

May 2022

Dialogic Language as Digital Ethos: an Analysis of Language Used in the Anti-Vaccine Conversation on Twitter

Jeffery A. Sternstein
University of Wisconsin-Milwaukee

Follow this and additional works at: <https://dc.uwm.edu/etd>



Part of the [Rhetoric Commons](#)

Recommended Citation

Sternstein, Jeffery A., "Dialogic Language as Digital Ethos: an Analysis of Language Used in the Anti-Vaccine Conversation on Twitter" (2022). *Theses and Dissertations*. 2953.
<https://dc.uwm.edu/etd/2953>

This Dissertation is brought to you for free and open access by UWM Digital Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of UWM Digital Commons. For more information, please contact scholarlycommunicationteam-group@uwm.edu.

DIALOGIC LANGUAGE AS DIGITAL ETHOS: AN ANALYSIS OF
LANGUAGE USED IN THE ANTI-VACCINE CONVERSATION ON TWITTER

by

Jeffery Sternstein

A Dissertation Submitted in
Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
in English

at

The University of Wisconsin-Milwaukee

May 2022

ABSTRACT

DIALOGIC LANGUAGE AS DIGITAL ETHOS: AN ANALYSIS OF LANGUAGE USED IN THE ANTI-VACCINE CONVERSATION ON TWITTER

by

Jeffery Sternstein

The University of Wisconsin-Milwaukee, 2022
Under the Supervision of Professor David Clark

Many scholars attribute social media's influence with a rise in distrust of expert advice. These scholars have suggested that people are turning to non-experts for advice because those non-experts seem to be more willing to openly discuss medical issues while also providing empathy, as opposed to the experts who have been trained to speak with detached authority. For this dissertation, I have done a study to find evidence supporting these theories. To do this, I looked at the Twitter conversation which has been focusing on anti-vaccination themes. Drawing on tweets from within that conversation, I conducted an inter-rater reliability test to categorize 1,000 tweets as either using a more empathetic and conversational tone versus those with the authoritative tone traditionally favored by experts. I then used those evaluations to conduct machine learning to evaluate over 50,000 additional tweets from the anti-vaccination conversation. I evaluated the relative success of tweets those tweets which used "authoritative" language compared to those that used "dialogic" language. Through this research, I was able to find a correlation between the degree to which the language within a tweet seemed to express empathy and encourage give-and-take forms of conversation and with engagement rates achieved by those tweets. Analysis suggests that the amount of influence this language use has on engagement rates is relatively minor, with tweets using

stronger levels of dialogic language earning approximately one additional like for every 5,000 followers an account may have over tweets using primarily authoritative language. This study was done with the intention of considering how an audience's preference for dialogic language might influence the way we prioritize authoritative voice in academic writing. As the data only marginally confirms this preference, this study shifts focus to ways of teaching students to be more responsible as readers in lieu of relying on experts using a more empathetic voice.

© Copyright by Jeffery Sternstein, 2022
All Rights Reserved

TABLE OF CONTENTS

List of Figures	vi
List of Tables	viii
Acknowledgements	ix
1 Introduction	1
2 Literature Review	5
2.1 Social Media and Post-Human Identity	8
2.2 Beyond the Point of Identity Formation	12
2.3 Studies on Online Credibility	14
2.4 Digital Ethos	19
2.5 Marketing Advice	21
2.5 Marketing for Medical Communities	24
2.6 Enter 2020	26
3 Methods	28
3.1 Measuring Favorable Reactions	30
3.2 Gathering Data	31
3.3 Evaluating the Language of the Tweets	36
3.4 Verifying Inter-Rater Reliability	40
3.6 Gathering Data from Machine Learning	44
4 Data and Results	52
3.8 Data from Preliminary Tweets My Readers Evaluated	52
3.9 Data from the Full Dataset Run	58
5 Tweet-Level Analysis	67
5.1 Current Status of the Anti-Vaccination Conversation	68
5.2 Observations from Preliminary Data	87
5.3 Observations from Full Dataset Run	91
5.4 Noteworthy Dialogic Tweets from the Preliminary Data	93
5.5 General Observations on Tweet Content	97
5.6 Observations on Other Ways of Being Heard	101
6 Implications and Conclusions	106
6.1 Limitations of this Study and Possibilities for Further Research	109
6.2 Concepts in Media Literacy	118
6.3 Recent Studies on the Efficacy of Media-Related Literacies	121
6.4 Conclusions	123
References	124
Appendix: Machine Learning Code	132

LIST OF FIGURES

Figure Description	Page
Figure 1: Average engagement for dialogic tweets	60
Figure 2: Average engagement for authoritative tweets	61
Figure 3: Average certainty ratings for engagement rates	63
Figure 4: @Justiceforevee – Sep 19	70
Figure 5: @picphysicians – Jan 24	71
Figure 6: @ChildrensHD – Sep 9	72
Figure 7: @ChildrensHD – Sep 17	72
Figure 8: @JenniferMarguli – Jul 30	74
Figure 9: @JenniferMarguli – Sep 2	75
Figure 10: @DrFrankLipman – Sep 4	76
Figure 11: @mercola – Aug 19	77
Figure 12: @truthlovingdr – Aug 19	78
Figure 13: “Know the facts” admin message	79
Figure 14: @BusyDrT “account suspended” admin message	83
Figure 15: Holocaust-style “No vax” patch	84
Figure 16: @delbigtree – Sep 10	84
Figure 17: @Vaxxed2 – Apr 6	86
Figure 18: @michellemalkin “Redacted” admin message	87
Figure 19: @michellemalkin – Sep 25	87
Figure 20: Number of dialogic tweets	88

Figure 21: @delbigtree – Jun 9	93
Figure 22: Responses to @delbigtree – Jun 9	94
Figure 23: @DrFrankLipman – Sep 17	96
Figure 24: “Doing research” meme – created by Jeffery Sternstein	98
Figure 25: @JenniferMarguli – Aug 10	98
Figure 26: @truthlovingdr – Aug 13	99

LIST OF TABLES

Table Description	Page
Common Twitter metrics	30
Inter-rater reliability instructions	42
Inter-rater reliability results	43
Binary sort examples	46
Dialogic tweeting breakdown by account (preliminary)	53
Average engagement rates for preliminary data	56
Average engagement rates for full dataset	59
Average engagement by dialogic rating	60
Average engagement by authoritative rating	61
Average certainty ratings by engagement rate	62
Dialogic tweeting breakdown by account (full dataset)	114

ACKNOWLEDGEMENTS

First off, I would like to thank my advisor, Dr. David Clark, not only for his guidance and encouragement, but also for handling the Python programming for the machine learning used in this dissertation. For their additional guidance and support, I would also like to thank the other members of my committee.

I would like to thank my uncle, Dr. Martin Sternstein, for helping me with the statistical analysis of my data. Additionally, I'd like to thank Trevor Sprague, Jeremy Maule, and Sara Maule for being my inter-rater reliability readers and for helping me categorize the initial set of tweets.

Finally, I would like to thank my parents and my sister for supporting me and encouraging me through this entire process, with a special thank-you to my father for making the option of graduate school a viable option in the first place.

Introduction

We are currently seeing groups of people openly denying scientific evidence and defying expert advice over large issues facing us today. This problem seems particularly concerning in health care and in medical issues, such as with the current COVID pandemic. While the distrust of authority figures is not a new phenomenon, many scholars attribute social media's influence with a rise in the visibility of that distrust (McComisky, McIntyre, Nichols). Tom Nichols, however, cautions us that blaming Facebook and Twitter is too simplistic of an explanation (6). It is reasonable to believe that a person's trust in someone's advice would be influenced by how they build their case and communicate that advice. I am interested in looking at how experts and laypeople are delivering their arguments and advice. More specifically, I would like to examine forms of *ethos* which may be better suited to engendering trust over a medium like social media.

In "Did Media Literacy Backfire?" danah boyd, new-media communications expert and researcher for Microsoft, points out that, while doctors and scientists may have valid information, there seems to be a personal connection that is missing from the way they deliver that information. According to boyd, many people are turning to non-experts for their information because, as she says, "Strangers on the Internet are willing to listen, empathize, and compare notes." Those strangers may not have definitive answers, but they seem to converse in an empathetic, human voice which, to some audiences, makes them seem more credible than the authoritative experts.

Along similar lines, Howard Gardner, professor of cognition and education at Harvard University, states that there has been a shift in the very nature of truth in today's media-rich environment, saying, "authority and objectivity have been supplemented – or even supplanted – by authenticity and transparency" (30). While Gardner never explicitly defines authenticity, he uses it to describe language which conveys a sense of being natural, uncontrived, unedited, and representative of the writer's character or spirit. He explains that media consumers no longer put their belief in a source based on that source's status, training, or expertise. Instead, Gardner says that they put their trust into those sources who *seem* candid and authentic.

Theorists from other disciplines make similar observations. As early as 1999, seven years before Facebook and Twitter were opened to the general public, the *Cluetrain Manifesto* announced that, over the internet, the "[h]omogenized 'voice' of business – the sound of mission statements and brochures – will seem... contrived and artificial," to consumers. The creators said that businesses will need to develop a genuine human voice, using language that is "Natural, open, honest, direct." Dave Kerpen says, "Marketing in a social media and Facebook world is not about broadcasting your message... It's about tapping into the conversation, listening, engaging, and empowering" (9). Even political satirist, Stephen Colbert, announced, "I don't trust books. They're all fact, no heart. And, that's exactly what's pulling our country apart today... We are divided between those who think with their head and those who know with their heart."

These theories have important implications for how doctors and scientists communicate with skeptical or hostile audiences. While the theories are compelling, the scholars discussing them have not done studies on the effectiveness of these forms of open, authentic, or

empathetic communication. For this dissertation, I have done a study to find evidence supporting or challenging these theories. To do this, I looked at the Twitter conversation which has been focusing on anti-vaccination themes. Drawing on tweets from within that conversation, I evaluated the relative success of tweets with a more empathetic tone versus those with the authoritative tone traditionally favored by experts.

Through this research, I was able to find a correlation between the degree to which the language within a tweet seemed to express empathy and encourage give-and-take forms of conversation (the combination of which I will refer to “dialogic” in this dissertation) and with engagement rates (a commonly used Twitter metric used to measure the success of a tweet) achieved by those tweets. However, the data suggests that the amount of influence this language use has on engagement rates is relatively minor. As I will discuss in more detail in what is to follow, a large increase in the use of dialogic language predicts an Increase of only one additional like for every 5,000 to 7,000 followers a Twitter account may have. Thus, the data seems to confirm these hypotheses about audience preferences for dialogic voice but only marginally so.

My ultimate goal in performing this research was to consider any implications these theories on preferences for empathetic *ethos* may have on the teaching of college composition. If portions of seemingly anti-intellectual audiences are responding to the empathetic voices boyd describes, developing an empathetic voice in writing could be a more important tool for delivering an intellectual message than we typically acknowledge in first-year composition classes. Teaching our students to communicate with a more empathetic style may open up avenues of communication for having meaningful conversations. Concerned audiences would

then get valid information AND the empathy they need in order to feel empowered by that information. If an empathetic tone had been found to be significantly more effective, using that tone may go a long way in opening up the dialogue on issues like the use of vaccinations. As a writing teacher, I also feel that a better understanding of the types of *ethos* readers are responding to in popular digital environments, like Twitter, may have implications for the type of voice we should be teaching our students to use in their writing.

Traditionally, we have taught our students to write in an authoritative voice and encouraged them – as they become experts in their fields – to continue to use that authoritative voice. I began this dissertation considering the question of how we might want to rethink what types of *ethos* we prioritize in academic writing if we were seeing a strong preference for empathetic voices in digital spaces – if teaching students to write more effectively may also mean teaching them how to utilize different forms of *ethos* not traditionally prioritized in academic writing. As the data I collected only marginally confirms this preference, I consider other opportunities for study to get better picture of the relationship between dialogic language and tweet success. In the absence of conclusive data on ways to equip our students as writers in conversations like these, I also discuss current thinking on ways of teaching students to be more responsible as readers.

Literature Review

Over the past couple of decades, there has been much discussion on how the changing digital landscape has complicated our concept of literacy. Our approaches to reading and writing have become more nuanced with regard to the medium through which that reading and writing takes place. With specific reference to media literacy, authors like Amber Buck and Collin Gifford Brooke highlight the importance of skills in evaluating new rhetorical situations of digital spaces and adapting to genre conventions of specific media platforms. However, other authors believe that recent events pertaining to the public's consumption of media suggest that we have misread some the important factors in developing media literacy.

In "Did Media Literacy Backfire?" danah boyd, new-media communications expert and researcher for Microsoft, argues that the ability to evaluate the credibility of online sources is a skill that should be considered vital to well-developed media literacy. However, it is a skill that most people have not paid enough attention to until recent events revealed how far behind most people are in developing that skill. She points primarily towards widespread acceptance of fake news and other misinformation revolving around the election cycle of 2016 as evidence of this ability to judge credibility.

"Children... are taught that they are the sole proprietors of knowledge. All they have to do is "do the research" for themselves and they will know better than anyone what is real" (3) boyd explains. However, "doing the research," she fears, has just become a matter of pulling the first result off of a Google search or scrolling through results of a search until finding an entry that says exactly what we were hoping to find. Boyd feels that most people now have learned to "trust their gut" when evaluating information.

Trusting their gut can, on the other hand, get readers into trouble because of the affordances social media gives authors and how those authors are encouraged to make use of those affordances. Tom Nichols, in his book *The Death of Expertise* points out social media and the internet have allowed many new voices to be heard – including those less-than-credible voices which would not have previously had a platform to broadcast their views and opinions. More importantly, the internet, he says, “allows people to mimic intellectual accomplishment by indulging in an illusion of expertise provided by a limitless supply of facts” (106). Nichols is quick to point out that having citable facts is not the same as having knowledge. However, Ola Erstad explains that, in her views of media literacy, a skilled author is one “who can act with authority across a series of domains and who is accustomed to forms of collaboration, genuine challenge, experimentation, risk-taking, curiosity and expressivity” (91) – the very type of person who can exploit the illusion of expertise Nichols mentioned, or, worse yet, make use of non-facts while still maintaining the act of authority.

The dark side to this shift in concepts of authority or expertise is what has led to a rise in post-truth rhetorics in the media today, primarily, but not exclusively, in social media. “Post-truth,” Bruce McComisky explains, “signifies a state in which language lacks any reference to facts, truth, and realities” (6). He distinguishes this from lying, where one is deliberately misrepresenting facts, by stressing that post-truth rhetorics are characterized by an indifference to the facts and one who will rattle off facts without even caring if they are true or not. In this case, McComisky worries that, “When language has no reference to facts, truths, or realities, it becomes a purely strategic medium” (6).

This is what then allows for science denialism – a concern specifically addressed by danah boyd in her previously mentioned essay. Lee McIntyre tells us that, especially in today’s political climate, “laypersons feel it is in their interest to question both the motives and the competence of scientists” (18). They will frequently reject the tested and supported information presented by scientists in favor of unsubstantiated claims from others. Nichols take the explanation for science denialism a little further in saying, “Americans now believe that having equal rights in a political system also means that each person’s opinion about anything must be accepted as equal to anyone else’s” (5) (echoing Issac Assimov’s quote from 1980, “the false notion that democracy means that 'my ignorance is just as good as your knowledge’”). Nichols continues, “Ignorance has become hip, with some Americans now wearing their rejection of expert advice as a badge of cultural sophistication” (21).

One way of looking at what is underlying the issue here is to consider new ways in which people are evaluating sources of information. Until recently, the spread of information was controlled by some form of gatekeeper – publishers, editors, etc. However, without those gatekeepers to control what information can be spread through the internet, readers have had to make the judgement on who to trust based on their own understanding or intuition. Just like in any conversation, much of that intuition may be a reaction to how the author they are reading portrays themselves in writing. As I will discuss shortly, the immersive nature of our interaction with digital writing pushes us to see that digital writing as a true expression of identity, and we have grown accustomed to making judgements of a person based only on our exposure to their social media presence. Early theories of post-humanism, specifically, concern themselves with the intersection between identity and digital spaces. Post-human concepts of

digital representation of the self raises questions about authenticity that were not an issue when media were controlled by editors and other gatekeeper experts.

Social Media and Post-Human Identity

If people are making judgements about an author's authority and authenticity based solely on their social media presence, we need to understand the link between digital writing and our perceptions of identity. Understanding the interplay between how we represent ourselves in social media spaces and identity involves seeing the self as a posthuman construct. According to N. Katherine Hayles, "In the posthuman, there are no essential differences or absolute demarcations between bodily existence and computer simulation, cybernetic mechanism and biological organism, robot teleology and human goals" (3). As a site of connection between bodily existence and computer simulation, we must consider the writing done for social media as an outlet for identity formation and identity play. Only by fully appreciating the writer's drive to embody him or herself within a social media environment can we gain a deeper understanding of micro-blogging.

According to Hayles, the posthuman view is characterized by the following assumptions:

1. The posthuman view privileges informational pattern over material instantiation;
2. The posthuman view considers consciousness as an epiphenomenon;
3. The posthuman view thinks of the physical body as only the first of a potential line of prostheses which may be replaced;
4. The posthuman view configures human beings so that they can be seamlessly articulated with intelligent machines (3).

In short, Hayles says, “In the posthuman, there are no essential differences or absolute demarcations between bodily existence and computer simulation, cybernetic mechanism and biological organism, robot teleology and human goals” (3). When studying social media, considering this view of human identity, we must consider the writing done for social media as primarily being an outlet for identity formation and identity play. Common reasoning assumes that there is a distinction between a person’s physical self and the representation of that self through language. With this in mind, the key takeaway I have from the work of Hayles and others is that we need to understand that there is no difference between what is commonly seen as the persona or identity that an author may craft on social media and that which is centered in a physical form. In the posthuman reasoning, identity is identity regardless of how or where it is performed.

Anne Wysocki discusses modern media in saying, “We come to be always already embedded – embodied – in mediation... We therefore need to consider our engagements with our media if we and the people in our classes are to learn about our embodiment and so what we consider ourselves to be and to be able to do in words” (4). Through personal profiles and public networking, social media becomes a site of identity performance. Richard Gilbert elaborates on the connection between identity and computer simulation along with our embodiment in media, saying, “In this conception, consciousness and aspects of the self (while ultimately still embodied within the human driver) will be increasingly externalized and distributed into digital forms... Within this new model, the source of identity remains internal... but the expression or enactment of this consciousness becomes increasingly external, disembodied, and distributed” (232). In Hayles’ terms, the external enactment of

consciousness is the immaterial informational pattern of information which represents the self. Thus, when we look at forms of writing on social media, we must consider that writing to be attempts to fully realize an incarnation of the self through informational pattern in order to fully understand how authors communicate via social media.

