



UvA-DARE (Digital Academic Repository)

How Do Individuals in a Radical Echo Chamber React to Opposing Views? Evidence from a Content Analysis of Stormfront

Bright, J.; Marchal, N.; Ganesh, B.; Rudinac, S.

DOI

[10.1093/hcr/hqab020](https://doi.org/10.1093/hcr/hqab020)

Publication date

2022

Document Version

Final published version

Published in

Human Communication Research

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/in-the-netherlands/you-share-we-take-care>)

[Link to publication](#)

Citation for published version (APA):

Bright, J., Marchal, N., Ganesh, B., & Rudinac, S. (2022). How Do Individuals in a Radical Echo Chamber React to Opposing Views? Evidence from a Content Analysis of Stormfront. *Human Communication Research*, 48(1), 116-146. <https://doi.org/10.1093/hcr/hqab020>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

ORIGINAL RESEARCH

How Do Individuals in a Radical Echo Chamber React to Opposing Views? Evidence from a Content Analysis of Stormfront

Jonathan Bright ^{1,2,*}, Nahema Marchal^{3,4}, Bharath Ganesh^{5,6} and Stevan Rudinac^{7,8}

1 Oxford Internet Institute, University of Oxford, UK

2 PhD Social and Political Sciences, European University Institute (2012), Italy

3 Institut für Politikwissenschaft, University of Zurich, Switzerland

4 PhD Social Science of the Internet, Oxford Internet Institute, University of Oxford (2021), UK

5 Centre for Media and Journalism Studies, University of Groningen, Netherlands

6 PhD Geography, University College London (2017), UK

7 Amsterdam Business School, University of Amsterdam, The Netherlands

8 PhD Computer Science, Delft University of Technology (2013), The Netherlands

Calls to “break up” radical echo chambers by injecting them with alternative viewpoints are common. Yet, thus far there is little evidence about the impact of such counter-messaging. To what extent and how do individuals who inhabit a radical echo chamber engage with messages that challenge their core beliefs? Drawing on data from the radical right forum Stormfront we address this question with a large-scale content and longitudinal analysis of users’ posting behavior, which analyses more than 35,000 English language contributions to the forum spanning 2011 through 2013. Our findings show that engaging with oppositional views is actually a core practice among Stormfront users which invites active participation and encourages engagement. Indeed, many “echoes” in the echo chamber we studied were not core beliefs being restated, but the sound of opposing viewpoints being undermined and marginalized. These findings underscore the limited potential for counter-messages to undermine radical echo chambers.

Keywords: Echo Chambers, Far Right, Cognitive Dissonance

<https://doi.org/10.1093/hcr/hqab020>

Over the last decade, the public conversation around the implications of digital technology for democracy has been dominated by fears about ‘echo chambers’ (Guess et al., 2018). The concept of echo chamber is not specific to digital media (Jamieson & Capella, 2010) but is commonly used today to describe online

*Corresponding author: Jonathan Bright; e-mail: jonathan.bright@oii.ox.ac.uk

communities where like-minded individuals “preferentially connect with each other” and echo each other’s opinions and beliefs (Bruns, 2019, p. 48). The relative isolation from outside perspectives that is characteristic of echo chambers has led scholars to qualify them as both *ideologically homogenous* and relatively *radical*, in the sense that opinions expressed in these spaces reflect views which might diverge from broader societal norms (Sunstein, 2001, Van Alstyne & Brynjolfsson, 2005).

While empirical research shows that exposure to diverse political news and opinion is common online (Bakshy *et al.*, 2015; Bright, 2018; Dubois & Blank, 2018; Fletcher & Nielsen, 2018; Röchert *et al.*, 2020; Scharkow *et al.*, 2020) numerous studies nevertheless document online spaces that exhibit echo chamber-like characteristics (Agathangelou *et al.*, 2017). This includes groups where highly misogynistic (Marwick & Caplan, 2018) and racist (Cleland, 2014) ideas are reified, for example, or where grave doubts are cast on ideas such as vaccinations (Coleman, 2017) and climate change (Bloomfield & Tillery, 2019; Oswald & Bright, 2021) which otherwise attract large-scale scientific consensus. Recently, the formation of these echo chambers, their social dynamics, and their impact on the beliefs and actions of those who inhabit them have generated a growing stream of research (Baele *et al.*, 2019; Bliuc *et al.*, 2018; Oswald & Bright, 2021; Regehr, 2020; Sasahara *et al.*, 2020; van Eck *et al.*, 2021).

Our study focuses on the online white nationalist forum Stormfront. Often dubbed as the first “hate Web site” (Meddaugh & Kay, 2009), Stormfront is widely seen as one of most influential far-right communities on the Internet of the last 20 years (Hatewatch, 2017), with thousands of members across the world. Stormfront members describe the forum as an “online refuge” (De Koster & Houtman, 2008); a place where they connect and actively engage with others who share their racist and white supremacist beliefs (Bowman-Grieve, 2009, Hartzell, 2020). It is, in other words, a forum characterized both by radical views and relative opinion homogeneity.

Ideologically, homogenous online communities like Stormfront have been decried as catalysts for extremist radicalization (Wojcieszak, 2010). Studies show that the insulation from outside perspectives and endless opinion reinforcement characteristic of echo chambers (Brugnoli, Cinelli *et al.*, 2019) distort individuals’ perception of reality through confirmation bias and facilitate the spread of misinformation (Choi *et al.*, 2020; Törnberg, 2018). This, in turn, has been shown to breed attitude extremity (Iyengar & Hahn, 2009; Stroud, 2010), intolerance to other people’s views (Garrett *et al.*, 2016), and even violence (Ganesh & Bright, 2020; Melton & Motyl, 2018; Wojcieszak, 2010). Stormfront itself is a bastion and breeding ground for violent extremism: Over the years, members of the site have been linked to a number of terrorists attacks, including the 2011 Norway attacks and the 2015 Charleston church shootings (Beirich, 2015). Recent longitudinal research on Stormfront has also found that users of the forum steadily become more radical in their discussion of anti-LGBTQ, anti-Black, and anti-Semitic narratives over time (Scrivens *et al.*, 2020).

While the impact of radical echo chambers is well-researched, how to best counter the threat they pose to society is a source of ongoing debate (Braddock & Horgan, 2016; Ganesh & Bright, 2019). Recently, significant efforts have been deployed by government, civil society, and activists to break up radical echo chambers by injecting them with opposing viewpoints and narratives that challenge their members' beliefs or offer alternative perspectives (Ganesh & Bright, 2019; Lee, 2020; Rosand & Winterbotham, 2019). Despite being a widespread strategic communication technique, research assessing the efficacy of counter-attitudinal messaging is still sparse and inconclusive (Ferguson, 2016), with some studies suggesting that such tactics may even be counter-productive (Castro & Hemmingsen, 2017; Schmitt et al., 2018). In particular, at this point it is still unclear *to what extent and how* members of these communities engage with views that challenge their core beliefs (Neubaum & Krämer, 2017). Our core interest in this article is thus to answer the following question: *To what extent and how do individuals within a radical echo chamber engage with opposing viewpoints?* Shedding light on this question will inform current debates about the usefulness of counter-messaging techniques as a way to break up radical echo chambers.

The rest of the article is structured in the following way. In the next section, we review existing literature on responses to opposing viewpoints, building on theories of cognitive dissonance and social identity. Then, we describe the nature of our data and measures. We selected Stormfront as a case study to address this question as it possesses all the characteristics of an echo chamber—namely, homophily associations and ideological homogeneity—yet also frequently hosts users who want to argue with or challenge white nationalist ideas. Finally, in the last section, we present our results. We show that, far from being an anathema to them, Stormfront members actively engage with counter-messages, working to weaken and reframe these views such that they bolster the overall ideology of white nationalism espoused by the forum.

Our article contributes to the literature in several ways. In theoretical terms, we address the question of whether dissonant material generates a response at all—an underexplored question; and, second, by qualifying the type of response elicited. In empirical terms, meanwhile, we test theories of dissonance and social identity in an underexplored empirical setting (namely, Stormfront). These results are then discussed in terms of their consequences for the study of echo-chambers, and efforts to inject more cross-cutting engagement into online life. Specifically, our research suggests that counter-messaging may not be an adequate strategy to deconstruct echo chambers and may even exacerbate the problem it aims to solve.

Engagement with Opposing Viewpoints

We first begin by surveying existing literature on whether and how individuals engage with counter-attitudinal information and opposing viewpoints. Then, we look

at the effect of engaging with these messages on user behavior over time, with special attention to the type of content they seek out afterwards.

We will begin by looking at the extent of engagement. Cognitive theories, such as selective exposure and cognitive dissonance theory (Festinger, 1957; Garrett, 2009; Stroud, 2014), posit that individuals are more inclined to seek out information they find agreeable or that matches their existing worldview. In the event of inconsistency between their own beliefs and other attitudes, however, people become “psychologically uncomfortable” (1957, p. 3) (see also Elliot and Devine, 1994; Harmon-Jones *et al.*, 1996) and work to reduce the dissonance they are experiencing (Gibbons *et al.*, 1997; Losch & Cacioppo, 1990) by avoiding situations that produce it (see Festinger 1957, pp. 129–131). This is one of the main reasons individuals may seek out “echo chambers” in the first place.

Research on online behavior supports this notion: although individuals rarely take active steps to completely avoid dissonant viewpoints (Garrett, 2009; Knobloch-Westerwick & Meng, 2009; Munson & Resnick, 2010)—except perhaps during heightened political times such as elections (John & Dvir-Gvirsman, 2015; Yang *et al.*, 2017)—when given the choice, internet users primarily seek out congenial content (for reviews see Knobloch-Westerwick, 2015; O’Keefe, 2002; Stroud, 2014). Cross-cutting content exposure is also known to generate ambivalence and reduce political engagement (Mutz, 2002). In the context of our study, these findings could suggest that users encountering opposing viewpoints might simply wish to selectively ignore them in order to avoid cognitive dissonance.

Many online discussion spaces, including Stormfront, have a post and comment structure (Gonzalez-Bailon *et al.*, 2010). That is to say, individuals using a forum create new posts containing original content, to which other users of the forum respond with comments, forming discussion “threads.” It may be therefore that counter messages expressing dissonant viewpoints receive less comments, and hence generate shorter discussion threads, than consonant ones.

