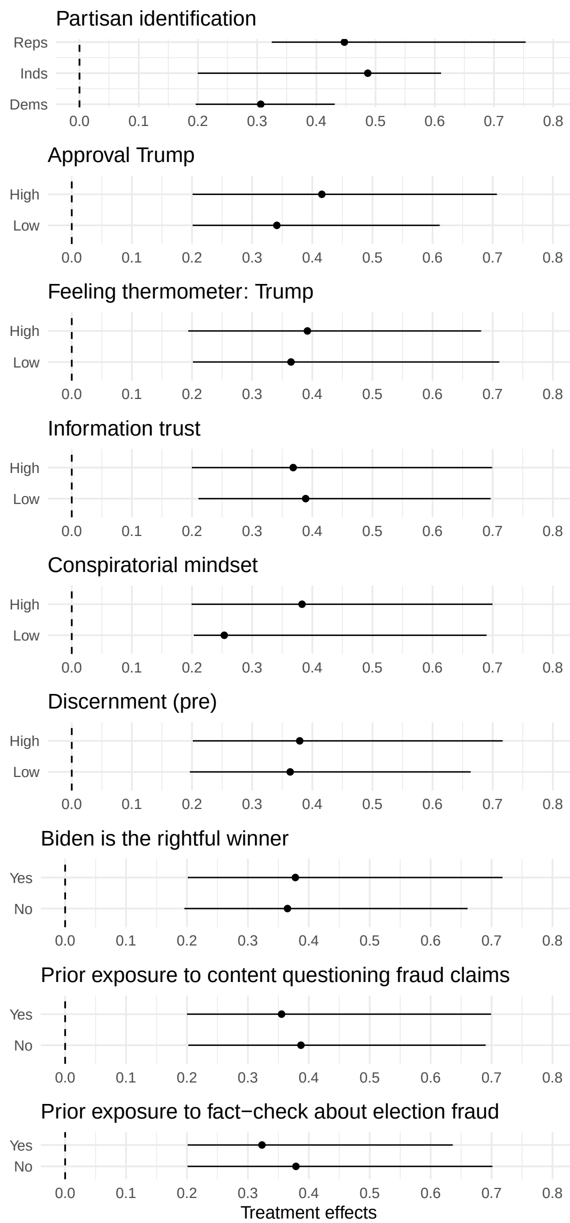

Heterogeneous treatment effects based on Bayesian Causal Forest models. Median effect by group shown with 95% confidence intervals (2.5th and 97.5th percentiles). Restricted to Pulse participants active for the majority of months during the study period. Includes the following exploratory moderators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

Figure S22: Treatment effects of exposure to fact-check article on the perceived truthfulness of non-targeted claims among active Pulse participants

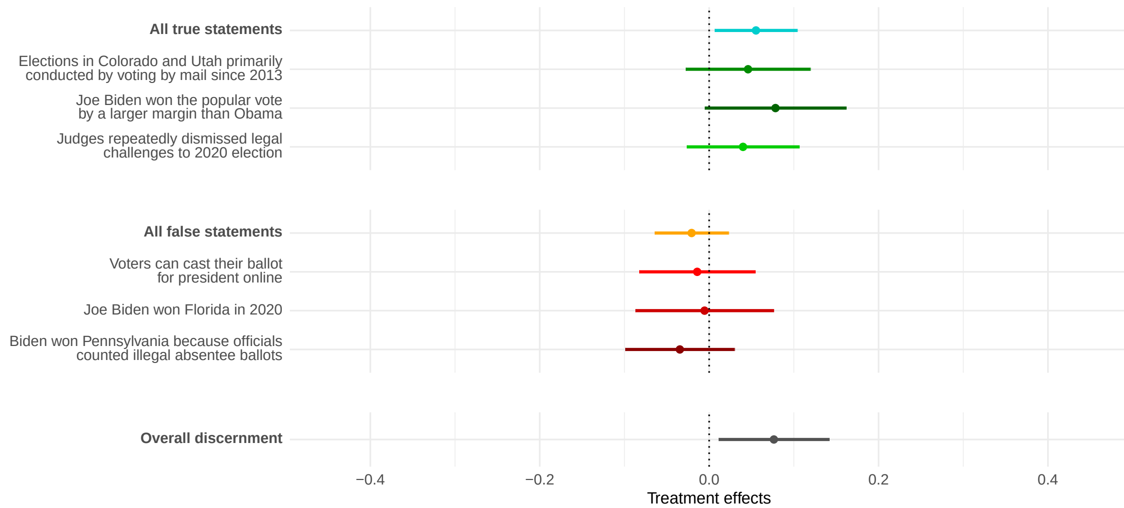

Sample average treatment effects of exposure to a fact-checking article on the perceived truthfulness of non-targeted statements. OLS regression coefficients shown with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted to Pulse participants active for the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