In discussing social media profiles and the external expression of consciousness, David Kreps goes back to a metaphor of the self from Deleuze, who likened identity to a series of masks. The commonly believed fallacy, he explained, is the assumption that there is a 'true' face behind the masks. The masks, in this view, are the only true expression of identity. Kreps connects this to social media, saying, "The profile is but one of its creator's many masks, and the representative burden lifts, becomes more playful, and perhaps even more revealing (of the) differential nature of the identity/question that created it" (112). The social media profile becomes the external mask that social networking readers see an author through. In this way, the expression of identity is externalized to digital media. Anne Wysocki further connects this to the writing process in pointing out, "We see ourselves in what we produce. We can look at what we produce to ask, "Is that who I (at least in part) am? Is that who I want to be? Is that a position through which I want to be seen?" (25). From this perspective, then, the primary purpose of the social media profile is to serve as a 'face' through which to present ourselves. All writing choices in the construction of the profile stem from the need to make this mask as full and expressive as possible.

To further elaborate on this idea, and to tie it back to modern, digital media, Zeynep Tufekci states, "The fundamental duality of being human: we are at once embodied and symbolic. Some technologies allow us to separate those two aspects... words without bodies"

(Whitehead, 34). In considering the symbolic representation of the self through language, sociolinguists Mary Bucholtz and Kira Hall outline five principles of linguistic identity formation which are useful in understanding the relationship between social media writing and posthuman theories and how authors may go about constructing themselves through informational patterns. While their principles are not unique to digital spaces, they do form the basis for which readers will judge an author's portrayal of themselves. For the purpose of understanding posthuman representations of the self in social media spaces, we need to keep three of those principles in mind as we go forward. Those principles of identity formation are:

1. The emergence principle – “Identity is best viewed as the emergent... and therefore as fundamentally a social and cultural phenomenon” (588);
3. The indexicality principle – “Identity relations emerge in interaction through several related indexical processes, including... the use of linguistic structures and systems that are ideologically associated with specific personas and groups” (594);
5. The partialness principle – “Any given construction of identity may be in part deliberate and intentional, in part habitual... in part an outcome of others' perceptions and representations” (606).

Combining these principles of linguistic identity formation with posthuman theories, which wed human and technology, provide the best perspective for the understanding of what goes on in social media writing. As such, we cannot approach social media writing with a concept of a pre-existing self in mind. As Bucholtz explained, we need to see identity as emergent through the writing done on social media. This, in turn, can shed more light on what we are seeing when we look at social media writing and why social media platforms have

become so popular and widely used. Part of what this will help to explain is the emphasis placed on social media writers' drive to develop an authentic sounding voice.

If, as suggested by posthumanist theory, the writer is actively embodying his or her profile - that the author is inhabiting the digital space in the same way that they inhabit the physical - establishing an *ethos* of authenticity would serve as external confirmation that the author has done so successfully. Exclamations of, "Oh, you sound so real" become validations of the user's identity, essentially saying, "You've successfully inserted yourself into this digital environment, and you've breathed life into the online persona you've crafted, so we can hang out together here, in this digital space." Comments like this would show, to the author, that he or she has successfully made the transition from bodily existence to living computer simulation. This, in turn, allows the reader to become fully immersed in the digital interaction – seeing the author's digital representation of themselves as a complete and fully-actualized individual. Understanding these theories on identity can lend credence to the observations made by boyd, McComisky and Gardner on the nature of an authentic, human voice seeming more credible in digital spaces.

Beyond the Point of Identity Formation

With an appreciation of the immersive nature of digital performances of identity, we can turn our attention back to issues of authority and post-truth rhetorics. Considering this complex convergence of concerns with media literacy, post truth rhetorics, post-humanism, and science denialism, I feel that we need to know more about why people are trusting the sources they do. Scholars who have addressed this question seem to look towards one of three

areas: message content, at-a-glance profile information, and compositional elements. Those looking at message content reduce the situation to the ideas being conveyed. The at-a-glance social media features include, primarily, profile information like user name, profile image, and the number of followers an account has. I however am more interested in the compositional elements like style, word choice, tone, etc.

Those studying message content typically gravitate towards two main influences: confirmation bias and the psychological appeal of conspiracy theories. Many authors, Nichols, McComisky and McIntyre included, quickly point to confirmation bias. It seems generally accepted that we tend to trust authors whose ideas fit with what we already believe to be true. There has also been much research coming out of the social sciences on the appeal of conspiracy theories, and why, no matter how outlandish, people are enthusiastic about buying into them. The motivation to believe in conspiracy theories can develop for several possible reasons, but one of the most cited is that conspiracy theories can bring order to an otherwise chaotic world – they provide a villain for people to blame for random, frequently tragic, misfortune. While I am sure that message content is a major component of determining credibility, I feel that stopping there sets up a bit of an impasse – it doesn't leave us with any ways of building a productive dialogue nor will it really help us “arm” our scientists and experts to battle this wave of non-truths and manipulations. That being the case, I will not spend much time discussing confirmation bias and the appeal of conspiracy theories here, though it is a topic I will return to in later chapters.

As a rhetoric student and a writing teacher, I feel it would be much more productive to study the at-a-glance features and compositional elements of a message. I still want to believe

that how we say something will have some significance in addition to what we are saying.

There have been recent studies looking at several at-a-glance features which influence a reader's impressions of credibility on social media. However, most of the work on how compositional elements contribute to credibility is more theoretical. Trying to test some of these theories is where I would like to focus my efforts.

Studies on On-line Credibility

Miriam Metzger and Andrew Flanagin theorized that readers on the internet, not having the time or capacity to evaluate information systematically, would use a collection of roughly six heuristics to evaluate the credibility of the information they found. Those heuristics, discussed in "Credibility and Trust of Information in Online Environments: The Use of Cognitive Heuristics," are:

- *reputation*
- *endorsement*
- *consistency*
- *self-confirmation*
- *expectancy violation*
- *persuasive intent*

Metzger describes the reputation heuristic as a form of name recognition. An example of the endorsement heuristic would be trusting a post based on the number of likes or retweets it has received. The consistency heuristic, she says, is based on finding information consistent across multiple sources. The self-confirmation heuristic, also known as confirmation-bias, suggests that readers will trust information that fits with what they already believe to be true. Metzger

gives the example of noticeable typos and improper grammar for the expectancy heuristic where information is judged based on the post fails to meet certain expectations of the reader. She identifies the persuasive intent heuristic as a judgement by the reader on if the source seems to be biased or manipulative. Metzger developed her theories on these heuristics by looking at several other studies but does not go so far as to discuss specific ways in which these heuristics are employed.

In a survey, Babajide Osatuyi sought to find where, within a tweet, people preferred to look for indicators of credibility. In “Information Sharing on Social Media Sites” he outlines his findings. The survey he designed gave participants a short list of tweet elements, and he asked his participants to rank them in the order of how important they were when determining their opinion of credibility. The survey results showed that “topic of interest” was considered to be the top credibility indicator. However, much like my initial thoughts on looking towards message content, Osatuyi seemed a little disappointed in this result, noting the obvious that, “topic of interest is typically an antecedent of most discussions” (2626), and, therefore, doesn’t leave much to be studied. Further results from his survey were more useful, as he found that providing links to other sources was seen as the next best indicator of credibility, followed by the embedding of videos within a post.

Looking more specifically at the effect network size may have on perceptions of credibility David Westerman, et al. supposed that cues on a user’s profile, like number of followers, should be useful to a reader in evaluating credibility. In their article, “A Social Network as Information: The Effect of system Generated Reports of Credibility on Twitter”, Westerman tells us his experiment in showing mock Twitter pages to readers and getting their

impressions of how credible those pages might be. He found that basic network size – the number of followers an author has – had no effect on perceptions of that author’s credibility. What he did find to make a difference was the ratio between the number of followers an author had and the number of accounts that author followed. Westerman explains that this ratio number of followers of an author and the number of accounts that author followed did not impact a reader’s perception of trustworthiness or goodwill, but a large gap between those numbers greatly affected a reader’s impression of competence. Those authors with a narrow gap were seen as much more competent than authors who has a large following but followed relatively very few other accounts.

In their article, “Tweeting is Believing? Understanding Microblog Credibility Perceptions”, Meredith Ringel Morris, et al. explained that, through a survey, they had determined that more readers use features which are visible at-a-glance to determine credibility, rather than features which may be obscured in the user interface, such as the details of an author’s bio and information studied by Westerman. Based on that, they conducted a controlled experiment on how some at-a-glance tweet features – specifically: message topic, user name, and user profile image - affected a tweet’s credibility. They found that different topics and styles of user names did influence a reader’s perceptions of credibility. Science tweets in their study were seen as more credible than tweets about politics or entertainment. Authors with topical user names (those that specifically refer to the topic of discussion - i.e. RhetoricStudent) were seen as more credible than authors with traditional names (i.e. JohnSmith), which were still more credible than those with internet-style names (those refer to random interests or personality traits – usually followed by a number - i.e.

ILoveCats74). Surprisingly, they found that the type of user profile image had no significant effect on perceptions of credibility, despite people identifying profile image as a source for making that judgement on the original survey Morris conducted. Authors who had no profile image at all, on the other hand, were seen as less credible.

Jiang Yang performed a study comparing how readers from the United States judged tweet credibility and how readers from China judged credibility which both confirmed some of Morris' work and contributed new findings. After showing a series of tweets to readers and then having them respond to a survey on those tweets, Yang ultimately did find some significant differences in how American audiences rated tweet credibility and how Chinese audiences did. For the purposes of this dissertation, I'll focus on what she found about readers from the United States. First, she confirmed what Morris found about user names and profile images – that authors with topical names were seen as more credible than those with internet style names, and authors with profile photos were found to be more credible than those with generic images. In addition, she found that tweets authored by men were seen as more credible than those authored by women. Yang also seemed to find evidence that people tweeting from locations with liberal-leaning populations were viewed as more credible than those tweeting from locations with more conservative populations. However, Yang admitted that 89% of her sample readers were from liberal backgrounds, so this could have been a natural bias of her readers.

Possibly shedding a little more light on the geographic location question, in subsequent study by Shafiza Mohd Shariff, et al. as described in "On the Credibility Perception of News on Twitter: Readers, Topics and Features", she looked at the relationship between the

demographics of a sample set readers on the demographics of a set of tweeters to see how that would impact a reader's impression of the credibility of the author. She found that the more similar a reader's education level and geographic location were to that of the author, the more credible the reader found that author to be. Her observations, though, gave no evidence for a correlation between gender or age and perceived credibility (as I read it, this does not necessarily contradict Yang's findings on gender, it only shows that the gender of the reader - in relation to the gender of the writer – does not affect perceptions of credibility). Similar to what Morris examined, Shariff also looked at how tweet topic factored into credibility. Like Morris, she found evidence of a reluctance to believe political news, however, she tied this reluctance specifically to female readers. Going beyond that, Shariff found that all Twitter readers from her sample found breaking news and news on natural disasters to be the most credible.

While most of this research points to the significance of the at-a-glance elements within a tweet, all of these authors recognize that a number of factors go into helping a reader make an assessment of the credibility of the information within a tweet. One author I found did manage to isolate a few compositional elements and measure their impact of credibility. Kyungsik Han, in his article "How Do You Perceive This Author? Understanding and Modeling Authors' Communication Quality in Social Media", did more of a comprehensive analysis of factors contributing to a reader's perception of author credibility on Twitter.

Han began with the assumption that a reader, being unable to investigate every claim on Twitter, would need to judge credibility based on cues within the tweets. For his study, he showed readers collections of 10-15 tweets from several different authors, and he asked those readers to give their impressions of the author's credibility. Han's analysis of his findings led

him to four factors linked to higher perceptions of credibility. First, he found that an author's use of more words longer than six letters (referred to as sixltr words by some social media theorists) lead to higher credibility ratings. He attributed this to the idea that the frequent use of sixltr words is suggestive of better education and higher social status. Secondly, Han found that having more articles in a tweet (a, an, the, etc) was linked to more positive perceptions of credibility. Han believed this indicated "that the use of concrete nouns or interest in objects and things leads to greater communication quality" (11). He also found that the expression of positive emotions in a tweet improved a reader's sense of author credibility, and, finally, the inclusion of more URLs and specific numbers also inspired more of a sense of credibility.

However, other than the word choice analysis done by Han, the literature does not seem to include much analysis of stylistic and compositional elements. In the same essay that I began this chapter with, danah boyd suggests that there are important stylistic elements which are winning over followers. She believes that a certain conversational approach may be an important key.

Digital Ethos

Looking at the health care industry, boyd observed that, while doctors and scientists may have valid information, there seems to be a personal connection that is missing from the way they deliver that information. According to boyd, many people are turning to non-experts for their information because, as she says, "Strangers on the Internet are willing to listen, empathize, and compare notes." Those strangers may not have definitive answers, but they seem to converse in an empathetic, human voice which, to some audiences, makes them seem more credible than the authoritative experts. From this view, it's the author's *ethos* which

becomes more important than the knowledge (or facts) they may or may not have. Boyd further notes that many of the anti-vaccination spokespeople do not necessarily even claim to have definitive information, yet they still may be perceived as credible.

This observation resonates with a more general observation McComisky made about post truth rhetorics. McComisky observed that, in post truth arguments, *ethos* and *pathos* function at the expense of logic. He says, “*Ethos* and *pathos* have themselves become effective sources of argument” (20). Thus, even without claiming to have valid information, anti-vaccine spokespeople can use this “empathetic, human voice” to craft an *ethos* to sway peoples’ opinions, making the facts irrelevant.

According to McComisky, “*Ethos*... describes the rhetorical effect (in terms of credibility) that one personality has on another personality’s willingness or capacity to be persuaded” (21). Aristotle identifies *ethos* as one of the three types of proofs a rhetorician may use for persuasion. This is usually translated simply as “the character of the speaker.” Cited as potentially the most powerful tool available to a rhetorician, Aristotle stresses that an audience will give more credence to a speaker they perceive to be goodwilled.

While most forms of contemporary literacy assume that a person’s credibility should be determined by their level of knowledge, based on the concerns that Boyd was pointing to, it seems that a trustworthy character, to some modern audiences, is not necessarily one concerned with correct information, but one who seems authentic and empathetic. Similarly, Howard Gardner, professor of cognition and education at Harvard University, states that “there has been a seismic shift... authority and objectivity have been supplemented – or even supplanted – by authenticity and transparency” (30). While Gardner never explicitly defines

authenticity, he uses it to describe language which conveys a sense of being natural, uncontrived, unedited, and representative of the writer's character or spirit. He explains that media consumers no longer put their belief in a source based on that source's status, training or expertise. Instead, Gardner says that they put their trust into those sources who seem candid and authentic. These theories linking an authentic voice in digital spaces to higher perceptions of credibility seem quite natural and compelling when thinking back to post-human views of identity.

Even outside of the field of rhetoric, many people saw this shift coming. Some of the theorizing on persuasion that has been done outside of the academic community very closely mirrors these theories on developing authentic human voices. Take, for example, advice given by marketing specialist when discussing approaches to social media.

Marketing Advice

Marketing specialists have long advised businesses and advertisers to approach social media differently than they approached the one-way communication of television, radio and newsprint ads. Social media is seen as a place where your message needs to seem more genuine, heartfelt and open to dialogue. As early as 1999, seven years before social media giants Facebook and Twitter were created, the *Cluetrain Manifesto* announced that, over the internet, the "[h]omogenized 'voice' of business – the sound of mission statements and brochures – will seem... contrived and artificial," to consumers. The creators say that businesses will need to develop a genuine human voice, using language that is "Natural, open, honest, direct." This type of advice continues to be repeated today.

Dave Kerpen, in *Likeable Social Media* (3rd edition published in 2019), says, “Marketing in a social media and Facebook world is not about broadcasting your message... It’s about tapping into the conversation, listening, engaging, and empowering” (9). Kerpen believes that audiences are responding to the personalization opportunities afforded by social media, and that they are becoming exceedingly wary of marketing ploys. He pushes the idea that spokespeople on social media need to develop a more authentic voice – that they need to be human and show personality in their social media interactions. He compares this to meeting someone at a cocktail party and knowing right away if that person is being sincere or fake. Social media readers, he believes, can read those they are interacting with in much the same way.

In an essay that was meant as a response to the Cluetrain Manifesto, “Markets are Conversations”, Doc Searls reinforces this idea of creating a persona with an authentic, human voice in saying, “If you’re going to join, don’t do it as a legal entity or wearing your cloak of officialdom. Join it as a person with a name, a point of view, a sense of humor, and passion” (113). He points out that our society’s first markets were all about people getting together, talking, arguing, bartering and sharing ideas. It’s only been in the past few years, with the influence of radio and television, that “market” has become a verb, and that businesses have been focused on crafting a message rather than a conversation. Platforms like Twitter, he suggests, are bringing us back to our earlier sense of what a market should be.

Following a similar theme, Shama Kabani, in her book *The Zen of Social Media Marketing*, says that social media advertising is all about forming relationships. She talks about taking an interest in your followers, asking them genuine questions, and providing them with

real value in your communication. In talking about Twitter, specifically, she says, “It’s about building a new form of community. It’s about learning. It’s about support, inspiration, and daily motivation” (110).

Dom Sagolla, a co-founder of Twitter, continues to give similar advice. While his book, *140 Characters*, is geared more towards the average user, he touches on these same themes of marketing, saying, “Communication and consumption must change... traditional media is a totalitarian aristocracy, subject to the political whims of the corporate few with power” (4), whereas communication via Twitter should be more like public speech where one should create a persona with an authentic voice. He also urges marketers to avoid focusing on numeric measurements of growth and reach and, instead, to ask themselves how they would rank themselves as members of a community.

We can better understand differences between the notion of the contrived voice of business and that of an authentic human voice by considering Erving Goffman’s work in *The Presentation of the Self in Everyday Life*. Goffman explains, “We tend to see real performances as something not purposefully put together at all, being an unintentional product of the individual’s unselfconscious... Contrived performances we tend to see as something painstakingly put together, one false item on another” (70). Prepared speech acts, he tells us, seem crafted out of context. They become removed from the moment of inspiration and, therefore, their claims to validity are weakened. Spontaneous speech acts, on the other hand, are made in context and show a connectedness that suggests something more authentic.

John Jones observed something similar in saying of microblogging, “The Twitter stream implied “raw conversation,” or unfiltered information... and therefore seemed to be a more

authentic record” (82). Jannis Androutsopoulos echoes this from the other side, saying that more formal language is “understood as strategically planned and staged, therefore supposedly ‘inauthentic’” (Androutsopolous 76). Mary Bucholtz summarized this idea in saying, “The gold standard of authenticity is the more vernacular speaker at his most casual and unself-conscious” (Bucholtz 414). Considering this, the advice given by the marketing experts above stem from the idea that Twitter’s conversation-like feel and in-the-moment status updates demands a sense of authenticity.