However, work from social identity and informational utility theories also provides reasons to expect that dissonant viewpoints might generate more engagement (Matz & Wood, 2005). Individuals may seek to engage with counter-messages because they contain information that users are curious about or find useful, regardless of whether it is consonant or not (Knobloch-Westerwick, 2008). Social dynamics and conformity processes within groups are also important in this respect (Neubaum *et al.*, 2018). As Festinger puts it, social groups are “. . . a major vehicle for eliminating and reducing [cognitive dissonance]” (1957, p. 177). When members of the same group collectively encounter dissonant beliefs in a social environment, they seek to contradict and undermine them; a process referred to as “trivialization” by Simon *et al.* (1995) or “denial” by Abelson (1959, pp. 344–345). This trivialization can involve direct counter-argument, undermining the logical basis on which the dissonant cognition rests, such that the total line of reasoning is recast as invalid (Festinger 1957, p. 135). Likewise, social identity theory suggests that members of the same ideological in-group will strive to maintain their beliefs in the

face of opposition (Mackie et al., 2000). Confirming this tendency, studies have shown that exposure to politically disagreeable news and information via social media motivates greater political participation (Min & Wohn, 2018; Valenzuela et al., 2012), and uncivil responses (Gervais, 2015). Both of these lines of thinking expect that dissonant material in an echo chamber should be met with vigorous reaction (rather than silence). On the basis of these diverging expectations, we thus propose the following as a research question:

RQ1: Which type of posts, dissonant or consonant, generate the most engagement?

With respect to the volume of engagement, it is also worth highlighting that algorithmic filtering of content might be expected to play a role: many social forums make use of this information to decide how to present content to users, and this can create a rich-get-richer dynamic whereby content which is initially popular becomes ever more recommended to users. We will return to this point in our methods section.

In addition to the amount of responses generated, we might also expect different types of people to engage with dissonant material. Individuals who go online with the explicit goal of discussing politics only form a small minority of Internet users, and these tend to be more educated, politically interested and knowledgeable than the average population (for a review see Stromer-Galley, 2014). Hence, there are good reasons to believe they may react more strongly to counter-attitudinal information. For social identity theorists, long-term membership and active participation in an online community such as a political forum is testament to an individual's ideological dispositions and to the strength of their political identity (Carr, 2017). There is also solid evidence that politically extreme individuals with strong attachment to ideological beliefs systems are more likely to maintain homogenous social networks (Boutyline & Willer, 2017; Jost et al., 2018), and develop a heightened sense of in-group identity (Hogg et al., 2004). Recent work on Stormfront itself seems to confirm this, showing that users in the 99th percentile of number of posts created account for 38.87% of overall posts in the forum, and are also those that use the most extreme language and spend the most time there (Kleinberg et al., 2021).

In our particular case, we might therefore expect dissonant material to prompt a stronger response from members who are very active on the forum, compared with those who do not engage very much, as these are more likely to have formed strong attachments to the forum's political identity and the views it represents, and therefore experience a need to defend and maintain its positive distinctiveness (Brewer & Kramer, 1985). We thus elaborate our first hypothesis:

H1: Highly active users of Stormfront will be more likely to engage with oppositional material than those who are less active.

We will now move on to discussing the behavioral consequences of having engaged with dissonant material. Here, cognitive dissonance and social identity theories provide slightly different perspectives. One possible behavioral response to

exposure to dissonant material would be to cease participation altogether, as exposure with disagreeable contributions creates an uncomfortable amount of cognitive dissonance (see Festinger 1957, p. 130). In this perspective, over the long term, the frequent appearance of oppositional voices could lead forum members to gradually leave the conversation, perhaps migrating to other parts of the forum with more congenial views. Supporting evidence for this idea is present in a number of studies: Matz & Wood demonstrate that people change social groups when faced with intra-group dissonance (2005, pp. 34–35), for example, while Torcal and Maldonado (2014) and Mutz (2002, 2006) show that exposure to political disagreement and argument can lead people to disengage with politics. Research on social media also confirms this: In the US, for example a great majority of users report finding it “stressful” to talk politics online with people they disagree with (Duggan & Smith, 2016), and many actively ignore or move on from dissonant content when they encounter it (Bode, 2016; Sibona, 2014).

In contrast, social identity theories have argued that dissonance effects are contingent upon the influence of group membership and group norms. Robertson and Reicher (1997) and McKimmie et al. (2009) argue that group norms may reduce dissonance when they validate counter-attitudinal behavior (here, engaging with opposing views). In that case, exposure to dissonant material would not lead to avoidance, but to further engagement as this demonstrates norm adherence. Indeed, engaging with dissonant material (and forcefully arguing back) may become a group norm in and of itself, and a way of displaying and performing group identity. These diverging lines of work lead us to pose our second research question:

RQ2: Does engagement with dissonant views make further participation more or less likely?

A final possible behavioral response, suggested by cognitive dissonance theory, is that users may seek to change their internal information balance by seeking out confirmatory views after having encountered an opposing one (Cohen, 1960; Eagly and Chaiken, 1998; Ehrlich et al., 1957; Festinger, 1957, p. 22; Frey, 1986). Group processes may be especially important for adding consonant cognitions, because these supportive pieces of information are boosted by the social support provided by group members (Festinger, 1957, p. 179). Indeed, research has shown that membership of radical groups allows individuals to reduce the dissonance involved in holding radical views (McKimmie et al., 2003). This is because even though these views might be generally in the minority, in the local context of an echo chamber they form the majority (Cooper & Stone, 2000).

Research on Stormfront supports this, arguing that it provides an important space for interaction between members whose views are often met with “resistance and condemnation in social settings” (Wong et al., 2015, p. 44). For that reason, Stormfront is often seen as a “second home” (de Koster & Houtman, 2008) by committed members of the “white race” (Daniels, 2009), who want to speak what they see as “truths”. Perry and Scrivens also argue that “white supremacists” use of the

Internet is not passive; rather, participants actively and discursively construct collective identities” (2016, p. 76). Thus, we would expect that individuals coming into contact with dissonant content would seek to engage with consonant content that supports the construction of white supremacist collective identities shortly afterwards, perhaps by engaging with consonant material or seeking out consonant areas of the site. This leads us to our final hypothesis:

H2: Individuals will be more likely to engage with attitude consonant material after engaging with dissonant material.

Methods

Data

We draw our data from the website Stormfront, which as described above is a major online white nationalist forum with a long history of hosting extremist discussion (Caren et al., 2012; De Koster and Houtman, 2008; Hara and Estrada, 2003; Meddaugh and Kay, 2009). Stormfront is divided into a number of thematic sub-forums, and the discussion itself takes a post-comment format: people using the site can propose new discussion topics (“posts”), and then other people comment on these posts.

We chose this forum for study for two reasons. First, the site fits the concept of echo chamber which we have defined above. It is a place for users with radical (white nationalist) viewpoints, which are marginalized in mainstream society but which are in a clear majority on the forum (see Bowman-Grieve, 2009). Yet, the site also has a sub-forum dedicated to hosting “opposing views”,¹ where people can post messages which are dissonant with the idea of white nationalism. The existence of this part of the forum allows us to address our overarching question: to what extent and how do individuals within a radical echo chamber engage with opposing viewpoints? The data used for the article was harvested automatically from the forum using a web scraper,² and covers the period from 1st of January 2011 to 31st of December 2013.³ The scraper collected both the text of posts contained in the forum, and any comments that were made in response. Hence each observation in the raw dataset is a submission to the forum: either a post or a comment (a graphical description of the data collection and analysis pipeline is available in [Appendix A1.7](#)).

The data used in the study was harvested from seven sub-forums: The opposing views sub-forum as described above, and six other sub-forums which are oriented towards other topics. All of these forums host English language content only (there are other parts of Stormfront that host comments in other languages, but we did not address these in our study). The six other subforums were selected to provide a range of different types of discussion topic. Stormfront hosts a wide range of discussions: some focusing on the political goals of the white nationalist

movement, and others which are essentially apolitical (focusing, for example, on dating or leisure activities) but which allow members to discuss such topics only with other white nationalists. We wanted to select a range of subforums, some political and others less so, whilst also keeping a balance such that the subforums we did select were (in sum) approximately equal in size to the opposing views forum. In these other sub-forums, oppositional discussion is not permitted, and posts containing oppositional information are deleted. Full details on our choices here are set out in [Appendix A1.1](#).

Data was only collected from open access parts of the forum, where no account was required to access the material (Stormfront also contains a number of subforums that can only be accessed by registered users, a point which we will return to in the limitations section). Further notes on sub-forums and data are available in [Appendix A1.1](#). Post-hoc verification of the data collected by the web scraper suggested it was approximately 93.2% complete (see [Appendix A1.2](#) for details of the verification process).

It is worth noting that, to the best of our knowledge, the technology powering Stormfront (vBulletin) did not “personalize” the content which was displayed to users during our period of study. Rather, any user visiting the forum would see new posts displayed in order of recency. However, users have the option to reorder posts by the amount of comments received. This option to reorder posts could contribute to a rich-get-richer dynamics where posts which were initially popular receive ever more attention. Further details on this are available in [Appendix A1.3](#).

It is also worth noting that our choice of observational data, rather than an experimental paradigm, necessarily means the strength of causal claims we can make is limited (a point to which we will also return in the limitations section). Users of Stormfront are not selected randomly, nor are they randomly exposed to oppositional messages—rather, they have to make a conscious choice to engage with the forum. However, we also believe the observational paradigm has important strengths in this context. First, individuals with white nationalist beliefs are a minority in the overall population: engaging them in a laboratory experiment would be complex. Second, and perhaps more importantly, as we argued above dissonance experiences, and their responses, are also a group phenomenon: being confronted by opposition will (we believe) have different impacts if a person feels that they are in a space with other members of their ingroup, as compared to people who feel they are in an unfamiliar space without any social support, such as a lab setting.