Table S14: 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*	−0.021	0.076*
	(0.025)	(0.022)	(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 S15: 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	0.078	0.040	−0.014	−0.005	−0.035
	(0.037)	(0.043)	(0.034)	(0.035)	(0.042)	(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 S23: Treatment effects of exposure to fact-check article on broader attitudes about the 2020 election among active Pulse participants

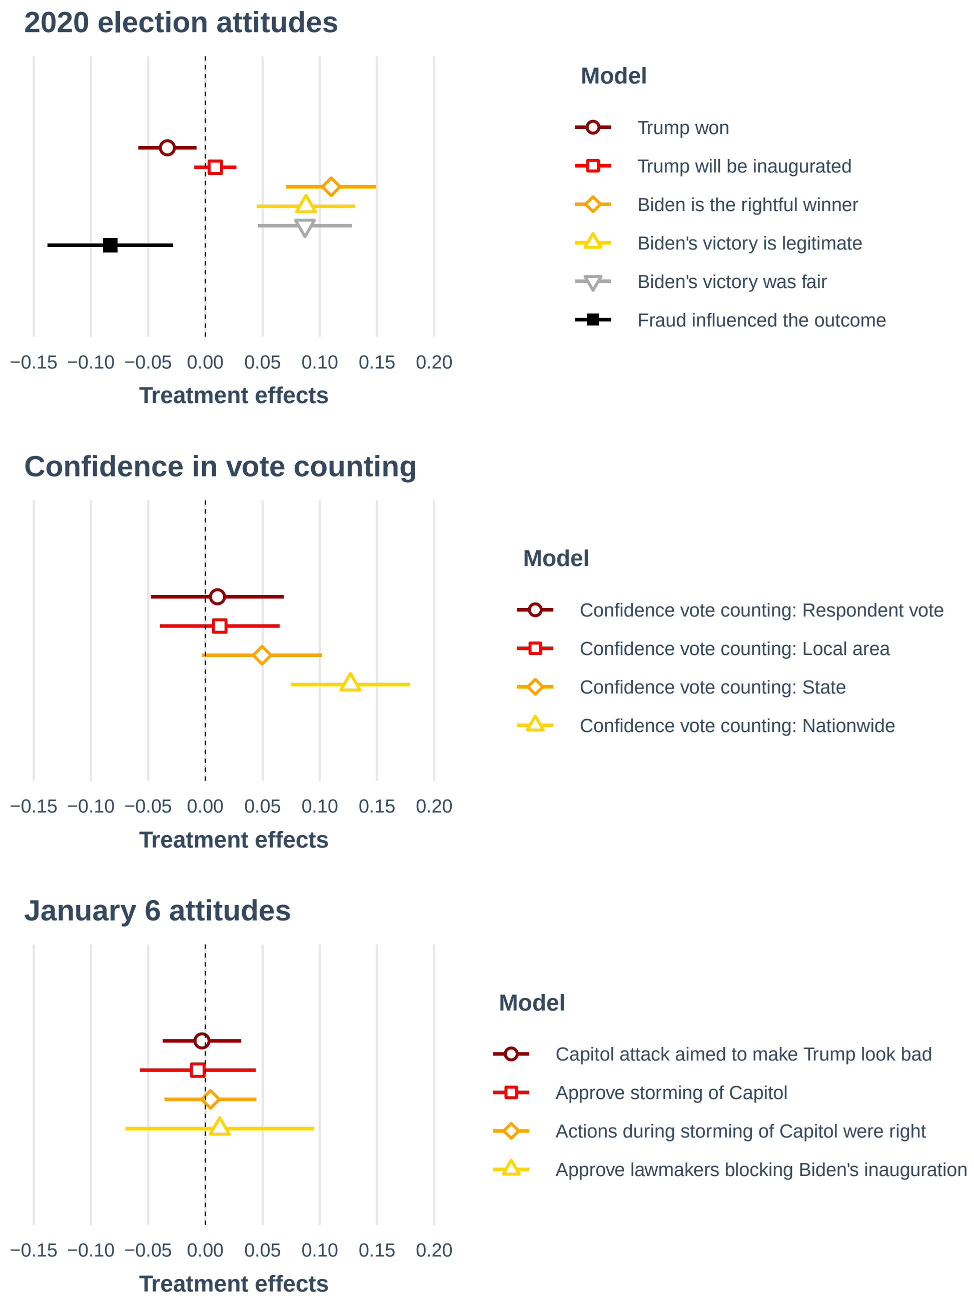

Sample average treatment effects of exposure to a fact-checking article on broader attitudes about the legitimacy of the 2020 election and the January 6 insurrection. OLS regression coefficients shown with 95% confidence intervals; estimated using pre-treatment covariates selected by lasso. Restricted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

Table S16: 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*	0.009	0.110***	0.088***	0.087***	−0.083**
	(0.013)	(0.009)	(0.020)	(0.022)	(0.021)	(0.028)
Lasso controls	✓	✓	✓	✓	✓	✓
Num.Obs.	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 S17: 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	0.013	0.050	0.127***
	(0.030)	(0.027)	(0.027)	(0.027)
Lasso controls	✓	✓	✓	✓
Num.Obs.	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 S18: 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	-0.007	0.004	0.013
	(0.017)	(0.026)	(0.020)	(0.042)
Lasso controls	✓	✓	✓	✓
Num.Obs.	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 (Trump supporters, Insurrection right, Block certification), ideology (Trump supporters, Block certification), information trust (Trump supporters, Block certification), political interest (Block certification), and education (Trump supporters, Block certification).

S4 Online behavior data

S4.1 Pulse data

The data includes 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 S24 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 election periods, with no statistically significant differences (t-tests) between Trump and Biden supporters. 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 election period (1,596 participants in the 2020 election data and 1,518 participants in the 2022 election data). Doing so does not meaningfully change the composition of the samples (see Table S2).