The change that the Cluetrain Manifesto and the previously mentioned marketing specialists are referring to is not limited to corporate marketing. If we turn our attention back to medical communities and the health care industry, we can see similar advice being given there.

Marketing for Medical Communities

In a move which echoes the sentiments from the Cluetrain Manifesto, but as recently as 2019, Lauren Vogel argues that doctors and health care organizations are missing a vital opportunity by being resistant to engaging with their patients on social media. Citing recent trending hashtags like #DoctorsAreDickheads, Vogel points out that people are venting their frustrations about health care over social media. She says, “In the past, patients had very few opportunities to connect... and limited recourse when unhappy about their care. But the balance of power has shifted as social media has enabled conversations and comparisons across social and geographic divides” (E87). She feels that, even in their online spaces, health

care providers are sticking to traditional one-way communication because they see the growing empowerment of patients as a threat.

Similar to the marketing advice above, Jeffery Stevens and Melanie Ross encourage hospitals and other health care organizations to develop their social media presence. In their article, “Social Media: Helping Health Systems Build Empathy and Engagement” they observe that over 40% of health care seekers would turn to social media for help finding a doctor, deciding if they need a second opinion, or gathering advice on how to treat chronic health problems. However, where most marketers focused on developing this authentic voice, the advice Stevens and Ross give to healthcare professionals seems to echo Kabani’s advice on making a connection with their audience. However, for the medical community, developing that natural, open, and honest voice is less about developing authenticity and more about demonstrating empathy for their patients.

To the extent that empathy in online medical communities has been studied, two different studies have linked higher perceptions of empathy with higher levels of participation in those online medical communities. Priya Namsbian observed that those who contribute more to conversations within online health communities are perceived as being more empathetic. In a somewhat circular direction, and more relevant to what I’ve been discussing here, Jing Zhao, et al. showed that higher levels of perceived empathy will encourage higher levels of participation in online conversations about health care.

Bringing this back, full circle, to dana boyd’s concerns, several authors argue that health care providers need to develop their social media presence specifically to combat anti-vaccination rhetoric. Zhongyi Gu *et al.* studied how information was spread about a specific

incident of the mistreatment of a batch of vaccines in China's Shandong province. They found that 78% of the people they surveyed had learned about the incident exclusively from social media and had learned it long before the incident was reported in any professional media source. The anti-vaccination community in China was able to use this spread of information to their advantage through social media. Gu argues that the medical community needs to put more effort into establishing two-way communication between health organizations and the public through social media.

Tying the anti-vaccination rhetoric back to the difference between the “homogenized voice of business” and a more authentic human voice, Neil Johnson *et al.* conducted a study in the spread of different vaccination viewpoints. While most of their study focused on network data and how given networks develop and behave, they made one relevant observation about content; Johnson observed that anti-vaccination messages provided “a wide range of potentially attractive narratives that blended topics such as safety concerns, conspiracy theories and alternative health and medicine” (2). Whereas, messages supporting vaccinations tended to be more monothematic and dry.

Enter 2020

When boyd first published her essay in 2017, the anti-vaccine commotion was already beginning to quiet down just a little. However, as I write this, the movement is seeing a resurgence of activity. The anti-vaccine rhetoric has a great appeal to the anti-mask crowds fighting to keep their personal freedoms prioritized over consideration for the health and safety of their neighbors. Additionally, there has been talk of making the COVID-19 vaccinations

mandatory, and renewed interest in conspiracy theories of government-controlled microchip tracking implants in the vaccines is spreading. The question of how our world's experts – the doctors, scientists, and other professionals – can build their credibility in the face of waves of science denialism remains relevant.

While the theory about an authentic, human *ethos* is compelling, it has yet to be tested. With this in mind, I took a look at tweets from various spokespeople within the anti-vaccine community to see if I could find evidence in support of the theory.

Methods

For this study, I was looking to see if I could find support for the statements made by boyd, Gardner, and others about the shift in how modern audiences were perceiving authority. More specifically, I was going to look for evidence to suggest that social media audiences were reacting more favorably to conversational, authentic styles of writing than they were to authoritative styles of writing within vaccine-related conversations.

In my initial concept for this study, I had thought about writing my own series of microblog statements, showing those statements to a group of survey participants, and asking those participants to rate each tweet based on various statements about the tweet's level of credibility. The statements would have been written to convey similar messages but be written in the different styles I was looking to examine. This approach would have helped isolate the one variable (style) I wanted to look at, and it potentially would have allowed for in-depth, follow-up interviews with the survey respondents.

However, after much consideration, divorcing the microblogging from the richer context of social media platforms seemed too artificial. As established by the studies I discussed in the previous chapter, the at-a-glance elements, like profile information, profile picture, and account name, and other elements like network overlap, reputation, and the greater Twitter conversation context all clearly exerting some degree of influence. If I were to try to simulate those elements, and if I were to find a correlation between language style and credibility, I would not know if that correlation represented a significant influence on a reader's perception of credibility. Isolating the style variable in this way seemed too unrealistic to draw any reliable conclusions from. If I were to try to find this language naturally occurring on Twitter, I would be

in a better position to avoid chasing any ‘false positives.’ I ultimately decided that it would be best to study actual posts within real-world contexts and find a way to gauge reactions from live readers.

The basic steps which I took to design a study to look at social media postings and the favorable or unfavorable reader reactions to those postings, which I will be discussing in this chapter, were:

1. Determining what social media platform I wanted to gather data from.
2. Determining how to measure favorable reactions from readers within the chosen social media platform.
3. Defining the group of authors or microbloggers I wanted to look at.
4. Gathering a set of tweets from those authors or microbloggers.
5. Determining how to distinguish authentic-style tweets from authoritative-style tweets.
6. Categorizing the social media posts based on the writing style.
7. Comparing the reactions of the readers (#2) to the different styles.

A few of the essays I discussed in the previous chapter looked at empathetic language within online medical communities. I’m more interested in what initially may draw people into these groups. If a certain type of language is being used to bring people into science-denialist movements, like the anti-vaccination movement, it’s that first stage of communication I’m interested in. That is why I’ve chosen to look at Twitter. Messages posted through Twitter have the potential to act as a broadcast to people outside of the anti-vaccination movement rather than other forms of word-of-mouth spread through closed medical communities with an already-devoted group of readers.

Measuring Favorable Reactions

In looking at Twitter data or at a tweet itself, there are a few measurements that are clearly visible: likes, retweets and number of comments. Any one of those metrics, by itself, can tell us pieces of how readers are responding to a tweet, but they do not give a complete picture. For a more nuanced look at how audiences are responding to tweets, marketing analysts usually talk about impressions, reach, and engagement rate.

Impressions	Number of times a tweet has been viewed
Reach	Number of unique viewers to have looked at a tweet
Engagement Rate	Number of likes and retweets divided by impressions
Sentiment	Estimation of whether respondent language is positive or negative
“the ratio”	Relationship between the number of likes and retweets compared to the number of comments

Unfortunately, impressions and reach are metrics that are only provided to the author of those tweets, and not a measurement that I could collect. Additionally, there are tools that can perform sentiment analysis on tweets (estimations of positivity and negativity expressed in a tweet), However, I was not able to find any tools which could gather the comments on the set of tweets I was intending on looking at. There are, nevertheless, other ways to approximate these measurements.

For the purposes of this study, I primarily based my judgement of a tweet’s performance on a modified measure of engagement rate. As I did not have access to the number of followers the account had at the time of posting – Twitter does not store that data – I had to

devise a way to approximate engagement rate. To derive an approximation of engagement rate, I looked at the number of likes and retweets divided by the number of followers that account had at this time I gathered my data. Using the number of followers the account had at the time of the data collection would make for a valid approximation of the actual engagement rate as the number of followers for each account should not have changed much within the timeframe I ultimately gathered data from.

As a follow-up measurement, I wanted to look at “The Ratio” (a.k.a. looking to see if a tweet “got ratioed”). “The Ratio” can approximate sentiment analysis by comparing the number of comments to the number of likes and retweets. The theory behind this measurement is that, if a tweet gets exceedingly more comments than likes and retweets, those comments are probably negative. Ratioed tweets are generally those that spark a lot of argument and anger, thus encouraging people to respond without ‘liking.’ While a ratioed tweet can generate a lot of conversation and publicity, I am still looking to see what stylistic elements may make a message more palatable to a potentially hostile audience.

Gathering Data

Since danah boyd specifically mentioned the anti-vaccination movement, I wanted to keep my focus on Twitter users who were active participants within that conversation. In order to compile a set of tweets to look at for this study, I initially wanted to pull tweets from a wide range of Twitter users based on their incorporation of anti-vaccination hashtags. I used a program called Chorus (chorusanalytics.co.uk) to track these hashtags. Chorus was able to

collect data on hashtag use within a three-month lookback period. I began with the following hashtags commonly associated with the anti-vaccination conversation:

- #antivaxx (225 occurrences within the Chorus lookback period)
- #vaxxed (156 occurrences)
- #antivax (152 occurrences)
- #nomandatoryvx (95 occurrences)
- #learntherisk (61 occurrences)
- #vaccineinjury (37 occurrences)
- (I also found that #wakeup was a popular anti-vaccination hashtag, but its use was widespread through multiple non-vaccine-related conversations, so I could not really use it)

In the interest of being thorough, I took a look at the following hashtags which were frequently used within the anti-vaccination and pro-vaccination conversations:

- #vaccines (164 recent occurrences)
- #MMR (80 recent occurrences)

Finally, more as a point of interest or follow-up research, I also looked briefly at these pro-vaccination hashtags:

- #vaccineswork (144 recent occurrences)
- #vaxxhappened (83 recent occurrences) (primarily used in satirical posts)

At the time I'm writing this, the hashtags #COVIDIOTS and #SHEEPL are being heavily used by the anti-vaccination community, but they had not come into such widespread use when I was initially gathering this information.

After searching through the tweets related to all of the anti-vaccine related hashtags listed above, I quickly found that the overwhelming majority of these tweets seemed to be one-off rants from people who did not routinely participate in the anti-vaccination conversation. Without that participation, it was difficult to see ways in which these one-off tweets were either establishing themselves as part of the anti-vaccination conversation or responding to others who were already a part of it. While these one-off rants are certainly worth study, my

focus was intended to be on those considered to be regular spokespeople for the anti-vaccine movement – those who accumulated followers through regular participation in the conversation and who were in a position to more-regularly convert people to being anti-vaccine sympathizers. So, for the purposes of this study, I wanted to avoid those authors with only fleeting involvement with anti-vaccine issues and tweets.

My objective shifted from looking for vaccine-related hashtags to one of looking for users who have been primarily focused on vaccine-related conversations. Focusing on the users who were highly engaged with the conversation would help keep the tweets, themselves, relevant while also allowing me to look at the style of an author across a series of tweets. I began by looking for a few of the big names I knew from the history of the anti-vaccination movement. Andrew Wakefield, the doctor who wrote the paper linking vaccines to autism, did not have his own Twitter account (he was associated with the Vaxxed and Vaxxed2 accounts – those dedicated to promoting the movies of the same titles – but I did not discover those accounts until later in my searching). I then looked at the account of Jenny McCarthy, celebrity spokeswoman for the movement, only to find that her current tweeting is devoted to her reality show, *The Masked Singer*. Jim Carrey's¹ account and Robert F. Kennedy Jr's account were a little better in that they both address political issues in their tweets, but the vaccine issue did not seem to be among them.

Having recently seen the PBS *Frontline* episode, "The Vaccine War," I searched for the people highlighted in that documentary. This led me to the accounts for Jennifer Margulis,

¹ Jim Carrey was dating Jenny McCarthy at the time that her son from a previous relationship, Evan, began showing symptoms of autism. Carrey, at the time, was also a strong voice in the anti-vaccine movement, but since the couple broke up five years later, he has been less vocal about his views.

Barb Loe, and the Generation Rescue organization. I gathered a few more accounts by looking at the people following or being followed by those initial three. I tried to stick to accounts with over 1,000 followers, as that seemed to be about the minimum number of followers to still suggest the author would have any significant influence. However, as I still wanted a wide range of accounts to look at, I gathered many accounts ranging from 1,000 followers to 75,000 (plus two outlier accounts with significantly higher followings - one with 290,000 and one with over two million followers). From there, I started looking at the network overlap of all of these accounts – again, looking for commonalities in who this larger group was following and being followed by. This, along with recommendations from a program called Twittonomy (twittonomy.com)² and, of course, recommendations from Twitter itself, led me to some of the people with larger followings, like Del Bigtree, Sherri Tenpenny and Michelle Malkin.

With a large collection of names and accounts, I reduced the number of accounts I would use for my analysis by eliminating those tweeters who had not posted anything within the past year. I also then eliminated anyone who was not following or being followed by anyone else in the collection. The combination of recent activity plus some degree of network overlap, I felt, would indicate a close involvement with the anti-vaccination community on Twitter and active participation in the anti-vaccination conversation (the only exception to the recent activity rule were the accounts for the movies, *Vaxxed* and *Vaxxed II*, both of which were directed by Andrew Wakefield and produced by Del Bigtree – I felt those were significant enough to keep, even though they have been inactive for a while). I also tried to keep people

² Twittonomy is a program primarily designed to provide various analytics on the followers of an account, mentions, and hashtag performance. One of its secondary functions provided me with suggestions based on my following list.

from a variety of different backgrounds – published authors, political advocates, and concerned parents.

A few of the key figures who were active on Twitter were:

- Michelle Malkin (2.2 million followers) – Fox news contributor and right-wing blogger
- Dr. Joseph Mercola (290k followers) – Osteopathic physician, author, and (according to his own Twitter profile) founder of the #1 natural health site (point of interest: during the time I’ve been looking at these tweets, Mercola has been censored by Twitter multiple times)
- Del Bigtree (75k followers) – Producer of the movies Vaxxed and Vaxxed II and producer and host of The HighWire talk show
- Dr. Sherri Tenpenny (44k followers) – Osteopathic physician and author
- Children’s Health Defense (31k followers) – Organization created by Robert F Kennedy Jr.

Initially, I gathered 37 anti-vaccination tweeters. After eliminating those accounts which were currently inactive, had posted the fewest tweets pertaining to the anti-vaccination conversation, or were suspended by Twitter for violations of Twitter’s code of conduct during the course of my research, I settled on 25 accounts to gather data from (see appendix).

Having selected these accounts for study, I used the Twitter-scraping program, Mozdeh, to compile the latest tweets from those accounts, up to a maximum of 3,200 tweets (a limit imposed by Twitter capacities, not Mozdeh’s). For some of these accounts, the data I collected goes as far back as 2014. However, I’ve been paying the closest attention to the tweets since March, 2020, when talk of making the COVID-19 vaccination mandatory lit a fire under this conversation.

The people and organizations I have chosen to follow here are, of course, dead set against making vaccines mandatory and are rushing to rally people against it. It’s this current,

renewed conversation that I wanted to focus my research on. While boyd's original comments were made in reference to earlier anti-vaccination conversations, flashpoints like this one can become major recruitment drives for the anti-vaxxers. This is where we're most likely to see people being drawn into these movements – certainly those who are potentially opening themselves up to recruitment due to the perceived inconvenience of wearing a mask. Plus, the resurgence of this conversation and the high stakes of the COVID pandemic make current application of this research vital.

Evaluating the Language of the Tweets

Since I was looking to find evidence for boyd's theories about empathy in the medical realm, I initially wanted to sort the tweets I found into the categories of "authoritative" and "empathetic." I recognized that authoritative and empathetic were not necessarily opposite ends of the same spectrum. It is conceivable that a tweet could simultaneously be authoritative and empathetic – something along the lines of, "I know the clear-cut truth, but I still feel your pain." However, I wanted to see what I actually found before making adjustments for that sort of crossover. As it turns out, I needn't have worried.

I knew that in the empathetic group, I wanted to include tweets which asked questions or utilized second person narration or spoke primarily about feelings. However, I would need more than those vague concepts in order to be able to sort through large numbers of tweets. I turned to sociolinguistic studies, looking for additional linguistic markers of empathy I might be able to identify tweets with. Sadly, I was not able to find any concrete markers of empathy. The closest I was able to find was Stirling and Manderson's analysis of the generalized "you" as

used in expressions of structural knowledge as opposed to the conveying of personal narratives. Their analysis, however, was highly dependent on context and did not allow for the use of “you” as a marker without consideration for context.

In a series of three articles from 1998 and 1999, Jenny Preece boasted of a technique for classifying empathetic statements through which she achieved an inter-researcher reliability rating of over 95% agreement. Unfortunately, she never specified what that technique entailed (I even tried emailing her directly, but she never responded). She did mention using three basic qualities of empathetic messages which, ultimately, I traced back to a study by William Ickles (via Levenson and Reuf, 1992). Ickes, in his article, “Empathic Accuracy,” measured subjects’ perceptions of their empathic connections to others. What Preece and Levenson found useful in Ickes work was his division of empathic understanding, empathic expression, and empathic communication (591). For my purposes, this division would be less useful as I’m primarily focused on empathic communication. In addition, his system of having subjects rate their own perceptions of their own levels of empathy felt too arbitrary to use for large numbers of tweets.

What I was able to find, more closely related to the authoritative classification, was a healthy amount of literature on the linguistic markers of certainty and markers of epistemic hedging. Using primarily the work on stance from Gray and Biber, along with the work from Englebretson, with a few minor additions by Kärkkäinen, Izadi, and Pérez-Paredes, I was able to put together a list of linguistic markers that could signify epistemic stance without reliance on context or arbitrary impressions. This struck me as useful as relating to boyd’s comment, “anti-vaxxers aren’t arguing that vaccinations definitively cause autism. They are arguing that we don’t know.” Additionally, epistemic hedging is certainly more of a true opposite end of the

authoritative spectrum than empathy is. I imagined that measuring reactions to tweets that present themselves with certainty versus those that present themselves with hedging could be a logical first step before looking at empathetic and non-empathetic tweets.

Based on the literature, I put together the following lists of linguistic markers:

Markers of Epistemic Certainty

- Adverbials
 - Obviously
 - Certainly
 - Definitely
 - Really
 - Actually
 - Surely
 - Factually
 - Typically
 - In Fact
- Compliment clauses
 - Certain
 - Sure
 - Know
 - Confirmed

• Modals

- Should
- Never
- Always
- Must

• Judgmental absolutes

- Truth
- Lie
- Deception
- Bullshit

• Copulas (when not used as auxiliary verbs in questions)

- Is
- Does
- Are
- Was
- Will (be)

Markers of Epistemic Hedging

- Explicit softening
 - Kind of
 - Sort of
- Adverbials
 - Perhaps
 - Possibly
 - Probably
 - Likely

• Compliment Clauses

- Don't know
- Not sure
- Think
- Doubt
- Wonder
- Expect

• Modals

- Might
- May
- Could
- Can
- Seems
- Would

Using these lists, I took a selection of 50 tweets from my collection, and I made note of all of the markers I found within them. After finding these markers and looking at how they were used, it became apparent that their usage was highly dependent on context, and I could not classify tweets as certain or hedging based solely on their presence within a tweet. This led me back to content analysis as the next best option, despite my initial hesitancy.

