Online hostility towards local election officials surged in 2020

Joelle Gross Samuel Baltz Mara Suttmann-Lea

Thessalia Merivaki Charles Stewart III

August 2023*

*IRB exempt through the MIT Committee On the Use of Humans as Experimental Subjects. Work on this project was supported by grants from Democracy Fund and the William and Flora Hewlett Foundation. We gratefully acknowledge feedback from Dari Tran, Joseph Anthony, Paul Gronke, Patrick Wu, and Matthew Brook O'Donnell.

Introduction

In recent years, more and more local election officials (LEOs) have incorporated social media into their voter communication strategies as a cost-effective and dynamic way to share election-related information and updates about the election process. When LEOs actively engage on social media, voters benefit by having the necessary tools to successfully complete key voting processes like registration and casting a mail ballot (Merivaki and Suttmann-Lea 2022; Suttmann-Lea and Merivaki 2023a). Considering that the public relies on social media for election-related news and information, the presence and activity of LEOs on social media increases voters' access to much-needed information about how to participate in elections, how elections are conducted, and the security of elections. The rise of misinformation has also transformed LEOs' communication strategies, with many actively stepping in to "set the record straight" and ensure that voters have accurate and reliable information.

There is significant variation in which election processes LEOs choose to emphasize and discuss on social media platforms (Merivaki and Suttmann-Lea 2022; Suttmann-Lea and Merivaki 2023a). There are also differences in how much voters actually engage with these posts, with some LEOs receiving a high number of reactions (likes, shares, comments) compared to others (Suttmann-Lea and Merivaki 2021). Researchers have started to focus on how voters interact with their election officials in these online spaces, especially when responding to posts shared by state and local election officials during an election cycle. One notable trend is an increase in harsh rhetoric targeted against state election officials. This rhetoric largely revolves around allegations that elections are not secure (Gross, Baltz, and Stewart III 2023). It reflects a growing challenge in election administration: a hostile working environment for election officials. Such hostility is cause for concern, given that it has prompted many election officials with years of institutional knowledge to resign from their jobs (Bidgood 2022). Such hostility also raises the concern that it may further drive election officials away from communication and outreach with voters, both online and offline.

In this chapter, we measure the social media environment within which local election

officials operate to inform voters about elections, build trust in election integrity, and insulate voters from election misinformation. We examines the content and sentiment of responses to over 250,000 tweets by local election officials over the course of a decade, from 2012 through the end of 2022. Beyond concerns about LEO turnover and impact on engagement with constituents, negativity towards LEOs on social media is of particular importance, partly because of the notable patterns in the demographics of LEOs that are targeted. Women and people of color may face distinctive types of online hostility, and negativity is especially focused on officials from jurisdictions or states where claims of voter fraud during the 2020 election were more prominent.

Focusing on content shared on Twitter, we find that LEOs have dramatically increased the number of tweets posted over the last ten years, with a surge around the 2020 election in particular. As the volume of LEO tweets has increased, the replies those tweets receive has also grown more negative. We take a close look at the few LEO accounts that receive the most replies, which are in Arizona, California, Colorado, Kansas, Texas, and New York. We find that very large numbers of negative replies have been responses to a very small number of tweets, and that these few tweets were mostly targeted by repliers on the right of the ideological spectrum. First we count the presence of relevant keywords in our corpus of 264,905 tweets and 133,199 replies, and we show that discussions of election fraud are concentrated in certain jurisdictions. We then estimate the sentiment of all the replies, and we find that local election officials are experiencing increasing negativity, a trend which mirrors previous findings about state election officials (Gross, Baltz, and Stewart III 2023).

Of course, not all local election officials are active on social media. Only about 30% have an active Facebook account, only 9% use Twitter, and fewer than 2% are active on Instagram (Suttmann-Lea and Merivaki 2023b). While the small percentage of LEOs active on these platforms may suggest their presence is not consequential for election administration, we argue that it is precisely because their presence is so spotty that harassment and threats against them online are concerning. Such actions may drive them to limit or eliminate both

their online presence as well as their willingness to engage in voter education and outreach under the specter of attacks, and disincentivize "offline" LEOs from communicating with voters in online spaces through social media (Brennan Center 2022). This is especially true since many local election officials report having limited capacity and resources to engage in voter education and outreach (Gronke and Mason 2022). It is therefore vital to understand what drives negativity towards these election officials, and areas where they may be more prone to be the target of online harassment.

Local Election Officials and Online Rhetoric

In the aftermath of the 2020 election, the work and lives of many election officials across the nation were upended. As a result of former President Donald Trump's harsh tirades against the integrity of American elections, election officials experienced an influx of violent and even criminal threats. One Reuters report set out to document the kind of threats these officials received, relying on 850 hostile messages compiled by Eisler et al. (2021) which shows the vitriol that many election officials receive on a near-daily basis (Reuters 2022). The threats persisted even after the election, and in some cases used explicitly sexualized or racist language. As a result of the growing danger against them, many states passed laws which mandated legal restrictions against threatening election officials (PBS 2022).

To some extent, these targeted attacks against election officials can be observed on social media platforms such as Twitter and Facebook. Social media is one of the most utilized tools for generating negative rhetoric towards public officials, and researchers find that this issue has been growing worse over time (Gorrell et al. 2019; Kuperberg 2021; Southern and Harmer 2019; Southern and Harmer 2021). Survey research on election officials confirms these trends: in 2022, the Brennan Center interviewed local election officials across the nation to understand how their working environment has changed as a direct result of threats and harassment. They found that LEOs overwhelmingly reported that they were concerned

about their experiences online and the possibility of an increase in threats on social media platforms, particularly Twitter (Brennan Center 2022).

The experience of election officials on social media mirrors harassment through other communication channels, such as text messages, emails, and phone calls, as detailed in Reuters (2022). A census of all the replies to *state*-level election officials' posts on Twitter identified changes in the usage of charged keywords related to topics like election fraud as well as the overall sentiment of the replies (Gross, Baltz, and Stewart III 2023). By both measures, the situation has grown worse for state election officials. Specifically, the usage of fraud-related keywords in Twitter replies increased over time, and spiked around the 2020 election, while the sentiments of those replies has become increasingly negative. Republican officials who received a large volume of negative replies tended to receive them from left-leaning users, and the reverse was true for Democratic officials. However, there were a few major exceptions, which largely featured Republican officials receiving a steady drumbeat of negative replies from right-leaning repliers.

Harsh language used against election officials offline and online can have important real-world effects: constant exposure to negative language can make elected officials feel unsafe (Korsell et al. 2020), lower their job satisfaction and create a hostile work environment (Wagner 2020), add extra and different types of labor to their jobs (Clark, Barsky, and Bustinza 2022), and harm their overall mental health (Every-Palmer, Barry-Walsh, and Pathé 2015). A direct result of the growing incidence of threats is that election officials across the country are fleeing their jobs, opting for a working environment free from harassment. With this mass resignation, a dearth of expertise and experience remains in their place, imperiling election offices across the nation (Bidgood 2022).