Measures

In this section, we outline our measures (descriptive statistics for all variables are available in [Appendix A1.6](#)). Our first measurement task was to characterize the type of engagements users had during their time on the forum. We define three such engagement types. First, users could engage in discussions by commenting in one of the ideologically consonant sub-forums (i.e. any sub-forum apart from the

opposing views sub-forum). These sub-forums are exclusively for discussion amongst white nationalists, and dissonant material is actively deleted by forum moderators. Second, users could engage with “consonant arguments” inside the opposing views sub-forum. These are discussion threads in the opposing views sub-forum where the initial post was in support of white nationalism. Finally, users could engage in dissonant arguments in the opposing views sub-forum: discussion threads where the initial post was against white nationalism.

We distinguished between consonant and dissonant arguments within the opposing views sub-forum by manually coding each post within this forum — that is, each original post to the forum that started a new discussion thread — into one of three categories: consonant, dissonant, and miscellaneous. Consonant posts were ones that supported white nationalism (our precise interpretation of this ideology is described in [Appendix A1.4](#)). Dissonant posts were ones that contradicted this belief system, and everything else was in the miscellaneous category. Krippendorff's alpha was 0.814 for the coding exercise, with a percentage agreement of 88%.

The posts were lengthy and required considerable time to read and code appropriately. Hence, we chose to code only contributions from the year 2012 in our data (though data from 2011 and 2013 was still used to calculate further measures as described below). We decided to use a continuous time period rather than a random sample of data because this facilitated the gap time and Markov chain analyses we chose to employ (these analyses are described further below). In total 1,468 posts were coded, which together attracted 34,368 comments. These posts and comments were made by 4,332 distinct users. Further details on the coding process are provided in [Appendix A1.5](#).

A second measurement task was to identify who counts as a “member” of the Stormfront echo chamber, especially when the forum contains contributions from people with views opposed to white nationalism. Clearly, our hypothesized effects only apply to those who already hold white nationalist beliefs. We approached this task in three ways. First, we discounted “guest” users, and considered only those people who had registered an account on the site. There were 5,073 contributions from guest users in our data (14% of total contributions): we unfortunately do not have any way of knowing how many different guests these contributions represent. Second, we discounted people who had never made a contribution outside of the opposing views sub-forum: having contributed in one of the other ideologically consonant sub-forums is, we would argue, a good indicator of being a member of Stormfront. Finally, we discounted anyone who made a dissonant post within the opposing views sub-forum itself. Overall, 3,324 of the 4,332 users in our data (77%) were classified as members in this way: the rest of the users were not used in the calculation of the results reported below. While we believe these adjustments give us the most accurate subset of users, we also checked the sensitivity of our models to this process of filtering by producing further models which did not filter out these people. The results were almost identical.

We also collected a number of other pieces of information. We measured the prior levels of engagement of each user in our dataset, based on the amount of previous contributions they had made to the forum before their current engagement. We use our data from 2011 in this calculation as well. We measured the length of the post, defined as the number of characters, to account for any potential confounding effects this might have on our main relationship of interest (for example, dissonant viewpoints might be longer ones, and longer ones might inherently be more likely to generate responses). And, for each engagement, we observed the location of a user’s next engagement, using data from 2013 if necessary. Descriptive statistics on all measures collected can be found in [Appendix A1.6](#).

Results

We begin our analysis by assessing the amount of engagement different types of post prompt from members of the forum (RQ1), and testing whether this engagement varies according to whether people are more active members of the forum (H1). We address this question and hypothesis in a series of multilevel models reported in [Table 1](#). Each observation in these models is a post made to the forum in the year 2012. We include the anonymized username of the initial poster, the day of the week and the month of the year as random effects, to account for any systematic variation in the amount of responses generated by different users and at

Table 1. Multilevel Linear Models of Post Response Levels and Previous Levels of Contributions

	Amount of responses (log)			Previous contribution levels of responders (log)		
	M1.0	M1.1	M1.2	M2.0	M2.1	M2.2
OV Forum Post	1.49			3.68 ***		
OV Consonant		1.21			3.71 ***	
OV Dissonant		2.74 +	2.29 ***		3.70 ***	1.05
OV Misc		0.94	0.78 **		3.52 ***	0.90
Content length (log10)	1.12***	1.09***	1.04	1.09 ***	1.09***	0.99
Marginal R ²	0.03	0.06	0.09	0.22	0.22	0.00
Conditional R ²	0.22	0.20	0.23	0.24	0.24	0.19
Conditional ICC	0.198	0.157	0.142	0.018	0.018	0.207
Observations	4,264	4,264	1,468	3,353	3,353	1,198

*Note: $p < 0.05$;

** $p < 0.01$;

*** $p < 0.001$

OV: Opposing Views

different times. Estimates of statistical significance were computed using the Kenward–Roger approximation for degrees of freedom (Halekoh & Højsgaard, 2014). Marginal and conditional R^2 were calculated using the method proposed by Nakagawa and Schielzeth (2013), and we also provide the conditional Intra-Class Correlation Coefficient (Nakagawa et al., 2017). Diagnostics show no reason to doubt the results reported in the models (see Appendix A2.1 for a list of diagnostic checks conducted), though suggested the dependent variables we used should be log transformed (as noted in the text below).

Models 1.0 to 1.2 in Table 1 address RQ1, looking at whether dissonant material receives a different amount of responses than consonant material. Here the dependent variable is simply the number of comments each individual post receives. As the number of comments exhibited a skewed distribution, this variable was incremented by one and log transformed, returning it to a good approximation of normality. Model 1.0 compares all material found in the opposing views [OV] forum to all material found outside it. As is clear from the results reported in Table 1, while opposing views do seem to generate more responses, the difference is not significant. Model 1.1 breaks down the opposing views [OV] content into three categories as defined above: consonant arguments, dissonant arguments and miscellaneous. Dissonant material again generates more responses than content outside the opposing views forum, but the result is only borderline significant. Model 1.2, finally, focusses only on material within the opposing views sub-forum. Here there is a significant effect: dissonant arguments receive more than twice the amount of responses that consonant arguments do. Hence overall there is some evidence that dissonant material generates more reaction than consonant material. Importantly, there is no evidence that it generates less reaction.

Models 2.0 to 2.2 in Table 1 address H1, concerning the prior levels of engagement of individuals on the forum. In these models, each observation is a post and the dependent variable is the average levels of past engagement of the people replying to that post, measured as the number of contributions they have previously made (again, this variable is incremented by one and log transformed). There is clear evidence for H1. Model 2.0 of Table 1 shows that people responding within the opposing views forum have been more than three times as active as those responding outside. Model 2.1 shows that this is true for all types of opposing views forum posts. Model 2.2 shows that, however, the differences *between* different types of opposing views post are not significant.

We will now look at our second research question (RQ2), concerning whether engagement with opposing views makes further engagement more or less likely. In Table 2, we address the question making use of a PWP Gap Time model (see Amorim & Cai, 2015; Prentice et al., 1981). Such an analysis is suited to our case because we do not observe whether (or when) people definitively decide to stop contributing to a forum, just the amount of time between each post. The technique also handles censoring in our data.

Table 2. PWP Gap Time Models Estimating the Probability of Posting Another Contribution, given the Location of the Previous Post

	Probability of posting again	
	M3.0	M3.1
OV Forum Post	1.18***	
OV Consonant		1.20***
OV Dissonant		1.14***
OV Misc		1.28***
Number of Replies (log10)	1.16***	1.16***
R ²	0.012	0.013
Users	3,324	3,324
Events	57,625	57,625

*Note: $p < 0.05$;

** $p < 0.01$;

*** $p < 0.001$

Model 3.0 of [Table 2](#) looks at the impact of engagements within the opposing views forum compared to engagements outside of it. There is good evidence to show that individuals exposed to counter-messages are not deterred by them: people using the opposing views forum are about 17% more likely to return than those engaging outside of it. Model 3.1 breaks this down by the type of opposing engagement. It shows that while all types of engagement have a positive effect, engaging with consonant arguments has more of an effect than engaging with dissonant arguments. Diagnostics show that the proportional hazards assumption of both models is violated, indicating the result changes over time, with the strongest part of the effect coming in the first 80 minutes after the initial post (see [Appendix A2.2](#) for more details on these diagnostics). Hence encounters with dissonant viewpoints stimulate a brief but immediate likelihood of returning to the forum.

We will now move on to our final hypothesis (H2) tackling re-confirmation. Do people seek consonant material following engagement with dissonance? We address this question using a discrete time Markov chain model. In the model, we treat engagement with the different categories of content (consonant arguments, dissonant arguments, and discussions outside of the opposing views forum) as different states between which people transition, with the transitions indicated by their engagements with the forum. These records of interactions are used to infer the transition probabilities in the model (following the method of [Yalamanchi and Spedicato, 2015](#)).

The inferred transition probabilities are shown in [Table 3](#). Note that we only show a subset of all the results here, see [Appendix A2.3](#) for the full transition probabilities, including a graphical representation. The results support H2: there is evidence that people engaging with dissonant viewpoints are around 2.5 percentage

Table 3. Markov Chain Transition Probability Estimations for the Likelihood of Engaging Outside of the Opposing Views Forum

Having seen a . . .	Probability of transition outside of opposing views forum	
	Estimate	95% Confidence Interval
Consonant viewpoint	20.45%	19.54–21.35%
Dissonant viewpoint	22.99%	22.08–23.89%
Number of users	3,324	
Number of state transitions	57,625	

Note: Probabilities are only presented for two initial states (consonant and dissonant) and one type of transition (move outside of the opposing views forum). Full transition probabilities for all initial states are available in Appendix 2.3.

points more likely to transition outside of the Opposing Views [OV] forum for their next engagement than people who engage with consonant material. The 95% confidence intervals reported in the table do not overlap, indicating that the difference reported here can be considered statistically significant at conventional levels.

As a final analytical step, we engaged in an inductive coding of responses to dissonant posts, following a similar process to Yardi and boyd (2010), to provide an idea of the types of discussion that they provoked. This process was conducted without prior expectations: all themes and counter-argument strategies identified were grounded in the data itself. Our overall finding was that responses to dissonant posts were overwhelmingly counter arguments defending the existing ideology of white nationalism on the forum. We identified two particularly prominent types of counter argument.