What my foray into working with epistemic hedging did do for me was that it convinced me that I don't really want to restrict myself to talking about empathetic as the alternative to authoritative. The lack of definitive stance now seemed just as important for analyzing stylistic elements that may make a tweet appealing to potentially hostile audiences, along with encouraging participation and that authentic voice that so many others were talking about. Faced with the allowances content analysis would afford me, I started to rethink what that opposite end of the spectrum should look like and how I could get at the true spirit of what has been discussed while incorporating all of the necessary pieces. I began thinking more along the lines of language that would set up more of a conversational tone – one that encouraged the sharing of thoughts and ideas while still sounding authentic. I finally settled on using the distinction between authoritative and “dialogic.”

Here is how I ultimately defined those terms:

- Authoritative - The language being used suggests the author is certain of their position, there is only one valid viewpoint, and the author is the only legitimate source of the viewpoint. The author assumes that viewpoint applies to all people equally.

- Idealized examples of my concept of this type of language might be – Vaccines work.

You have no reason to be concerned about their safety. Everyone should get the vaccination regardless of past experiences or issues.

- Actual examples from my data might include: “Do we need vaccines? No, we do not. Get an in-depth understanding of how and why our immune system can handle infectious disease” (@stopvaccinating Sept 17, 2020), or “Universal mask mandates are a sham. No more fines. No more arrests. No more orders” (@michellemalkin Sept 25, 2020).
- Dialogic - The language being used elicits, encourages, or rewards others' participation in a discussion, fosters the give-and-take of ideas, and recognizes contingencies and complexities of an issue, suggesting that what is correct for one may not be correct for others.
 - Possible, idealized examples of this type of language could be – Vaccines should be safe for most people, but if you have any concerns we should discuss them. Thank you for sharing your personal stories.
 - Actual examples from my data may include: “Maybe unsafe and ineffective? 🤔 no one knows” (@eileeniorio Sept 29, 2020), or “Sorry for your loss sister. Thank you for sharing your story. Blessings to you” (@uTobian Sept 25, 2020).

The authoritative/dialogic classification best captured the key spirit of an authentic voice, empathetic expression, a conversational approach, epistemic hedging and their collective distinction from authoritative language. This two-category approach helped me capture the primary distinction I was concerned with while not over-complicating the sorting process.

Were I to have found evidence suggesting differences in how favorably readers were responding to those two approaches, I could break them down further in follow-up studies.

Verifying Inter-Rater Reliability

I began with 1000 tweets (40 most recent tweets from each of the 25 accounts I was following), but I had to throw 43 of the tweets out due to issues with the tweets themselves (3

were in French and the 40 tweets from GenRescue were all devoted to promoting their charity poker tournament - My readers were all in total disagreement as to how to handle those), leaving me with 957 tweets.

When I went through the tweets myself, I was primarily concerned with the tweets I've been referring to as "authoritative" (1) and "dialogic" (2). In going through these tweets, it became apparent that I had to include two additional categories for tweets that could not be legitimately classified as either authoritative or dialogic. Those categories became "personal" (3) – shout-outs to friends or "what I had for breakfast" types of tweets - and tweets which only consisted of a link with little or no commentary (4).

After categorizing the list of tweets myself, I asked three other people to go through the exercise of sorting the same tweets into those four categories in order to verify if these categories could be applied to the language of the tweets with a reasonable amount of consistency by different readers. In order to verify that the categories and language I was using were clear on their own and that this type of sorting could be repeatable, I kept my instructions to my readers to a minimum. I only provided them with the exact definitions of the four categories I have above, and did not provide any examples. This is exactly what was given to them for instructions:

	Directions: For each tweet, decide if the language being used seems authoritative or dialogic. Mark your evaluation in column E
1	Authoritative = The language being used suggests the author is certain of their position, there is only one valid viewpoint, and the author is the only legitimate source of the viewpoint. Assumes that viewpoint applies to all people equally.
2	Dialogic = The language being used elicits, encourages, or rewards others' participation in a discussion, fosters the give-and-take of ideas, and recognizes contingencies and complexities of an issue (not one-size-fits-all)
3	Personal comment - does not seem to take an authoritative stand nor does it encourage a discussion.
4	A link with little or no commentary by the person sharing it.

Before getting into the results, there were a couple problems I need to mention. First, I noticed a problem with two of my readers overusing categories 3 and 4. For the first 200 tweets, my 1st reader assigned a 4 to any tweet which included a link – even when there was ample commentary by the person sharing the link. After that first 200, I spoke to reader #1 and clarified what that category was meant for. I did not ask him to go back and reevaluate those 200, but I made sure he would proceed with a better understanding. From that point forward, reader #1 seemed to use category 3 as a catch-all, even for many tweets that my other readers classified as 1s or 2s. Reader #2 seemed to develop a similar problem with overusing category 4, primarily in the 2nd half of the list.

Here is a quick summary of the results: Of the 957 tweets, all four readers (myself included) agreed on 46% of them. Of that 46%, 358 tweets were authoritative, 16 were dialogic, 24 were personal, and 44 were just a link with little or no commentary. An additional 33% were tweets that three readers agreed on with one reader disagreeing. In that 33%, the majority group found 241 authoritative tweets, 26 dialogic tweets, 31 personal tweets, and 17

link-only tweets. Beyond that, we had a 2 and 2 reader split for 7% of the tweets, a 3-way split for 13%, and a 4-way split for less than a quarter of a percent.

	Total	Subset	Percentage
Full agreement	957	442	0.46186
Authoritative	442	358	0.809955
Dialogic	442	16	0.036199
Personal	442	24	0.054299
Link only	442	44	0.099548
3 Reader agreement	957	315	0.329154
Reader 1 disagreed	315	168	0.533333
Reader 1 disagreed (w/o 3s & 4s)	315	58	0.184127
Reader 2 disagreed	315	82	0.260317
Reader 2 disagreed (w/o 4s)	315	44	0.139683
Reader 3 disagreed	315	45	0.142857
I disagreed	315	20	0.063492
Authoritative	315	241	0.765079
Dialogic	315	26	0.08254
Personal	315	31	0.098413
Link only	315	17	0.053968
2 and 2 split	957	69	0.0721
3 way split	957	129	0.134796
4 way split	957	2	0.00209

Due to the complexity of the language I'm looking at and the potential overlap of the categories, it would have been unreasonable to expect 100% agreement between the four readers. So, I counted every time we had agreement between three out of the four readers as a valid categorization. In other words, if three of the readers rated a tweet as a 1 (authoritative), then I will be accepting that categorization. So, that gave me valid categorizations for 80% of the tweets I had the readers look at – 757 tweets in total. Based on this 80% agreement rate, I felt secure in considering “authoritative” and “dialogic” to be valid and workable categories for this study.

Getting Data from Machine Learning

At this point, I was able to look at the data from those first 957 tweets, and I will discuss some of these preliminary observations in the next chapter. However, my ultimate goal was still to get ratings for my entire list of the 57,533 tweets I had collected. In order to do that, my advisor, Dave Clark, and I were going to use the 957 tweets (599 authoritative, 42 dialogic, 316 “other” – where “other” was a combination of link only tweets, personal tweets, and the tweets I did not get agreement on) to teach a computer to rate the tweets. Dave initially recommended using a “Bag of Words” (BoW) method he was familiar with, where the computer would learn to recognize the distinction between authoritative and dialogic tweets based on associations with common words it found in the sample we gave it.

To set this up, we used LogisticRegression from the scikit-learn library. With this BoW method, we would feed the 957 tweets rated by my readers into the computer along with the classifications assigned to them by my readers. The machine learning system would break down the tweets and organize the data based on what significant words it found in each of the categories along with the frequency with which those words were used in each of the categories (more on how we defined significant in a moment). The machine learning would then use that breakdown to create an algorithm it would be able to apply to future tweets we would give it in order to determine which category the new tweet would most likely fall into. It’s important to note here that the BoW method only looks for the presence of these significant words; it does not consider word positions within a sentence or relationships to other words. When looking at a new tweet, the computer would compare the words it found in the new tweet to those in groups it learned from. Applying the algorithm it created from the

learning set of 957 tweets, the system would determine the likelihood that the new tweet would be classified as an authoritative tweet or a dialogic tweet.

The results of this calculation would be given as a certainty rating on how strongly it associated the wording of the new tweet with authoritative and dialogic tweets. In other words, it would give each new tweet an authoritative score and a dialogic score. These certainty ratings (or scores) would be on a scale between 0 and 1 – a 0 showing no similarities to that category of tweets and a 1 showing great similarities (the authoritative certainty rating and the dialogic certainty rating would not necessarily add up to an even 1, but I found that to be a good way to conceptualize how these scores might work). Based on this, we might classify any tweet with a higher authoritative certainty rating than a dialogic certainty rating as authoritative and vice versa. However, this also allowed me to look closely at any tweet receiving a high dialogic rating, even if that rating was not higher than the tweet's authoritative rating – we would effectively have access to a graduated scale of authoritative and dialogic strength rather than a simple binary. Some examples that would eventually come from my data looked like this:

- The removal of stop words
- The removal of capitalization
- The removal of punctuation

First, we removed all user names from the tweet data. This was so the machine would not automatically base an evaluation with the association of a user name. So, if @eileeniorio wrote 6 dialogic tweets and 29 authoritative tweets in our learning set, the system would not automatically associate all of her tweets with the authoritative set (as I will discuss in the next chapter, the overwhelming majority of everyone’s tweets were authoritative – had we included user names, all of the tweets we had the computer rate would automatically be high on the authoritative scale).

We also converted all emoji’s into text, so they could be used in assigning ratings. The assumption there was that something like the heart emoji (translated to “â¸ [sic]”) could be commonly used in dialogic expressions, whereas other emojis, like the clapping emoji (translated to “ðŸ [sic]”) could be commonly used in authoritative expressions – especially when punctuating words like @sallykp’s tweet, “NOT👏ABOUT👏A👏VIRUS Wake up world” (Sept 27, 2020).

Then, we removed all function words. These words are referred to as “stop” words in the field of language processing for machine learning. This group consists of commonly used determiners, coordinating conjunctions, and prepositions. The system would then ignore words like “an,” “the,” and “and,” in order to maintain the focus on more significant words.

I debated a little about the removal of capitalizations. While the use of all capital letters in a tweet, as a commonly accepted textual representation of yelling (as in @sallykp’s tweet

above), would be a clear indication of an authoritative approach, I felt that having the computer identify “virus,” “Virus,” and “VIRUS” as three distinctive words would be overly complicated in a way that might reduce the relative importance of the presence of those words.

Finally, we eliminated punctuation. While questions, themselves, may be suggestive of dialogic statements, and exclamations may be suggestive of authoritative statements, we felt the simple use of a question mark or an exclamation point should not necessarily lead to a rating one way or the other. For example, many of the tweets my readers classified as dialogic included phrases like, “Thanks” or “Reach out to me if...” followed by an exclamation mark. Conversely, there were tweets like one from @stopvaccinating on September 17, 2020 saying, “Do we need vaccines? No, we do not,” where the use of a question was determined to be done in an authoritative way. We wanted the computer to focus on how the question or exclamation was phrased and not on the punctuation itself.

Having programmed in these considerations, we were ready to run a test set. Our first attempt at this did not provide any usable results. We fed the 957 tweets in and then ran a trial sorting with another 1000 tweets for the computer to rate. For the results we did get back, the computer rated everything as primarily authoritative – it returned no tweets for which its dialogic certainty rating was higher than its authoritative certainty rating. Since, in the initial set of 957 tweets, only 42 of them were examples of dialogic tweets, this did not provide a sample large enough for the computer to really learn what a dialogic tweet looked like. So, I went ahead and hand-coded 2,500 additional tweets. In that batch, I identified an additional 109 dialogic tweets. Having 151 examples of dialogic tweets in this second run gave the system

more to work with than in the first run with only 42. As we discovered, this difference was enough to allow the machine learning to better identify dialogic tweets.

Our second run, based on the set of nearly 3,500 tweets (the initial 957 plus the 2,500 newly rated tweets), gave us much better results. In a test of 2,500 new tweets, the computer returned 239 tweets for which the dialogic certainty rating was higher than the authoritative certainty rating. As an alternative way of looking at the data, it returned 130 tweets to which it gave a dialogic strength rating above 0.5. Based on those numbers alone, this seemed promising, but I was sure to take a closer look at the data to be sure. I carefully reviewed the 130 tweets that received a dialogic certainty rating above a 0.5, and I saw some promising results. I did not feel the need to review the highly rated authoritative tweets as we already believed the overwhelming majority of the tweets would be authoritative.

Before I get to what made the results promising, I'll say that the computer rating was not perfect. I did want to take note of some of its quirks or apparent inaccuracies as issues to be aware of. The only real inaccuracy I saw was something that I had mentioned also tripped up my reliability readers – rhetorical questions. For example, on September 7th, @stopvaccinating tweeted, “What’s really driving society to vaccinate, vaccinate, vaccinate? One thing is clear: it’s not for our health.” The machine learning classified this as dialogic by a wide margin; it gave this tweet a 0.945 dialogic certainty rating and 0.155 authoritative certainty rating. The first sentence in that tweet certainly sets it up as one that should encourage a response. However, the follow-up of, “it’s not for our health” belies any apparent invitation for discussion. @stopvaccinating clearly believes that society is vaccinating for economical and political reasons. They only phrased this tweet as a question to show an

apparent concern for health as a lie. Similarly, On August 25th, @HighWireTalk asked, “Is the #Massachusetts flu shot mandate a foreshadowing of what’s to come if a COVID vaccine becomes available?” This was in reference to Massachusetts’ mandate from August 19, 2020, that all students in child care, K-12, and postsecondary schools get the flu vaccine (mass.gov). Again, this tweet may have been phrased in the form of a question, but Del Bigtree, host of the High Wire talkshow, seems to be expressing more of a definitive opinion than inviting discussion. In this case, the computer gave this tweet a dialogic certainty rating of 0.903 and an authoritative certainty rating of 0.385. Considering this issue of rhetorical questions was an issue for human readers as well, I don’t believe it would be realistic to expect any other forms of computer learning to be able to distinguish a rhetorical question from a question actually intended to encourage an open dialogue.

The issue with rhetorical questioning aside, I was encouraged in that I found several tweets which I felt were strong examples of dialogic language which had been captured by the machine learning. Most of these were ones which included the phrases, “I’m so sorry” or, “Thanks for sharing.” Another example would be like @stopvaccinating’s tweet from August 30th, asking, “Has a pediatrician ever bullied, harassed, or demeaned you for asking about vaccine safety or for refusing to vaccinate?” This tweet both seems to invite participation and involvement by encouraging a response and carries within it a sense of empathy in expressing concern over a situation many of their followers may have been upset by. By doing both of those things, a tweet like that fits right into what I was looking to capture with the dialogic category, so it was reassuring to see the BoW method picking it out. In our test run, this tweet scored a dialogic strength of 0.909 and an authoritative strength of 0.260.

Based on these test results, I was confident in using the BoW method. There were a couple other vector approaches we could have looked at – those that looked at other elements of sentence construction. I was curious about those methods, and we began working on putting a couple of them together. However, after a few complications with programming and time delays, we decided to proceed with gathering the data through the BoW method. We ran the 53,908 tweets we had not used for the learning through the computer and got ratings back for all of them. I will discuss these results in the next chapter.

Data and Results

Data from the Preliminary Tweets My Readers Evaluated

Upon receiving the first 957 tweets from my readers, I did some analysis as an estimation on what I might find through the rest of the data I took. In doing so, I found a few noteworthy things that I knew I would have to look for when analyzing the rest of the data.

The first thing that I noticed was the large volume of authoritative tweets compared to a surprisingly limited number of dialogic tweets. Of the 957 tweets my readers looked at, based on their responses, I was able to categorize 641 of them as either authoritative or dialogic (116 fell into a 'neither' category which included personal comments and shared links without any original comments, and the remaining 200 we did not get a consensus on). Of those 641 tweets, my readers categorized 599 of them as authoritative and only 42 of them as dialogic. Based on this, the spokespeople for the anti-vaccination movement utilize authoritative language on Twitter quite a lot more than they use dialogic language. I was immediately struck by the evidence that, if we are seeing this "seismic shift" in displays of trustworthy *ethos*, where authority and objectivity have been replaced by authenticity and transparency, it's either that Twitter is not a place where we're seeing it, or that the key members of the anti-vaccination conversation are just not putting expert advice into practice on their tweeting (this is on the assumption that they're being exposed to professional marketing advice). Either way, this already seems to suggest that boyd's theory that dialogic language is a large factor in attracting members to the anti-vaccination community is either unfounded or is happening in places

other than Twitter, which I thought of as a natural platform for the initial contact with larger audiences.

In looking at who was responsible for the 42 dialogic tweets, there seemed to be a fairly even spread of dialogic language throughout the accounts I was following. The detailed breakdown is as follows:

Name	Followers	#Dialogic in 50	#Authoritative in 50	Notes
Michelle Malkin	2,200,000	1	25	Fox News contributor - Censored by Twitter during research
Dr Joseph Mercola	290,500	1	34	Commercial site - Censored by Twitter
Frank Lipman MD	75,400	0	37	Author
Del Bigtree	47,600	1	28	Producer
Dr Sherri Tenpenny	43,800	1	21	Osteopathic Dr, Author, Speaker
LotusOak	40,100	0	40	Account suspended by Twitter during research
The HighWire	37,800	1	28	Created by Del Bigtree
Childrens Health Defense	31,000	1	29	Created by Robert F Kennedy Jr
Toby Rogers PhD, MPP	19,400	5	20	Conspiracy theorist - Account suspended during research
Generation Rescue	18,500	0	0	Frontline - Non-profit – I removed this from reader results due to self-promotional tweets
Barb Loe, NVIC	16,300	3	21	Frontline - Non-profit - Account suspended by Twitter during research
Larry Cook	14,300	3	21	"Healthy Lifestyle Advocate" - Account suspended by Twitter during research
sally	11,400	3	13	Account suspended by Twitter during research
Physicians for Info	11,400	1	24	Non-profit
Inside Vaccines	10,900	3	28	Link to MeWe - Account suspended by Twitter during research
Vaxxed II: The People's Truth	10,700	0	21	Produced by Del Bigtree and Andrew Wakefield
Jefferey Jaxen	10,300	0	32	Journalist

Eileen Iorio	10,100	6	29	Author
Catie	7,500	1	18	Grieving mother
Noforcedvaccination	5,700	0	18	
Jennifer Margulis	4,900	1	27	Frontline - Journalist, Author
Vaxxed_Supporter	3,100	1	20	Grieving mother
Truth Lover	3,000	3	27	
Wayne Rohde	2,900	3	20	Author
One Pissed Off Mom	2,800	3	13	

There do seem to be minor concentrations of the number of dialogic tweets in accounts with fewer followers. The accounts with more than 20,000 followers (the top 8) tweeted 11 of the dialogic tweets. The accounts with between 10,200 and 20,000 followers (the middle 8 – not including Generation Rescue) were responsible for 13 of the dialogic tweets. The accounts with fewer than 10,200 followers (the bottom 8) posted 18 of the dialogic tweets. If this trend remains consistent or more pronounced throughout the remainder of the data, it could suggest that authoritative authors might gather larger followings or it could suggest a natural evolution in the type of language authors use as their following grows.