What is defined as "the great resignation" in election administration has disproportionately impacted local election officials, many of whom report unwillingness to serve in the next presidential election (Brennan Center 2022). These LEOs, as the Brennan Center report documented, cite increasing harassment and threats on social media as the primary

causes of hostile working conditions (Brennan Center 2022). As of now, there is limited insight as to what a hostile social media environment looks like for local election officials, and whether it is similar, or worse, compared to what *statewide* election officials are experiencing. Whereas state election officials are the federally and state designated officials in charge of elections, they are not in charge of the day-to-day election-related procedures as local election officials are. Due to their broader name recognition, voters tend to primarily associate state election officials with the logistics of the elections they participate in. However, local election officials are often most people's first point of contact for election-related information (Merivaki and Suttmann-Lea 2022). Moreover, voters' attitudes about the election process are conditioned on how they view their local election officials (Suttmann-Lea and Merivaki 2023c). Taken together, constituents' relationship with their local election officials and the rise in online threats is concerning given constituents' easier access to act on threats towards these proximate officials, and officials' limited resources for physical protection against such threats.

There is reason to believe that local officials and state officials receive similar levels of targeted negativity. In August 2022, a report regarding the findings of a joint investigation by the Department of Justice and Federal Bureau of Investigations discussed 1,000 contact reports that were made to the Threats Against Election Officials Task Force. Among many findings, the task force demonstrated that of the reports received, 48% were directed at state officials, and 48% were directed at local officials, with the remaining 4% directed at vendors or liaisons (Office of Private Sector 2022). In the next section we describe how we complement these survey-based measures with analyses of publicly available data.

Methodological Approach and Data Collection Process

We will investigate whether the state-level results found by Gross, Baltz, and Stewart III (2023) also apply at local levels of election administration by replicating the methodology

used in that work to measure the social media environment in responses to state election officials. We measure the extent of negative rhetoric that local election officials experience on Twitter in four steps. *First*, we collected all tweets shared by local election officials over the span of the last decade, as well as all of the replies to those tweets. *Second*, we counted the number of times that certain keywords related to fraud or accusations of wrongdoing appeared in the replies. *Third*, we estimated and compared the sentiment of the replies to officials' tweets. *Finally*, we estimated the ideologies of a sample of the repliers to understand if negative replies to tweets are more likely to come from conservative or liberal-leaning users on Twitter, focusing on jurisdictions with the highest variation in replies to LEOs, and whose substantive relevance in the discussion about the conduct of elections is significant: Maricopa County, Arizona; Los Angeles County, California; Denver County, Colorado; Johnson County, Kansas; Harris County, Texas; and New York City, New York. For analysis on the ideologies of repliers, due to the time-intensiveness of the process to generate ideology scores for users on Twitter, we chose only two disparate local jurisdictions from the above list to focus on: Maricopa County, Arizona and New York City, New York.

Our dataset consists of the entire universe of tweets from accounts associated with county-level election officials and election offices, and every reply that any user made to those tweets. We collected tweets and replies from these accounts that were made between January 1st, 2012 and January 1st, 2023. In total, we collected replies and tweets for 258 accounts nationwide (Suttmann-Lea 2022). From 2012 to the end of 2022 these accounts tweeted 264,905 times, and those tweets attracted 133,199 replies.

Two major decisions shaped the contents of this dataset: the date range, and the types of tweets collected. We chose 2012 as the beginning of the time period of interest for two reasons. First, the 2012 US presidential election is widely viewed as the first election after the widespread adoption of Twitter by the country's political and media elites. Second, this starting point also provides a basis of comparison for the particularly unusual elections of 2016 and 2020, which both featured a presidential candidate (and, in the latter case, a

sitting president) who focused substantial negative rhetoric on election administration in an attempt to make the logistics of elections into a salient political issue.

Regarding the types of tweets collected, we focus on replies to tweets because our interest is in posts that are directed *at* local election officials, and which are at least ostensibly made with the intention of being seen by the local election official. Replies are the only type of public information that fits these criteria. Other ways of directing messages to local election officials, like Twitter direct messages or messages sent by means other than social media, cannot be measured by the public. On the other hand, other types of public tweets that might be considered responsive to local election officials' posts, like retweets of their tweets, are at least ostensibly meant to be read by a different audience, as they begin a new chain of conversation on a different subpage of the website (the user's own account page).

Keywords

The harsh language directed at election officials in recent years is often motivated by perceptions of officials' involvement in illegitimate electoral activity. Therefore, a central question in our analysis is how often the topic of election legitimacy appears in replies to local election officials. One particularly direct measure of that phenomenon is to count the number of times that words related to electoral fraud appear in the replies to local election officials.[1] However, there is no standard dictionary of fraud-related keywords, so we constructed a list of keywords that relate to discussion of undemocratic activities. This is not an exhaustive list of words that might correspond to that topic, and neither does the appearance of any one of these keywords imply that the replier is accusing the original poster of illegitimate electoral activities — the replier may instead be alleging that some other actor engaged in fraudulent activity, they might merely be discussing the idea of fraud, or they may even be pushing back against accusations that fraud occurred. However, as the usage of these

[1] In the context of election administration, electoral fraud is distinct from voter fraud, which includes actions such as voter impersonation and double voting. "Electoral fraud is understood as intervention of election officials in the conduct of elections, especially in regard to accurately counting votes" (Stewart III 2021).

words increases, taken together, it is reasonable to believe that discussion of election fraud is becoming more prominent, and that is the phenomenon that we are directly measuring.

The keywords we use are: accomplice, attack, bribe, cheat, complicit, criminal, faked, fascist, forged, forgery, fraud, hoax, illegal, jail, lawsuit, liar, lies, lock up, lying, prison, racist, restrict, RINO, russia, sedition, steal, stole, suppression, traitor, treason, and unconstitutional. Nearly all of these are common nouns or verbs closely connected to types of illegitimate electoral activities, while a few (such as "racist", "fascist", and "RINO") are intended to capture the sorts of labels that might have been applied at any point in the ten-year period to election officials who are seen as administering elections in an illegitimate way for politically motivated reasons.[2] Our outcome measure in the keyword analysis is the raw count of the number of times that these words appeared in the officials' replies. We focus on the raw number of keywords (rather than, say, keywords as a proportion of all words used) because the theoretical quantity of interest is how exposed election officials are to negative rhetoric, and an official who has more fraud-related keywords in their replies now than they did a few years ago is indeed being subjected to more discussion of fraud than they were previously, regardless of the volume of tweets they also receive on other topics.

Sentiments

Because of the many limitations of generating a list of specific keywords and counting their appearances, we also pursue a much more general approach to understanding the qualities of the rhetoric directed at local election officials. We estimate the negativity or positivity of these replies using a standard method for scoring the sentiment of texts on social media: the Python package VADER (Hutto and Gilbert 2014). Using VADER to generate sentiment scores for individual words, we score every word to every reply to a tweet created by a local election official, and then we compute the mean sentiment score for all of the replies to a

[2] In the case of state election officials, a leave-k-out analysis showed that none of the keywords was load-bearing in driving the overall trends; even if one of these keywords does not credibly or always represent a discussion of election fraud, it is not alone driving the result.

given tweet.[3]

We use these scores to examine trends in sentiment over time and across space, by analyzing which localities had the lowest mean sentiment score and when. We fit Ordinary Least Squares (OLS) trend lines to the tweet-level average sentiment scores, both for each account individually and binned together, and we use the intercept and slope of these best fit lines to compare the trend in negativity across accounts and over time.