Figure S24: Months active in Pulse data by 2020 presidential vote choice

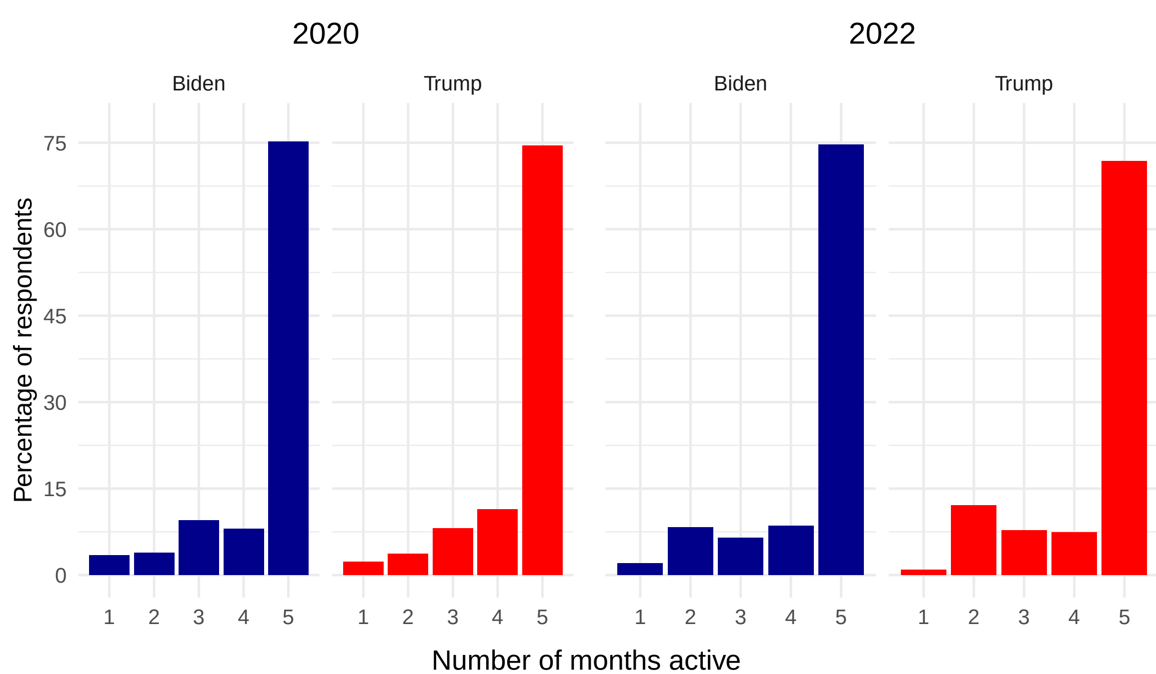

Months active measured from September 2020–January 2021 and September 2022 (starting on August 31)–January 2023. Participants are considered active when they visit at least one website during a calendar month.

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 7,109 domains and Lin et al. (2023) rates 11,520 domains.

- 6,840 domains are rated by both NewsGuard and Lin et al.
- 269 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`, `kroger.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 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 nomenclatures in February 2023, after the end of our study period. Sources with a score of 100 are now coded 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 `cutpointr` (Thiele and Hirschfeld 2021) package in R to find the optimal cutpoint by maximizing the Youden Index (= sensitivity + specificity - 1). We determined that the

[S1] Lin et al. (2023) compare NewsGuard ratings to these five domain ratings, but do not include it as part of their publicly available principal component scores to satisfy NewsGuard's data publication requirements.

optimal cutpoint value for mapping Lin et al.'s ratings score onto NewsGuard's binary ratings was 0.5 (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 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 election fraud content, we included a maximum of three visits to a given election fraud article. This process left 1,127,500 unique URLs from 4,255 rated domains for 2020 and 806,478 unique URLs from 3,748 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, 2,354 out of 2,537 (93%) have the same rating in our 2020 election data and 2,134 out of 2,268 (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 of 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 of 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.

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 concurrent requests.

We manually examined the unscraped URLs for domains with 100 or more that were unscraped. 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 346,533 unique URLs across 3,766 rated domains for 2020, and from 309,495 unique URLs across 3,513 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`), and scraped articles containing less than 25 characters (about five words). This left 326,942 URLs for 2020 and 285,864 URLs for 2022.

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 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, a proxy for text not specific to the article of interest.[S2] 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.

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.

Finally, 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 election campaigns (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

[S2]Complete sentences or paragraphs started appearing below that threshold, making it difficult to determine whether sections were related to the main article or not.

[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.

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)”
- “(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”

After removing duplicates, we were left with 32,969 articles containing election fraud keywords for 2020 and 10,945 articles for 2022.

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 for whether (1) they indeed mention election fraud (“yes” or “no”) and, if “yes”, whether (2) they questioned or contradicted election fraud claims (“yes” or “no”). 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) for the four human coders for the final set of 100 articles was 0.754 for mentions of election fraud and 0.813 for whether the article questioned or contradicted fraud claims. The final version of the LLM prompt is presented in Table S19.

Table S19: Prompt

As an expert annotator, you will be asked to annotate a sample of news articles about U.S. politics to determine (1) if they discuss election fraud or not and (2) if they include a statement questioning or contradicting claims that election fraud is widespread or could change the outcome of one or more elections.

"Election fraud" refers to illegal interference with the process of casting, tallying, and certifying votes in an election. Election fraud can take different forms—including but not limited to in-person election fraud, fraudulent activity involving absentee or mail ballots, and illegal actions to prevent eligible voters from casting ballots or preventing those ballots from being counted—and can occur at different points of the election process, from voter registration to the tallying of ballots to interference with the recording and certification of the results.

General scope:

- The articles you code may include extraneous text that was captured during the web scraping process. Any irrelevant text should be ignored; do not answer Yes for any question below based on text that appears to be extraneous to the article (e.g., a headline for another article).
- Hacking an email system, releasing confidential information, investigating political groups, committing campaign finance violations, etc. are NOT included.
- International actors may try to interfere with elections in various ways, including some that are illegal; only claims about interference with the process of casting and tallying votes are relevant here (e.g., hacking into election systems).
- Identification requirements or election roll purges that have not been determined by a court to be illegal are not considered to be within the definition above (these may be referred to by critics as "voter suppression").
- Mentions of protests against election results or associated events such as what happened on January 6, 2021 are not considered to be within the definition above unless they specifically mention fraud (e.g., "attempt to overthrow the American government by the loser of the election" on January 6 would not count).
- If a scraped article includes multiple distinct articles or blog posts (e.g., on Yahoo), only code the first article or post for each of the two questions below (i.e., disregard any text after the first distinct article or post).
- Any comments by readers on the article (e.g., that appear afterward) should be excluded.
- Satire articles (e.g., The Onion, Borowitz Report, Babylon Bee) should be excluded, as should references to real or alleged fraud in other countries that do not mention fraud in the context of U.S. elections.
- Letters to the editor are excluded. References to "election integrity" are not, by themselves, considered a reference to election fraud.

For each article, please code the following based on the definitions and scope conditions provided above:

Question 1. Does the following article mention election fraud? The mention can be brief and does not have to relate to the 2020 election specifically. General references to election fraud that are hypothetical or unspecific to a particular election also count. These mentions can include articles that promote, discuss, question, or contradict election fraud claims. (Yes/No)

Table S20: Prompt (continued)

Scope of Question 1:

- In the context of articles that concern the 2020 presidential election and its aftermath, a statement that someone is "refusing to concede" or "refusing to accept" the election results is sufficient to be considered to be a reference to fraud claims (and thus to code Question 1 as "Yes"), as are references to efforts to "overturn" or "steal" the election, general references to "election denier" or "election denial," or descriptions of the election as "invalid." If you find language like the terms listed in this scope condition, you should code the article as "Yes" even if the article itself is not about election fraud or does not mention election fraud explicitly.
- Code articles as "Yes" on Question 1 if any dimension of the process of casting, tallying, and certifying votes is described as being vulnerable or prone to fraud.

If the answer to question 1 is "No," code question 2 as "No."

If the answer to question 1 is "Yes," code question 2 per below:

Question 2. Does the following article include a statement questioning or contradicting claims that election fraud is widespread or could change the outcome of one or more elections under current or past practices? (i.e., not claims about fraud if a change in the rules were enacted in the future). These statements can be made by the author of the article or sources they cite, quote, etc. and do not necessarily have to be endorsed by the author of the article. (Yes/No)

Scope of Question 2:

- In the context of the 2020 election, code articles as "Yes" on Question 2 if they refer to "election denier" or "election denial" or say someone "refused to accept" election results. ("Refused to concede," by contrast, should not be treated as evidence for coding Question 2 as a "Yes.")
- Code articles as "Yes" on Question 2 if they say someone is lying in the process of claiming the election was "stolen."
- Code articles as "Yes" on Question 2 even if the reference is brief and no evidence is provided (e.g., makes reference to "baseless," "unfounded," "false," "unsupported," or "absurd" election fraud claims).
- Code articles as "No" on Question 2 if they refer only to efforts to "overturn" the 2020 election without further questioning fraud claims. If they question these claims directly (e.g., saying the person was lying in trying to convince people to overturn the election), code Question 2 as "Yes."
- Code articles as "Yes" on Question 2 if they describe courts rejecting or denying lawsuits making claims that election fraud is widespread or changed the outcome of one or more elections.
- Code articles as "Yes" on Question 2 if they include a statement questioning or contradicting the legality or legitimacy of efforts to change or interfere with the certification of election results.
- Per the previous bullet, a court rejecting or denying legal arguments that seek to change or interfere with the certification of election results should be coded as "Yes" as well.
- Statements in which Donald Trump, his allies, or his supporters question or contradict the legality or legitimacy of efforts to certify the results of the 2020 election should not be treated as evidence for coding Question 2 as a "Yes."

Provide your answer in JSON in the following format, without providing any justification for your choice: 1:Yes/No, 2:Yes/No.

S4.8 Validation

To validate the LLM results and compare the performance of different models, we took a random sample of 100 articles containing election fraud keywords (50 articles from 2020 and 50 articles from 2022); articles coded for developing the LLM prompt were excluded. Each article was coded by four human coders and separately coded with GPT-4o and GPT-4o mini. We resolved all disagreements between the human coders before comparing their results with the results from the two GPT models.

Table S21 shows that there is a high level of agreement between GPT-4o mini and GPT-4o outputs.

