Negativity bias in the spread of voter fraud conspiracy theory tweets during the 2020 US election

Mason Youngblood^{*1,2,3}, Joseph M. Stubbersfield⁴, Olivier Morin^{3,5}, Ryan Glassman⁶, Alberto Acerbi⁷

¹Dept. of Psychology, Graduate Center, City University of New York, USA

²Dept. of Biology, Queens College, City University of New York, USA

³Minds and Traditions Group, Max Planck Institute for Geoanthropology, DE

⁴Dept. of Psychology, University of Winchester, UK

⁵Institut Jean Nicod, ENS, EHESS, PSL University, CNRS, FR

⁶IBM Watson, 51 Astor Place, New York, NY, USA

⁷Centre for Culture and Evolution, Brunel University London, UK

*youngblood@shh.mpg.de

Abstract

During the 2020 US presidential election, conspiracy theories about large-scale voter fraud were widely circulated on social media platforms. Given their scale, persistence, and impact, it is critically important to understand the mechanisms that caused these theories to spread. The aim of this study was to investigate whether retweet frequencies among proponents of voter fraud conspiracy theories on Twitter during the 2020 US election are consistent with frequency bias and/or content bias. To do this, we conducted generative inference using an agent-based model of cultural transmission on Twitter and the *VoterFraud2020* dataset. The results show that the observed retweet distribution is consistent with a strong content bias causing users to preferentially retweet tweets with negative emotional valence. Frequency information appears to be largely irrelevant to future retweet count. Follower count strongly predicts retweet count in a simpler linear model, but does not appear to drive the overall retweet distribution after temporal dynamics are accounted for. Future studies could apply our methodology in a comparative framework to assess whether content bias for emotional valence in conspiracy theory messages differs from other forms of information on social media.

Keywords: voter fraud, conspiracy theory, social media, transmission bias, negative valence

Introduction

Allegations of malign acts, carried out in secret by powerful groups, have been offered as explanations for major events throughout history, from ancient Rome, through the medieval period, to the present day (1–3). People across the globe have shared these conspiracy theories (4,5). Conspiracy theories have been a part of North American culture since the colonial period, with beliefs about conspiring, “un-American” groups of witches, enslaved Africans, Masons, Catholics, and Jews dominating early versions, before shifting to the US government itself as the source of conspiring agents in the nineteenth and twentieth centuries (6,7).

Conspiracy theories are typically defined as explanations of important events which allege secret plots by powerful actors as salient causes (8–10). Belief is not inherently irrational (as conspiracies do occur; see (11,12)), but conspiracy *theories* (as opposed to simply conspiracies) are allegations that survive and spread despite a lack of reliable evidence (8,10). Belief in conspiracy theories is associated with reduced engagement with mainstream politics (13,14), increased support for political violence and extremism (13,15), and increased prejudice towards minority groups (16,17).

A range of recent and ongoing conspiracy theories allege that the result of the 2020 US presidential election was achieved through large-scale electoral fraud (18). Building on allegations of voter fraud made prior to the 2016 election (19) and years of Republican messaging about electoral fraud and illegal voting (20), these conspiracy theories were widely circulated on social media platforms like Twitter. Major political and public figures, including US President Donald Trump, boosted these theories using hashtags like #stopthesteal (21) and eventually had their accounts suspended for incitement of violence following the January 6th attack on the US Capitol (22). More specific claims, such as hacked voting machines being programmed in favor of then-Presidential Candidate Joe Biden and large numbers of ballots being thrown out in trash bags (23,24) have been used to justify election audits and tighter voting laws in states like Arizona (25) and Georgia (26). The Justice Department has found “no evidence of widespread voter fraud” (27), and the Cybersecurity and Infrastructure Security Agency concluded that 2020 was “the most secure election ever” (28). Despite this, polls suggest that up to a third of Americans (29) and the majority of Republicans (30) believe that Biden won the election illegitimately through voter fraud. Exposure to such claims has been shown to reduce confidence in democratic institutions (31) and is thought to have contributed to motivating the US Capitol attack (32). Given the scale, persistence, and impact of voter fraud conspiracy theories, it is critically important to understand the mechanisms that caused them to spread.

While conspiracy theories, as everything else, are disseminated through social media, the nature of the association between social media usage and conspiracy theory belief is an open question (18,33–35). Social media does provide, however, a source of data that can be used to test theories about their spreading. Most studies of the spread of conspiracy theory messages on social media have focused on the content of posts in general, highlighting the importance of negative content (36), emotional content (37), or out-group derogation (38,39). However, content is only one of the possible features that influence the success of a social media post. In what follows, we use a framework inspired by cultural evolution that allows us to distinguish among various features and assess their relative importance. It is important to highlight up-front that this is not a comparative study, so we are unable to make conclusions about whether detected processes are unique to conspiracy theories relative to other forms of information. Instead, our goal is simply to characterize what transmission processes are present in a particular high-profile case of conspiracy theory spread on social media.

Broadly, cultural evolution adopts an evolutionary framework to research the stability, change and diffusion of cultural traits (40). Transmission biases—biases in social learning that cause individuals to adopt some cultural variants over others—are thought to be some of the most important factors driving cultural evolutionary patterns (41). According to this perspective, the probability that a behavior will be adopted is influenced by various cues. Frequency bias, which includes conformity and novelty bias, is when the frequency of a variant in the population disproportionately affects its probability of adoption. Content bias is when the inherent characteristics of a variant affects its probability of adoption. Demonstrator bias is when some characteristic of the individuals expressing a variant affects its probability of adoption (see review in (41)). Importantly, transmission biases can lead to discernible changes in the cultural frequency distributions of populations (42). For example, in the context of Twitter, a positive frequency bias (i.e. conformity) would cause users to be more likely to retweet content that has already been heavily retweeted by other users, thus increasing the right skew of the overall retweet distribution. This framework allows us to consider both individual susceptibility and the influence of social context on wider population level patterns.

Using generative inference, it is possible to infer the underlying cognitive biases of individuals in a population from the cultural frequency distribution that they generate. Generative inference is a statistical procedure in which a model is run many times with varying parameter values to generate large quantities of simulated data. This simulated data is then compared to real data using approximate Bayesian computation (ABC) to infer the parameter values that likely generated it (43). Carrignon et al. (44) recently applied generative inference to the spread of confirmed and debunked information on Twitter and found that the retweet distributions of both were more consistent with random copying than with conformity. However, their model did not include parameters for the influence of follower count and did not explore the influence of content bias due to computational limitations (44).

The aim of this study is to investigate whether retweet frequencies among proponents of voter fraud conspiracy theories on Twitter during the 2020 US election are consistent with frequency bias and/or content bias. To do this, we conducted generative inference using an agent-based model (ABM) of cultural transmission on Twitter that combines elements from Carrignon et al. (44), Lachlan et al. (42), and Youngblood and Lahti (45) (see Methods for details). Our ABM simulates a population of Twitter users with follower counts, activity levels, and probabilities of tweeting original tweets from real users in the observed data. Every six hours, a subset of users become active and either compose a new tweet or retweet an existing tweet. The probability of an existing tweet being retweeted is based on four factors: (1) the attractiveness of the content in the tweet (e.g. emotional valence and/or intensity), (2) the follower count of the user who tweeted it, (3) how many times it has already been retweeted, and (4) the age of the tweet. The influence of each of these factors on retweet probability is controlled by separate parameters which correspond to content bias (c), follower influence (d), frequency bias (a), and age dependency (g), which we fitted to real data using the random forest version of ABC (46). This ABM follows Carrignon et al. (44) in assuming a fully-connected population so that, under neutral conditions, retweet probability is independent of follower count. The parameter d simulates departure from this baseline, where follower count has increasing influence on retweet probability. Our original intent was to use d to assess demonstrator bias, but a more parsimonious explanation for an effect of follower count is network structure—tweets from people with more followers appear on more users’ timelines. Without complete details of Twitter’s proprietary timeline algorithm it is very difficult to disentangle these two factors, so we will describe d with the more general term “follower influence”.