Moving on from these initial observations, I calculated the approximate engagement rate for all of the 957 tweets in this initial sample (the number of likes and retweets the tweet received divided by the number of followers the account had). As a reminder, the engagement rate, as I was calculating it, was the number of likes a tweet received added to the number of times that tweet was retweeted, and that total was divided by the number of followers the account had. I say this is estimated because I'm using the number of followers the account had at the time I pulled my data as opposed to the number of followers the account had at the time the tweet was posted. For example, on September 27, 2020, @uTobian tweeted, "The U.S.

vaccine schedule is savage and barbaric. It is not based in science. It has nothing to do with health.” This tweet received 2,985 likes and was retweeted 1,053 times. @uTobian (Toby Rogers) has 28,821 followers. So, this tweet has an estimated engagement rate of 0.1401 ($2,985 + 1,053 = 4,038$, and $4,038 / 28,821 = 0.1401$). Just for reference, this is well above the average engagement rate for this batch of tweets (0.016) and was one of the top performing tweets of this group.

In performing those calculations, I immediately noticed a small number of tweets with a ratio considerably larger than one. This means that they would have gotten more likes and retweets than they have followers. Ratios slightly above one are certainly possible because a single reader may like and retweet a post (so they would be double counted), and followers of followers may see and respond to a tweet (and would not be counted as a follower of the author). However, some of these clear outliers went to a ratio as high as 21. I looked more carefully at some of these tweets and noticed that they were all actually retweets and Mozdeh was counting the number of likes and retweets the original tweet received – sometimes accounting for many more than the number of followers the retweeter had. For example, on April 28th, @truthvaxwarrior tweeted, “Under no circumstances will I or my family be getting a #coronavirus #vaccine. We have suffered enough from the damage they have inflicted upon us. #exvaxxer,” and my data showed that this tweet had an extremely high engagement rating of 21.1691. When I looked at that tweet on Twitter, it turned out that it had been a retweet from Candace Owens (@RealCandaceO), a writer for the New York Times with three million followers, that she tweeted the day before. Her original tweet received nearly 75 thousand likes and retweets. When compared to Owens’ 3 million followers, this would give us an

engagement rate of about 0.025, whereas, when compared with @truthvaxwarrior's 3,512 followers, it presents an inaccurate picture. The retweet itself was retweeted again three times and received no additional likes. I could not determine why Mozdeh pulled in data from the original tweet instead of just the retweet. After looking at many of the retweeted posts, I found five posts where this seemed to be an issue – ones with engagement rates between 1.5 and 21 – and I removed them as outliers.

Once I had reliable engagement rate approximations, I calculated averages for the authoritative tweets and the dialogic tweets. For the authoritative tweets, I found an average engagement rate of 0.011753799 with a standard deviation of 0.034042563 and a standard error of 0.001396783. For the dialogic tweets, I found an average engagement rate of 0.004517618 with a standard deviation of 0.012322261 and a standard error of 0.001901366.

Population	Sample Size	Mean Engagement	Standard Deviation	Standard Error
Auth - Engagement	594	0.011753799	0.034042563	0.001396783
Dialog - Engagement	42	0.004517618	0.012322261	0.001901366

A two-sample *t*-test was run as follows:

H_0 : The mean engagement rate for Auth = the mean engagement rate for Dialog

H_a : The mean engagement rate for Auth > the mean engagement rate for Dialog

Conditions checked:

- Two independent random samples were obtained.
- Both samples were large enough for the Central Limit Theorem to apply

Statistics obtained from the data:

- $n_{Auth} = 594$, $\bar{x}_{Auth} = 0.011753799$, $s_{Auth} = 0.034042563$
- $n_{Dialog} = 42$, $\bar{x}_{Dialog} = 0.004517618$, $s_{Dialog} = 0.012322261$

Calculations obtained from the two-sample t -test:

- $df = 95.28$
- $t = 3.067$
- $P = 0.0014$

With this small a P -value, 0.0014, there is very strong evidence to reject the null hypothesis that the average engagement rate for authoritative tweets is the same as the average engagement rate for dialogic tweets. Instead, there is convincing evidence that the true mean engagement rate for Authoritative tweets is GREATER than the true mean engagement rate for Dialogic tweets.

One issue I noticed that may have been a factor in the low mean for the dialogic tweets was that, of the 42 dialogic tweets, 14 of them had engagement rates of 0 – meaning they received no likes or retweets at all – whereas of the 594 authoritative tweets, only 30 had engagement rates of 0 – a considerably lower ratio. In looking closer at some of those zero-engagement rated tweets, I saw that nearly all of them began with the @ symbol – meaning they were replies or were otherwise directed at a specific person. This made sense as fewer people may have taken note of those tweets, and person-to-person communications may naturally be more dialogic in nature. An example of this may be @1pissedoffmom1's tweet from September 26, 2020, "@RelevantMom Where do you think she stands on v mandates?" This tweet was directed primarily to @RelevantMom and may not have been seen by others. To see how much this may have affected my calculations, I recalculated the mean engagement rate for the dialogic tweets without including the zeroes. Without those zero-rated tweets, the mean engagement rate for the remaining dialogic tweets was 0.006776428. This would

certainly be better than the original 0.004517618, but not nearly enough of a difference to bridge the gap between the dialogic engagement rate and the 0.011753799 authoritative engagement rate.

As a follow up, I wanted to look at the relative engagement rates for authoritative and dialogic tweets from just a single author. I chose to look at Eileen Iorio because she had the most non-zero-rated dialogic tweets (4). She also had 24 non-zero-rated authoritative tweets. The mean engagement rate for her authoritative tweets was 0.020725968, and the mean engagement rate for her dialogic tweets was 0.00021179. So, even just looking at the one user who had the most dialogic tweets within this small starting sample of my data, her authoritative tweets greatly outperformed her dialogic tweets.

Data From the Full Dataset Run.

Once my full dataset of 53,908 tweets had been run through the machine learning, it provided some interesting results. Before I made any calculations, however, I removed the outlier tweets – those that had had engagement ratings high enough to suggest that Mozdeh, the program I was using to collect the data, was giving engagement credit to a retweet which was rightfully due to an originating account with a much larger following – as I had done with my preliminary set. In this case, I wanted to be just a little more conservative, so I removed 372 tweets with ratings above 1.25. Especially for the smaller accounts, ratings at or above 1.0 were still possible, if unrealistic. I considered dropping my definition of an outlier down to 0.5. There were only an additional 198 tweets with engagement ratings between 0.5 and 1.25. I

decided to keep them in to make sure I was giving proper consideration to the smaller accounts which were most likely the accounts those ratings were attached to.

While I was excited to now have a more detailed scale of the levels of dialogic and authoritative language, I wanted to start out looking at the numbers in the way that I looked at the initial set of tweets my readers first evaluated. To do this, I sorted all of the tweets into the simple binary of dialogic or authoritative – any tweet which received a higher dialogic certainty rating than authoritative rating was labeled as “dialogic” and any tweet which received a higher authoritative certainty rating than dialogic rating was labeled as “authoritative.” After sorting them like this, I ended up with 46,226 authoritative tweets and 7,309. So, much like I found in the preliminary set, the overwhelming majority of the tweets seemed to use more authoritative language than they used dialogic language. However, when I calculated the average engagement rates for those groups, it told a very different story from what I saw in my preliminary set. The average engagement rating for the authoritative group was 0.009363, and the average engagement rating for the dialogic group was 0.009754. So, the dialogic set averaged 0.000391 higher than the authoritative set. This suggests that, as a group, the dialogic set performed better than the authoritative set.

Population	Sample Size	Mean Engagement	Standard Deviation	Standard Error
Auth - Engagement	46226	0.00936303	0.0557853	0.00000121
Dialog - Engagement	7309	0.009754139	0.064549578	0.0000088315

From there, I wanted to make use of the certainty rating as a scale, rather than just use the simple authoritative/dialogic binary. I could not simply create a scatter graph, as with over 50,000 points of data, it would just show as a solid block. So, I took the average engagement

rate for all tweets which scored a dialogic certainty rating between 0 and 0.1, between 0.1 and 0.2, and so on. Here's what I found:

RatingBase - Dialogic	RatingCap - Dialogic	# of Tweets	Avg Engagement
0	0.1	30550	0.008673318
0.1	0.2	13071	0.010933905
0.2	0.3	3551	0.009411627
0.3	0.4	1964	0.010627836
0.4	0.5	1523	0.009090056
0.5	0.6	861	0.011118189
0.6	0.7	710	0.00839628
0.7	0.8	589	0.009948697
0.8	0.9	449	0.009578728
0.9	1	268	0.008948454

Plotted out on a graph, the data looks like this:




Figure 1

Looking at this information, it did not seem as though the dialogic certainty rating showed a correlation with engagement rating. I then repeated this process based on the authoritative certainty rating. Here is how those numbers turned out:

RatingBase - Authoritative	RatingCap - Authoritative	# of Tweets	Avg Engagement
0	0.1	2162	0.006876904
0.1	0.2	8113	0.012919839
0.2	0.3	4943	0.008719509
0.3	0.4	4307	0.010969693
0.4	0.5	3851	0.007638796
0.5	0.6	3492	0.007992396
0.6	0.7	3763	0.007910633
0.7	0.8	4027	0.008104289
0.8	0.9	5355	0.007877681
0.9	1	13523	0.009773655

Plotted out on a graph, the data looks like this:




Figure 2

Once again, just at a glance, the data does not seem to suggest any real correlation between authoritative strength rating and engagement rate.

For one final way of looking at the data visually, I sorted the tweets based on engagement rating. I then calculated the average authoritative and dialogic certainty ratings for those groups. For all of my tweets combined, they had an average engagement rate of

0.00941642579789089 and a standard deviation of 0.00325568968699088. I rounded these to 0.009 and 0.003 and sorted the tweets into groups based on standard deviations away from the mean. For engagement rates below the mean, I created groups of up to 1 standard deviation below the mean, up to 2 standard deviations below the mean, and up to three standard deviations below the mean. This brought me down to zero. For groups above the mean, I tried to keep them as uniform as possible, creating the groups of up to one standard deviation above the mean, up to two standard deviations from the mean, 3 to 4 standard deviations from the mean, 5 to 7 standard deviations from the mean, 8 to 17 standard deviations from the mean, 18 to 50 standard deviations from the mean, and then anything above 50 standard deviations from the mean. Broken down in this way, this is the data I had:

Group	RatingBase - Engagement	RatingCap - Engagement	# of Tweets	Avg Auth	Avg Dialog
-3 SV	0	0.003	41207	0.553232924	0.140842915
-2 SV	0.003	0.006	4677	0.677777144	0.101788707
-1 SV	0.006	0.009	1876	0.687114997	0.100833834
1 SV	0.009	0.012	1158	0.694974994	0.108212208
2 SV	0.012	0.015	733	0.676054526	0.115334202
3 - 4 SV	0.015	0.021	851	0.676475383	0.117309734
5 - 7 SV	0.021	0.03	765	0.653425862	0.117512684
8 - 17 SV	0.03	0.06	898	0.656796524	0.114204028
18 - 50 SV	0.06	0.159	751	0.614753716	0.138663988
50+ SV	0.159	1.25	620	0.505086921	0.151336703

Plotted out on a graph, the data looks like this:




Figure 3

This graph does suggest a couple points worth noting. First off, we should keep in mind that this does not demonstrate that authoritative tweets performed better, it mostly shows that all tweets tended to have higher authoritative certainty ratings than they did dialogic ratings. Having already determined that 46,226 of these tweets had higher authoritative ratings than dialogic ratings, this was no surprise. What did seem significant to me was how the two lines seemed to converge below two standard deviations below the mean and then again, more slowly, at values above the mean, converging more quickly at values far above the mean. Looking at the data this way suggests that there could be correlations between tweet performance and the certainty ratings – as we move up the scale of engagement rate, the dialogic certainty ratings seemed to increase slightly and the authoritative certainty ratings seemed to decrease slightly. This convergence in the certainty ratings was also true for the tweets with engagement ratings approaching zero.

In order to make better sense of the data I was seeing, I shared these numbers with Dr. Martin Sternstein, a mathematics professor at Ithaca College in New York (and my uncle). He

took the calculations further and performed linear regression t -tests as done with the preliminary data above, and he found that the data did show a small but definite correlation between authoritative and dialogic certainty ratings and the tweet's engagement rate.

In the following analysis,

- A = Authoritative Certainty Rating
- B = Dialogic Certainty Rating
- C = Engagement Rate

Here is the analysis he provided me with, in his words:

Using your data and running a linear regression t -test of C against A yields:

- Predicted $C = 0.01054 - 0.0001963A$
- Correlation $r = -0.342$
- Coefficient of determination $r^2 = 11.7\%$
- P -value = 0.000
- 95% confidence interval for the slope: $(-0.00020, -0.00019)$

With this small a P -value, there is strong evidence of a linear relationship between C and A . However, only 11.7% of the variation in C can be explained by this linear model. Also remember that a linear relationship does not imply causation.

Using your data and running a linear regression t -test of C against B yields:

- Predicted $C = 0.00913 + 0.0001904B$
- Correlation $r = +0.320$
- Coefficient of determination $r^2 = 10.2\%$
- P -value = 0.000
- 95% confidence interval for the slope: $(0.000185, -0.000195)$

With this small a *P*-value, there is strong evidence of a linear relationship between *C* and *B*. However, only 10.2% of the variation in *C* can be explained by this linear model. Also remember that a linear relationship does not imply causation.

It is interesting to note that the correlation between *C* and *A* is negative, while the correlation between *C* and *B* is positive. That is, a one-unit increase in *A* predicts a 0.0001963 decrease in average *C*, while a one-unit increase in *B* predicts a 0.0001904 increase in average *C*. [Or since the *A* and *B* values are always between 0 and 1, you could say that a 0.1 increase in *A* predicts a 0.00001963 decrease in average *C*, while a 0.1 increase in *B* predicts a 0.00001904 increase in average *C* (sent to me via email on 12/10/21).

The key observations here are that there is a positive correlation between dialogic certainty rating and engagement rating and a negative correlation between authoritative certainty rating and engagement rate, however this correlation only accounts for a small portion of a tweet's performance - approximately 10 and 12 percent, respectively.

With a coefficient of determination of about 10% for each correlation, this tells us that, at most, the use of dialogic or authoritative language can only account for 10% in observable differences in engagement rate. That does not, however, mean that dialogic tweets should be earning 10% higher engagement ratings than authoritative tweets. More accurately, using the slope of the linear relationship, as explained above, "a 0.1 increase in *B* predicts a 0.00001904 increase in average *C*." Since *C*, engagement rate, represents the number of likes and retweets divided by the number of followers an account has, this means that a 0.1 increase in dialogic certainty should be accompanied by 0.00001904 of a like for every follower an account has.

To put that into perspective, for an account with 30,000 followers (a rough average of the accounts I was looking at, if we ignore Michelle Malkin who, with 2.2 million followers, has more than double the number of followers of all other accounts I was following combined), that would be a little more than half of a like. If such an account posted a tweet that would be rated a 1.0 in dialogic certainty (highest possible) and another tweet which would have been rated a 0.0 in dialogic certainty (lowest possible), the more dialogic tweet could be expected to receive 6 more likes than the other. Since most tweets scored in the mid-ranges of dialogic or authoritative certainty, a deliberate effort to make a tweet more dialogic would more likely result in a 0.6 or 0.7 increase in dialogic certainty. That much of an increase would only correspond to 3 or 4 additional likes. For reference, the average number of likes and retweets all for the tweets I collected (not including the outlier retweets) was 510 per tweet. So, a difference of 3 or 4 hardly seems significant.

In short, these findings suggest a positive but arguably insignificant relationship between the use of dialogic language and tweet performance.

Tweet-Level Analysis

The anti-vaccine movement is certainly not new. The crusade against the practice of inoculation even precedes the first official vaccination. Azhar Hussain cites the writings of Reverend John Williams in 1721 and a sermon by Reverend Edmund Massey in 1772 as both speaking out against the practice of inoculation, depicting it as an attempt to defy God's will and avoid the divine punishment of disease. Both of these instances predate the world's first vaccination by Doctor Edward Jenner of eight-year-old James Phipps in 1796 (Stern 611).

Anti-vaccination sentiment first moved into the political realm in London when compulsory vaccination laws were passed in 1821 (Stern 617). Britain's working class saw this as a violation of their privacy and of their bodily integrity, so the Anti-Vaccination League was formed (Hussain 2). Many years later, in 1898 the League won a major victory when the British Parliament was forced to remove all penalties for non-compliance with the vaccination laws (Hussain 2).

More recently, the anti-vaccine movement was revitalized by Andrew Wakefield's infamous paper in *The Lancet* which linked the MMR (measles, mumps, and rubella) vaccine to the development of autism. Wakefield's theory was that the onset of symptoms of autism does tend to coincide with the recommended timing for the MMR vaccine, so many parents who noticed that correlation latched onto Wakefield's findings. One such parenting team was celebrity couple Jenny McCarthy and Jim Carrey. Using their celebrity status as a platform, McCarthy and Carrey easily spread mistrust of vaccination recommendations. Subsequent studies disproved any causal relationship between the MMR vaccine and the onset of autism (Hussain 2), and on February 2, 2010 *The Lancet* retracted Wakefield's paper when it was

discovered that Wakefield received funding and subject referrals from lawyers who were involved in litigation against vaccine manufacturers (Palfreman) (Hussain 2). This, however, did not mark the end of the of the anti-vaccine movement.