First, we identify what we term *repeated rebuttal*. Dissonant viewpoints that appear on the site, especially those that strongly criticize white nationalism on the grounds of either racism, hypocrisy or ignorance (pointing, for example, to the fact that white nationalism seems to marginalize the issue of slavery), attract stock rebuttal responses that had been previously featured on the site many times. We saw repetition of “the mantra”: that critiques of white nationalists as “racists” were merely criticisms of white people in disguise⁴. We also saw repetition of the idea that many white people were also slaves, and that mainstream arguments about slavery overlooked this. What is critical here, in our view, is not the arguments themselves but the fact that they were repeated over and over again by the different members of the site. Indeed, forum moderators at times stated that they had allowed opposing views onto the site purely so that common arguments against them could be rehearsed. Arguably one of the most common “echoes” in this echo chamber is the sound of opposing viewpoints being undermined and marginalized.

Second, we identified *selective evidencing* as a further type of counter-argument. This involves incorporating decontextualized news reports, government statistics, and academic articles in order to support some of the “scientific” claims of white

nationalism. Posts about differential rates of sickness between races were a strong theme, for example highlighting a report suggesting that black people were more likely to have AIDS⁵ or contract skin cancer.⁶ Sometimes this practice of selective evidencing involved reporting on social problems and suggesting that they are related to these genetic differences: for example, statistics showing apparent disproportionate receipt of social security checks by racial minorities,⁷ or a study showing that multicultural neighborhoods have lower trust.⁸ What we found striking was that much of the evidence used was drawn from what might be regarded as “mainstream” sources (such as major news outlets, government statistics portals and academic institutions), rather than from other white nationalist sites. Simply by omitting crucial pieces of context, or slightly misreading the intent of the article, members of Stormfront were able to reformat mainstream information in a way that made it consistent with their belief system. This raises serious questions about the value of measuring mere exposure to diverse sources as an indicator of homophilous communities (supporting the work of [Haller and Holt, 2019](#)).

Discussion

This article has investigated how users on the radical “echo chamber” Stormfront react to counter-messages that challenge their shared ideology. At odds with the literature on cognitive dissonance, our analysis demonstrates that opposing views are not systematically avoided by Stormfront members, and do not cause them to cease participation. Instead, these posts stimulate discussion, and encourage users to remain active on the site overall. They are also frequently engaged with by the most active members of the forum. We also identified the many ways in which these opposing views are managed to minimize the amount of dissonance they might provoke: by traducing, decontextualizing and reinterpreting them; and through the provision of spaces with purely consonant cognitions that people can seek out soon after having engaged with counter-messages.

The main consequence of our findings is that counter-messaging might not be an adequate strategy to dismantle a radical echo chamber such as Stormfront. As counter-messaging and strategic communication are increasingly touted as a valid approach to combating extremism online by social media platforms and governments (see [Archetti, 2019](#)), our research stresses the limited potential for counter-messages and suggests the possibility that counter-messaging can exacerbate the very problem it seeks to solve by provoking further engagement in the forums it arrives in.

Today, of course, the ecosystem of far-right extremism makes use of a variety of platforms including web forums and imageboards such as 4Chan and 8Chan, mainstream social media platforms, and bespoke platforms for hate such as Gab.ai and Parler ([Baele et al., 2020](#)). Building on our findings, we suggest that communication solutions like counter-messaging are likely to have little effect on forums and platforms dedicated to hate and white supremacy. Following recent attempts by

platforms such as Facebook, Twitter, and YouTube to deplatform extremist users (Rogers, 2020), user migration to other sites including Gab, Parler, and encrypted messaging applications such as Telegram makes our findings increasingly relevant to the formulation of policies to break up extremist echo chambers.

The relevance of our findings is not limited to forums and platforms such as Stormfront. Research on mainstream social media platforms, where a significant amount of counter-messaging activity currently takes place (see Ganesh, 2019), shows that far right users continue to engage in forms of collective identity formation, even in environments that are more porous to dissonance than Stormfront (e.g. Deem, 2019; Gaudette et al., 2020). These findings also echo existing work on echo chambers in traditional media. In the same way that Jamieson and Cappella conceptualized the American conservative media “echo chamber” as an epistemic community built on the systematic marginalization of liberal figures and viewpoints (2008, p. 163), far-right online echo chambers seem to thrive on the marginalization and ridicule of their opponents.

Of course, beyond the behavioral data we observe, we can only speculate about whether exposure to counter attitudinal information has any effect on the beliefs and opinions of those on Stormfront. Some existing literature suggests that such exposure might strengthen users’ beliefs in the superiority of their views (Karlsen et al., 2017), by allowing them to maintain the ideological boundary between their in-group and a despised “other” (Graham, 2019; Meddaugh & Kay, 2009). This connects to work that has shown how arguments are more persuasive if they are presented alongside the opposing view (see e.g. O’Keefe, 2002). It also links back to Festinger’s early claim that a mass experience of dissonance could lead to beliefs becoming stronger (see discussion in Abelson, 1959, p. 346; see also Tormala and Petty, 2004; McKimmie et al., 2003, 2009). The constant exposure to relatively low levels of dissonance may contribute to beliefs becoming progressively stronger on the forum, supporting work which has found that consistent users of the forum become progressively more radical in their beliefs (Scrivens et al., 2020). However, we cannot demonstrate this with the data in our study.

It is also worth acknowledging the other limitations to our study which are inherent in the data we have collected. One limitation is our focus on just one website, and one time period. This narrow focus was necessitated by the fact that the qualitative coding task was nuanced and time consuming, and also that a time series of observations was necessary to address our hypotheses. However, further work could usefully address whether our findings generalize to other contexts. A second issue is our reliance on observational scraped data. This means that we focus only on *active* engagements in the Stormfront forum. We cannot know how many people actually read the contributions we are interested in, nor what the impact of reading (as opposed to actively engaging through commenting) was. We also cannot analyse the behavior of guest users. This limits the application of dissonance theory that we make use of here. Future work could usefully extend this analysis by using a combination of survey and digital trace data.

Relatedly, there is the issue of moderation. Although in theory anyone is free to create an account and to create content on Stormfront, we know that in practice many

comments are deleted. Therefore, the oppositional viewpoints which make it onto the forum do not necessarily represent the full range of all possible opposition, but rather the ones that the moderators have selected for inclusion. We know that the moderators are strategic in this behavior, and often mentioned that they are “allowing” opposing viewpoints. This relates to another issue, which is our focus on observational data, which necessarily limits the strength of the causal claims we can make. A final point, related to the above, is that oppositional viewpoints are corralled into only one part of the overall Stormfront site. We do not know what impact oppositional viewpoints would have if they appeared in other parts of the site, and it may be that their impact is minimized precisely because they appear in an area dedicated to opposition (even if we are able to observe the difference between dissonant and consonant viewpoints within this one subforum). It may well also be that different types of users are attracted to different areas of the forum. For all of these reasons, further tests of our findings in a field experimental setting, where the number and type of counter-messages injected in the echo chamber is controlled, would be highly valuable.

Funding

The work was financed by the VOX-Pol Network, which was funded under the EU 7th Framework Programme (grant number 312827).

Notes

1. This section of the forum is described as being “For all our opponents who want to argue with White Nationalists.”
2. This scraper was called the “Dark Web Monitor”. Further details on the monitor are available on the website: <https://dws.pm/monitor/>
3. Data and replication scripts supporting the article are available from the following location: <https://ora.ox.ac.uk/>
4. For further details see: <https://www.adl.org/education/references/hate-symbols/anti-racist-is-a-code-for-anti-white>
5. See <http://edition.cnn.com/2008/HEALTH/conditions/07/29/black.aids.report/index.html>
6. See <http://news.bbc.co.uk/1/hi/health/5219752.stm>
7. See http://www.census.gov/aprd/www/statbrief/sb95_22.pdf
8. See: <https://www.ft.com/content/c4ac4a74-570f-11db-9110-0000779e2340>
9. See <https://archive.org/web/>
10. See vbulletin.com
11. Available from https://www.vbulletin.com/docs/html/main/sitebuilder_modules_activitystream
12. See <https://archive.org/web/>

References

- Abelson, R. P. (1959). Modes of resolution of belief dilemmas. *Journal of Conflict Resolution*, 3(4), 343–352. <https://doi.org/10.1177%2F002200275900300403>
- Agathangelou, P., Katakis, I., Rori, L., Gunopulos, D., & Richards, B. (2017). Understanding online political networks: The case of the far-right and far-left in Greece. *Lecture Notes in Computer Science*, 10539 LNCS, 162–177. https://doi.org/10.1007/978-3-319-67217-5_11
- Amorim, L. D., & Cai, J. (2015). Modelling recurrent events: A tutorial for analysis in epidemiology. *International Journal of Epidemiology*, 44(1), 324–333. <https://doi.org/10.1093/ije/dyu222>
- Archetti, C. (2019). The unbearable thinness of strategic communication. In C. Bjola and J. Pammont (Eds.), *Countering online propaganda and extremism: The dark side of digital diplomacy* (pp. 81–96). London: Routledge.
- Baele, S. J., Brace, L., & Coan, T. G. (2019). From “incel” to “saint”: Analyzing the violent worldview behind the 2018 Toronto attack. *Terrorism and Political Violence*, 1–25. <https://doi.org/10.1080/09546553.2019.1638256>
- Baele, S. J., Brace, L., & Coan, T. G. (2020) Uncovering the far-right online ecosystem: An analytical framework and research agenda. *Studies in Conflict & Terrorism*. <https://doi.org/10.1080/1057610X.2020.1862895>
- Bakshy, E., Messing, S., & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239), 1130–1132. <https://doi.org/10.1126/science.aaa1160>
- Beirich, H. (2015). White homicide worldwide. Retrieved from: https://www.splcenter.org/sites/default/files/d6_legacy_files/downloads/publication/white-homicide-worldwide.pdf
- Bliuc, A.-M., Faulkner, N., Jakubowicz, A., & McGarty, C. (2018). Online networks of racial hate: A systematic review of 10 years of research on cyber-racism. *Computers in Human Behavior*, 87, 75–86. <https://doi.org/10.1016/j.chb.2018.05.026>
- Bloomfield, E. F., & Tillery, D. (2019). The circulation of climate change denial online: Rhetorical and networking strategies on Facebook. *Environmental Communication*, 13(1), 23–34. <https://doi.org/10.1080/17524032.2018.1527378>
- Bode, L. (2016). Pruning the news feed: Unfriending and unfollowing political content on social media. *Research & Politics*, 3(3). <https://doi.org/10.1177/2053168016661873>
- Boutyline, A., & Willer, R. (2017). The social structure of political echo chambers: Variation in ideological homophily in online networks. *Political Psychology*, 38(3). <https://doi.org/10.1111/pops.12337>
- Bowman-Grieve, L. (2009). Exploring “Stormfront”: A virtual community of the radical right. *Studies in Conflict & Terrorism*, 32(11), 989–1007. <https://doi.org/10.1080/10576100903259951>
- Braddock, K., & Horgan, J. (2016). Towards a guide for constructing and disseminating counternarratives to reduce support for terrorism. *Studies in Conflict & Terrorism*, 39(5), 381–404. <https://doi.org/10.1080/1057610X.2015.1116277>
- Brewer, M. B., & Kramer, R. M. (1985). The psychology of intergroup attitudes and behavior. *Annual Review of Psychology*, 36(1), 219–243. <https://doi.org/10.1146/annurev.ps.36.020185.001251>