The data used in this study comes from a team of researchers at Cornell Tech, who retrieved millions of tweets and retweets relating to voter fraud conspiracy theories between October 23 and December 16 of 2020 (47). After iteratively building a set of search terms from the seeds “voter fraud” and #voterfraud and using them to collect data in real-time, they estimate that they collected ~60% of tweets about voter fraud conspiracy theories during that period. An anonymized version of the *VoterFraud2020* dataset is public (<https://github.com/sTechLab/VoterFraud2020>), and Abilov et al. (47) generously provided us with access to their full disambiguated dataset. Importantly, this dataset includes tweets from users who were “purged” from Twitter following the US Capitol attack (48). Beyond the quality of the data and its intrinsic historical importance, the choice of this dataset was motivated by its exceptional scale and

focus. Few other datasets gather such a large number of social media actors engaging in one conversation topic over a period of several months.

Additionally, we conducted secondary analyses with general linear mixed models (GLMM) to assess the potential targets of content bias, while accounting for other factors such as follower count and whether the account holder was verified. Twitter verifies some accounts to make sure they are authorized by the person they claim to represent, but only undertakes this costly verification for high-profile accounts whose status is signaled by a “blue check mark” icon. The emotional valence of tweets was measured using the valence aware dictionary and sentiment reasoner (VADER), a sentiment analysis model trained for use with Twitter and other social media data (49). The output of VADER includes a compound score of the overall emotional valence and intensity from strongly negative to positive, in addition to the proportion of words that are negative, positive, and neutral. A large body of research suggests that content with negative valence has an advantage over content with positive valence across several domains (50,51). In digital media, evidence of negativity bias has been suggested within online “echo chambers” (52,53) and for tweets about political events both from individual users (36) and institutions (54). False rumors, on the other hand, seem to be subject to a positivity bias (55), despite the fact that “fake news” articles tend to contain more negative language (56). Other studies have suggested that just the strength of emotion influences the transmission of content on social media (37,57) (but see critiques in (58)), and, in two experimental studies, the appeal of conspiracy theories was associated with the intensity of emotion evoked, rather than the valence of that emotion (59).

Based on this research, if content bias is detected, then we hypothesize that it will be targeted towards more emotional content (i.e. more positive or negative according to VADER, which collapses valence and intensity into one indicator), but we remain agnostic as to the direction of valence.

Results

The *VoterFraud2020* dataset is divided into several sub-communities, including both detractors and proponents of the conspiracy theories. We chose to focus on cluster #2, the “proponent” community that tweets and retweets content in English and does not have significant connections to the “detractor” community (see Methods). After subsetting the *VoterFraud2020* data to only include user and tweet data from cluster #2, we ended up with 3,982,990 tweets from 341,676 users. Note that we calculated the number of users as all unique users that either tweeted or retweeted content from cluster #2. The agent-based model was initialized with a population size (N) of 341,676. The model and methods were preregistered in advance of the analysis (see Methods).

	M	95% CI	NMAE
c : content bias	4.612	[3.479, 5.679]	0.530
d : follower influence	0.362	[0.007, 1.663]	0.741
a : frequency bias	0.295	[0.004, 0.916]	1.458
g : age dependency	2.043	[0.019, 7.824]	2.016

Table 1. The median, 95% credible interval, and out-of-bag normalized mean absolute error of the posterior distribution for each dynamic parameter in the agent-based model.

The posterior distributions for content bias (c), follower influence (d), frequency bias (a), and age dependency (g) can be seen in Figure 1 and Table 1. Higher values of c , d , and g are indicative of stronger effects of those parameters, where 0 is neutrality. Values of a that are lower and higher than 1 are indicative of novelty and conformity bias, respectively, where 1 is neutrality. The median estimate for content bias is 4.612, with a 95% credible interval (CI) that spans from 3.479 to 5.679. This indicates that content bias plays a significant and strong role in driving retweet frequencies. Follower influence has a median estimate of 0.362 with a wide 95% CI and a right-skewed posterior that is converging towards zero. This indicates that follower count has a weak effect on retweet frequencies that is either difficult to estimate or varies among users. The median estimate for frequency bias is 0.295 with a similarly wide and right-skewed posterior. If we assume that neutrality in frequency bias is $a = 1$, where retweet probability is perfectly proportional to the number of times a tweet has already been shared, then this result is indicative of an extremely strong novelty bias. However, we feel that a novelty bias of this magnitude is unrealistic, and that this result instead suggests that retweet probability is mostly decoupled from the number of times a tweet has already been shared (so neutrality is $a = 0$). This is much more consistent with our personal experiences on Twitter, where the timeline includes a balance of new tweets from followers and trending tweets that have already been heavily retweeted. According to this interpretation, the wide 95% CI for a could reflect departures from neutrality resulting from conformity bias, variation in user behavior, or how Twitter’s timeline algorithm weights trending tweets. The posterior distribution for age dependency is also converging towards zero, but has high out-of-bag error and probability mass across the entire prior. We chose to make no conclusions about this parameter. Chaotic posterior distributions such as this one can sometimes

result from misspecified priors, but we feel that this is unlikely to be the case here as higher values of g generate retweet distributions that are further away from the observed data (Figure S5).

Figure 1. The prior (dotted lines) and posterior (solid lines) distributions for each of the four dynamic parameters from the ABM that were estimated using ABC.

Importantly, the details of Twitter’s timeline algorithm are not publicly available, which means that it is not statistically possible for us to differentiate between a bias produced by the algorithm or user behavior. This issue is most relevant for the results related to follower influence and frequency bias, as follower and retweet count are the two cues most likely utilized by Twitter’s timeline algorithm (60).

	IRR	95% CI
<i>Tweet length</i>	1.478	[1.477, 1.480]
<i>Presence of media</i>	1.515	[1.507, 1.523]
<i>Follower count</i>	1.505	[1.496, 1.515]
<i>Compound score</i>	0.959	[0.958, 0.960]

Table 2. The incidence rate ratio (IRR) and 95% confidence interval for each predictor in the best fitting model. IRR, the exponentiated beta estimate, is interpreted as the rate at which the outcome variable is expected to change per unit increase in a predictor (one standard deviation for scaled and centered predictors). Wald confidence intervals were used due to the high sample size.

Based on the null models for the GLMM, user appears to be the only grouping variable that explains a high level of variance in the data ($ICC_{\text{user}} = 0.610$, $ICC_{\text{date}} = 0.109$, $ICC_{\text{hour}} = 0.039$). As such, we chose to include user as a random effect in our base model. Adding tweet length as a control variable improved model fit ($\Delta AIC = 40803$; LRT: $\chi^2 = 40805$, $p < 0.0001$). Adding the presence of media (i.e. images or videos) as a control variable improved model fit ($\Delta AIC = 510$; LRT: $\chi^2 = 511$, $p < 0.0001$). Both follower count and verification status further improved model fit ($\Delta AIC > 2$), but the model with follower account was significantly better ($\Delta AIC = 6204$; LRT: $\chi^2 = 6205$, $p < 0.0001$) so we updated our base model accordingly. All three content measures further improved model fit ($\Delta AIC > 2$), but the model with the compound score was significantly better than the models with the proportion of negative words ($\Delta AIC = 831$; LRT: $\chi^2 = 830$, $p < 0.0001$) or positive words ($\Delta AIC = 1276$; LRT: $\chi^2 = 1276$, $p < 0.0001$). All model specifications and AIC values for the primary GLMM are in Table S2, along with a partial specification curve analysis in Figure S7. The best fitting model included user as a random intercept, and tweet length, presence of media, follower count, and the

compound score as fixed effects (see Table 2). Note that we chose to not report p -values for our GLMM results, given the unreliability of p -values at higher sample sizes (61). Instead, we focused on the effect sizes (calculated here as incidence rate ratios) and 95% confidence intervals, where intervals that do not overlap with 1 are interpreted as evidence for a significant effect.

Tweet length, presence of media, follower count, and the compound score all have significant effects on retweet frequency (Table 2). The incidence rate ratio (IRR) for tweet length is 1.478, indicating that if a tweet is one standard deviation (SD) longer then it is 47.8% more likely to be retweeted. The IRR for presence of media is 1.515, so tweets with images or videos are 51.5% more likely to be retweeted. Follower count has a similarly strong effect, where tweets from users with one SD more followers are 50.5% more likely to be retweeted. The IRR for the compound score is much lower but still significant. Tweets with a compound score that is one SD lower (more strongly negative) are 5.7% more likely to be retweeted. The models that include the proportion of words that are negative or positive instead of the compound score, fit during model choice, show similar effects: tweets with one SD more negative words are 8.0% more likely to be retweeted, and tweets with one SD more positive words are 6.4% less likely to be retweeted. Pseudo R^2 values, calculated using log-normal approximation, indicate that the fixed effects alone account for about 10% of the variance in the data ($R^2 = 0.101$), whereas the fixed effects and random intercept together account for about 68% of the variance in the data ($R^2 = 0.676$) (62). A variance inflation factor test indicates that there are no significant issues with multicollinearity between predictors (VIFs < 2). Residual diagnostics for the best fitting model indicate that, while there are some extreme low and high outliers, the Poisson family is appropriate and there are no significant problems with overdispersion (see Figure S8).