The Current Status of the Anti-Vaccination Conversation

As discussed in chapter 2, issues with the COVID-19 lockdowns and mandates have led to a flashpoint within the anti-vaccination conversation. As would be expected, much of the talk I saw on Twitter revolved around skepticism over the effectiveness of the COVID vaccines, speculation over ulterior motives for those pushing the vaccines, and outrage over people or organizations requiring vaccination. In addition to this, I saw many tweets from the group I was following taking an anti-mask stance. As I will discuss later in this chapter, the connection between the anti-vaccination argument and an anti-mask argument may not be surprising, but neither is it a rationally logical connection. Regardless, this did provide for a lot of heated postings.

In analyzing these accounts, I've found that they mainly fall into one of three categories:

- Accounts of grieving mothers,
- Accounts of representatives of non-profit organizations (primarily those pursuing anti-vaccine agendas or those offering support to vaccine injured patients).
- Accounts belonging to authors who have published books on anti-vaccination literature.

In addition, there were many accounts that could have fallen into one of these three categories but are also tied together in a way that makes me want to discuss them as special categories. I will get to those shortly.

Grieving Mothers

Of the accounts I was following, a few of them were accounts from grieving mothers – those who were not authors or spokespeople but were parents who believed that their children were vaccine-injured. One such mother, known to me only as @1pissedoffmom1, regularly tweeted things like, “You know how those front loading washing machines develop mold. the cure is to leave the door open. yeah... that... with masks. <https://t.co/6ossQrn5Dd>” (Aug 21, 2020)³. The embedded link in her tweet leads to an article from DailyMail.co.uk about two dentists in New York (two out of the 14,893 registered dentists in New York in 2021, according to op.nysed.gov) who claim to be seeing an increase in cavities and gum disease due to “mask mouth.” The headline for the link would have read, “Dentists declare 'mask mouth' a new phenomenon as they see an explosion in patients suffering from tooth decay and gum disease after wearing face coverings” (Court). While The Daily Mail has largely been discredited as unreliable and sensationalist (Jackson), it’s these types of headlines that fuel posts by @1pissedoffmom1.

Other noteworthy grieving mothers include one I only know as “Vaxxed_Supporter”⁴ (@truthvaxwarrior). Her profile page describes her as, “Mother of vaccine injured child. Vaccines can and do cause autism.” Vaxxed_Supporter does not post many original tweets, but she posts a lot of retweets from public figures, like Robert F. Kennedy Jr. Following her may be useful to those who want to stay aware of what those public figures are saying about vaccines

³ The account @1pissedoffmom1 has since been blocked by Twitter and these tweets could not be retrieved for a screenshot

⁴ The banner on her profile page is a promotional image from the movie, *Vaxxed*, so I think it’s safe to assume that her name is a reference to contributions she has made to the movie

without, themselves, sifting through all of the non-vaccine-related tweets posted by those figures. However, her account did not provide me with much original tweeting to work with.

With slightly more original tweeting is another grieving mother, Catie (@justiceforevee). Catie's daughter, Eeve, died 36 hours after receiving a series of vaccines, and Catie



Figure 4

launched an anti-vaccine, non-profit named “Justice for Eeve” (justiceforevee.org). Catie is very well immersed in the anti-vaccine conversation, with a followers to following ratio of 4 (7,500 followers and 1,800 following) and a high rate of replies to other people’s tweets. As such, she was valuable to me in gathering subjects, and her empathetic replies to people sharing their own stories made me hopeful to find dialogic tweeting. However, most of her original tweeting involved promoting her own organization.

Non-Profit Organizations

The next group of accounts are those belonging to non-profit organizations. One of those accounts with 11,400 followers belongs to Physicians for Informed Consent (@picphysicians). PIC is an organization led by 30 M.D.s, D.O.s, and Ph.D.s. Their website (physiciansforinformedconsent.org) says that they fully support anyone who chooses to get vaccinated and to wear masks, but the literature they provide on the site suggests that they don't believe in either. Like most of the other non-profits I will discuss, most of their tweets seem written primarily for self-

promotion. When I first looked at the data I gathered, I was hopeful to find dialogic tweeting because many of their tweets began with the phrase, "Did you know." Ultimately, however, my readers saw this as an insincere rhetorical tactic, and the way they present their information after that phrase placed them in the authoritative category.



Figure 5

The next two non-profits on my list are Barb Loe Fisher's National Vaccine Information Center (@NVICLoeDown) with 16,300 followers and Generation Rescue (@GenRescue) with 18,500 followers. I will discuss NVIC shortly, as it falls into the special group of suspended accounts. Generation Rescue's tweeting was a source of frustration for me, as they were the group that happened to be holding their celebrity poker tournament fundraiser at the time I

was gathering my data, and their Twitter stream, at the time, was just a deluge of tweets advertising the tournament. As these tweets had little to do with the vaccine conversation, and my readers were unsure how to classify them, these were the ones I removed from my group of initial tweets.

Children's Health Defense (@ChildrensHD), with 31,000 followers, is the non-profit organization that was founded by Robert F. Kennedy Jr. Children's Health Defense primarily used their tweeting to spread news articles about

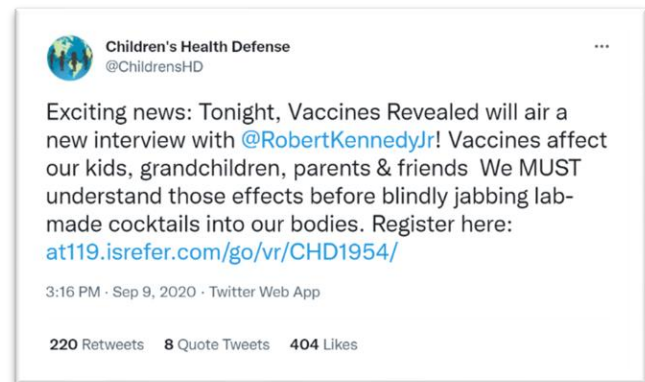


Figure 7

vaccines. For example, on September 17, 2020, they tweeted, "Merck fast-tracked Gardasil by presenting misleading data to the FDA & fabricating a health crisis. They claimed they were "filling an unmet medical" need but the only thing Merck wanted to fill was the \$6 billion hole

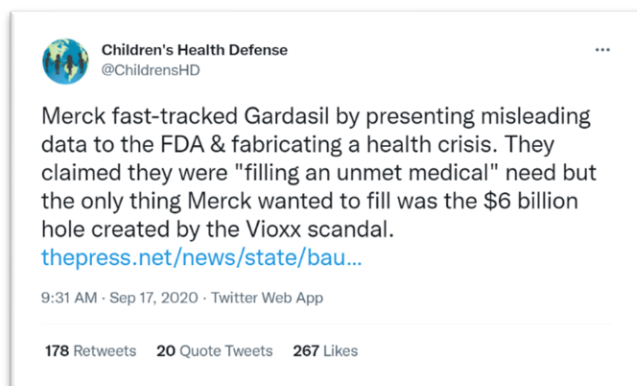


Figure 6

created by the Vioxx scandal," and they provided a link to a small, local paper (the article they linked to has since been removed from the site). Beyond that, they did not do a lot of self-promotion for the organization, itself. However, they did do a

lot of tweeting to promote Kennedy's vaccine-related activities, including speeches and interviews, such as on September 9th, 2020, when they tweeted, "Exciting news: Tonight, Vaccines Revealed will air a new interview with @RobertKennedyJr! Vaccines affect our kids,

grandchildren, parents & friends We MUST understand those effects before blindly jabbing lab-made cocktails into our bodies.”

Published Authors of Anti-Vaccine Literature

There were a few accounts I was following, with a low number of followers, which seemed to have a few good examples of tweets which use the dialogic language discussed by boyd and Gardner. One of these, with 2,900 followers, was from author Wayne Rohde. Rohde published a book on the vaccine courts (the government organization responsible for administering the National Vaccine Injury Compensation Program established in 1986) and hosts a podcast on the subject. At first, I was hopeful when my initial readers returned three dialogic tweets from the first 1,000 tweets of my dataset. However, upon reviewing those tweets, a few were not about vaccine-related issues at all but about The Bears and opinions of their use of Trubisky in their starting lineup.

The author with the most dialogic tweets in my initial collection was Eileen Iorio (@eileeniorio). Iorio wrote a book on the HPV vaccine and had 10,100 followers at the time that I pulled my data. What makes Iorio’s tweeting particularly interesting for some of what I’m looking at in this dissertation, beyond the number of potentially dialogic tweets, was an incident with her tweeting on September 26th, 2020. On that date, Iorio posted a series of tweets about Amy Barrett (at the time, nominated to the supreme court by Donald Trump) and Barrett’s support for *Jacobson V Mass.*, which Iorio cited as “the 1905 precedent for forced vaccination, eugenics, forced sterilization and mandatory masks and lockdowns.”

Through a discussion with others on this thread, it became clear that Iorio had some of her facts incorrect and was tweeting false and misleading information. Despite realizing this,

lorio refused to correct or remove her posts. In a response to @HMcbadger, @annaroo1021, @RandyEBarnett, and the official account for the president of the United States (Donald Trump, at the time), on the same day, she defends her choice, saying, “I mean that I won’t delete the post even though I didn’t quite get it right. I’m not a lawyer. It’s still an issue that Jacobson is being used at all. It should be overturned.” After further discussion, still on the same day, she adds, “It’s up too long to remove it.” So, with a clear disregard for the truth of the situation, she expressed a clear preference for spreading misinformation rather than retracting her Tweets. At one point, she mentioned adding a disclaimer to the original tweet, but I was unable to find any disclaimer. The entire thread has since been removed from the Twitter record.

Moving further up the list based on number of followers, is a key figure from Frontline’s “The Vaccine War” story, Jennifer Margulis (@JenniferMarguli). An author of several books, including one with the title, *Your Baby, Your Way*, Margulis was first drawn into the anti-vaccine movement over concerns about why doctors wanted to vaccinate her newborn baby against Hepatitis B, a disease that that Margulis

reasoned could not be a threat to her daughter until her daughter was sexually active. Since then, she has fashioned herself as a champion for the cause, trying to rebrand it as “the medical



Figure 8



Figure 9

safety movement.” More recently, she has come out against mask mandates and enforced social distancing, tweeting things like, “The erosion of freedom is far more threatening to our lives than a virus” (Sept. 15, 2020), and, “Being around a sick person doesn't make everyone sick, folks! I've breathed the same air as 2 confirmed #COVID19 patients, 1 a family member (shared toilet). I didn't get as much as a sniffle. A strong immune system beats

illness. So why isn't public health talking about this? (Sept. 5, 2020).

Nearing the top my list of authors, with 75,000 followers, is Dr. Frank Lipman (@DrFrankLipman). An M.D. of functional Medicine, Lipman’s tweeting actually seemed the most well-rounded and tempered of the group I had put together. In addition to his anti-vaccine tweets, many of his other tweets focused heavily on the benefits to getting good sleep and currently trendy diets like low-carb and intermittent fasting. Lipman’s reputation as an anti-vaccination supporter seems to come primarily from a time when he spoke out against the swine flu vaccination, specifically (Lipson). Since then, it seems that he may have tempered his perspectives.

Compared to the other people on my list, Lipman did relatively little tweeting about COVID 19 and even less tweeting about vaccines. He did give advice about maintaining healthy

vitamin D levels to help fight off COVID, and he warned about maintaining mental health during a quarantine. However, within the first 120 of his tweets in my dataset, I only found two tweets with a loosely anti-vaccine stance: on September 19, 2020, he tweeted, “COVID-19 Vaccine Makers Keep Safety Details Quiet, Alarming Scientists,” and on September 4, 2020, he tweeted, “Rising speculation that President Trump may pressure the FDA to approve a Covid-19 vaccine before testing has been fully completed.” Surprisingly, on September 17th, he even sent out a tweet supporting social distancing practices. His followers reacted to that tweet with anger and outrage, but I will discuss that tweet, specifically, a little later in this chapter.



Figure 10

The account with the 2nd highest number of followers on my list belongs to Dr. Joseph Mercola (@mercola). With 290,500 followers, Mercola has almost four times as many followers as the 3rd highest account on my list (Frank Lipman). Sheera Frenkel of the New York Times calls Mercola the single most influential spreader of Coronavirus misinformation. His profile page on Twitter proudly announces that he is the founder of the #1 natural health site, Mercola.com⁵.

While I’ve included Mercola with my group of authors, his primary source of income from the spread of his anti-vaccine message comes from that website. As Frenkel explains, “As

⁵ A quick Google search for “top natural health sites” turns up a few results that place Mercola.com within the top three. Other results didn’t mention Mercola.com.

his popularity grew, Dr. Mercola began a cycle. It starts with making unproven and sometimes far-fetched health claims, such as that spring mattresses amplify harmful radiation, and then selling products online — from vitamin supplements to organic yogurt — that he promotes as alternative treatments.” Recently, Mercola has begun a new practice of deleting all of his website’s content after 48 hours, because he is afraid of censorship, the “dark money forces of billionaires,” and “a new era of authoritarian technocracy” (Mercola).

Focusing on his tweeting, this first thing I noticed was how he tags many of his posts, saying, “⚠️ Twitter has CENSORED my website! ✅ Visit **Mercola.com** to read this and more news articles,” this driving traffic from his Twitter account to his website.

Some of the unsubstantiated claims

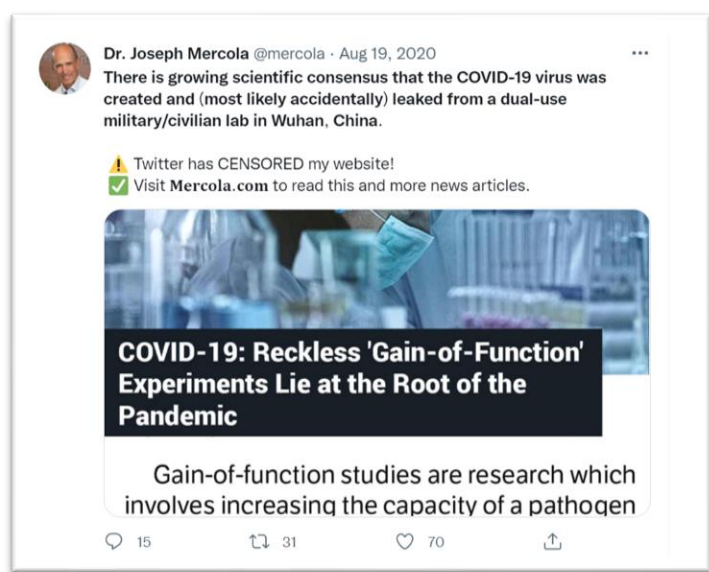


Figure 11

he has made include, “Today, we also have something no previous tyrant has had, namely the technology to track, trace, control and manipulate individuals wherever they live” (September 19, 2020) and, “In the 20 years that vaccine makers have tried to develop a coronavirus vaccine, efforts have failed due to dangerous, many times lethal, side effects” (September 21, 2020).

Possible Fourth Category: Doctors Who do not Derive an Income from Anti-Vaccination

Sources

There was one anonymous account with 3,059 followers, titled “Truth Lover” (@truthlovingdr). The type of doctor that Truth Lover may or may not be is never specified on the account. As an anonymous account, I feel that it is safe to say Truth Lover is not on Twitter to promote anything that they may have



Figure 12

published, however I also cannot confirm their qualifications to this fourth possible category. Additionally, considering that this was the only account that may have fallen into this category, I can't say I'm sure if this really constitutes a category or if it really was just a unique occurrence. Either way, their tweets were still of interest to me because Truth Lover did seem to write a few dialogic tweets. While Truth Lover does frequently engage directly with other users and encourages other users to engage back, their posts do not seem to get many likes, retweets, or responses.

Special Mention Category: Banned Accounts

There were several accounts, that happened to be in mid-range size of the accounts I was looking at, which were all suspended by Twitter during my research. Presumably they were suspended for their tweets containing “misleading or disputed information that could lead to harm” (according to <https://help.twitter.com/en/rules-and-policies/notices-on-twitter>). I feel this is a safe assumption because, when I search for all of those accounts now, I get a special administrative message from Twitter saying, “Know the facts. To make sure you get the best information on the coronavirus (COVID-19) and coronavirus vaccines, resources are available from public agencies.” The message then also includes links to [vaccines.gov](https://www.vaccines.gov)



Figure 13

and the CDC’s Twitter account (searching for other suspended accounts, such as @realDonaldTrump, the administrative message simply says, “Account suspended. Twitter suspends the accounts which violate the Twitter rules.”). These suspended accounts all had between approximately 11,000 followers and 45,000. Whether it is coincidence that they seemed to be grouped this way or if there is actually some account size range that Twitter considers large enough to be threatening but small enough to not cause an uproar when suspended, which just happens to match my mid-range accounts, I cannot say.

The owner of at least one of these accounts, that belonging to Toby Rogers (@uTobian), is no stranger to being suspended. On September 24th, 2020, Rogers replied to @jdelugach, saying, “Yep. I get kicked off of here pretty regularly. The key is not to panic (and) to just give it

a day or two. More often than not things go back to normal.” In this case, however, the suspension seems to be sticking as his account was suspended shortly after I gathered my data and has continued to be so for the entire time that I’ve been writing this dissertation. Rogers’ tweeting was some of the more entertaining, because he frequently got into complete with name-calling and personal attacks unrelated to the topic of discussion (I’ll talk more about that later in this chapter). And, yet, he sent a lot of love out to his followers, sending the empathetic tweets I originally looking for, such as a reply he sent to Catie, on September 25th, 2020, saying, “@justiceforevee Sorry for your loss sister. Thank you for sharing your story. Blessings to you. 🙏” Beyond that, with a Ph.D. in economics, he primarily approaches the issue from that standpoint, saying things like, “I have examined financial conflicts of interest for hundreds of studies. The answer is Pharma spends billions to distort the science while parents’ groups hold bake sales to fund a few studies” (September 28th, 2020).

The next noteworthy suspended account belongs to Barb Loe Fisher (@NVICLoeDown), founder of National Vaccine Information Center. Also someone featured on Frontline’s *The Vaccine War*, most of Loe Fisher’s posts were promotional posts for the NVIC. The majority of those were phrased, “Join us for an enlightened conversation about...” I was immediately curious about how well received Loe Fisher’s tweeting would be, as she has one of the worst follower to following ratio, by far, of all of the accounts I was looking at. She is only following 6 accounts. According to Westerman’s findings, which I discussed in chapter 1, this should suggest that Loe Fisher would have very little credibility, as it suggests she’s not really following the conversation on Twitter. However, she did have over 16,000 followers – a significant number for someone who should not be seen as credible – bringing her ratio to 2,717:1

followers to following (one account, which I will discuss shortly, had a worse ratio but substantially more accounts being followed, but the next worse ratio was 635:1). That being said, she did have a low engagement rate for her tweets: an average of 0.0017.