We conducted this analysis for all of the 258 local jurisdictions which have active accounts identified by Suttmann-Lea (2022). To give a sense for the overall tendency in the whole universe of replies, we present their sentiments together in one figure with a line of best fit. Because this is a summary of hundreds of thousands of replies to hundreds of accounts, we then focus on a few of the counties and jurisdictions which both receive enough replies to be analyzable and also demonstrate more or less positive public reception. In many cases, we aim to highlight local jurisdictions within states that are frequently targeted by those who deny the legitimacy of recent American elections, but to avoid generalizing inappropriately from this highly distinctive group of states, we also include trends in jurisdictions of low-salience states for comparison.

Ideologies

We supplement our analysis of the topic and tone of tweets by estimating the ideology of the repliers who respond to the tweets by election officials. We also use a standard approach to this problem: the method, and associated R package, developed by Barberá (2017). This package uses the list of accounts a given user is following to generate an ideal point estimate of their left-right ideology, by treating the decision to follow (or not follow) an elite on a pre-coded list of elites with known ideologies as an indicator of a Twitter user's own ideology. Because the Twitter API has an access rate limit that only makes it possible to estimate the ideology of one user per minute on average, we can only sample the population of repliers:

[3]The results that we present throughout this chapter are robust to different types of aggregation, with the median, mean, and sum producing similar results.

for a given tweet by a local election official, we score the ideologies of ten unique users (or every replier in cases where fewer than ten people replied), sampling repliers uniformly and without replacement, and excluding the original poster in cases where they replied to their own tweet.

Findings: Rhetoric against LEOs on Social Media

In this section, we delve into the keywords, sentiment, and content of the replies to local election officials on Twitter. We first consider the corpus of tweets as a whole, describing the number of accounts included, when those accounts were created, and how many tweets they have posted over the last decade. We plot the number of replies over time and the average sentiment of the words in those replies, and then characterize the overall trend in sentiment across every reply to any one of these accounts. We then narrow our focus to a handful of local jurisdictions, all of which more or less match the broader trend of increasing negativity over time, but which highlight the variation in trends that we witness across these accounts. After checking the sentiment of replies to accounts from these jurisdictions, we focus on the few accounts that receive the most replies to address the question of whether repliers tend to be conservative or liberal-leaning Twitter users.

Overall Trends in Replies

We begin by considering some descriptive questions about the number and trend in replies to local election officials: how many tweets do LEOs make, how many replies do those tweets receive, and how has the sentiment of those replies trended?

To be included in this corpus, we impose two restrictions to ensure that each account engages in some minimum amount of activity and receives some level of engagement. First, the account must have posted at least 10 tweets. Second, the account must have received at least 10 replies total across those 10 tweets. Among the 258 LEO Twitter accounts

identified by Suttmann-Lea and Merivaki (2023b), 154 satisfy these minimal requirements. Those 154 accounts tweeted 255,119 times (compared to 264,905 tweets in the whole corpus) and received 133,006 replies (compared to 133,199 in the whole corpus).

Figure 1a shows the number of Twitter accounts in the corpus over time.[4] The figure shows that about half the population of LEO accounts was already present and active on Twitter by 2014, and the vast majority of active LEO accounts had made at least one post before the 2018 election.[5] Figure 1b shows the actual activity level of those accounts, plotting the total number of tweets by any LEO account in each month from January 1, 2012 to January 1, 2023. The plot shows structured cycles of activity throughout the year, with accounts more active towards the end of the year than the beginning, and particularly inactive in the summer. The number of tweets surges in October and November of even-numbered general election years, and the total number of tweets in each of these election months steadily increased from 2014 through 2020. The 2020 election dramatically stands out: October 2020 in particular saw a spike in tweets by LEO accounts that was nearly double the number of tweets posted in any other month.

Since over 95% of LEOs' tweets were posted in 2016 or later, we now focus on that period in Figure 2.[6] Figure 2 shows how the volume of tweets and the sentiment of users' replies to them varied during the period between January 1, 2016, and December 31, 2022. In Figure 2, darker regions represent a larger number of replies. It is immediately visible that LEOs have been sharing more content on Twitter since the 2016 presidential election, which is partially explained by the increase in the number of LEO Twitter accounts over time, but also due to increased usage of Twitter as a tool to communicate with voters (Suttmann-Lea and Merivaki 2023c).

[4]Specifically: the number of Twitter accounts that had made at least one tweet between January 1, 2012 and the time on the x-axis, and went on to satisfy the activity requirements for inclusion in the corpus.

[5]Note that the figure represents accounts that still existed at the time of data collection in 2023. Some accounts may have been created in this period and subsequently deleted.

[6]This is really just for the sake of readability. All the results we describe in the span from 2016 to 2022 also apply to the whole decade.

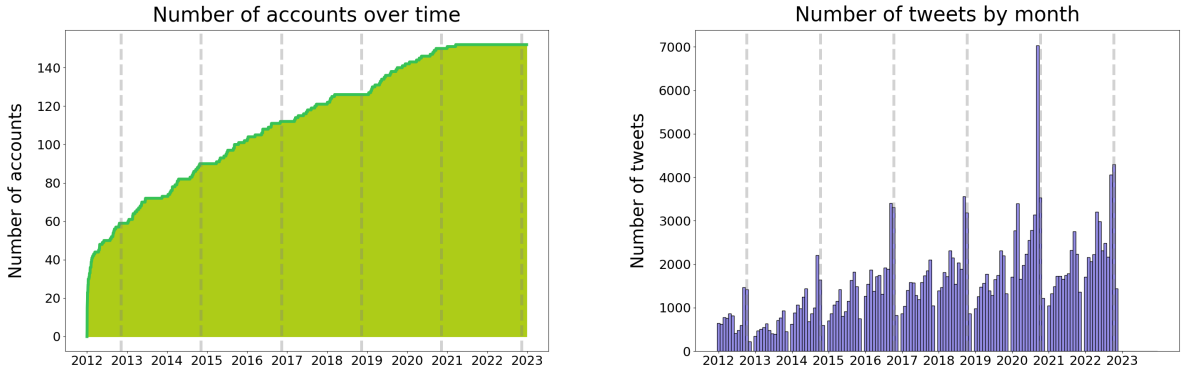

(a) Number of LEO accounts

(b) Number of LEO tweets

Figure 1: Subfigure a) shows the number of LEO accounts over time, with an account added to the corpus the first time it tweets after January 1, 2012. Subfigure b) shows the total number of tweets by those accounts, grouped by month, with a gap at the start of each year.