An additional GLMM found that quote tweets tend to have reduced negative valence relative to the original tweets that they are quoting (see SI). This lends support to our generative inference results, given our decision to treat quote tweets like original tweets when computing retweet distributions (see Methods). If quote tweets tend to be less attractive than original tweets, then our estimate for content bias is likely more conservative than it would have been if we had treated quote tweets like retweets instead of original tweets.

Discussion

Based on the results of generative inference, the observed retweet distribution is consistent with a strong content bias. Tweets with a lower compound score are more likely to be retweeted, suggesting that this bias is targeted towards negative emotional valence. Follower count, on the other hand, has a weak effect in the agent-based model and a strong effect in the GLMM. This discrepancy is likely due to the fact that the former accounts for key temporal dynamics that cannot be included in the latter, such as activity levels over time, the effect of current retweet count on future retweet probability, and age dependency. For frequency bias, the results of generative inference could be interpreted as evidence for either an extremely strong novelty bias or for frequency information being irrelevant. We feel that a novelty bias of this magnitude is unrealistic, and we interpret this result as evidence that retweet probability is mostly decoupled from the number of times a tweet has already been retweeted. Interestingly, we also found that quote tweets tend to contain less negative emotional valence than their targets. This means that users do not tend to amplify negativity when commenting on a retweet, despite having a content bias for negative valence. Quote tweets are often thought to reflect criticism, so this reduced negativity may be unique to discussions between like-minded users.

Importantly, the results related to follower influence and frequency bias are difficult to separate from the influence of Twitter’s timeline algorithm, which is heavily based on user characteristics and engagement (60). In other words, we do not know how much the algorithm prioritizes tweets that are popular or from followed users. In the case of follower count, this makes it impossible to disentangle the effect of network structure from a demonstrator bias. We were unable to collect the follower network because of account suspensions following the January 6th attack (47), but even if we had access to the network structure the relevance of it would depend upon user behavior. If users primarily share information that they see passively on their timeline, then network structure is much more important, whereas if users are searching with keywords and hashtags, then it is less so. In this case, we think the absence of an observed effect of follower count is due to the latter. Proponents of conspiracy theories may actively seek out messages using keywords and hashtags, which are known to be important in the spread of disinformation (63), whereas more generalist Twitter users may rely more on the timeline and be more impacted by network structure (52). People may also vary in how selectively they retweet particular users or groups (e.g. only retweeting former President Trump). Based on the confounding effects of the algorithm and network structure, we are most confident in concluding that the spread of voter fraud claims among proponents of voter fraud conspiracy theories on Twitter during and after the 2020 US election was partly driven by a content bias causing users to preferentially retweet tweets with more negative emotional valence.

Our results are consistent with previous work suggesting that emotionally negative content has an advantage on social media across a variety of domains, including news coverage and political discourse (36,52–54). Other studies, though, have shown that positive messages spread more slowly but reach more people (64), that exposure to both

positive and negative tweets increases the probability of a user tweeting content with similar emotional valence (65), and that tweets with greater emotional intensity (independent of valence) are more likely to be retweeted (37,57). In such cases, there seems to be significant variation across domains and individuals. Messages about same-sex marriage, for example, are more likely to be retweeted if they use positive language, whereas messages about climate change are more likely to be retweeted if they use negative language (37). We suspect that conspiracy theory content generally falls into the latter category. Similarly, another study found that there is variation in how users respond to emotional content, where some “highly susceptible” users are more likely to be influenced by positive messages (65).

Most recently, Pröllochs et al. (55) found that false rumors with more positive language are more likely to go viral, despite the fact that articles espousing false rumors tend to have more negative language (56). The discrepancy between our findings and this evidence for positivity bias in “fake news” is puzzling, but we think it is due to differences in granularity and methodology. Pröllochs et al. worked with unique news-related rumors spreading through a general Twitter population (55,66) whereas we analyzed sub-conversations about the same topic within a single conspiracy theory community. Content biases for emotional valence may vary when messages are shared between like-minded individuals or embedded within a single conversation. Additionally, while Pröllochs et al.’s methodology is based on a well-established and cross-culturally validated emotion model, their primary valence measure has a low correlation with human raters when applied to social media data ($r = 0.11$) (55) relative to VADER ($r = 0.88$) (49) and other methods like SentiStrength (>60% accuracy) (65).

That being said, the sentiment analysis in this study has several limitations that should be highlighted. First, some new phrases that emerged during the 2020 election (e.g. “stop the steal”) may differ in emotional valence compared to their constituent words in VADER’s lexicon (49). Second, several of the common words and phrases used to create the *VoterFraud2020* dataset (e.g. “fraud”) have negative valence in VADER’s lexicon. Luckily, this only appears to introduce a slight skew in our valence measure (Figure S6), which is common in studies of emotion and sentiment on social media (55,58,64,65) and is unlikely to affect our results as GLMMs make no distributional assumptions about predictors (67). To improve the robustness of modeling of emotional contagion on social media, Burton et al. (58) recently came up with three recommendations for future studies: going beyond correlational evidence, analyzing the effect of specification decisions on model estimates, and preregistration. We fully agree with these recommendations, and we hope that we adequately addressed them by using a preregistered generative inference framework to ensure that the data was consistent with transmission bias before conducting GLMM and ensuring that our estimates were robust across a reasonable range of modeling specifications.

Regarding the spread of conspiracy theories, previous research has proposed “herd behavior”, in which rational individuals with limited information defer to the beliefs of the majority, to be a potential explanation (68,69). Our study addresses the sharing of conspiracist tweets among proponents, who presumably already believed some voter fraud claims before the election took place, but our lack of clear evidence for a frequency bias suggests that a disproportionate tendency to “follow the herd” may not be the primary driver of the spread of conspiracy theory messages. Rather, our study suggests that the content of conspiracy theory messages were, in the case we studied, more salient cues for cultural transmission. That being said, we were unable to incorporate many of the contextual and individual-level factors that have been associated with conspiracy theory belief in survey-based and experimental studies, such as anxiety (70), age (71), demographic diversity among social ties (35), and morbid curiosity (72). Such data are difficult, if not impossible, to infer from social media corpora, but could be included in future analyses based on more detailed profiles of individuals who have adopted conspiracy theory beliefs (as has been done in the extremism literature (73,74)).

While recognized as important, the transmission processes involved in the spread of conspiracy theories have received relatively little attention in research and are not well understood (75). This study demonstrates the value of cultural evolutionary approaches for understanding the transmission processes at play on social networks, and it highlights the importance of considering the roles of both the content of conspiracy theories and the context in which they are shared. Identifying and characterizing the biases influencing the transmission of conspiracy theories can help us to generate potential methods for countering the spread of harmful conspiracy theories and promoting the spread of genuine information (76).

A previous study using generative inference to investigate behavior on Twitter found that retweet patterns of both confirmed and debunked information were more consistent with unbiased random copying than with conformity (44). At first glance our study seems to contradict this result, but Carrignon et al. (44) did not include a parameter for follower influence in their agent-based model, and they assumed neutrality for content bias due to computational limitations. When we ran our agent-based model with neutral values for both content bias and follower influence, we too found that the model best fit the observed data when copying was unbiased by frequency (see Figure S4). The discrepancy between the results when parameters are estimated together instead of individually highlights the importance of considering equifinality—the fact that different processes can lead to similar patterns at the population level (77). If different processes lead to only subtle differences in retweet frequencies, then the effect of one could be mistakenly attributed to another if both are not considered simultaneously. For example, in our study we found that the observed

retweet distribution was consistent with content bias when content bias was estimated alongside other parameters (Figure 1 and Table 1), but not when it was estimated in isolation (Figure S4).