- Bright, J. (2018). Explaining the emergence of political fragmentation on social media: The role of ideology and extremism. *Journal of Computer-Mediated Communication*, 23(1), 17–33. <https://doi.org/10.2139/ssrn.2839728>
- Brugnoli, E., Cinelli, M., Quattrociocchi, W., & Scala, A. (2019). Recursive patterns in online echo chambers. *Scientific Reports*, 9(1), 20118. <https://doi.org/10.1038/s41598-019-56191-7>
- Bruns, A. (2019). *Are filter bubbles real?* Cambridge: Polity.
- Caren, N., Jowers, K., & Gaby, S. (2012). A social movement online community: Stormfront and the white nationalist movement. *Research in Social Movements, Conflicts and Change*, 33, 163–193. [https://doi.org/10.1108/S0163-786X\(2012\)0000033010](https://doi.org/10.1108/S0163-786X(2012)0000033010)
- Carr, C. T. (2017). Social media and intergroup communication. In: *Oxford research encyclopedia of communication*. <https://doi.org/10.1093/acrefore/9780190228613.013.460>
- Castro, K. I., & Hemmingsen, A.-S. (2017). The trouble with counter narratives. Retrieved from: <https://www.econstor.eu/bitstream/10419/197640/1/880712619.pdf>
- Choi, D., Chun, S., Oh, H., Han, J., & Kwon, T. “Taekyoung.” (2020). Rumor propagation is amplified by echo chambers in social media. *Scientific Reports*, 10(1), 310. <https://doi.org/10.1038/s41598-019-57272-3>
- Cleland, J. (2014). Racism, football fans, and online message boards: How social media has added a new dimension to racist discourse in English football. *Journal of Sport and Social Issues*, 38(5), 415–431. <https://doi.org/10.1177/0193723513499922>
- Cohen, A. R. (1960). Attitudinal consequences of induced discrepancies between cognitions and behavior. *Public Opinion Quarterly*, 24(2), 297–318. <https://psycnet.apa.org/doi/10.1086/266950>
- Coleman, M. C. (2017). Rhetorical logic bombs and fragmented online publics of vaccine science. *Journal of Contemporary Rhetoric*, 7(4), 203–216. http://contemporaryrhetoric.com/wp-content/uploads/2018/02/Coleman7_4_1.pdf
- Cooper, J., & Stone, J. (2000). Cognitive dissonance and the social group. In D. J. Terry & M. A. Hogg (Eds.), *Applied social research. Attitudes, behavior, and social context: The role of norms and group membership* (pp. 227–244). Lawrence Erlbaum Associates, Publishers.
- Daniels, J. (2009). *Cyber racism: White supremacy online and the new attack on civil rights*. Rowman & Littlefield Publishers.
- De Koster, W., & Houtman, D. (2008). Stormfront is like a second home to me. *Information Communication and Society*, 11(8), 1155–1176. <https://doi.org/10.1080/13691180802266665>
- Deem, A. (2019). Extreme speech: The digital traces of #whitegenocide and alt-right affective economies of transgression. *International Journal of Communication*, 13, 3183–3203. Retrieved from <https://ijoc.org/index.php/ijoc/article/view/9631/2717>.
- Dubois, E., & Blank, G. (2018). The echo chamber is overstated: The moderating effect of political interest and diverse media. *Information, Communication & Society*, 21(5), 729–745. <https://doi.org/10.1080/1369118X.2018.1428656>
- Duggan, M., & Smith, A. (2016). The political environment on social media. Retrieved from: <https://www.pewresearch.org/internet/2016/10/25/the-tone-of-social-media-discussions-around-politics/>
- Eagly, A. H., & Chaiken, S. (1998). Attitude structure and function. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., pp. 269–322). Oxford University Press.

- Ehrlich, D., Guttman, I., Schönbach, P., & Mills, J. (1957). Postdecision exposure to relevant information. *Journal of Abnormal and Social Psychology*, 54(1), 98–102. <https://psycnet.apa.org/doi/10.1037/h0042740>
- Elliot, A. J., & Devine, P. G. (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology*, 67(3), 382–394. <https://doi.org/10.1037/0022-3514.67.3.382>
- Ferguson, K. (2016). Countering violent extremism through media and communication strategies. Retrieved from <https://www.paccsresearch.org.uk/wp-content/uploads/2016/03/Countering-Violent-Extremism-Through-Media-and-Communication-Strategies-.pdf>
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford University Press.
- Fletcher, R., & Nielsen, R. K. (2018). Are people incidentally exposed to news on social media? A comparative analysis. *New Media and Society*, 20(7), 2450–2468. <https://doi.org/10.1177/1461444817724170>
- Frey, D. (1986). Recent research on selective exposure to information. *Advances in Experimental Social Psychology*, 19, 41–80. [https://doi.org/10.1016/S0065-2601\(08\)60212-9](https://doi.org/10.1016/S0065-2601(08)60212-9)
- Ganesh, B. (2019). Evaluating the Promise of Formal Counter-Narratives. In B. Ganesh and J. Bright (Eds.), *Extreme digital speech: Contexts, responses and solutions* (pp. 89–98). VOX-Pol Network of Excellence.
- Ganesh, B., & Bright, J. (2019). *Extreme digital speech: Contexts, responses and solutions*. VOX-Pol Network of Excellence.
- Ganesh, B., & Bright, J. (2020). Countering Extremists on Social Media: Challenges for Strategic Communication and Content Moderation. *Policy & Internet*, 12(1), 6–19. <https://doi.org/10.1002/poi3.236>
- Garrett, K. (2009). Echo chambers online? Politically motivated selective exposure among internet news users. *Journal of Computer-Mediated Communication*, 14(2), 265–285. <https://doi.org/10.1111/j.1083-6101.2009.01440.x>
- Garrett, K., Weeks, B., & Neo, R. (2016). Driving a wedge between evidence and beliefs: How online ideological news exposure promotes political misperceptions. *Journal of Computer-Mediated Communication*, 21(5), 331–348. <https://doi.org/10.1111/jcc4.12164>
- Gaudette, T., Scrivens, R., Davies, G., & Frank, R. (2020). Upvoting extremism: Collective identity formation and the extreme right on Reddit. *New Media & Society*, 1–18. <https://doi.org/10.1177/1461444820958123>
- Gervais, B. T. (2015). Incivility online: Affective and behavioral reactions to uncivil political posts in a web-based experiment. *Journal of Information Technology and Politics*, 12(2), 167–185. <https://doi.org/10.1080/19331681.2014.997416>
- Gibbons, F. X., Eggleston, T. J., & Benthin, A. C. (1997). Cognitive reactions to smoking relapse: The reciprocal relation between dissonance and self-esteem. *Journal of Personality and Social Psychology*, 72(1), 184–195. <https://psycnet.apa.org/doi/10.1037/0022-3514.72.1.184>
- Gonzalez-Bailon, S., Kaltenbrunner, A., & Banchs, R. E. (2010). The structure of political discussion networks: A model for the analysis of online deliberation. *Journal of Information Technology*, 25(2), 230–243. <https://doi.org/10.1057/jit.2010.2>
- Graham, E. (2019). Boundary maintenance and the origins of trolling. *New Media & Society*, 21(9), 2029–2047. <https://doi.org/10.1177/1461444819837561>