Larry Cook (@stopvaccinating), “healthy lifestyle advocate,” is also no stranger to being banned from social media. On September 25th, 2020, he tweeted, “In the last 2 weeks I went Live 5 times. Facebook deleted them all. Why do you suppose that is?” In looking at his website (larrydcook.com), most of his work seems to be dedicated to promoting a vegan diet, but his tweeting definitely shows a more frightening side to areas where his advocacy crosses over into more controversial topics and opinions. On September 28th, 2020, he sent a tweet addressed to Governor Cuomo, saying “Reject satanic laws.” I do not know, specifically, what this comment was in reference to, but the comment itself shows us his perspective on politics. Worse than that, on September 23rd, he shared a video of a black man knocking over tables and chairs in a restaurant patio area (<https://t.co/SpfF7koJsA>). The caption on the video reads, “BLM rioters are already targeting businesses.” Cook’s tweet, in which this video was embedded, tells us, “Buy your guns now”. While he maintains that he’s concerned about vaccines because he believes they cause autism and fibromyalgia (September 26th, 2020), this crossover of issues speaks more about a generalized hostility towards progressive ideals, but I will speak more about the crossover of issues later in this chapter.

A woman known only as sally (@sallyKP), maintained another of the suspended accounts I was following. Sally’s primary topic of discussion in her tweets focused on her view that the mandates being put into place were just blatant systems of control and had very little to do with actually preventing the spread of the virus. With her tweeting, she announces,

“Remove your masks” (September 27th, 2020), and “Reclaim your freedom” (September 28th, 2020), and she backs this up with tweets like, “From the UK... “No extended eye contact” NOT👏ABOUT👏A👏VIRUS Wake up world” (September 27th, 2020).

Two more accounts I was following that were banned, which, in hindsight, should not have been a surprise, were Inside Vaccines and LotusOak. First off was Inside Vaccines). Inside Vaccines (@InsideVaccines) seemed to like to tweet fairly outlandish things like, “I think that people in Europe during the 1300s coped better with plague than we are coping with COVID,” and, “The existence of human beings may soon be a conspiracy theory” (both on September 28th, 2020). However, what I immediately noticed when I first started following Inside Vaccines, was that their profile immediately directed you to their MeWe account. MeWe.com is a largely unmoderated social media site that has become popular for some of the more extreme conspiracy theorists primarily because of the light moderation. If Inside Vaccines was using their Twitter account to funnel followers to their MeWe account, then some of their more extreme content would likely filter through at times.

The other banned account that should not have been a surprise, with 40,100 followers, was known to me as LotusOak (@lotusOak2). LotusOak primarily did a lot of retweeting and sharing of external links. After they disappeared, I discovered a few things about their history with suspensions. According to an article on The Conversation by Filippo Menczer and Pik-Mai Hui, @LotusOak was an account that was suspended in late 2018 or early 2019. That account listed the name Vira Burnayeva, and was cited by Menczer as one of the dominating anti-vaccine accounts. Later, Menczer found an account @ViraBurnayeva which listed the name LotusOak and also tweeted out anti-vaccination messages. Menczer concluded that this was a case of a

single source circumventing Twitter's attempts to silence them. The account @lotusOak2 was likely just another link in the same chain. However, with over 40,000 followers, repeatedly recreating their account does not seem to have slowed them down too much⁶. I looked for @lotusOak3, but it does not exist... yet.

The largest of the suspended accounts I was following belonged to Dr. Sherri Tenpenny (@BusyDrT). With 43,800 followers, Tenpenny is an osteopathic doctor and an author. According to Anagha Srikanth, in an article published in 2021 (after my data collection) named Tenpenny (along with Dr.



Figure 14

Joseph Mercola, who I discussed previously) as one of just a few people who have been responsible for spreading the majority of vaccine misinformation. In one of her more astounding recent appearances, in June 2021 she testified before an Ohio court, as an expert witness, claiming that the COVID-19 vaccine magnetizes people and contains particles which interface with 5G relay towers (Bischoff). As far as her Tweeting activity prior to her suspension

⁶ As a sidenote on Westerman's findings on followers to following ratio, I feel like there must be a limit to that logic. As a reminder, Westerman showed that accounts with large followings that also had a large following lists were seen as more credible because they appeared to be well immersed into the Twitter community, whereas accounts that were only tweeting out and not following others seemed disconnected from the community. LotusOak has one of the best ratios of the accounts I was following while also being one of the largest accounts I was following. They had 40,100 followers and were following 33,700 accounts, giving them a ratio of nearly 1:1. Personally, I feel that the most accounts someone can legitimately monitor and pay attention to has to be somewhere around 200 to 300 (I, myself, struggle in keeping up with 100 friends on Facebook). As we know LotusOak has gone through several accounts and seems to build themselves back up quickly. I believe it probable that they had a follower list from an incarnation of a previous account. When they created the @lotusOak2 account, they likely started following all of the accounts that had previously followed them, simply as a way of announcing their return to Twitter. If that is the case, it seems like quite an effective tactic.

goes, she would say things like, “Thanks to some of the vaccines on the CDC schedule - many of our kids actually won’t be able to become parents” (September 25, 2020) and, later that day, “They’ve got diapers with chips in them too!” Neither of these tweets were supported with a link.

Special Mention Category: The Bigtree Family

I then have three accounts I was following that were all tied, in one way or another, to one individual: @delbigtree, @HighWireTalk, and @vaxxed2. Del Bigtree is a movie producer, talk-show host, and the CEO of the non-profit advocacy group ICAN (Informed Consent Action Network). In one of his more infamous recent moves, while giving a speech at “Parents Call the Shots” in Austin, Texas, in March of 2019, Bigtree equated the struggle of members of the anti-vaccine movement to the persecution and plight of the Jewish people in the Holocaust, pinning a Holocaust-style, yellow star of David to his lapel during his speech, comparing his desire to not get a shot from his doctor to the systematic murder of six million people.



Figure 15 - Retrieved from ADL.com

Bigtree’s tweeting from his personal account (@delbigtree) was split between attacks



Figure 16

on Anthony Fauci, indirect attacks on people willing to wear masks, and support for Donald Trump. On September 10, 2020, he attacked parents who were simply preparing

for a quarantine, saying, “If you find yourself packing extra clothes so your CHILD is prepared to be QUARANTINED then I think you should be on the list of WORLD’S WORST PARENTS!!” And, in May of 2020, he tweeted, “If I had an employee that made wrong prediction after wrong prediction I’d eventually fire them. What credibility does Fauci have left? #COVIDIOT #BeBrave.” There were a few occasions where he tweeted out questions and calls-to-arms to his followers that matched the dialogic style I was initially looking for, and I will discuss those in some detail later in this chapter.

There are then also Twitter accounts devoted to internet talkshow that Bigtree hosts, The HighWire (@HighWireTalk), and his movie, *Vaxxed 2: The People’s Truth* (@vaxxed2). The Twitter profile page for The HighWire boasts, “High above the circus of mainstream media spin, death-defying talk without the corporate safety net.” The tweets coming from this account are not as aggressive and confrontational as those coming from Bigtree’s personal account. While there is some self-promotion and show topic announcements, there is actually less of it than I would have expected. The majority of the tweets coming from this account seems to be news-sharing and re-tweets from other sources.

The movie, *Vaxxed: From Cover-Up to Catastrophy*, was produced by Del Bigtree, written by Andrew Wakefield and Del Bigtree, and directed by Andrew Wakefield. The Twitter account “We Are Vaxxed”⁷ (@vaxxed 2), represents the sequel, *Vaxxed 2: The People’s Truth*. *Vaxxed 2* was produced by Del Bigtree and Robert F. Kennedy Jr and was directed by Brian Burrowes but still starred Andrew Wakefield. When I collected my data, this account had been inactive since May of 2020, however I felt obligated to keep it in my set because it was the closest thing I had to an account connected to Wakefield, the original spokesperson for the movement’s current resurgence. The overwhelming majority of their tweets were announcements about how and when the movie could be purchased. As an interesting note, the account became active again early in 2021 but has been primarily retweeting content from Peeps TV – the streaming channel currently hosting both movies.



Figure 17

Special Mention Category: 2.2 Million Followers

The final account I was following belongs to Fox news contributor, Michelle Malkin. With 2.2 million followers, Malkin’s following dwarfs all other accounts on my list. Since the time that I gathered my data, Malkin has come under fire for associations with white nationalists, neo-Nazis, Holocaust deniers, and Groypers (ADL2). I did capture some tweets of

⁷ Yes, this seems like an ironic name for this page considering they are not vaccinated.




Figure 18

hers in my dataset which clearly show these leanings. For example, on September 23, 2020, she tweeted,

“God bless the #ProudBoys,” and, the next day, she tweeted, “Black Lives Matter is a domestic terrorist organization.” Twitter has “redacted” many of her posts that I have tried to go back to, but I have not observed her having any problems with having her account suspended. Her aggressive and confrontational tweeting style does persist into COVID-related topics. On September 25, 2020, she tweeted, “Universal mask mandates are a sham. No more fines. No more arrests. No more orders. NO MORE TASERING of citizens yearning to breathe free.”⁸



Figure 19

Observations from preliminary data

As I mentioned in the previous chapter, the first thing that I noticed when looking at the preliminary data was the imbalance between the number of authoritative tweets and the number of dialogic tweets. In the initial 1000 tweets that my readers looked at, they identified 599 tweets as authoritative and only 42 of them as dialogic. Based on the disparity between these numbers, I feel that we cannot make sweeping claims about widespread use of a dialogic

⁸ Astoundingly offensive misappropriation of the phrase “yearning to breathe free” considering her position on immigration – evidenced by her own tweets, one example being on July 14, 2020, when she said, “Stop importing foreign workers... #ExpandTheBan #ExtendTheBan #ImmigrationMoratorium.”

ethos on Twitter. More specifically, we cannot say that it's an open, authentically human voice that is drawing people into sympathizing with the anti-vaccination movement on Twitter.

Another finding I mentioned in the previous chapter was how the Twitter users with fewer followers seemed to be responsible for more of the dialogic tweeting. The one-third of the accounts with the most followers on my list were responsible for 26% of the dialogic tweets my readers identified. The middle-third of the accounts that I was looking at were responsible for 31% of the dialogic tweets. The third of the

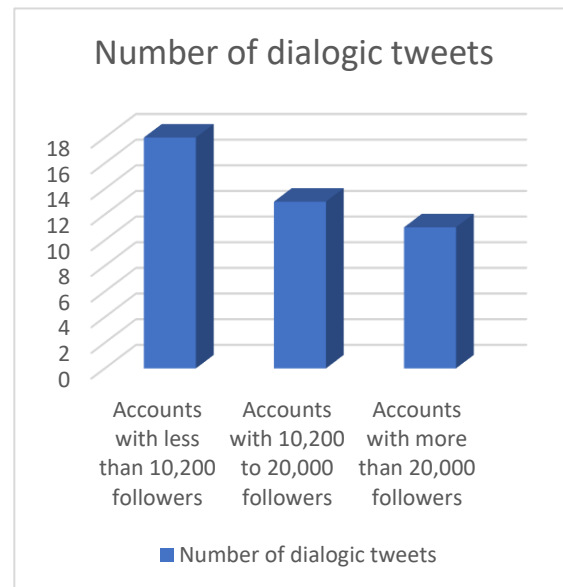


Figure 20

accounts with the fewest followers were responsible for 42% of the dialogic tweets. This trend seems to suggest one of two possible conclusions about the use of a dialogic *ethos* on twitter.

The first possibility, quite simply, is that authoritative authors attract larger followings. This would suggest the opposite of what this dissertation sought evidence for – that audiences were reacting more favorably to dialogic tweets than they were to authoritative tweets. If authoritative tweeters are attracting more followers, then that is clearly the type of language that the anti-vaccination sympathizers on Twitter are responding more favorably to. For as useful of an indication as likes and retweets may be in looking at audience reactions, the key question here is still one about what is attracting people to causes like the anti-vaccination movement. Therefore, the amassing of followers still lies at the heart of the issue and would be

worth looking at had I the ability to follow the progress of these accounts over their lifetimes on Twitter.

The second possible conclusion we might draw from the trend that the accounts with fewer followers utilized the most dialogic language is that the shift from a dialogic voice to a more confident, authoritative voice could be a natural evolution of writing style as one accumulates more followers. It seems logical that users with few followers may be more tentative with their tweets and may use more markers of epistemic hedging (as discussed in chapter 2). As, however, they build a following, they may become much more confident in their own opinions and in their writing. The building of followers would be a validation of their perspective which would, in turn, give them the confidence to be more authoritative.

At a few points, Toby Rogers (@uTobian), one of the accounts I was following, got into a few arguments with other twitter users in which he attacked their characters just based on the number of followers they had. On September 24th, responding to an argument with a user calling himself Roger Roger (@canjetsfan), Rogers said, “3,000+ tweets and only 15 followers!? That's gotta be some kind of record for futility.” In a similar argument earlier that day, he also attacked a user with the name Max(@OriginalName99), pointing out that Max only had 7 followers.⁹ This does reinforce the idea that, regardless of what the facts might be, some Twitter users believe that a user with more followers should have the greater authority - a form of *argumentum ad hominem*. Additionally, they seem to be of the opinion that the simple fact that a user has more followers makes that user’s opinions more valid than a user with fewer followers.

⁹ The account @uTobian has since been blocked by Twitter and these tweets could not be retrieved for a screenshot

In order to determine which of these conclusions should be drawn, we would need to conduct a study to follow different Twitter accounts from their inception and through their growth and development over the course of five to ten years – something well beyond the scope of this dissertation, but a possibility for a future, follow up study. As an illustration of what I'm thinking, I would begin by identifying newly created accounts with a low number of followers which seem to have potential for long-term involvement in the anti-vaccine conversation (those that tweet regularly, remain focused on anti-vaccine issues, and follow other anti-vaccine accounts). I would then track the evolution of the language they use and compare that with their rate of growth. If the beginning sample was large enough, a few of those accounts should grow well over the next five to ten years while showing a progression from using dialogic language to using authoritative language. I would then try to see if I could determine if their following grew after the evolution of their language or if their language evolved after their following grew. This, however, remains a question for a later date.

On top of this, the data from the full run of all 57,533 tweets in my full dataset did not necessarily support this distribution of dialogic tweets. That breakdown suggests a more complicated relationship between the size of an account's following and the frequency of dialogic tweets, but I will discuss that further in my next chapter.

Observations from Full Dataset Run

As a reminder, after running my full dataset of 57,533 tweets through the machine learning and analyzing the results, the data did suggest a linear relationship between dialogic/authoritative certainty rating and average engagement rate. More specifically, a 0.1 increase in authoritative certainty predicts a 0.00001963 decrease in average engagement rates, while a 0.1 increase in dialogic certainty predicts a 0.00001904 increase in average engagement rates. While we need to remember that correlation does not necessitate causation, this information does show that tweets with higher dialogic certainty ratings did seem to perform better than tweets with higher authoritative certainty ratings. This, then, could support the theory that some audiences may respond well to expressions of authenticity and empathy in digital environments like Twitter.

The real question this information leaves us with is how significant of a difference this makes. As mentioned, the data shows that an increase of 0.1 in dialogic certainty predicts a 0.00001904 increase in average engagement rates. This means that, with an increase of 0.1 in dialogic strength, we would predict 0.00001904 of a like for every follower an account has (remembering that engagement rate is represented by the number of likes (and retweets) divided by the number of followers that account has). In other words, a tweet which would receive a 1.0 dialogic certainty (highest possible score) could be predicted to receive one more like (or retweet) for every 5,000 followers that the originating account had when compared to a tweet that scored a 0.0 (lowest possible score) in dialogic certainty.

The average dialogic certainty rating for the entire set of 57,533 tweets was 0.134013911. So, for a deliberately crafted dialogic tweet, we are more realistically looking at a

difference of a 0.7 or 0.8 increase in dialogic certainty. That would equate to a 0.0001428 increase in engagement rate, or approximately one like or retweet for every 7,000 followers and account had.

Considering how small that increase actually seems to be, I feel it is worth recognizing and doing additional research on, however I don't believe it justifies danah boyd's comments that strangers on the internet being "willing to listen, empathize, and compare notes" is a main reason for people to be getting their medical information from their social networks, at least as far as Twitter is concerned. There seems to be a very limited amount of dialogic tweeting occurring, and, to the extent that it does occur, its influence seems to be rather limited.

That being said, a correlation does suggest there is more here that is worth looking at – we shouldn't ignore the influence of a dialogic approach just because it isn't the main influencer of a tweet's success. It has certainly always been important to know your audience and how to best communicate with that audience. More specific audiences may respond better to dialogic language than those on Twitter. Additionally, looking at these tweets 'in the wild' could have presented certain disadvantages. As I mentioned, the average dialogic certainty of the tweets I was looking at was 0.134. If I had looked at tweets which were carefully crafted to utilize dialogic language, then the data may have shown a greater difference in predicted engagement rates. That, however, would be a question for another study and one which I will discuss further in my next chapter. In the meantime, we can still look at some noteworthy tweets from this dataset in order to consider how a carefully crafted dialogic tweet might sound.

Noteworthy Dialogic Tweets from the Preliminary Data

There were a few tweets within the group of the preliminary 1000 tweets that I was able to find that seemed to be particularly good examples of how I thought a dialogic tweet should function. As it turns out, they were some of the more successful of the dialogic tweets in the preliminary group of 1000 tweets.

On June 9th of 2020, Del Bigtree (@delbigtree) tweeted out the question, “Why do YOU wear a mask?” We must keep in mind that Bigtree does not support the mandated use of masks. Much like his stance on vaccines, he encourages his followers to resist accepting the guidance of medical authority. I find it curious why he phrased his question in this way – seemingly allowing his followers to provide their own reasoning for the justification of masks.

Regardless of intent, the text alone (which was all my readers could see) suggests this tweet to be a well crafted dialogic prompt – it encourages participation and (again, just based on the text) differing viewpoints.

However, the cartoon accompanying the text does belie the intent presented by his question. The reasons for wearing masks presented by the people shown in the cartoon



Figure 21

are clearly phrased in a way to make the use of masks seem ridiculous. His hashtag, #BeBrave, clearly is not intended for those choosing to wear a mask. While the prompt for participation still makes me want to classify this as a dialogic tweet, when considered with the cartoon, I can't give him credit for encouraging different viewpoints – hence my curiosity on why he phrased the question in the way he did.