One of our biggest takeaways from this study is the importance of algorithmic transparency and accountability (78,79). Without access to detailed information about recommendation algorithms, researchers will continue to face difficulty in constructing realistic null models and making inferences about behavior on social media. Given that Twitter’s timeline algorithm uses deep learning (60) it is likely impossible for them to simply make the details of it public. Instead, the company could try to infer how the algorithm boosts different kinds of content by running natural experiments on the platform, as has been recently done for racial and gender bias in the image cropping algorithm (80) and right-leaning political bias in the timeline algorithm (81). For example, Twitter could publish the results of a model in which simulated users randomly share information from simulated timelines constructed by the algorithm to see how different kinds of content spread under neutral conditions. Luckily, algorithmic transparency and accountability are increasingly being viewed as public policy priorities by governments around the world (82,83). We hope that future studies can take advantage of improved transparency to develop more effective policy recommendations for fighting the spread of conspiracy theories and disinformation on social media platforms.

In conclusion, our methodology, based on a cultural evolution framework, allowed us to weigh the relative importance of different features influencing the spread of voter fraud claims among conspiracy theorists on Twitter. Most importantly, we found that retweet frequencies of voter fraud messages posted during and after the 2020 US election are consistent with a content bias for tweets with more negative emotional valence. While previous research focused a priori on the role of tweets’ content, without evaluating other possibilities, we were able to show that content is indeed central when compared with other possible mechanisms of social influence. Future studies could apply our methodology in a comparative framework to assess whether content bias for emotional valence in conspiracy theory messages differs from other forms of information on social media.

Methods

The data for this study comes from the *VoterFraud2020* dataset, collected between October 23 and December 16 of 2020 by Abilov et al. (47). This dataset includes 7.6 million tweets and 25.6 million retweets that were collected in real time using Twitter’s streaming API. The *VoterFraud2020* dataset was collected according to Twitter’s Terms of Service and is consistent with established academic guidelines for ethical social media data use (47). Abilov et al. (47) started out with a set of keywords and hashtags that co-occurred with “voter fraud” and #voterfraud between July 21 and October 22, and expanded their search with additional keywords and hashtags as they emerged (e.g. #discardedballots and #stopthesteal). They estimate that their dataset includes at least 60% of tweets that included their search terms. Abilov et al. (47) also applied the infomap clustering algorithm to the directed retweet network to identify different communities that engaged with voter fraud conspiracy theories. We ran our analysis using only the user and tweet data from cluster #2, the “proponent” community that tweets primarily in English and does not have significant connections to members of the “detractor” community. We restricted our analysis to cluster #2 so that retweets would be indicative of the spread of the conspiracy theories among proponents, as opposed to discourse and debate between both proponents and detractors. The activity levels and original tweet probabilities from these data only reflect users’ interactions with conspiracy theory content. By using these as input to our agent-based model, we are only simulating the subset of users’ behavior that pertains to the voter fraud conspiracy theory and is relevant for producing the observed retweet distribution.

The agent-based model (ABM) we used has elements from Carrignon et al. (44), Lachlan et al. (42), and Youngblood and Lahti (45), and is available as an R package on GitHub (<https://github.com/masonyoungblood/TwitterABM>). The ABM is initialized with a fully-connected population of N users and is run for 216 timesteps, each of which correspond to a six-hour interval in the real dataset (the highest resolution possible given computational limits). Each simulated user is assigned a follower count (T), an activity level (r), and a probability of tweeting an original tweet (μ) from a real user in the observed data. In this way, we retain the correlation structure of follower count, activity, and retweet probability to reflect real variation in users’ behavior on Twitter (e.g. users with few followers who exclusively retweet). μ is the proportion of a user’s total tweets and retweets that are original tweets (population mean of ~ 0.45). T is scaled with a mean of 1 and a standard deviation of 1. The ABM is also initialized with a set of tweets with retweet frequencies drawn randomly from the first timestep in the observed data. Each tweet is assigned an attractiveness (M). At the start of each timestep, a pseudo-random subset of users becomes active (weighted by their values of r) and tweets according to the observed overall level of activity in the same timestep. Each active user either tweets original tweets or retweets existing tweets based on their unique value of μ . New original tweets are assigned an attractiveness of M , while retweets occur with probability $P(x)$:

$$P(x) = \frac{F_x^a \cdot T_x^d \cdot M_x^c \cdot \frac{1}{age_x^g}}{\sum_{x=1}^n F_x^a \cdot T_x^d \cdot M_x^c \cdot \frac{1}{age_x^g}}$$

The denominator simply normalizes the probability of retweeting x against the probability of retweeting all other tweets in the population. F is the number of times that a tweet has been previously retweeted, and is raised by the level of frequency bias (a). a is the same across all agents, where values > 1 simulate conformity bias and values < 1 simulate novelty bias. T is raised by the level of follower influence (d). d is the same across all agents, where values of 0 simulate neutrality by removing variation in follower count and values > 0 simulate increasing levels of follower influence. M is the attractiveness of the tweet, and is drawn from a truncated normal distribution with a mean of 1, a standard deviation of 1, and a lower bound of 0. M is raised by the level of content bias (c). c is the same across all agents, where values of 0 simulate neutrality by removing variation in the attractiveness of content and values > 0 simulate increasing levels of content bias. The final term simulates the decreasing probability that a tweet is retweeted as it ages, where g controls the rate of decay. Once the active users are done each tweet increases in age by 1 and the next timestep begins. Lastly, we should note that we chose to exclude “top n ” dynamics (e.g. trending topics) from our ABM, because they did not improve the fit of neutral models of cultural transmission on Twitter in Carrignon et al. (44) and hashtags/keywords related to voter fraud rarely made it into the top trending topics during our study period (<https://www.exportdata.io/trends/united-states>).

In summary, the following are the dynamic parameters in this ABM that we estimated using approximate Bayesian computation (ABC):

- c - variation in the salience of the attractiveness of content
- d - variation in the salience of follower count
- a - level of frequency bias
- g - rate of decay in tweet aging

All other parameters in the ABM were assigned static values based on the real dataset. The output of this ABM is a distribution of retweet frequencies (see Figure 2), which was used to calculate the following summary statistics: (1) the proportion of tweets that only appear once, (2) the proportion of the most common tweet, (3) the Hill number when $q = 1$ (which emphasizes more rare tweets), and (4) the Hill number when $q = 2$ (which emphasizes more common tweets). We used Hill numbers rather than their traditional diversity index counterparts (Shannon’s and Simpson’s diversity) because they are measured on the same scale and better account for relative abundance (84,85).

The same summary statistics were calculated from the observed retweet distribution of the real dataset. For purposes of the summary statistic calculations quote tweets were treated like original tweets, as they themselves can be retweeted. Then, the random forest version of ABC (46) was conducted with the following steps:

- 500,000 iterations of the ABM were run to generate simulated summary statistics for different values of the parameters: c , d , a , and g .
- The output of these simulations were combined into a reference table with the simulated summary statistics as predictor variables, and the parameter values as outcome variables. The parameter values were logit-transformed prior to inference, as is recommended for bounded outcome variables in non-linear regression-based ABC (86,87).
- A random forest of 1,000 regression trees was constructed for each of the four parameters using bootstrap samples from the reference table. The number of sampled summary statistics and the optimal minimum node size were tuned to minimize prediction error using the *tuneRanger* package in R (88) (see Table S1).
- Each trained forest was provided with the observed summary statistics, and each regression tree was used to predict the parameter values that likely generated the data.

Uniform prior distributions were used for all four of the dynamic parameters: $c = \{0-12\}$, $d = \{0-4\}$, $a = \{0-2\}$, $g = \{0-8\}$. We plotted the output from 10,000 iterations to ensure that we were capturing enough of the parameter space before running the full analysis (see Figure S1). We conducted posterior checks by running the agent-based model with parameter values drawn from the posterior distributions to see how closely the output matched the original data (see Figure S1 and Figure S2). We also repeated the analysis using the basic rejection form of ABC to ensure that our main conclusions are robust to the choice of the random forest algorithm (see Figure S3). Finally, we ran four additional rounds of the ABM, each with only a single term from the probability function included, to investigate the behavior of each parameter in isolation (see Figure S4 and Figure S5).

Sentiment analysis was conducted using VADER from the natural language toolkit in Python, a model that performs similarly to human raters when applied to social media posts from platforms like Twitter ($r = 0.88$) (49). VADER assigns a valence score to each word (and emoji or emoticon) in a tweet, and weights those scores according to

a set of rules (e.g. negation, capitalization, punctuation). The main output of VADER is a compound score that sums and normalizes the weighted valences of the words in a tweet to give an overall score of emotional valence and intensity between -1 (strongly negative) and +1 (strongly positive) (Figure S6). VADER also outputs the proportion of words in a tweet that are identified as neutral, positive, or negative. VADER was specifically trained to handle emojis, URLs, hashtags, and tagged users during sentiment analysis so we did not remove those from our dataset. Up-to-date details about VADER can be found in the GitHub repository (<https://github.com/cjhutto/vaderSentiment>).