Table S21: Comparison of GPT-4o and GPT-4o mini coding

Model	Metric	(1) Mention	(2) Question
Overall			
LLMs	Agreement	93%	92%
	Krippendorff's alpha	0.85	0.84
2020			
LLMs	Agreement	94%	94%
	Krippendorff's alpha	0.87	0.88
2022			
LLMs	Agreement	92%	90%
	Krippendorff's alpha	0.84	0.80

Table S22 presents the percent agreement and Krippendorff's alpha between the consensus human coder judgments and the GPT-4o and GPT-4o mini models.

Table S22: Comparison of LLM output with human coders

Model	Metric	(1) Mention	(2) Question
Overall			
GPT-4o mini	Agreement	87%	89%
	Krippendorff's alpha	0.72	0.78
GPT-4o	Agreement	90%	93%
	Krippendorff's alpha	0.78	0.86
2020			
GPT-4o mini	Agreement	94%	92%
	Krippendorff's alpha	0.87	0.84
GPT-4o	Agreement	92%	98%
	Krippendorff's alpha	0.82	0.96
2022			
GPT-4o mini	Agreement	80%	86%
	Krippendorff's alpha	0.58	0.72
GPT-4o	Agreement	88%	88%
	Krippendorff's alpha	0.74	0.76

The LLM models perform better in coding articles from the 2020 election than the 2022 election, which negatively impacts Krippendorff's alpha when weighing each article and year equally.

Table S23 adds weights based on the total number of views and articles from each year in the Pulse data to account for the fact that the vast majority of views and articles were from the 2020 election.

Table S23: Comparison of LLM output with human coders

Model	Metric	(1) Mention	(2) Question
Raw			
GPT-4o mini	Agreement	87%	89%
	Krippendorff's alpha	0.72	0.78
GPT-4o	Agreement	90%	93%
	Krippendorff's alpha	0.78	0.86
Weighted based on number of views			
GPT-4o mini	Agreement	91%	91%
	Krippendorff's alpha	0.81	0.82
GPT-4o	Agreement	91%	96%
	Krippendorff's alpha	0.80	0.92
Weighted based on number of articles			
GPT-4o mini	Agreement	91%	91%
	Krippendorff's alpha	0.79	0.81
GPT-4o	Agreement	91%	96%
	Krippendorff's alpha	0.80	0.91

S4.9 Coding articles using LLM

The results presented in this paper are based on 43,901 articles (combining the 2020 and 2022 data) containing election fraud keywords that were coded using GPT-4o mini via OpenAI's batch API (implemented via the `OpenAI` package in Python). We used GPT-4o mini given that it achieves acceptable reliability, as shown in Table S23, at a significantly lower cost than the GPT-4o model. We set the temperature to 0 to increase replicability and allowed a maximum of 25 tokens for the JSON output to prevent deviation from the format defined in the prompt.

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 election period, we scraped 3,123 unique fact-checking articles, of which 435 were coded by GPT as mentioning election fraud. For 2022, we scraped 618 unique fact-checking articles, of which 22 were coded by GPT as related to election fraud.

S4.11 Descriptive statistics

Tables S24 and S25 include descriptive statistics about exposure to different types of election fraud content overall and by vote choice during the 2020 and 2022 election periods.

Table S24: Distribution of content exposure during the 2020 U.S. election (Pulse data)