- Guess, A., Nyhan, B., Lyons, B., & Reifler, J. (2018). Avoiding the echo chamber about echo chambers. Retrieved from https://kf-site-production.s3.amazonaws.com/media_elements/files/000/000/133/original/Topos_KF_White-Paper_Nyhan_V1.pdf
- Halekoh, U., & Højsgaard, S. (2014). A Kenward-Roger approximation and parametric bootstrap methods for tests in linear mixed models-the R Package Pbkrttest. *Journal of Statistical Software*, 59(9), 1–32. <https://doi.org/10.18637/jss.v059.i09>
- Haller, A., & Holt, K. (2019). Paradoxical populism: How pegida relates to mainstream and alternative media. *Information Communication and Society*, 22(12), 1665–1680. <https://doi.org/10.1080/1369118X.2018.1449882>
- Hara, N., & Estrada, Z. (2003). Hate and peace in a connected world: Comparing moveon and Stormfront. *First Monday*, 8(12). <https://doi.org/10.5210/fm.v8i12.1104>
- Harmon-Jones, E., Brehm, J. W., Greenberg, J., Simon, L., & Nelson, D. E. (1996). Evidence that the production of aversive consequences is not necessary to create cognitive dissonance. *Journal of Personality and Social Psychology*, 70(1), 5–16. <http://dx.doi.org/10.1037/0022-3514.70.1.5>
- Hartzell, S. L. (2020). Whiteness feels good here: interrogating white nationalist rhetoric on Stormfront. *Communication and Critical/Cultural Studies*, 17(2), 129–148. <https://doi.org/10.1080/14791420.2020.1745858>
- Hatewatch. (2017). Waning storm: Stormfront.org loses its domain. Retrieved from Southern Poverty Law Center website: <https://www.splcenter.org/hatewatch/2017/08/29/waning-storm-stormfrontorg-loses-its-domain>
- Hogg, M. A., Abrams, D., Otten, S., & Hinkle, S. (2004). The social identity perspective: Intergroup relations, self-conception, and small groups. *Small Group Research*, 35(3), 246–276. <https://doi.org/10.1177/1046496404263424>
- Iyengar, S., & Hahn, K. S. (2009). Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1), 19–39. <https://doi.org/10.1111/j.1460-2466.2008.01402.x>
- Jamieson, K. H., & Capella, J. (2010). *Echo chamber: Rush Limbaugh and the conservative media establishment*. Oxford University Press.
- John, N. A., & Dvir-Gvirsman, S. (2015). I don't like you any more: Facebook unfriending by Israelis during the Israel-Gaza conflict of 2014. *Journal of Communication*, 65(6), 953–974. <https://doi.org/10.1111/jcom.12188>
- Jost, J. T., van der Linden, S., Panagopoulos, C., & Hardin, C. D. (2018). Ideological asymmetries in conformity, desire for shared reality, and the spread of misinformation. *Current Opinion in Psychology*, 23, 77–83. <https://doi.org/10.1016/j.copsyc.2018.01.003>
- Karlsen, R., Steen-Johnsen, K., Wollebæk, D., & Enjolras, B. (2017). Echo chamber and trench warfare dynamics in online debates. *European Journal of Communication*, 32(3), 257–273. <https://doi.org/10.1177/0267323117695734>
- Kleinberg, D., van der Vegt, I., & Gill, P. (2021). The temporal evolution of a far-right forum. *Journal of Computational Social Science*, 4, 1–23. <https://doi.org/10.1007/s42001-020-00064-x>
- Knobloch-Westerwick, S. (2015). The selective exposure self- and affect-management (SESAM) model: Applications in the realms of race, politics, and health. *Communication Research*, 42(7), 959–985. <https://doi.org/10.1177/0093650214539173>

- Knobloch-Westerwick, S., & Meng, J. (2009). Looking the other way: Selective exposure to attitude-consistent and counterattitudinal political information. *Communication Research*, 36(3), 426–448. <https://doi.org/10.1177/0093650209333030>
- Knobloch-Westerwick, S. (2008). Information utility. In W. Donsbach (Ed.), *International Encyclopedia of Communication* (pp. 2273–2276). Blackwell.
- Lee, B. (2020). Countering violent extremism online: The experiences of informal counter messaging actors. *Policy & Internet*, 12(1), 66–87. <https://doi.org/10.1002/poi3.210>
- Losch, M. E., & Cacioppo, J. T. (1990). Cognitive dissonance may enhance sympathetic tonus, but attitudes are changed to reduce negative affect rather than arousal. *Journal of Experimental Social Psychology*, 26(4), 289–304. [https://psycnet.apa.org/doi/10.1016/0022-1031\(90\)90040-S](https://psycnet.apa.org/doi/10.1016/0022-1031(90)90040-S)
- Mackie, D. M., Devos, T., & Smith, E. R. (2000). Intergroup emotions: Explaining offensive action tendencies in an intergroup context. *Journal of Personality and Social Psychology*, 79(4), 602–616. <https://doi.org/10.1037/0022-3514.79.4.602>
- Marwick, A. E., & Caplan, R. (2018). Drinking male tears: Language, the manosphere, and networked harassment. *Feminist Media Studies*, 18(4), 543–559. <https://doi.org/10.1080/14680777.2018.1450568>
- Matz, D. C., & Wood, W. (2005). Cognitive dissonance in groups: The consequences of disagreement. *Journal of Personality and Social Psychology*, 88(1), 22–37. <https://doi.org/10.1037/0022-3514.88.1.22>
- McKimmie, B. M., Terry, D. J., & Hogg, M. A. (2009). Dissonance reduction in the context of group membership: The role of metaconsistency. *Group Dynamics*, 13(2), 103–119. <https://doi.org/10.1037/a0014554>
- McKimmie, B. M., Terry, D. J., Hogg, M. A., Manstead, A. S. R., Spears, R., & Doosje, B. (2003). I'm a hypocrite, but so is everyone else: Group support and the reduction of cognitive dissonance. *Group Dynamics*, 7(3), 214–224. <https://doi.org/10.1037/1089-2699.7.3.214>
- Meddaugh, P. M., & Kay, J. (2009). Hate speech or “reasonable racism?” The other in Stormfront. *Journal of Mass Media Ethics*, 24(4), 251–268. <https://doi.org/10.1080/08900520903320936>
- Melton, Z. J., & Motyl, M. (2018). Segregation and belief polarization as boundary conditions for when fusion leads to self-sacrifice. *The Behavioral and Brain Sciences*, 41(e211), 211. <https://doi.org/10.1017/S0140525X1800170X>
- Min, S. J., & Wohn, D. Y. (2018). All the news that you don't like: Cross-cutting exposure and political participation in the age of social media. *Computers in Human Behavior*, 83, 24–31. <https://doi.org/10.1016/j.chb.2018.01.015>
- Munson, S. A., & Resnick, P. (2010). Presenting diverse political opinions: How and how much. *Proceedings of CHI 2010*, 1457–1466. <https://doi.org/10.1145/1753326.1753543>
- Mutz, D. C. (2002). The consequences of cross-cutting networks for political participation. *American Journal of Political Science*, 46(4), 838. <https://doi.org/10.2307/3088437>
- Mutz, D. C. (2006). *Hearing the other side: Deliberation versus participatory democracy*. Cambridge University Press.
- 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>

- Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R^2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2), 133–142. <https://doi.org/10.1111/j.2041-210x.2012.00261.x>
- Neubaum, G., & Krämer, N. C. (2017). Opinion climates in social media: Blending mass and interpersonal communication. *Human Communication Research*, 43(4), 464–476. <https://doi.org/10.1111/hcre.12118>
- Neubaum, G., Rösner, L., Ganster, T., Hambach, K., & Krämer, N. C. (2018). United in the name of justice: How conformity processes in social media may influence online vigilantism. *Psychology of Popular Media Culture*, 7(2), 185–199. <https://doi.org/10.1037/ppm0000112>
- O’Keefe, D. J. (2002). *Persuasion: Theory and research*. SAGE Publications Ltd.
- Oswald, L., & Bright, J. (2021). How do climate change skeptics engage with opposing views? *arXiv preprint*. Available from: <https://arxiv.org/abs/2102.06516>
- Perry, B. and Scrivens, R. (2016). White pride worldwide. Constructing global identities online. In J. Scheppe and M. A. Walters. (Eds.), *The globalization of hate. Internationalizing hate crime* (pp. 65–78). Oxford Scholarship Online.
- Prentice, R. L., Williams, B. J., & Peterson, A. V. (1981). On the regression analysis of multivariate failure time data. *Biometrika*, 68(2), 373–379. <https://doi.org/10.2307/2335582>
- Regehr, K. (2020). In(cel)doctrination: How technologically facilitated misogyny moves violence off screens and on to streets. *New Media & Society*, 1–18. <https://doi.org/10.1177/1461444820959019>
- Röchert, D., Neubaum, G., Ross, B., Brachten, F., & Stieglitz, S. (2020). Opinion-based homogeneity on YouTube: Combining sentiment and social network analysis. *Computational Communication Research*, 2(1), 81–108. <https://doi.org/10.5117/CCR2020.1.004.ROCH>
- Rogers, R. (2020). Deplatforming: Following extreme Internet celebrities to Telegram and alternative social media. *European Journal of Communication*, 35(3), 213–229. <https://doi.org/10.1177%2F0267323120922066>
- Rosand, E., & Winterbotham, E. (2019, March 20). Do counter-narratives actually reduce violent extremism? The Brookings Institution. <https://www.brookings.edu/blog/order-from-chaos/2019/03/20/do-counter-narratives-actually-reduce-violent-extremism/>
- Sasahara, K., Chen, W., Peng, H., Ciampaglia, G. L., Flammini, A., & Menczer, F. (2020). Social influence and unfollowing accelerate the emergence of echo chambers. *Journal of Computational Social Science*. <https://doi.org/10.1007/s42001-020-00084-7>
- Scharkow, M., Mangold, F., Stier, S., & Breuer, J. (2020). How social network sites and other online intermediaries increase exposure to news. *Proceedings of the National Academy of Sciences*, 117(6), 2761–2763. <https://doi.org/10.1073/pnas.1918279117>
- Schmitt, J. B., Rieger, D., Rutkowski, O., & Ernst, J. (2018). Counter-messages as prevention or promotion of extremism?! The potential role of YouTube. *Journal of Communication*, 68(4), 780–808. <https://doi.org/10.1093/joc/jqy029>
- Scrivens, R., Davies, G., & Frank, R. (2020). Measuring the evolution of radical right-wing posting behaviors online. *Deviant Behaviour*, 41(2), 216–232. <https://doi.org/10.1080/01639625.2018.1556994>
- Sibona, C. (2014). Unfriending on Facebook: Context collapse and unfriending behaviors. *Proceedings of the Annual Hawaii International Conference on System Sciences*, 1676–1685. <https://doi.org/10.1109/HICSS.2014.214>