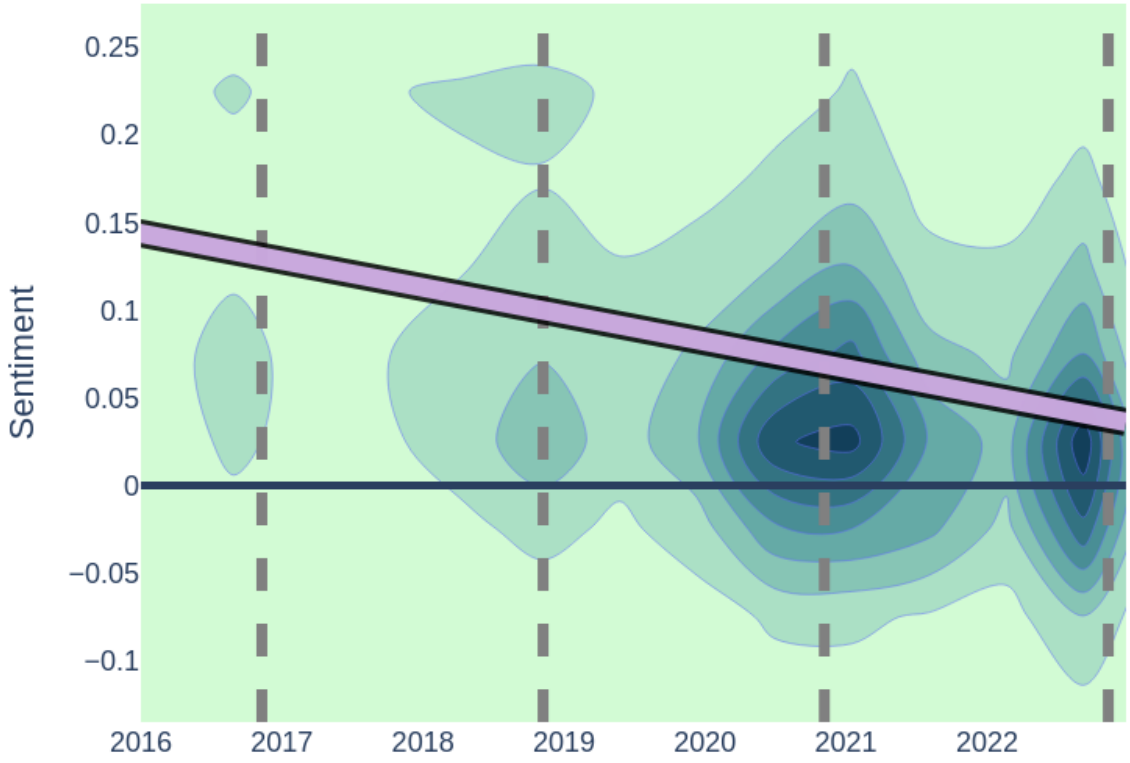

Figure 2: Darker regions represent a higher volume of tweets by election officials, with the time each tweet was posted on the x-axis and the sentiment of the replies to each tweet on the y-axis. A trend line shows the Ordinary Least Squares regression of sentiment on the number of days since January 1, 2012. General elections are represented by dashed vertical lines.

Figure 2 includes an Ordinary Least Squares fit, where the observations are each original tweet by a LEO account: the x-values are the time each tweet was posted, and the y-values are the average sentiment of the replies to each original tweet. The OLS line has a negative slope, so the sentiment of replies has been getting more negative over time across the whole corpus. That said, the measure of sentiment has no inherent units, so the extent of this decline is not directly interpretable, which is why we next discuss changes in the volume of negative rhetoric over time for individual accounts.

Replies and Fraud-Related Rhetoric in six Regions

We now narrow our attention to six key counties or local jurisdictions: Maricopa County, Arizona; Los Angeles County, California; Denver County, Colorado; Johnson County, Kansas; Harris County, Texas; and New York City, New York. We focus on these six localities for descriptive purposes, because a) they received a large enough volume of replies and thus yield high variation on the outcome measure, and b) they have varying levels of centrality in the popular narratives around American election administration. They also represent geographically distant parts of the country, have varied demographic compositions, and represent different levels of government. One of these accounts, Arizona's Maricopa County, was central in the narratives spread by then-President Donald Trump regarding the 2020 election.

The election administration accounts from the six localities in Figure 3 each portray a distinct story regarding what local officials are experiencing on Twitter. In New York City, New York, as seen in subfigure 3a, the election administration increased its Twitter presence during the 2020 election cycle. As a result, the account saw an increase in fraud-related keywords, and as many as 100 or more replies feature this type of negative rhetoric for a given tweet that they posted. Harris County, Texas, exemplifies a similar trend: as seen in subfigure 3b, this accounts sees the greatest increase of negativity throughout the 2020 election cycle and even into 2021 and 2022.

Tweets from the election administration in Johnson County, Kansas practically generate little to no fraud-related responses, as seen in subfigure 3c. Johnson County is not a prominent target in election fraud narratives, which provides a possible explanation for why few fraud-related keywords appear in these replies, and the same can be said for Denver, Colorado, as seen in subfigure 3d. California's Los Angeles County, however, experienced an uptick in discussions of fraud near the 2020 election, with some tweets receiving hundreds of fraud-related keywords, and that uptick persisted well into 2021 and even 2022, as seen in subfigure 3e.