Variable	Mean	SD	p10	p25	Median	p75	p90
Overall							
News content	1877.94	11333.46	13	64	305	1302	4193
Trustworthy news	1764.46	11275.56	11	61	271	1141	3943
Untrustworthy news	113.47	663.60	0	0	4	32	152
Election fraud content	18.30	65.67	0	0	0	6	39
” questioning fraud claims	15.00	57.94	0	0	0	5	30
” not questioning fraud claims	3.20	14.88	0	0	0	0	4
” from trustworthy sources	13.99	56.37	0	0	0	4	28
” from untrustworthy sources	4.32	23.68	0	0	0	0	3
” from fact-check	0.27	2.48	0	0	0	0	0
Biden voters							
News content	1473.87	3495.94	19	70	328	1174	3652
Trustworthy news	1407.27	3341.92	13	67	315	1083	3523
Untrustworthy news	66.60	645.72	0	0	3	15	58
Election fraud content	26.35	84.73	0	0	1	11	67
” questioning fraud claims	24.35	79.25	0	0	0	10	62
” not questioning fraud claims	1.95	7.89	0	0	0	0	3
” from trustworthy sources	23.95	78.12	0	0	0	10	56
” from untrustworthy sources	2.40	18.10	0	0	0	0	1
” from fact-check	0.41	2.89	0	0	0	0	0
Trump voters							
News content	1959.64	5115.23	15	83	368	1814	5124
Trustworthy news	1761.20	4978.39	14	72	306	1477	4491
Untrustworthy news	198.44	790.63	0	1	13	79	407
Election fraud content	13.35	45.77	0	0	0	4	28
” questioning fraud claims	7.56	26.80	0	0	0	3	17
” not questioning fraud claims	5.61	21.97	0	0	0	1	8
” from trustworthy sources	5.57	21.80	0	0	0	2	11
” from untrustworthy sources	7.78	31.99	0	0	0	1	10
” from fact-check	0.19	2.42	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 S25: Distribution of content exposure during the 2022 U.S. election (Pulse data)

Variable	Mean	SD	p10	p25	Median	p75	p90
Overall							
News content	1211.94	3392.30	13	47	148	703	3074
Trustworthy news	1127.07	3220.72	13	44	133	665	2894
Untrustworthy news	84.87	896.31	0	0	3	11	52
Election fraud content	3.61	18.89	0	0	0	0	5
” questioning fraud claims	3.08	16.92	0	0	0	0	4
” not questioning fraud claims	0.51	3.27	0	0	0	0	0
” from trustworthy sources	2.69	16.43	0	0	0	0	3
” from untrustworthy sources	0.92	7.76	0	0	0	0	0
” from fact-check	0.02	0.21	0	0	0	0	0
Biden voters							
News content	1250.13	3691.77	12	54	170	800	2770
Trustworthy news	1173.25	3468.62	12	47	163	752	2670
Untrustworthy news	76.88	1175.86	0	0	2	10	35
Election fraud content	5.95	26.60	0	0	0	1	8
” questioning fraud claims	5.47	24.77	0	0	0	0	8
” not questioning fraud claims	0.47	2.37	0	0	0	0	0
” from trustworthy sources	4.97	24.39	0	0	0	0	7
” from untrustworthy sources	0.98	9.20	0	0	0	0	0
” from fact-check	0.03	0.28	0	0	0	0	0
Trump voters							
News content	1753.93	3918.11	18	86	284	1492	4487
Trustworthy news	1607.19	3764.75	18	76	226	1167	4424
Untrustworthy news	146.74	794.34	0	1	7	32	158
Election fraud content	2.85	12.15	0	0	0	0	5
” questioning fraud claims	1.90	7.48	0	0	0	0	3
” not questioning fraud claims	0.92	4.98	0	0	0	0	1
” from trustworthy sources	1.36	5.70	0	0	0	0	2
” from untrustworthy sources	1.50	8.51	0	0	0	0	1
” from fact-check	0.01	0.17	0	0	0	0	0

Respondents are YouGov Pulse panel members who were active for the majority of months throughout the study period. Estimates calculated with survey weights.

Table S26: 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	2527.24	13349.74	28	117	603	2145	5612
Election fraud content	22.85	70.38	0	0	1	10	56
” questioning fraud claims	18.70	62.15	0	0	0	8	48
” not questioning fraud claims	4.03	16.48	0	0	0	1	7
2020 Mobile							
News content	507.15	1573.10	6	29	97	352	1036
Election fraud content	9.26	52.66	0	0	0	4	9
” questioning fraud claims	7.69	45.88	0	0	0	3	7
” not questioning fraud claims	1.51	9.74	0	0	0	0	2
2022 Laptop/Desktop							
News content	1940.58	4285.80	24	90	398	2057	5040
Election fraud content	5.49	23.64	0	0	0	1	9
” questioning fraud claims	4.69	21.30	0	0	0	1	8
” not questioning fraud claims	0.78	4.11	0	0	0	0	1
2022 Mobile							
News content	314.59	1024.14	8	32	79	211	583
Election fraud content	3.03	21.84	0	0	0	0	2
” questioning fraud claims	2.77	20.33	0	0	0	0	1
” not questioning fraud claims	0.26	1.84	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.