- Simon, L., Greenberg, J., & Brehm, J. (1995). Trivialization: The forgotten mode of dissonance reduction. *Journal of Personality and Social Psychology*, 68(2), 247–260. <https://psycnet.apa.org/doi/10.1037/0022-3514.68.2.247>
- Stromer-Galley, J. (2014). Political discussion and deliberation online. In K. Kenski & K. H. Jamieson (Eds.), *The Oxford handbook of political communication*. https://doi.org/10.1093/oxfordhb/9780199793471.013.015_update_001
- Stroud, N. J. (2010). Polarization and partisan selective exposure. *Journal of Communication*, 60(3), 556–576. <https://doi.org/10.1111/j.1460-2466.2010.01497.x>
- Stroud, N. J. (2014). Selective exposure theories. In K. Kenski & K. H. Jamieson (Eds.), *The Oxford handbook of political communication* (pp. 119–142). Oxford University Press.
- Sunstein, C. R. (2001). *Republic.com*. Princeton University Press.
- Torcal, M., & Maldonado, G. (2014). Revisiting the dark side of political deliberation: The effects of media and political discussion on political interest. *Public Opinion Quarterly*, 78(3), 679–706. <https://doi.org/10.1093/poq/nfu035>
- Tormala, Z. L., & Petty, R. E. (2004). Source credibility and attitude certainty: A metacognitive analysis of resistance to persuasion. *Journal of Consumer Psychology*, 14(4), 427–442. https://doi.org/10.1207/s15327663jcp1404_11
- Törnberg, P. (2018). Echo chambers and viral misinformation: Modelling fake news as complex contagion. *PLOS ONE*, 13(9), e0203958. <https://doi.org/10.1371/journal.pone.0203958>
- Valenzuela, S., Kim, Y., & Gil de Zuniga, H. (2012). Social networks that matter: Exploring the role of political discussion for online political participation. *International Journal of Public Opinion Research*, 24(2), 163–184. <https://doi.org/10.1093/ijpor/edr037>
- van Eck, C. W., Mulder, B. C., & van der Linden, S. (2021). Echo chamber effects in the climate change blogosphere. *Environmental Communication*, 15(2), 145–152. <https://doi.org/10.1080/17524032.2020.1861048>
- Wojcieszak, M. (2010). “Don’t talk to me”: Effects of ideologically homogeneous online groups and politically dissimilar offline ties on extremism. *New Media and Society*, 12(4), 637–655. <https://doi.org/10.1177/1461444809342775>
- Wong, M. A., Frank, R., & Allsup, R. (2015). The supremacy of online white supremacists – an analysis of online discussions by white supremacists. *Information & Communications Technology Law*, 24(1), 41–73. <https://doi.org/10.1080/13600834.2015.1011845>
- Yalamanchi, S. B., & Spedicato, G. A. (2015). Bayesian inference of first order Markov chains. Retrieved from <https://cran.microsoft.com/snapshot/2015-09-16/web/packages/markovchain/vignettes/BayesianInference.pdf>
- Yang, J. H., Barnidge, M., & Rojas, H. (2017). The politics of “unfriending”: User filtration in response to political disagreement on social media. *Computers in Human Behavior*, 70, 22–29. <https://doi.org/10.1016/j.chb.2016.12.079>
- Yardi, S., & boyd, d. (2010). Dynamic debates: An analysis of group polarization over time on Twitter. *Bulletin of Science, Technology & Society*, 30(5), 316–327. <https://doi.org/10.1177/0270467610380011>

Appendix A1: Description of Data Collection and Coding Procedures

A1.1 Sub-forum Description

A description of the sub-forums included in the study can be found in [Table A1](#). The opposing views sub-forum is the main area of interest, where individuals using the forum can encounter both dissonant and consonant material. The other six sub-forums were selected to provide a range of different types of conversation: some oriented towards the ideological and campaigning aspects of white nationalism, others more informal. These six were also selected as together they contained approximately the same amount of comments as the opposing views forum itself during the study period. The values reported in the table refer to 2012 only, which was the main year under study.

A1.2 Data Completeness Verification

Our data was scraped automatically from the Stormfront website, and covers the period 2011 to 2013 (though we only apply our coding to the year 2012). The automatic scraper sought to recursively download threads from the Stormfront website by

Table A1. Stormfront Sub-forums Included in the Study

Sub-forum name	Description (as shown on Stormfront)	Number of Posts Observed (2012 only)	Number of Comments Observed (2012 only)
Opposing Views Forum	“For all our opponents who want to argue with White Nationalists”	1,468	34,368
Events	“White Nationalist demonstrations, rallies, conferences, talk shows, media interviews.”	116	844
Local and Regional	“Contact information for those who want to work together in their communities.”	175	1,344
For Stormfront Ladies Only	“Sugar and spice, and everything nice.”	208	5,640
Politics & Continuing Crises	“Practical politics, including Ron Paul, Bugsters, Tea Parties, Occupy Wall Street, and, yes, the occasional conspiracy theory.”	1,922	23,995
Strategy and Tactics	“Promoting White Rights through local organization.”	139	1,077
Talk	“Meet other White Nationalists for romance or friendship.”	236	2,087

The descriptions are taken directly from the stormfront.org website (see stormfront.org/forum).

iterating through pages within these threads and extracting the comments and posts within them into a dataset.

While we believe the scraper performed reasonably well, we also know that intermittent server outages and malformed HTML on the site may have caused some data quality issues. Therefore, we sought to verify the completeness of the data by comparing it to the actual number of contributions made to our forums of interest during this time period.

We sought these numbers by making use of the Wayback Machine.⁹ This website archives snapshots of other websites from around the web. It allowed us to look at the state of the Stormfront forum at the start and end of our data capture window (we used Wayback Machine snapshots from the 30th of December 2010 and the 2nd of January 2014, which were the closest available dates to the start and end of our data capture period). In particular, we looked at records of the number of contributions made in each sub-forum of the site. The way these are counted on Stormfront is not completely clear, however in total for our forums of interest we observed 281,341 contributions whilst on the Wayback machine there were 290,691. So overall our total is 96.8% of that reported on Stormfront. A nuance to note is that we actually have slightly more contributions in some of our forums than are listed on the site. It wasn't quite clear why this is and may mean that the Stormfront forum is reporting slightly different statistics. If we disregard these "overfull" forums from our validation analysis our completeness figure becomes a more conservative 93.2%. This is the figure we have reported in the main text.

A1.3 Post Display Mechanisms

We also wanted to understand the way in which posts are displayed to users. Different social sites have different ways of ordering the content they host and presenting it to those visiting the site: some may base it on recency, others on post popularity, whilst others may try and personalize their offering to individual users. Clearly, the ordering of posts may make a difference to our outcome variables.

Stormfront, is powered by a forum hosting technology called vBulletin.¹⁰ We first reviewed the current description of how vBulletin works, and in particular how it proposes content to viewers.¹¹ This suggested that site administrators are able to choose how content is displayed, picking from a number of options (e.g., most recent first, most commented, etc.).

To check how the administrators of Stormfront had configured their site during the time period of study, we again made use of the Wayback Machine.¹² We visited the earliest available snapshot of the forum for each month in 2012, and recorded the order the posts were displayed in. In all of the months we checked, the posts were displayed in order of recency. However, it is worth noting that, having clicked through to a given subforum, users can choose to alter the order themselves, and display either by number of replies the posts have received or the number of views.

A1.4 Defining White Nationalism

A definition of white nationalism was critical for our coding procedure to take place. Our definition is based on a close reading of material contributed to the Stormfront website, though works by [Meddaugh and Kay \(2009\)](#) and [Bowman-Grieve \(2009\)](#) also offered considerable inspiration. We define the ideology as consisting of the following basic beliefs.

First, white nationalists believe that there are important differences between 'races' of humans, and that these differences emerge from genetic features. This does not necessarily equate with believing that different races are inferior to the "white race", though many white nationalists also believe other races are less intelligent and/or more prone to criminality than they are.

Second, white nationalists believe that different nations of the world should be ethnically homogeneous, and as such that white people should be provided with an "all white" or at least "majority white" nation in which to live. Opinions diverge considerably about justifiable means for achieving such an outcome, with some advocating peaceful solutions whilst others suggest force would be legitimate.

Finally, white nationalists tend to believe that the white race is in decline or being replaced by policies of multiculturalism and immigration. Some link this to ideas about a (Zionist) conspiracy of global elites.

Needless to say, we think that the beliefs outlined here show the "extremism" of white nationalist views, and their separation from mainstream society.

A1.5 Coding Process

We chose to code every post (these are original contributions which start discussion threads) in the opposing views forum in the year 2012. The coding was performed by two of the authors. Each author was assigned half of the 1,468 posts to code. The authors read the complete text of the post, and assigned it to one of three categories: "consonant," "dissonant," and "miscellaneous." The median post was around 540 characters in length, or approximately 100 words.

"Consonant" posts are ones that were broadly ideologically consonant with the belief system of white nationalism as described in [Appendix A1.3](#). Postings which expressed support for any of the beliefs we identify as relating to white nationalism were labelled as consonant. Such posts could still be about questions about the best way to live as a white nationalist, or questions about the detail of the belief system. However, they would not be ones which raised major questions about the faith.

"Dissonant" posts, by contrast, are ones which sought to attack or undermine white nationalism in some way. There were many different approaches taken by the authors of dissonant posts. Some sought to highlight what they perceived of as logical fallacies in the white nationalist argument. Others contradicted some of the scientific claims on which the idea of racial differences was based. Also common were postings which involved invective or ad hominem attacks on white nationalists themselves. The main thing these posts have in common is that a white nationalist would feel that their views were, in some way, being challenged.

"Miscellaneous" posts were ones which did not clearly fall into either of the above two categories. These might be posts which were off topic (i.e., did not engage in ideological debate), or ones in which the intent of the poster was not clear.

One hundred of these posts were also selected at random and coded by both of the authors, allowing us to produce an estimate of inter-coder reliability for the coding exercise. We observed an 88% agreement between the authors, which equated to a Krippendorff's Alpha of 0.814, which is usually considered a good level of inter-coder reliability.

The confusion matrix for the coding is shown in [Table A2](#). There was disagreement on 12 of the items coded, of which 6 were disagreements between the key categories of interest (i.e. posts where one coder marked "Consonant" and one marked "Dissonant"). The balance of sample sizes between all the three categories is reasonable. Overall we feel the results support the validity of the coding scheme.