Figure 2. The retweet distributions resulting from conformity, novelty, content, and follower influence using this ABM (100 iterations each), alongside the observed retweet distribution (in black). Biases were all modelled with a g of 0.25 and the following parameter values: $a = 1.4$ (conformity bias), $a = 0.6$ (novelty bias), $c = 1$ (content bias), and $d = 1$ (follower influence). The x-axis (the identity of each tweet) and the y-axis (the number of times each tweet was retweeted) have been log-transformed.

To determine the potential targets of content bias we conducted GLMM using the *lme4* package in R (89). Retweet frequency was used as the outcome variable. To determine which grouping variables would be suitable as random effects we ran separate null models with each and calculated the intraclass correlation coefficient (ICC), or the proportion of the variance in retweet frequency explained by the grouping levels of each variable. Once random effects were chosen we added predictor variables in three stages, using the Akaike information criterion (AIC) and likelihood-ratio test (LRT) to choose between competing models. First, we determined whether tweet length would be an appropriate control variable. Next, we determined whether a tweet includes media (i.e. image or video) would be an appropriate control variable. Then, we added follower count and verification status to see which measure of platform size best improves the model. The 1.9% of tweets from users with missing verification statuses and follower counts were assigned verification statuses of “false” and follower counts of 0. Finally, we added the compound score, the proportion of negative words, and the proportion of positive words to see which measure of content best improves the model. All continuous predictor variables were scaled and centered prior to analysis. Model choice and residual diagnostic tests were conducted using a random 10% of observations, but the best fitting model was run using the entire dataset. The Poisson family was used since our outcome variable was count data and did not appear to have over- or underdispersion issues (Figure S8).

To ensure that our decision to treat quote tweets like original tweets did not bias our results related to content, we did a second round of GLMM to determine whether quote tweets have different emotional valence compared to the original tweets that they are quoting. We refer to original tweets that are quoted as target tweets (i.e. targets), and the tweets that quote them as quote tweets (i.e. quotes). Here we only considered target and quote tweets from cluster #2. Whether a tweet was a target (0) or a quote (1) was used as the outcome variable, and the identity of each target tweet was used as a random effect. In other words, each target and all of its quotes were assigned the same random effect. Like above, we first added tweet length and the presence of media as control variables (for both targets and quotes). Then, we added the compound score, the absolute value of the compound score, the proportion of negative words, the proportion of positive words, and the proportion of neutral words as predictor variables to see which measure of content best improves the model. The absolute value of the compound score and the proportion of neutral words were included as indicators of a general reduction in the intensity of emotion independent of positive or negative valence. All continuous

predictor variables were scaled and centered prior to analysis, and model choice and residual diagnostic tests were conducted using all observations (Figure S9).

Acknowledgments

We would like to thank Anne Kandler for providing us with feedback on our agent-based modeling. This research was supported, in part, under National Science Foundation Grants CNS-0958379, CNS-0855217, ACI-1126113 and the City University of New York High Performance Computing Center at the College of Staten Island. This work has received funding from the “Frontiers in Cognition” EUR grant, ANR-17-EURE-0017 EUR.

Author Contributions Statement

M.Y. developed the agent-based model and conducted all statistical modeling. R.G. conducted the sentiment analysis. A.A. and O.M. provided feedback on the agent-based model and statistical modeling. M.Y., J.S., A.A., and O.M. contributed to writing the manuscript.

Data & Code Availability

The agent-based model, analysis code, and processed data used in this study can be found on GitHub: <https://github.com/masonyoungblood/TwitterABM>. The full anonymized *VoterFraud2020* dataset can be found on Abilov et al.’s website (<https://github.com/sTechLab/VoterFraud2020>), and the full disambiguated dataset with tweet text is available from Abilov et al. upon request.

Preregistration

Our complete methods, model, and predictions were preregistered in advance of data analysis (<https://osf.io/invyf/>), except for the post hoc comparison between tweets and quote tweets. Departures from our preregistration that came up during the course of peer review can be found in the supplementary information.

Competing Interests Statement

The authors declare no competing interests.

References

1. Brotherton R. *Suspicious Minds: Why We Believe Conspiracy Theories*. New York: Bloomsbury Sigma; 2015.
2. Pagán VE. Conspiracy theories in the Roman Empire. In: Butter M, Knight P, editors. *Routledge Handbook of Conspiracy Theories*. 1st ed. New York: Routledge; 2020.
3. Zwiernin C. Conspiracy theories in the middle ages and the early modern period. In: Butter M, Knight P, editors. *Routledge Handbook of Conspiracy Theories*. 1st ed. New York: Routledge; 2020. p. 542–54.
4. West HG, Sanders T. *Transparency and Conspiracy: Ethnographies of Suspicion in the New World Order*. Durham: Duke University Press; 2003.
5. Butter M, Knight P. General introduction. In: Butter M, Knight P, editors. *Routledge Handbook of Conspiracy Theories*. 1st ed. New York: Routledge; 2020. p. 1–8.
6. Olmsted K. Conspiracy theories in US history. In: Uscinski JE, editor. *Conspiracy Theories and the People Who Believe Them*. Oxford University Press; 2018. p. 285–97.
7. Goldberg RA. Conspiracy theories in America: a historical overview. In: *Conspiracy Theories in American History: An Encyclopedia*. Oxford University Press; 2003. p. 1–13.
8. Douglas KM, Uscinski JE, Sutton RM, Cichocka A, Nefes T, Ang CS, et al. Understanding conspiracy theories. *Polit Psychol* [Internet]. 2019 Feb 1;40(S1):3–35. Available from: <https://doi.org/10.1111/pops.12568>
9. Goertzel T. Belief in conspiracy theories. *Polit Psychol* [Internet]. 1994 Oct 21;15(4):731–42. Available from: <http://www.jstor.org/stable/3791630>
10. Keeley BL. Of conspiracy theories. *J Philos* [Internet]. 1999 Oct 21;96(3):109–26. Available from: <http://www.jstor.org/stable/2564659>
11. Dentith M. *The Philosophy of Conspiracy Theories*. Palgrave Macmillan; 2014.
12. Pigden C. Popper revisited, or what is wrong with conspiracy theories? *Philos Soc Sci* [Internet]. 1995 Mar 1;25(1):3–34. Available from: <https://doi.org/10.1177/004839319502500101>
13. Imhoff R, Dieterle L, Lamberty P. Resolving the puzzle of conspiracy worldview and political activism: belief in secret plots decreases normative but increases nonnormative political engagement. *Soc Psychol Personal Sci* [Internet]. 2020 Mar 16;12(1):71–9. Available from: <https://doi.org/10.1177/1948550619896491>
14. Jolley D, Douglas KM. The social consequences of conspiracism: exposure to conspiracy theories decreases intentions to engage in politics and to reduce one’s carbon footprint. *Br J Psychol* [Internet]. 2014 Feb 1;105(1):35–56. Available from: <https://doi.org/10.1111/bjop.12018>
15. Uscinski JE, Parent JM. *American Conspiracy Theories*. New York: Oxford University Press; 2014.
16. Jolley D, Meleady R, Douglas KM. Exposure to intergroup conspiracy theories promotes prejudice which spreads across groups. *Br J Psychol* [Internet]. 2020 Feb 1;111(1):17–35. Available from: <https://doi.org/10.1111/bjop.12385>
17. Kofta M, Soral W, Bilewicz M. What breeds conspiracy antisemitism? The role of political uncontrollability and uncertainty in the belief in Jewish conspiracy. *J Pers Soc Psychol*. 2020;118(5):900–18.
18. Enders AM, Uscinski JE, Klofstad CA, Premaratne K, Seelig MI, Wuchty S, et al. The 2020 presidential election and beliefs about fraud: continuity or change? *Elect Stud* [Internet]. 2021;72:102366. Available from: <https://www.sciencedirect.com/science/article/pii/S0261379421000858>
19. Cottrell D, Herron MC, Westwood SJ. An exploration of Donald Trump’s allegations of massive voter fraud in the 2016 General Election. *Elect Stud* [Internet]. 2018;51:123–42. Available from: <https://www.sciencedirect.com/science/article/pii/S026137941730166X>
20. Edelson J, Alduncin A, Krewson C, Sieja JA, Uscinski JE. The effect of conspiratorial thinking and motivated reasoning on belief in election fraud. *Polit Res Q* [Internet]. 2017 Jul 31;70(4):933–46. Available from: <https://doi.org/10.1177/1065912917721061>
21. Sardarizadeh S, Lussenhop J. The 65 days that led to chaos at the Capitol. *BBC News* [Internet]. 2021 Jan; Available from: <https://www.bbc.com/news/world-us-canada-55592332>
22. Conger K, Isaac M. Twitter permanently bans Trump, capping online revolt. *New York Times*. 2021 Jan;