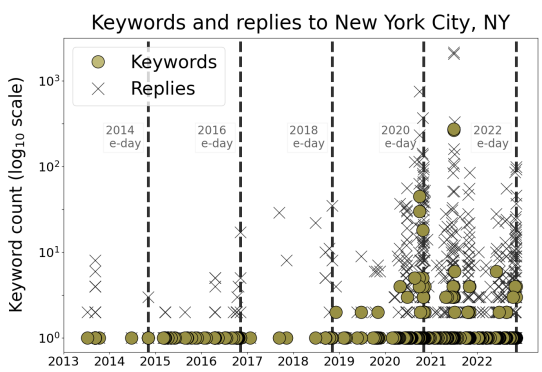

(a) New York City, NY

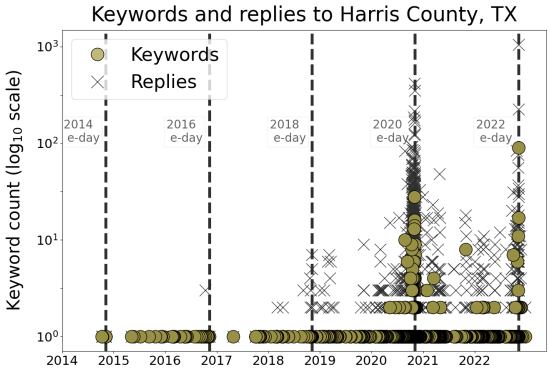

(b) Harris County, TX

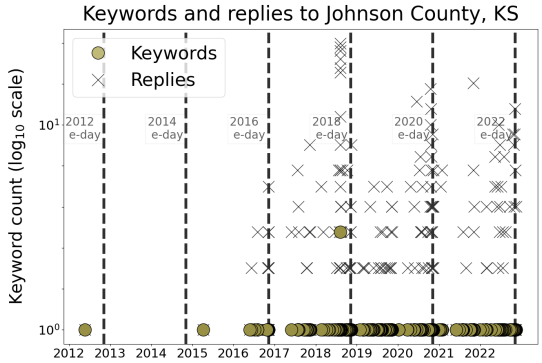

(c) Johnson County, KS

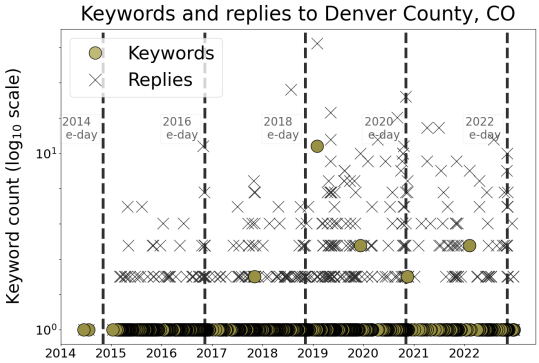

(d) Denver County, CO

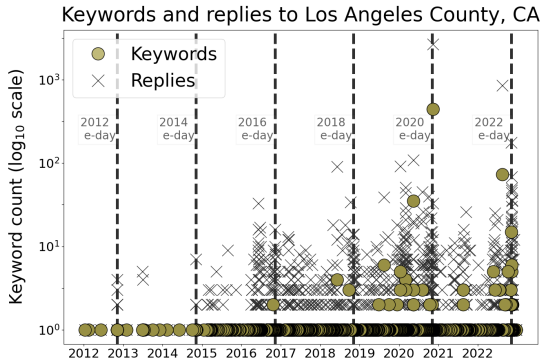

(e) Los Angeles County, CA

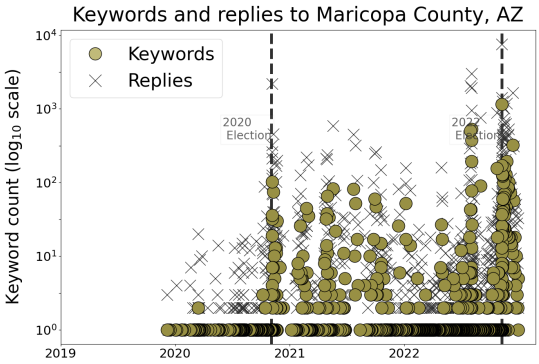

(f) Maricopa County, AZ

Figure 3: The total number of replies to tweets by election officials in six local jurisdictions with varying levels of salience, and the number of fraud-related keywords used in those replies.

Finally, Maricopa County in Arizona received by far the most keyword activity. The Maricopa County elections board did not have a Twitter account before 2020. During the

beginning of the county's first year on Twitter there was little fraud-related keyword activity in the replies, but after the 2020 election and nearing the beginning of 2021, the negative, fraud-related keywords start to pile on. Most notably, the keywords start at the beginning of 2021 and continue well into 2022. The activity of the replies to this Twitter account reflect the experience of the officials in this local jurisdiction: as seen in subfigure 3f, Maricopa County was elevated to a high level of national prominence as it continues to battle allegations of fraud and other falsehoods about the legitimacy of its elections.

Sentiments of Replies

While keywords help to capture the fraud-related conversations happening on these election officials' accounts, they do not completely explain the general trend towards negativity over time. Thus, we introduce sentiment analysis to plot trends in negativity over time for each of the six local jurisdictions in Figure 4. To capture sentiment, we generated a sentiment score for every word in the reply to a tweet using the VADER package in Python (Hutto and Gilbert 2014). Next, we took the average sentiment score for each reply and then the average of all of the replies to get an overall sentiment score for the replies.[7] Ultimately, we get a single value which roughly represents the average sentiment of the replies to a given tweet that an official made. Finally, we fit an OLS line to the plots to understand how this trend changes over the entire time period of this analysis.

As Figure 4 shows, five of the six jurisdictions we focus on have experienced increases in negativity over time; only New York City has a positive trend in the sentiment of its replies. Despite the negativity, however, our analysis finds that several of these jurisdictions receive replies that are on average more positive than negative: these include Denver County, Colorado (as seen in subfigure 4d) and Los Angeles County, California (subfigure 4e). And yet, even in a county such as Johnson County, Kansas, which is not particularly salient in

[7]Though there is no consensus method for distinguishing in general between human and automated activity on twitter, we drop some tweets where the only reply is obviously automated. This step only affects two of the six plots and it actually plays against our substantive claim, by making the OLS slopes more positive.

national narratives of election fraud, there is far more negativity in its replies during and since the period around the 2020 election.

Maricopa County is again the most extreme case. In Maricopa, shown in subfigure 4f, the replies have grown steadily more negative over time like the other jurisdictions, but what distinguishes this county is that the average sentiment of tweets has fallen to the point that it is now below zero (which represents the dividing point between a positive and negative post). Because Maricopa County only created its election administration Twitter account at the start of 2020, it is not possible to know what might have occurred beforehand. Yet, as seen on the plot, there is a clear change just from early 2020 to late 2020: the election is clearly a focal point for negativity in the jurisdiction as the barrage of negativity persists past the end of the year, continuing throughout 2020, all of 2021, and even into 2022.

A possible counterargument to the claim that election official Twitter replies are uniquely more negative over time is to say that all Twitter accounts are growing more negative over time. A separate work has examined whether this trend extends to American officials who do not oversee elections (Gross, Baltz, and Stewart III 2023). That paper analyzes the sentiment of the replies for accounts associated with state-wide Departments of Transportation. If negativity is getting worse more broadly, then we should expect to find similarly falling sentiments in the responses to the accounts of another branch of state government. The finding is that negativity towards Department of Transportation account has been extremely flat over time, which suggests that the trend in negativity that we observe in American election officials' accounts is evidence of public sentiment towards the people who administer American elections specifically, and not just one example of a broader phenomenon unrelated to American elections.