Table A2. Confusion Matrix

	Consonant	Dissonant	Misc
Consonant	40	2	3
Dissonant	4	28	2
Misc	1	0	20

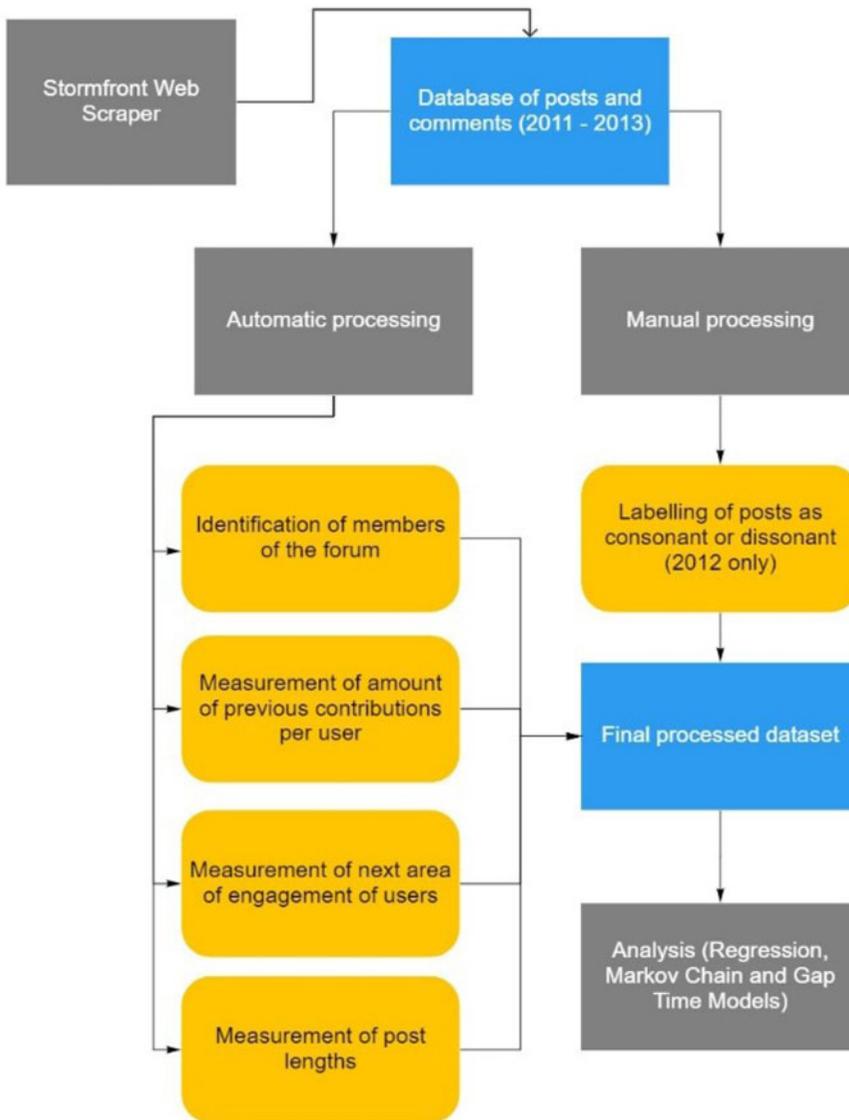
A1.6 Descriptive Statistics

[Table A3](#) presents descriptive statistics for the study.

Table A3. Descriptive Statistics

	Users	
Members	3,324	(77%)
Non-members	1,008	(23%)
Total users	4,332	
	Threads	Comments
Inside opposing views	1,468	34,368
Outside opposing views	2,796	34,987
Total	4,264	69,335
Consonant threads	702	
Dissonant threads	469	
Miscellaneous	297	
	Median	Mean
Seniority	155	200
Post length (characters)	538	1,111

A1.7 Graphical representation of the data collection and analysis pipeline



Appendix A2: Model Diagnostics

A2.1 Diagnostics for Post Response and Contributor Seniority Models

Models 1.0–2.2 are multilevel models. Data were grouped by the user creating the post, as we might expect different users to generate different typical reactions when posting threads, and by the month of posting, as we might expect response levels to fluctuate over time. We did not use the forum of posting as a grouping factor because variation here is already largely accounted for by the post type categorical variable.

The following diagnostic checks were performed. The normality of the dependent variable was inspected (leading it to be log transformed). Plots of residuals versus fitted values were inspected. These suggested no problems in the case of models 1.0–1.2, though some skew was observed in the case of models 2.0–2.2. This skew was caused by posts with few responses. The dependent variable of these models is the average seniority of people commenting on a post: clearly, posts which had fewer comments would be more likely to have higher variance in the average seniority of responders. Further models were produced with posts with less than ten responses removed. The fit was much improved, but the results were largely the same.

The normality of residuals was also inspected. This suggested some problems in the case of models 1.0–1.2 which were resolved by removing posts with zero responses (the results of these new models being substantially the same as the ones reported in the main text). These plots also suggested the same truncation of posts with few responses for models 2.0–2.2 which was already performed above. Finally, hat values were calculated to check for high leverage points. However, no hat values greater than the conventional cut-off of 0.4 were observed.

Overall the diagnostics provided no reason to doubt the results of the models reported in the main text.

A2.2 Gap Time Analysis Diagnostics

Models 3.0 and 3.1 are a type of proportional hazards model known as a recurrent events model (specifically a PWP Gap Time model). These models function as a normal event history analysis, except that there are multiple events per individual. The model hence clusters standard errors by individual, and also stratifies observations by the number of previous events and individual has had. There are three main diagnostic considerations for such a model. First, the number of individuals can be relatively low at higher strata, which can lead to instability in estimates (see Amorim & Cai, 2014). Hence, truncated models were produced that disregarded events after the 169th event (which was the largest strata with at least 50 individuals within it). The results of these models were substantially the same. Second, the models were checked for influential observations, measured using DF Beta scores. No observations above the conventional cutoff were observed (the cutoff being $2l < i > < i > n$).

Finally, the proportional hazards assumption was checked for both models. This was found to be violated, meaning that the ratio of hazards between groups is not constant over time. The evolution of the estimate, compared to the main effect reported in model 3.0 (the solid black line) is shown in [Figure A1](#) above. The figure shows that the effect is strongly positive for around the first hour after posting. After this time, it fluctuates around zero until around a week, when it turns negative. This indicates that the main “stimulus” effect of encountering an opposing view occurs shortly after having engaged with it. This is reported in the main text.

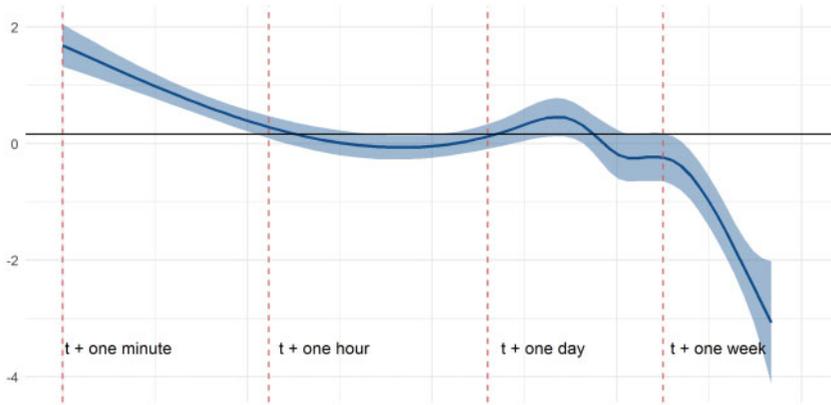


Figure A1 Trend line indicates a smoothed LOESS estimation of the evolution of the main effect from model 3.0 over time. The black horizontal line indicates the estimate reported in model 3.0.

A2.3 Markov Chain Full Transition Probabilities

Table A4 above provides the full transition probabilities for all possible pairs of states in our Markov chain model. The probabilities reflect our estimate of the next type of discussion thread someone is likely to contribute to, given the previous one they contributed to. For example, if an individual contributed to a consonant discussion

Table A4. Markov Chain Transition Probabilities

Previous contribution	Next contribution	Estimate	Lower bound	Upper bound
Consonant	Consonant	0.5011	0.487	0.5153
Consonant	Outside	0.2045	0.1954	0.2135
Consonant	Dissonant	0.2171	0.2078	0.2264
Consonant	Unknown	0.0773	0.0718	0.0829
Outside	Consonant	0.0539	0.0515	0.0563
Outside	Outside	0.8555	0.846	0.865
Outside	Dissonant	0.0665	0.0638	0.0691
Outside	Unknown	0.0241	0.0225	0.0257
Dissonant	Consonant	0.1991	0.1907	0.2076
Dissonant	Outside	0.2299	0.2208	0.2389
Dissonant	Dissonant	0.4948	0.4815	0.5081
Dissonant	Unknown	0.0762	0.071	0.0814
Unknown	Consonant	0.1661	0.1537	0.1784
Unknown	Outside	0.1984	0.1849	0.2119
Unknown	Dissonant	0.2194	0.2052	0.2336
Unknown	Unknown	0.4161	0.3966	0.4357

thread within the opposing views forum, we estimate a 0.5011 probability that their next contribution will also be a consonant discussion thread in the opposing views forum.

The transition probabilities from the Markov chain are also represented graphically in Figure A2.

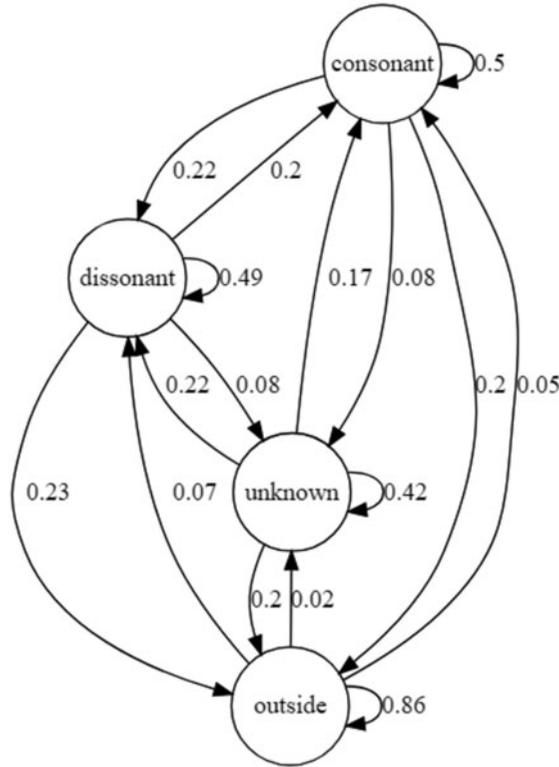


Figure A2 Fitted Markov chain estimates of transition probabilities between different areas in the forum.