23. Spring M. "Stop the steal": The deep roots of Trump's "voter fraud" strategy. BBC New. 2020 Nov;
24. Cohen L. 6 conspiracy theories about the 2020 election – debunked. CBS News. 2021 Jan;
25. Cooper JJ, Christie B. Election conspiracies live on with audit by Arizona GOP. Associated Press [Internet]. 2021 Apr; Available from: <https://bit.ly/3k4soQD>
26. Corasaniti N, Epstein RJ. What Georgia's voting law really does. New York Times. 2021 Apr;
27. Balsamo M. Disputing Trump, Barr says no widespread election fraud. Associated Press [Internet]. 2020 Dec; Available from: <https://apnews.com/article/barr-no-widespread-election-fraud-b1f1488796c9a98c4b1a9061a6c7f49d>
28. Tucker E, Bajak F. Repudiating Trump, officials say election "most secure." Associated Press. 2020 Nov;
29. Cillizza C. 1 in 3 Americans believe the "Big Lie." CNN. 2021 Jun;
30. Skelley G. Most Republicans still won't accept that Biden won. FiveThirtyEight [Internet]. 2021 May; Available from: <https://fivethirtyeight.com/features/most-republicans-still-wont-accept-that-biden-won/>
31. Albertson B, Guiler K. Conspiracy theories, election rigging, and support for democratic norms. Res Polit [Internet]. 2020 Jul 1;7(3):2053168020959859. Available from: <https://doi.org/10.1177/2053168020959859>
32. Beckett L. Millions of Americans think the election was stolen. How worried should we be about more violence? The Guardian. 2021 Apr;
33. Hall Jamieson K, Albarracín D. The relation between media consumption and misinformation at the outset of the SARS-CoV-2 pandemic in the US. Harvard Kennedy Sch Misinformation Rev. 2020;1.
34. Stempel C, Hargrove T, Stempel GH. Media use, social structure, and belief in 9/11 conspiracy theories. Journal Mass Commun Q [Internet]. 2007 Jun 1;84(2):353–72. Available from: <https://doi.org/10.1177/107769900708400210>
35. Min SJ. Who believes in conspiracy theories? Network diversity, political discussion, and conservative conspiracy theories on social media. Am Polit Res [Internet]. 2021 May 2;49(5):415–27. Available from: <https://doi.org/10.1177/1532673X211013526>
36. Schöne JP, Parkinson B, Goldenberg A. Negativity spreads more than positivity on Twitter after both positive and negative political situations. Affect Sci [Internet]. 2021; Available from: <https://doi.org/10.1007/s42761-021-00057-7>
37. Brady WJ, Wills JA, Jost JT, Tucker JA, Van Bavel JJ. Emotion shapes the diffusion of moralized content in social networks. Proc Natl Acad Sci [Internet]. 2017 Jul 11;114(28):7313 LP – 7318. Available from: <http://www.pnas.org/content/114/28/7313.abstract>
38. Osmundsen M, Bor A, Vahstrup PB, Bechmann A, Petersen MB. Partisan polarization is the primary psychological motivation behind political fake news sharing on Twitter. Am Polit Sci Rev [Internet]. 2021/05/06. 2021;115(3):999–1015. Available from: <https://www.cambridge.org/core/article/partisan-polarization-is-the-primary-psychological-motivation-behind-political-fake-news-sharing-on-twitter/3F7D2098CD87AE5501F7AD4A7FA83602>
39. Rathje S, Van Bavel JJ, van der Linden S. Out-group animosity drives engagement on social media. Proc Natl Acad Sci [Internet]. 2021 Jun 29;118(26):e2024292118. Available from: <http://www.pnas.org/content/118/26/e2024292118.abstract>
40. Mesoudi A. Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences [Internet]. University of Chicago Press; 2011. Available from: <https://doi.org/10.7208/9780226520452>
41. Kendal RL, Boogert NJ, Rendell L, Laland KN, Webster M, Jones PL. Social learning strategies: bridge-building between fields. Trends Cogn Sci [Internet]. 2018;22(7):651–65. Available from: <https://doi.org/10.1016/j.tics.2018.04.003>
42. Lachlan RF, Ratmann O, Nowicki S. Cultural conformity generates extremely stable traditions in bird song. Nat Commun [Internet]. 2018;9. Available from: <http://dx.doi.org/10.1038/s41467-018-04728-1>
43. Kandler A, Powell A. Generative inference for cultural evolution. Philos Trans R Soc B Biol Sci. 2018;373(1743).
44. Carrignon S, Bentley RA, Ruck D. Modelling rapid online cultural transmission: evaluating neutral models on Twitter data with approximate Bayesian computation. Palgrave Commun. 2019;5(83).
45. Youngblood M, Lahti D. Content bias in the cultural evolution of house finch song. Anim Behav. 2022;185:37–48.
46. Raynal L, Marin J-M, Pudlo P, Ribatet M, Robert CP, Estoup A. ABC random forests for Bayesian parameter inference. Bioinformatics [Internet]. 2019;35(10):1720–8. Available from: <http://arxiv.org/abs/1605.05537v0><http://dx.doi.org/10.24072/pci.evolbio.100036>
47. Abilov A, Hua Y, Matatov H, Amir O, Naaman M. VoterFraud2020: a multi-modal dataset of election fraud claims on Twitter. arXiv [Internet]. 2021; Available from: <http://arxiv.org/abs/2101.08210>
48. Romm T, Dwoskin E. Twitter purged more than 70,000 affiliated with QAnon following Capitol riot. Washington Post [Internet]. 2021 Jan; Available from: <https://www.washingtonpost.com/technology/2021/01/11/trump-twitter-ban/>
49. Hutto CJ, Gilbert E. VADER: a parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media. 2014. p. 216–25.
50. Baumeister RF, Bratslavsky E, Finkenauer C, Vohs KD. Bad is stronger than good. Rev Gen Psychol [Internet]. 2001 Dec 1;5(4):323–70. Available from: <https://doi.org/10.1037/1089-2680.5.4.323>
51. Rozin P, Royzman EB. Negativity bias, negativity dominance, and contagion. Personal Soc Psychol Rev [Internet]. 2001 Nov 1;5(4):296–320. Available from: https://doi.org/10.1207/S15327957PSPR0504_2
52. Asatani K, Yamano H, Sakaki T, Sakata I. Dense and influential core promotion of daily viral information spread in political echo chambers. Sci Rep [Internet]. 2021;11(1):7491. Available from: <https://doi.org/10.1038/s41598-021-86750-w>
53. Del Vicario M, Vivaldo G, Bessi A, Zollo F, Scala A, Caldarelli G, et al. Echo chambers: emotional contagion and group polarization on Facebook. Sci Rep [Internet]. 2016;6(1):37825. Available from: <https://doi.org/10.1038/srep37825>
54. Bellovary AK, Young NA, Goldenberg A. Left- and right-leaning news organizations use negative emotional content and elicit user engagement similarly. Affect Sci [Internet]. 2021; Available from: <https://doi.org/10.1007/s42761-021-00046-w>
55. Pröllochs N, Bär D, Feuerriegel S. Emotions explain differences in the diffusion of true vs. false social media rumors. Sci Rep. 2021;11:22721.
56. Acerbi A. Cognitive attraction and online misinformation. Palgrave Commun [Internet]. 2019;5(1):15. Available from: <https://doi.org/10.1057/s41599-019-0224-y>
57. Stieglitz S, Dang-Xuan L. Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. J Manag Inf Syst [Internet]. 2013 Apr 1;29(4):217–48. Available from: <https://doi.org/10.2753/MIS0742-1222290408>
58. Burton JW, Cruz N, Hahn U. Reconsidering evidence of moral contagion in online social networks. Nat Hum Behav. 2021;5(12):1629–35.
59. van Prooijen J-W, Ligthart J, Rosema S, Xu Y. The entertainment value of conspiracy theories. Br J Psychol [Internet]. 2021 Jul 14;n/a(n/a). Available from: <https://doi.org/10.1111/bjop.12522>
60. Koumchatzky N, Andryeyev A. Using deep learning at scale in Twitter's timelines. Twitter's Engineering Blog. 2017.
61. Halsey LG. The reign of the p-value is over: What alternative analyses could we employ to fill the power vacuum? Biol Lett. 2019;15(5).
62. Nakagawa S, Johnson PCD, Schielzeth H. The coefficient of determination R² and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. J R Soc Interface [Internet]. 2017 Sep 30;14(134):20170213. Available from: <https://doi.org/10.1098/rsif.2017.0213>