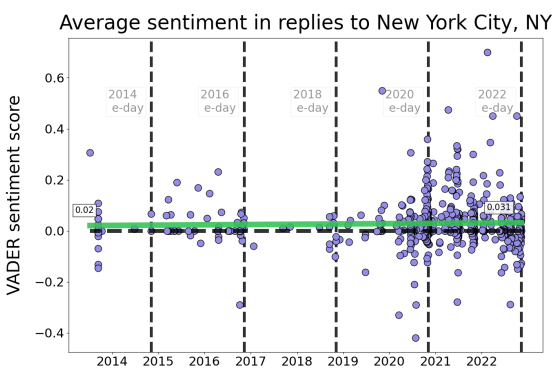

(a) New York City, NY

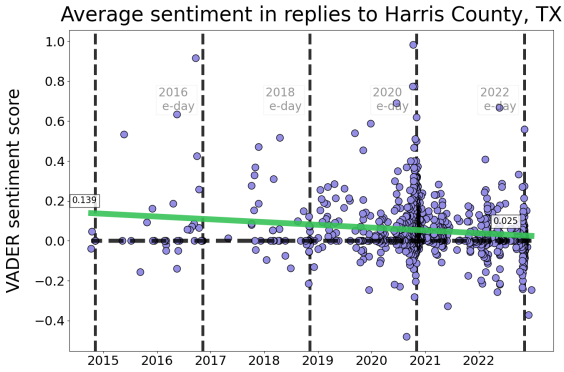

(b) Harris County, TX

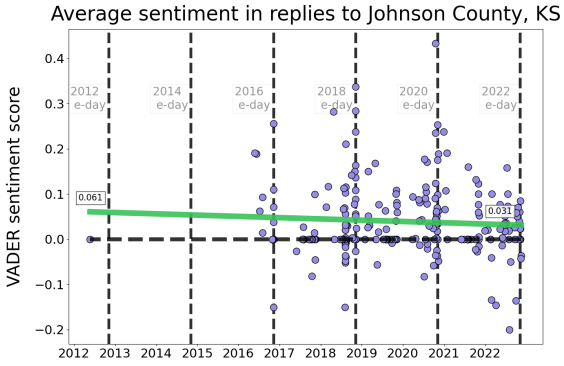

(c) Johnson County, KS

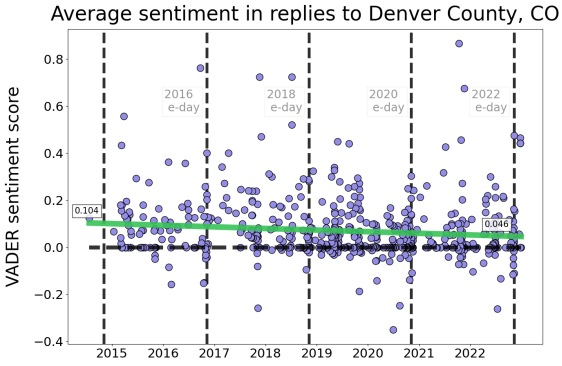

(d) Denver County, CO

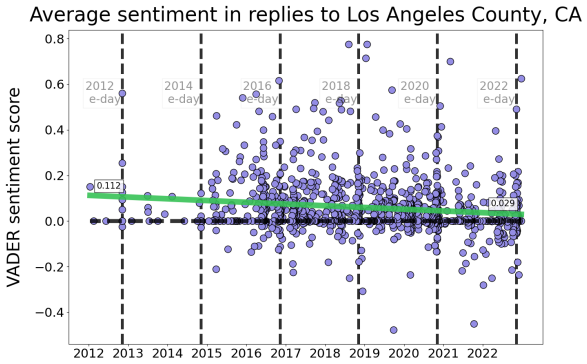

(e) Los Angeles County, CA

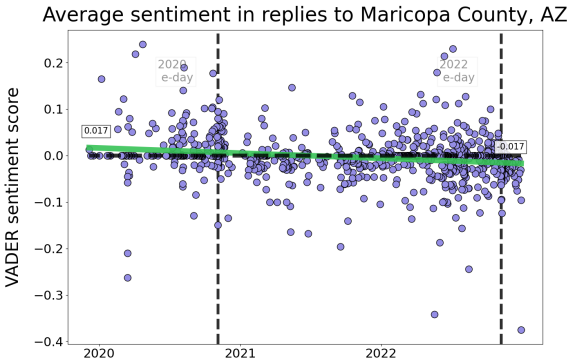

(f) Maricopa County, AZ

Figure 4: The mean sentiment score of replies to each tweet by six local-level election administration Twitter accounts, with each dot representing the average sentiment in replies to a single tweet. The average is computed by averaging the VADER sentiment scores of each word in each reply, and then taking the average across replies.

Ideologies of the Repliers

The previous sections highlight the differences in sentiment and keyword usage which occur across counties and local jurisdictions in various states. In certain counties like Maricopa County, Arizona, the negativity is increasing. The next question, then, is to ask who is responsible for increases in negativity in replies across counties. To determine the ideologies of users, we relied on Barberá's method for ideology estimation, which generates an ideal point estimate based on how many conservative or liberal elites a given user follows on Twitter (Barberá 2017).

In Figure 5 there is one point for each original tweet made by the Maricopa County LEO account. The y-axis position of a data point shows the average sentiments of replies to that tweet. The points are scaled according to the number of replies that the tweet received, with a larger circle corresponding to a post that received more replies. Squares represent tweets which had replies where the median ideology was liberal users and circles represent points where the median ideology of the repliers were conservative. The darkness of the point represents points which are more partisan. The plot shows that the respondents to this account are estimated to include both liberals and conservative users. However, the plot also shows that much of the positivity is in replies to tweets where most of the repliers are liberals, while tweets with replies of negative sentiment tend to have more conservative repliers.

Around the 2020 election, as we have noted, the Maricopa County account's tweets received many more replies than it ever had in the past. While the replies near the election ramp up, most of them are largely positive and from liberal users. It is evident, however, that the LEOs in Maricopa receive a mix of responses to their posts. For example, in a post made by the county on November 4^{th}, 2020, just a few days after the election but before the final results for presidential election had been called, they posted a cartoon photo with a ballot with a face and a clapper board with the caption "That's a wrap for Election Day! All Vote Centers are now reported. We'll update unofficial results daily until all valid ballots

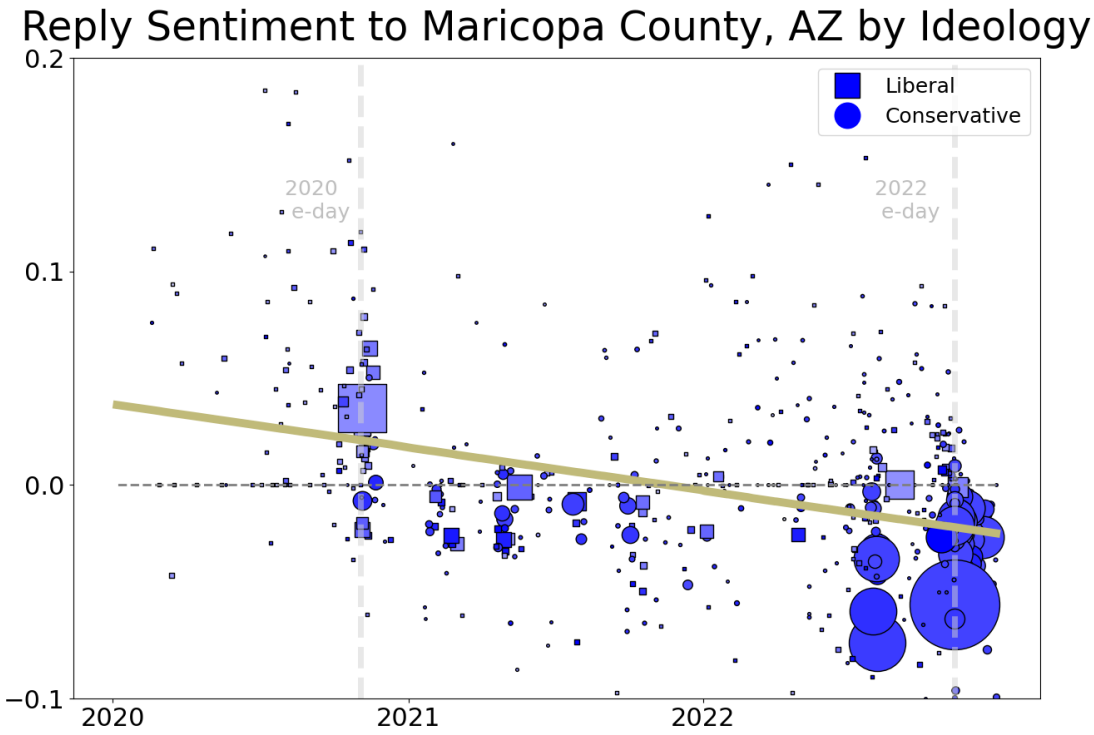

Figure 5: Maricopa County, AZ: : the mean sentiment scores of replies to each tweet by this account, sized by the number of replies, and shapes correspond to the median ideologies of a random sample of the repliers.

are counted."[8] The replies to the tweet are mixed, yet a substantial number of repliers commend the county for their hard work during a tumultuous election season. Many reply with statements like "Thank you for your hard work" and "Thanks for everything". Despite the praise, however, the account also receives negative responses alleging fraud. For example, one replier states in Maricopa "Voting fraud at its finest".

The account has been subject to more negativity as time has passed since the 2020 election. Most notably, this LEO account received substantial negativity during the summer and fall of 2022, which were the 2022 primary and midterm seasons. One tweet which drove negativity to the account was during the primary election. The Maricopa County account posted an image of a ballot drop off box with the caption, "We have emergency Vote Centers

[8]The tweet can be accessed here: https://twitter.com/MaricopaVote/status/1323919282334478336

open today for voters who have issues preventing them from voting on Election Day. But early voters can still drop off your ballot at any of our open Vote Centers or ballot drop boxes this weekend."[9] Repliers expressed skepticism and anger. Several repliers asked whether the ballot drop boxes would be monitored. One replier even stated: "If I was a patriot in AZ, I would camp out with my cellphone by one of these boxes." These responses not only express harsh sentiments, but they also reveal a deep mistrust in electoral legitimacy. And indeed, this one pledge that, if the user were in Arizona, they would camp out by the ballot boxes, reflects a broader fact: it is possible for individuals across the country — not only in these local jurisdictions themselves — to reply to these local election officials with negative, abrasive, and distrustful rhetoric.

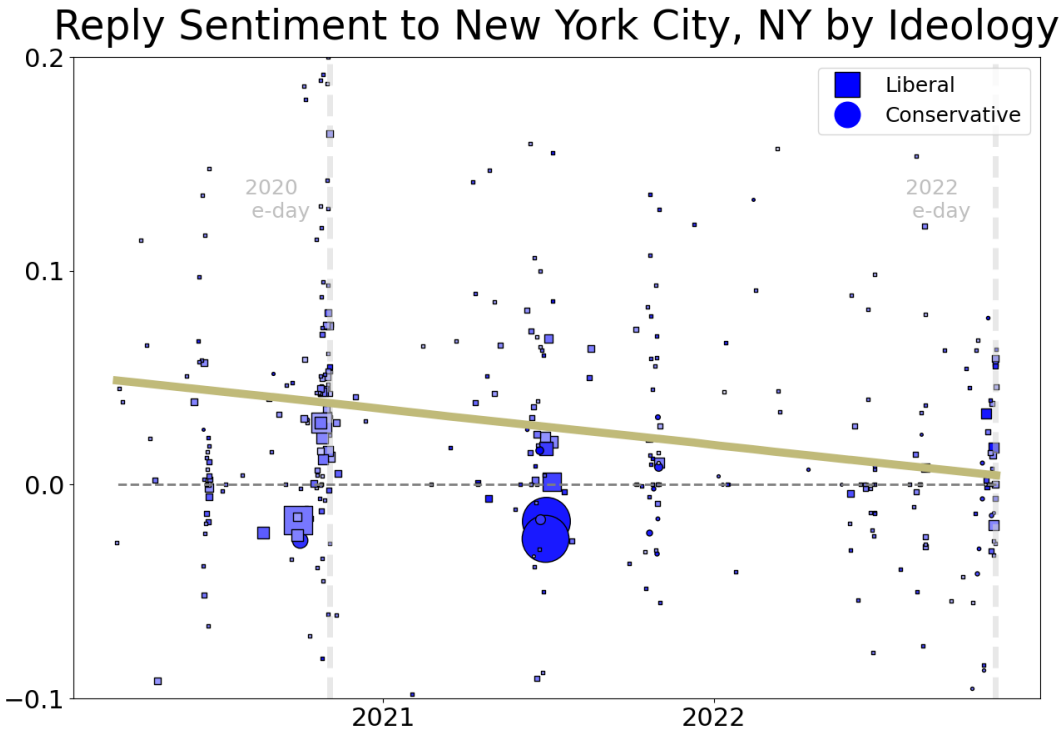

Figure 6: New York City, NY: the mean sentiment scores of replies to each tweet by this account, sized by the number of replies, and the shape of the point represents the median ideologies of a random sample of the repliers.

[9]Tweet available at: https://twitter.com/MaricopaVote/status/1553417904233652227

Is this trend limited to Maricopa County due to its particular position in the glare of elite-fueled conspiracy theories? Figure 6 shows that, for the LEO account for New York City, New York, like that of Maricopa County, there is increasing negativity in the replies they receive over time. In this section, we focus just on tweets since 2020, because this is the time period when the account received increasing negativity. While there are fewer tweets which receive quite as many replies as those in Maricopa, the ones which do receive replies show the same sort of trend we see in Arizona. Often, liberal users generate positivity towards their tweets and conservative users generate negativity. While negativity is not exclusive to conservative users, in this specific plot they drive the majority of the negative replies.

One such example of a tweet created by the officials in the county that generated a large number of replies occurred in October 2021. This tweet was posted in the month before New York's 2021 municipal election, which selected a new mayor of New York City. In this October 2021 tweet, the board of elections posted a graphic reminding people of important dates; it states: "The November General Election is right around the corner! Mark these important dates on your calendar."[10] The content of the replies are mixed, but here too, there are allegations of illegitimate electoral activity. Some replies allege that the county will send out ballots to deceased voters, while others question the mask policies at polling places and insist that voters need to "vote out everyone that protects these authoritarian mandates". Clearly the increasingly harsh rhetoric in the corpus overall is not limited to just the most salient case, but can be found in other jurisdictions as well, and the timing of this outlier further suggests that local (and not just national) political events can shape how people respond to LEOs.

Implications about the Ability of LEOs to Reach Voters

Previous work has established that *statewide* election officials are receiving increasing levels of negativity on social media (Gross, Baltz, and Stewart III 2023). In this chapter, we

[10]See that status here: https://twitter.com/BOENYC/status/1445491747744669696

extended that analysis by documenting and analyzing the experience of *local* officials on social media. We collected all the tweets and replies for 258 local jurisdictions which have active Twitter accounts for their election offices, collecting data dating as far back as 2012. After analyzing the prevalence of fraud-related keywords in the replies, the sentiments of the replies, and generating ideology estimates for some of the most negative replies, our analysis provides several key takeaways.