63. Hindman M, Barash V. Disinformation, “fake news” and influence campaigns on Twitter [Internet]. 2018. Available from: <https://knightfoundation.org/reports/disinformation-fake-news-and-influence-campaigns-on-twitter/%0Ahttps://knightfoundation.org/reports/disinformation-fake-news-and-influence-campaigns-on-twitter>
64. Ferrara E, Yang Z. Quantifying the effect of sentiment on information diffusion in social media. *PeerJ Comput Sci.* 2015;1(e26).
65. Ferrara E, Yang Z. Measuring emotional contagion in social media. *PLoS One.* 2015;10(11):1–14.
66. Vosoughi S, Roy D, Aral S. The spread of true and false news online. *Science* (80-). 2018;359:1146–51.
67. Stroup WW. *Generalized Linear Mixed Models: Modern Concepts, Methods and Applications*. Vol. 81, International Statistical Review. CRC Press; 2013.
68. Sunstein CR. *Conspiracy theories and other dangerous ideas*. New York: Simon & Schuster; 2014.
69. Sunstein CR. *On rumors: how falsehoods spread, why we believe them, and what can be done*. Princeton: Princeton University Press; 2014.
70. Radnitz S, Underwood P. Is belief in conspiracy theories pathological? A survey experiment on the cognitive roots of extreme suspicion. *Br J Polit Sci.* 2017;47(1):113–29.
71. Guess A, Nagler J, Tucker J. Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Sci Adv.* 2019;5:eau4586.
72. Scrivner C, Stubbersfield JM. Curious about threats: Morbid curiosity and interest in conspiracy theories. *OSF.* 2022;
73. Youngblood M. Extremist ideology as a complex contagion: the spread of far-right radicalization in the United States between 2005–2017. *Humanit Soc Sci Commun.* 2020;7(49).
74. Becker MH. When Extremists Become Violent: Examining the Association Between Social Control, Social Learning, and Engagement in Violent Extremism. *Stud Confl Terror* [Internet]. 2019;1–21. Available from: <https://doi.org/10.1080/1057610X.2019.1626093>
75. Bangarter A, Wagner-Egger P, Delouée S. How conspiracy theories spread. In: Butter M, Knight P, editors. *Routledge Handbook of Conspiracy Theories*. New York: Routledge; 2020. p. 206–18.
76. Salali GD, Uysal MS. Effective incentives for increasing COVID-19 vaccine uptake. *Psychol Med.* 2021/09/20. 2021;1–3.
77. Barrett BJ. Equifinality in empirical studies of cultural transmission. *Behav Processes.* 2019;161:129–38.
78. Matei SA, Russell MG, Bertino E. Transparency in social media: tools, methods and algorithms for mediating online interactions. Matei SA, Russell MG, Bertino E, editors. Springer; 2015.
79. Shah H. Algorithmic accountability. *Philos Trans R Soc A.* 2018;376(2128):20170362.
80. Agrawal P, Davis D. Transparency around image cropping and changes to come [Internet]. Twitter’s Product Blog. 2020. Available from: https://blog.twitter.com/en_us/topics/product/2020/transparency-image-cropping
81. Huszár F, Ktena SI, O’Brien C, Belli L, Schlaikjer A, Hardt M. Algorithmic amplification of politics on Twitter. Twitter [Internet]. 2021; Available from: <https://bit.ly/3niLy6V>
82. Markey EJ. Algorithmic Justice and Online Platform Transparency Act [Internet]. United States Senate; 2021. Available from: <https://www.congress.gov/bill/117th-congress/senate-bill/1896/text>
83. Koene A, Clifton C, Hatada Y, Webb H, Patel M, Machado C, et al. *A governance framework for algorithmic accountability and transparency*. 2019.
84. Chao A, Gotelli NJ, Hsieh TC, Sander EL, Ma KH, Colwell RK, et al. Rarefaction and extrapolation with Hill numbers: a framework for sampling and estimation in species diversity studies. *Ecol Monogr* [Internet]. 2014 Feb 1;84(1):45–67. Available from: <https://doi.org/10.1890/13-0133.1>
85. Roswell M, Dushoff J, Winfree R. A conceptual guide to measuring species diversity. *Oikos* [Internet]. 2021 Mar 1;130(3):321–38. Available from: <https://doi.org/10.1111/oik.07202>
86. Blum MGB, François O. Non-linear regression models for Approximate Bayesian Computation. *Stat Comput.* 2010;20:63–73.
87. Sisson SA, Fan Y, Beaumont MA. *Handbook of Approximate Bayesian Computation*. CRC Press; 2018.
88. Probst P, Wright M, Boulesteix A-L. Hyperparameters and tuning strategies for random forest. *Wiley Interdiscip Rev Data Min Knowl Discov.* 2018;
89. Bates D, Mächler M, Bolker B, Walker S. Fitting linear mixed-effects models using lme4. *J Stat Softw.* 2015;67(1):1–48.

Supplementary Information

	Sampled summary statistics	Min node size	MSE
α : content bias	4	65	1.809
δ : follower influence	4	109	1.076
α : frequency bias	4	74	0.843
g : age dependency	3	89	1.523

Table S1. The optimal number of sampled summary statistics and minimum node size for the random forests for each of the four dynamic parameters in the ABM.

Figure S1. The summary statistics from 10,000 simulations from the prior distributions (green) and posterior distributions (red), plotted against the observed data (black). The left panel shows the most important variables for generative inference, whereas the right panel shows the two least important variables. Variable importance was assessed using the Gini-based impurity method. We appear to be capturing more than enough of the parameter space to make inferences about our observed data, marked by the black star. Importantly, the posterior distributions appear to have converged towards parameter values that do a much better job of recreating the observed data.

Figure S2. The retweet distributions resulting from 1,000 simulations from the prior distributions (green) and posterior distributions (red), plotted against the observed data (black). Note that this shows the top 100,000 most retweeted tweets, which represent only 2.51% of the entire distribution. The posteriors do not recreate the distribution of the most common retweets, but they do exclude the more right-skewed outcomes that are consistent with conformity and demonstrator bias (Figure 2). Based on the fact that the two summary statistics most important for inference are those that are most sensitive to the long tail of the distribution (proportion of tweets that only appear once and the Hill number when $q = 1$, see Figure S3), it is unsurprising that the posteriors from these inferences are less sensitive to the distribution of the most common tweets.

Figure S3. The prior (dotted lines) and posterior (solid lines) distributions for each of the four dynamic parameters from the ABM, estimated using the basic rejection form of ABC across five tolerance levels. The results are broadly consistent with our main conclusions: a strong signal for content bias, convergence towards zero for follower influence and frequency bias, and uninterpretable results for age dependency.

Figure S4. The prior (dotted lines) and posterior (solid lines) distributions from four additional rounds of the ABM. Each round was run for 1,000 iterations with only a single term from the probability function included. For example, panel A shows the posterior distribution for c when the ABM was run with uniform prior for c (from 0 to 8) with fixed neutral values for d (0), a (1), and g (0). In other words, panel A shows the estimated parameter values for c that best recreate the real data when all other parameters are ignored. Random forest ABC was conducted exactly as described in the main text.

Figure S5. The retweet distributions from four additional rounds of the ABM (500 iterations each), alongside the observed retweet distribution (in solid black). The dotted black lines show the distribution resulting from the lowest value of that parameter sampled from the prior, whereas the dashed black lines show the distribution resulting from the highest value of that parameter sampled from the prior. Most importantly, frequency bias and age dependency (arguably the two processes most likely to confound one another) lead to different patterns in the retweet distributions.

Figure S6. The distribution of VADER's compound scores assigned to tweets in our dataset alongside the mean value (dashed line). The distribution skews negative, which makes sense given the subject matter, the full range of possible values between -1 and 1 are represented. VADER assigned 31.5% of tweets a compound score of 0.

Model specification	df	AIC
retweets ~ (1 user)	2	1121395
retweets ~ (1 day)	2	3907893
retweets ~ (1 hour)	2	3919423
retweets ~ scale(length) + (1 user)	3	1080592
retweets ~ has_media + scale(length) + (1 user)	4	1080082
retweets ~ scale(followers) + has_media + scale(length) + (1 user)	5	1073736
retweets ~ scale(verified) + has_media + scale(length) + (1 user)	5	1079940
retweets ~ scale(compound) + scale(followers) + has_media + scale(length) + (1 user)	6	1071799
retweets ~ scale(negative) + scale(followers) + has_media + scale(length) + (1 user)	6	1072630
retweets ~ scale(positive) + scale(followers) + has_media + scale(length) + (1 user)	6	1073075

Table S2. The model specification, degrees of freedom, and AIC for each candidate model for the primary GLMM (conducted with a random 10% of observations).

Figure S7. A partial specification curve analysis for the primary GLMM (conducted with a random 10% of observations). Panel A (on the left) shows the IRRs for the effect of the proportion of negative words on retweet probability under different combinations of predictors that appeared in the best fitting model. Panel B (on the right) shows the same for the effect of followers. Error bars represent Wald confidence intervals which were used due to high sample size. The IRR for followers from the simplest model in B is absent because the model did not converge. In all specifications, the proportion of negative words and the number of followers have significant effects in the same direction as in the best fitting model.

Figure S8. A Q-Q plot (left) and a standardized residual plot (right) for the main GLMM, constructed using the *DHARMA* package in R (Hartig, 2020). The model appears to be a good fit, and a dispersion test indicates that there is no significant dispersion in the data ($p = 0.73$). That being said, there is a low level of significant outliers among residuals (outliers = 0.98%, $p < 0.0001$), and a Kolmogorov-Smirnov test indicates that the residuals do significantly deviate from uniformity ($p < 0.0001$). Based on a visual inspection of the Q-Q plot, and the fact that the creator of the *DHARMA* package has suggested that slight departures from uniformity can be significant when sample sizes are extremely high¹, we conclude that the model is a good fit to the data. There is also a significant level of zero-inflation in the data ($p < 0.0001$), but the level of it is so low (2.2%) that we would prefer to use the Poisson family than to separately model zero and non-zero values.

¹ <https://github.com/florianhartig/DHARMA/issues/181>

Quote Tweet Analysis

After further subsetting the tweets for the quote tweet analysis, we ended up with 91,227 original tweets and 383,778 quote tweets from 102,227 users. The identity of each target tweet was included as a random effect (ICC = 0.146), and adding tweet length as a control variable significantly improved model fit ($\Delta\text{AIC} = 140971$; LRT: $\chi^2 = 140972$, $p < 0.0001$). Adding the presence of media as a control variable also significantly improved model fit ($\Delta\text{AIC} = 6684$; LRT: $\chi^2 = 6686$, $p < 0.0001$). All five content measures further improved model fit ($\Delta\text{AIC} > 2$), but the compound score improved fit significantly better than all of the content measures ($\Delta\text{AIC} = 15314$; LRT: $\chi^2 = 15317$, $p < 0.0001$). All model specifications and AIC values for the quote GLMM are in Table S4. The best fitting model for the quote tweet analysis included the identity of each quoting event as a random intercept, and tweet length, presence of media, and the compound score as fixed effects.

	OR	95% CI
<i>Tweet length</i>	0.163	[0.160, 0.165]
<i>Presence of media</i>	0.111	[0.106, 0.117]
<i>Compound score</i>	1.889	[1.869, 1.909]

Table S3. The odds ratio (OR) and 95% confidence interval for each predictor in the best fitting model. OR, the exponentiated beta estimate, is interpreted as the rate at which the outcome variable is expected to change per unit increase in a predictor (one standard deviation for scaled and centered predictors). Wald confidence intervals were used due to the high sample size.

Tweet length and the compound score both significantly predict whether a tweet is original or a quote. The odds ratio (OR) for tweet length is 0.163, indicating that if a tweet is one SD longer then it is 83.7% more likely to be an original tweet. Similarly, the odds ratio (OR) for presence of media is 0.111, indicating that if a tweet includes an image or video then it is 88.9% more likely to be an original tweet. In other words, quote tweets tend to be shorter than original tweets and not include media. The compound score, on the other hand, has a strong positive effect on the probability that a tweet is a quote. The OR of 1.889 indicates that tweets with a compound score that is one SD higher are 88.9% more likely to be a quote. In other words, quote tweets tend to be more positive than their targets that they are quoting. The next best fitting model, which included the proportion of words that are negative, outcompeted the models with the absolute value of the compound score and the proportion of words that are neutral ($\Delta\text{AIC} > 2$). This suggests that the shift in the tone of quote tweets is driven by a reduction in negative emotion rather than a general reduction of all emotion independently of valence. Pseudo R^2 values, calculated using the theoretical variances, indicate that the predictor variables alone account for about 50% of the variance in the data ($R^2 = 0.497$), whereas the predictor variables and random effects together account for about 61% of the variance in the data ($R^2 = 0.626$) (Nakagawa et al., 2017). A variance inflation factor test indicates that there are no significant issues with multicollinearity between predictors (VIFs < 2). Residual diagnostics for the best fitting model can be seen in Figure S9.

Model specification	df	AIC
retweets ~ (1 event)	2	420624
retweets ~ scale(length) + (1 event)	3	279654
retweets ~ has_media + scale(length) + (1 event)	4	272969
retweets ~ scale(compound) + has_media + scale(length) + (1 event)	5	257655
retweets ~ scale(abs(compound)) + has_media + scale(length) + (1 event)	5	265669
retweets ~ scale(negative) + has_media + scale(length) + (1 event)	5	261714
retweets ~ scale(positive) + has_media + scale(length) + (1 event)	5	270873
retweets ~ scale(neutral) + has_media + scale(length) + (1 event)	5	270060

Table S4. The model specification, degrees of freedom, and AIC for each candidate model for the quote GLMM (conducted with all observations).

Figure S9. A Q-Q plot (left) and a standardized residual plot (right) for the quote analysis GLMM, constructed using the *DHARMA* package in R (Hartig, 2020). The model appears to be a good fit, but there is a low level of significant outliers (outliers = 0.76%, $p = 0.0054$), and a Kolmogorov-Smirnov test indicates that the residuals do significantly deviate from uniformity ($p < 0.0001$). Based on a visual inspection of the Q-Q plot, and the fact that the creator of the *DHARMA* package has suggested that slight departures from uniformity can be significant when sample sizes are extremely high², we conclude that the model is a good fit to the data. We did not conduct a dispersion test for the quote analysis GLMM, as they are unreliable for logistic models with binary outcomes³.

² <https://github.com/florianhartig/DHARMA/issues/181>

³ <https://github.com/florianhartig/DHARMA/issues/79>

Departures From Preregistration

Several modifications to our framing and analysis came up during the course of peer review, so our final study departs from our original preregistration⁴ in the following ways:

1. The way that demonstrator bias was originally treated in our ABM was confounded with the basic effect of network structure, where people with more followers are more likely to be retweeted because their tweets appear on more timelines. The parameter d is now broadly interpreted as the effect of the influence of follower count rather than demonstrator bias. We have also removed the hypotheses and interpretation related to demonstrator bias.
2. We originally assigned follower counts (T) and activity levels (r) separately and imposed a uniform probability of tweeting an original tweet (μ) on agents. By doing so, we removed the correlations among T , r , and μ that capture variation in user behavior (e.g. users with few followers who exclusively retweet). Each agent is now assigned a unique value of T , r , and μ that come from a real user in the observed data.
3. Follower counts greater than or equal to 100,000 are no longer excluded from the ABM.
4. We increased the iterations of the ABM from 200,000 to improve statistical power during generative inference.
5. The presence of media (e.g. images or videos) was added as a control variable to the general linear mixed modeling.

References

- Hartig, F. (2020). *DHARMA: residual diagnostics for hierarchical regression models* (R package version 0.3.3.0). <https://cran.r-project.org/package=DHARMA>
- Nakagawa, S., Johnson, P. C. D., & Schielzeth, H. (2017). The coefficient of determination R^2 and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of The Royal Society Interface*, 14(134), 20170213. <https://doi.org/10.1098/rsif.2017.0213>

⁴ <https://osf.io/jnyvf/>