First, we find the incidence of tweets generated by local officials has grown over time. While local election officials have been on Twitter for the better part of a decade, the 2020 election was a turning point for their social media presence: they suddenly start tweeting much more than they ever had before. Next, when analyzing sentiment of the replies we find increasingly negative replies over time. While it is difficult to measure the degree of negativity, it is clear that in some places, like Maricopa County, the replies are getting far worse.

Keyword usage is concentrated in some contentious jurisdictions and not others; not every jurisdiction receives allegations of fraud but several of them receive a noticeable uptick after the 2020 election. Additionally, where we observe examples of a high concentration of negative replies, most (though not all) of that negativity comes from conservative users. Specifically, the tweets that we analyzed from accounts in Maricopa County, Arizona and New York City, New York which generated negative replies from conservatives were those that discussed what have become contentious issues in recent months, but were previously mundane topics, such as ballot drop box locations and the certification of election results.

Finally, to answer our initial motivating question of whether these trends are similar or different from those which described the experience of statewide officials, ultimately we largely see the same trend. Similar to their statewide counterparts, local election officials are not exempt from growing negativity and increasing allegations of fraud on social media. Local election officials are the backbone of election administration in our highly decentralized system; if they continue to experience threats and allegations of fraud, whether in person

or online, they may continue to flee these essential jobs. If LEOs reduce their activity on social media to avoid negativity there, voters may face challenges learning essential information about election administration and navigating an information environment where the distribution of false content is rampant.

Works Cited

Barberá, Pablo. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data". *Political Analysis*, vol. 23, no. 1, 2017, pp. 76–91.

Bidgood, Jess. America's Election Workers are Leaving in Droves. 2022. apps.bostonglobe.com/nation/politics/2022/10/democracy-under-siege/election-workers-leaving-in-droves/.

Brennan Center. Local Election Officials Survey. 2022. www.brennancenter.org/media/9197/download.

Clark, Amanda D., Christina S. Barsky, and Monica A. Bustinza. "Heavy Lifting: Emotional Labor and Election Administration". *Administration & Society*, 2022. https://doi.org/10.1177/0894439319865519.

Eisler, Peter, Jason Szep, Linda So, and Sam Hart. Anatomy of a death threat. 2021. graphics.reuters.com/USA-ELECTION/THREATS/mopanwmlkva/index.html.

Every-Palmer, Susanna, Justin Barry-Walsh, and Michele Pathé. "Harassment, stalking, threats and attacks targeting New Zealand politicians: A mental health issue". *Australian & New Zealand Journal of Psychiatry*, vol. 49, no. 7, 2015, pp. 634–41. https://doi.org/10.1177/0004867415583700.

Gorrell, Genevieve, Mehmet E. Bakir, Mark A. Greenwood, Ian Roberts, and Kalina Bontcheva. "Race and Religion in Online Abuse towards UK Politicians". *Working Paper*, 2019.

Gronke, Paul, and Paul Mason. "The State of Election Administration in 2022". *Democracy Fund*, 2022.

Gross, Joelle, Samuel Baltz, and Charles Stewart III. "What Happens When the President Calls You an "Enemy of the People?" Election Officials and Public Sentiment". *Election Law Journal: Rules, Politics, and Policy*, vol. 0, no. 0, 2023, null. https://doi.org/10.1089/elj.2022.0058.

Hutto, C.J., and Eric Gilbert. "VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text". *The International AAAI Conference on Web and Social Media*, 2014.

Korsell, Lars, Tomas Axelson, Liselotte Frisk, and Jonas Stier. "Harassment and Threats, Concern and Fear: The Experiences of Local Politicians in Ludvika, Sweden". *Journal for Deradicalization*, vol. 22, 2020, pp. 188–217.

Kuperberg, Rebecca. "Incongruous and illegitimate: Antisemitic and Islamophobic semiotic violence against women in politics in the United Kingdom". *Journal of Language Aggression and Conflict*, vol. 9, no. 1, 2021, pp. 100–126. https://doi.org/10.1075/jlac.00055.kup.

Merivaki, Thessalia, and Mara Suttmann-Lea. "Can electoral management bodies expand the pool of registered voters? Examining the effects of face-to-face, remote, traditional, and social media outreach". *Policy Studies*, 2022, pp. 1–31.

Office of Private Sector. Threats to Election Workers Likely More Prevalent in States with Recounts, Audits or Public Election Disputes. 2022.

PBS. Maine passes law protecting poll workers facing increased threats since the 2020 election. 2022. www.pbs.org/newshour/show/maine-passes-law-protecting-poll-workers-facing-increased-threats-since-the-2020-election.

Reuters. Campaign of Fear: The Trump world's assault on U.S. election workers. 2022. www.reuters.com/investigates/section/campaign-of-fear/.

Southern, Rosalynd, and Emily Harmer. "Othering Political Women: Online Misogyny, Racism and Ableism Towards Women in Public Life". *Online othering: Exploring digital violence and discrimination on the web*, ed. by Karen Lumsden and Emily Harmer, Palgrave Macmillan, 2019, pp. 187–210.

———. "Twitter, Incivility and "Everyday" Gendered Othering: An Analysis of Tweets Sent to UK Members of Parliament". *Social Science Computer Review*, vol. 39, no. 2, 2021, pp. 259–75. https://doi.org/10.1177/0894439319865519.

Stewart III, Charles. 2020 Survey of the Performance of American Elections. 2021. https://doi.org/10.7910/DVN/FSGX7Z.

Suttmann-Lea, Mara. Voter Education in the Digital Age: Local Election Official Use of Social Media and Webpages During the 2020 U.S. Presidential Election. 2022. mediawell.ssrc.org/expert-reflections/voter-education-in-the-digital-age-local-election-official-use-of-social-media-and-webpages-during-the-2020-u-s-presidential-election/.

Suttmann-Lea, Mara, and Thessalia Merivaki. "Don't Drown the Message: The Effects of Voter Education on Mail Ballot Acceptance in North Carolina". *Journal of Election Administration Research and Practice*, vol. 1, no. 2, 2023.

———. "Educating Voters Online: Local election Officials' use of Social Media During the 2020 U.S. Presidential Election". *Presented at the American Political Science Association conference*, 2021.

———. "How LEOs Educate Voters about Voting Requirements and Election Reform". *Local Election Administrators in the United States: The Frontline of Democracy*, ed. by Paul Gronke, Palgrave McMillan (Under Contract), 2023.

———. "The Impact of Voter Education on Voter Confidence: Evidence from the 2020 U.S. Presidential Elections Elections". *Election Law Journal*, vol. Online First, 2023.

Wagner, Angelia. "Tolerating the trolls? Gendered perceptions of online harassment of politicians in Canada". *Feminist Media Studies*, 2020, pp. 1–16. https://doi.org/10.1080/14680777.2020.1749691.