Bigtree did get a positive reaction out of this tweet, earning a 0.049 estimated engagement rate. He also did get a lot of participation in this thread, prompting 733 comments. As should have been predictable, the overwhelming majority of those comments were from people defending their choice to NOT wear a mask. A few examples of responses include comments like, “100% against sheeple

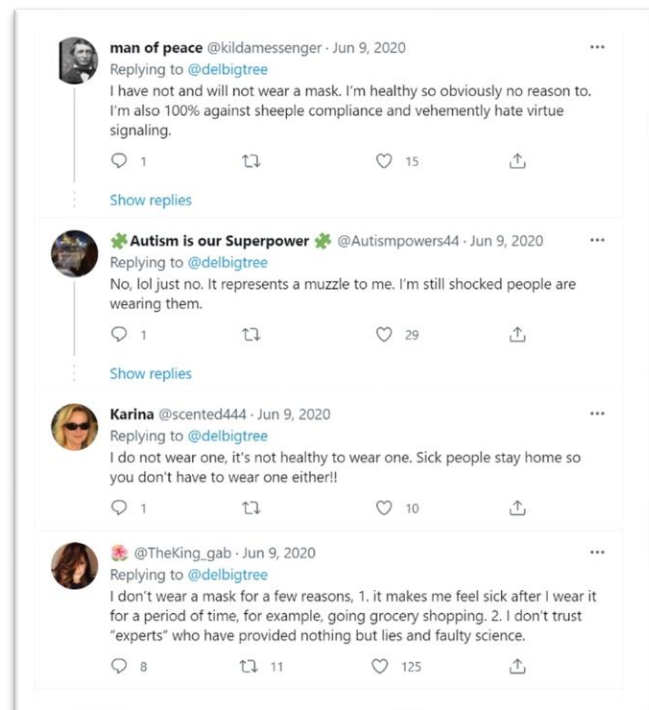


Figure 22

compliance,” “It represents a muzzle to me,” and “I don’t trust “experts.”” Regardless of intent or of the opinions of the responders, Bigtree’s dialogic approach did seem to act as a rallying cry to potential followers, though it still earned him an average number of likes for his tweets in the preliminary 1000 tweets. His actual average for those tweets was 1,962 likes, and this tweet earned 1,675 likes.

A similar call for participation was made by Larry Cook (@stopvaccinating) on September 29th, 2020, when he tweeted the question, “Has a pediatrician ever bullied,

harassed, or demeaned you for asking about vaccine safety or for refusing to vaccinate?”¹⁰ This tweet earned an estimated engagement rate of 0.013. This level of engagement is high compared to Cook’s other tweets. The average number of likes his tweets in the group of the preliminary 1000 tweets received was 67 likes, and this tweet received 164 likes.

Beyond those two examples, there were many good examples of tweets saying things like, ‘I’m so sorry this happened to you,’ or ‘thank you for sharing your story’ (@sallyKP, @uTobian, @justiceforevee, @truthvaxwarrior, @truthlovingdr), but very few of these tweets earned any likes or retweets. What is more likely the case is that, to the extent that they were effective, they may have earned likes and retweets for the original post that they were written in response to. While these tweets were considered dialogic because they expressed empathy or encouraged participation in a discussion, I did also find just a couple tweets that brought up some differing opinions.

¹⁰ The account @stopvaccinating has since been blocked by Twitter and this tweet could not be retrieved for a screenshot

On September 1st, 2020, Frank Lipman (@DrFrankLipman), someone who typically speaks out against vaccines and other advice given out by medical experts tweeted out the message, “COVID-19 study links strict social distancing to much lower chance of infection.” However, unlike the examples from Bigtree and Cook, this acknowledgement of different viewpoints was actually met with a lot of open hostility. Lipman’s followers responded with comments like, “Pure Propaganda,” “Stop spreading

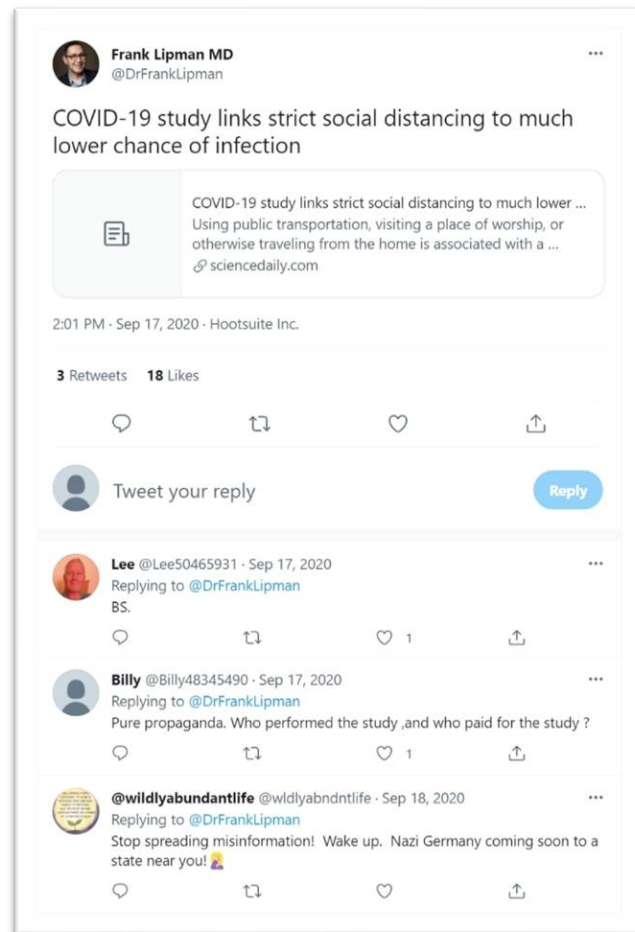


Figure 23

misinformation,” and, succinctly put, “BS.” The tweet also received very few likes – Lipman’s average in the preliminary 1000 tweets was 58 likes, and this tweet earned only 18 likes. This, again, seems to suggest that people following the anti-vaccine movement are certainly not doing it because they are looking for an open, honest dialogue.

Similarly, the account “Inside Vaccines” (@InsideVaccines), an account that typically speaks out against expert opinions, seemed to defend doctors on September 26th, 2020, when they tweeted, “In my first-hand experience with multiple doctors, most describe what they see and make a recommendation. They are quite willing to answer questions and equally willing for

the patient to make the final decision about treatment.”¹¹ Again, this tweet earned very few likes – 18 likes – when compared to their average of 54 likes for their tweets within the preliminary 1000.

General Observations on Tweet Content

While I found only a slight benefit to trying to appeal to readers through the use of dialogic language, I did make a few observations tangentially related to my original research question which I would like to note here. Originally, for this dissertation, I wanted to focus on stylistic elements, such as tone and word choice, rather than content – the actual ideas expressed. I wanted to see if the way a microblogger on Twitter presented their views could be as important as what they said. As I mentioned in chapter one, all of the theorists currently studying confirmation bias focus on content. However, Osatuyi, as we remember, points out that content is “an antecedent” of the conversation (2626). Agreeing with him, I felt that focusing on this antecedent sets up a bit of an impasse – it doesn’t leave us with any ways of building a productive dialogue nor will it really help us “arm” our scientists and experts to battle this wave of non-truths and manipulations. However, looking at the messages from the set of preliminary 1000 tweets, and knowing that crafting a dialogic *ethos* may not be as powerful of a tool as I and others thought it would be, I feel that there could be important takeaways from the content of these tweets.

¹¹ The account @InsideVaccines has since been blocked by Twitter and this tweet could not be retrieved for a screenshot



Figure 25

First, one trend I noticed, in particular, has me questioning what people within the anti-vaccine movement mean when they talk about “doing your research.” Now, every first year composition teacher struggles to teach their students that “doing research” should

not mean scanning for the one source that supports your opinion and ignoring the rest.

However, a post by Jennifer Margulis (@JenniferMarguli), a woman who has a Ph.D. in English, seems even worse. On August 10th, 2020, Margulis shared a Newsweek article with the headline, “Sweden, which never had lockdown, sees COVID-19 cases plummet.” Margulis’ accompanying commentary said, “Sweden is doing very well! So lockdowns and masking may not have been the right approach after all?”

After seeing this, I read the Newsweek article myself.



Figure 24

The article, written by Soo Kim, absolutely does not say that Sweden was doing very well, nor did it say that Sweden was not participating in a form of a lockdown. Kim does point

out that there has not been a government-mandated lockdown in Sweden. However, she reports that, “the Scandinavian nation ranks eighth among countries with the highest number of COVID-19 deaths per 100,000 people.” She goes on to discuss the fact that the population of Sweden, recognizing how vital social distancing has been in controlling the spread of COVID-19, collectively self-imposed lockdown practices across the country. That self-imposed lockdown was what lead to a dramatic decline in COVID cases. So, based on Margulis’ comments, it is clear that she (with her Ph.D. in English) did not read past the headline of the article before sharing it with all of her followers, citing it as an argument against lockdowns.

Looking at the comments to this post, there were a couple of the 25 comments which encouraged her to read past the headline. But, the majority of the comments supported Margulis’ view, and the post was retweeted 106 times. I would assume that a great many more people than that looked at this and accepted her “research” without any question. So, research, here, does not even entail finding an article that supports your opinion. It only entails finding a headline that can be interpreted as supporting your opinion. This reluctance to read past the headline and to make snap decisions on assumptions of what an article says does lend more weight to the influence of confirmation bias I was trying look past.



Figure 26

In scanning through the first 6,000 tweets in my list, I found 19 additional posts about Sweden and their apparent success without masks or lockdowns. Not all of them stemmed from the same

Newsweek article, but they still seemed to focus, selectively, on things said about Sweden's commitment to avoiding placing restrictions on people and ignored much of the other news coming out of the country.

Curious about this trend, I scanned the same group of 6,000 tweets for any mention of Israel – a country frequently cited as having the greatest success controlling the pandemic through their aggressive pursuit of mask use, social distancing, and delivering vaccinations. I found a total of four mentions of Israel, and none of them said anything positive.

@1pissedoffmom called news from the country as “Israeli propaganda” while others decried Israeli efforts as a threat to medical freedom. So, regardless of what was printed beyond the headlines of articles shared through Twitter, Sweden's lax approach was hailed as a success while Israel's actual success was hardly mentioned.

The other issue about the content of all of the tweets that I've looked at for this dissertation – an issue that suggests clear support for the influence of confirmation bias – is the link between the anti-vaccine conversation and the anti-mask conversation. As I discussed in Chapter 2, the 25 accounts I had chosen to follow were specifically ones on a crusade against vaccines. However, when the COVID pandemic began, they all quickly also took up the anti-mask and anti-lockdown fight. Logically speaking, each of these arguments have nothing to do with the others – people are really concerned about the presence of mercury in vaccines, but there is no mercury in my 100% organic cotton mask. It is possible that anti-vaccine spokespeople took up the anti-mask cause as a ploy to gather additional followers from a potentially sympathetic audience. However, more realistically, this points to a general mistrust of science and a powerful aversion to being told what to do.

Observations on Other Ways of Being Heard

Early in my research, I did come across a clip of Jenny McCarthy saying, “Is it mercury? Is it the schedule? Is there just too many? My answer to people, and what I’ve been telling them, is it’s all of the above. We don’t know for sure, which is why we keep saying, study it. But, they won’t” (Palfreman). This is, of course, is the exact language boyd echoed when she said, “Keep in mind that anti-vaxxers aren’t arguing that vaccines definitively cause autism. They are arguing that we don’t know” (boyd). However, other statements from McCarthy actually belie the idea that “we don’t know” is their true argument. In the “Green our Vaccines” rally on June 4, 2008, McCarthy gave a speech in which she said, “The ingredients, like the frikkin mercury... need to be removed immediately after we saw the devastating effects it took on our children.” Based on that speech, it sounded like she was definitively saying that mercury caused autism. McCarthy later saying that “We don’t know,” could just be a tactic commonly referred to as “moving the goalposts.” Maarten Boudry refers to this technique as an immunizing strategy¹² where, “A theory-in-crisis is often belatedly modified by its advocates so as to be less vulnerable to refutation, by introducing ad hoc elaborations and special clauses that explain away apparent failures and reduce the empirical content of the theory” (Boudry 147).

We cannot really say one way or the other if this was truly a case of epistemic hedging used to invite dialogue or if it was an attempt to inoculate herself against further attack. Since I’ve only been looking at Twitter, I certainly cannot make any broad claims as to dialogic language use across other forms of anti-vaccine rhetoric, but I would be careful about accepting

¹² No, the irony of anti-vaccine rhetoric making use of an “immunizing strategy” is not lost on me.

statements like McCarthy's at face value for any future research on the topic. That all being said, the lack of evidence of the use of dialogic language on Twitter does not discount the theory that people may be rallying to the anti-vaccine cause because they feel they are being heard in a way that they do not feel when talking to their doctors. Since, however, that sense may not be coming from deliberate dialogic engagement on the part of the spokespeople I've been looking at, some of my general observations about this conversation could suggest that it is still embedded in the way the arguments have been formed. Where I believe this is most evident in anti-vaccine tweets, and most neglected in advice from medical experts, is in recognition of the basic assumptions that these arguments stem from.

As Katarzyna Elliott-Maksymowicz points out in her article, "How much can you say in a tweet? An approach to political argumentation on Twitter," the 280 character limit on a tweet does not lend itself well to complex arguments. She notes that many people on Twitter use enthymemes as a way of condensing their arguments down to single speech act. She explains, "Because the enthymematic form allows much to go unsaid, it allows arguments to be made even by singular speech acts of few characters. This makes even singular speech acts a potent way of expressing even moral arguments in limited space such as character-restricted tweets" (3).

An enthymeme is a syllogism where one of the proofs, or even the conclusion, can go unsaid because it is a commonly held cultural assumption. When reading through the collection of tweets I've been looking at, I have noticed that many of them are based in common assumptions which are not shared by those arguing in favor of vaccines. If there is no evidence that suggests that utilizing dialogic language may be an effective tool for bridging the

gap between the sides of this conversation, then adopting the same, unspoken assumptions common in anti-vaccine arguments may be a way through which people with concerns may feel heard. In reading through my collection of tweets, I feel that I've seen two trends in assumptions being made: one on value judgements and one on motivations of the CDC.

The first common focus of enthymemes I've been seeing is on core values. I remember looking at the "The Mask Speaks" illustration tweeted by Del Bigtree which I included earlier in this chapter. Looking closely at that illustration, there is one person on the far left side of the middle row whose mask says, "I want safety, not freedom." I remember looking at that and being quite confused. I knew it was supposed to be a jab or an insult to that type of thinking, but I could not imagine actually taking that as an insult. I showed this to a few people in my own social circles, and most of them said, "You can't be free if you're dead." However, continuing to read through the tweets from this group, I can see how strongly they believe in the motto, "Live free or die." While, in its general use, the concept of freedom has become problematic, these enthymemes resonate with concepts of freedom shared by those participating in the anti-vaccine and anti-mask conversations.

While none of the tweets I captured actually used that motto, I did see many tweets which relied on an understanding of freedom and its value as the people participating in the anti-vaccine conversation understand it. An example of this can be seen in the tweet from Jennifer Margulis I quoted earlier in this chapter, where she said, "The erosion of freedom is far more threatening to our lives than a virus." Similarly, on July 4, 2020, Sally tweeted, "Take charge of your own life... Be brave and live free." The general sense I get from the collection of

their tweets is that they would be happy to knowingly expose themselves and others to danger so long as they can maintain their sense of what they perceive to be their freedom.

I would also probably include that this group also values freedom over social responsibility. I don't believe that, at this point, we can argue against the value of freedom, no matter how grave the danger may be and no matter how vital the role of social responsibility. So, that leaves us with trying to construct an argument which will lead to seeing the act of getting vaccinated as a way of exercising one's freedom and choice (I'm thinking something like the truth campaign which turned anti-smoking effort into a choice rather than a command).

The second common focus I've been seeing in these enthymemes includes two assumptions about the CDC: the first of these assumptions is that all government agencies will have ulterior motives that will not necessarily align with the public's best interests. Those motives would include, but not be limited to, the survival of the agency itself and securing future funding. The second basic assumption is that the CDC is an inextricable part of the pharmaceutical industrial complex (my words – not theirs). I believe these assumptions are implied largely by many of the tweets discussing the economics of the situation. One example I've already mentioned in this chapter was Toby Rogers' tweet about "Pharma" spending billions of dollars to "distort the science." Along similar lines, on September 12, 2020, Truth Lover tweeted, "The CDC has been manipulating the data on vaccine safety for years and the media is fully complicit." A few days later, on September 17, Truth Lover added, "They will do anything to protect their cash machine. They don't care at all about safety or protecting the public." And, also on September 12, 2020, Eileen Iorio tweeted, "Remember the CDC press kit

on how to increase demand for flu vaccines?? Create a panic in the media. Use fear mongering words in headlines, create anxiety and fear.”

I do believe the concerns behind these two assumptions about the CDC have some legitimacy and that we cannot ethically dismiss them out-of-hand. Therefore, in order to match a sense of dialogue through enthymeme, we need to construct an argument which begins with those two assumptions and still leads, logically, to the conclusion that we should listen to the advice of the CDC. As the effectiveness of such an argument lies outside the scope of this study, in my final chapter, I will look at what we can be doing now, from a pedagogical standpoint.

Implications and Conclusions

As I mentioned in the introduction to this dissertation, my ultimate question was asking what implications any findings from this study may have on how we approach the curriculum of a first-year composition class. Traditionally, we have taught our students to write in an authoritative voice and encouraged them – as they become experts in their fields – to continue to use that authoritative voice. However, I imagined that, if the hypothesis on a digital audience’s favoring of a more empathetic *ethos* proved true, then we might have wanted to rethink the approaches we take to teaching informative or persuasive writing. If, as boyd argued, strangers who are willing to listen and empathize are seeming more trustworthy and persuasive than impersonal advice from experts, it may be that we should be teaching and encouraging students to write with a voice of authenticity and transparency that Gardner linked to this type of approach to credibility. That would likely, in turn, lead to new perspectives on academic writing as well.

In approaching this topic, I had done a lot of imagining of what a dialogic, academic *ethos* might look like. I imagined three main components: more leniency for epistemic hedging, more self-referential language, and more occasions of speaking directly to one’s reader - possibly going so far as to encourage follow-up interaction after a reader has read the paper. The first two components are ones that seem natural for many students, and composition teachers have spent a lot of effort trying to break them of these habits. For example, any time a student would write “I think” in an essay, I would tell them that phrasing is a great way for organizing their thoughts in a rough draft but that they should remove the “I think” from

subsequent drafts and allow the statement to stand on its own. It's also worth noting that popular word-processing programs, like Grammarly and Microsoft Word, frequently make similar suggestions, prompting users to remove epistemic hedging words from their writing in order to make themselves sound more confident – more authoritative. This tendency towards epistemic hedging and self-referential language are key components to building an authentic and (seemingly) transparent voice in writing. They allow the author a way of being more present in their writing, and they give the reader that sense of a human voice being behind the writing. They suggest a closer, more personal connection between author and audience. Additionally, bringing this language into a first-year composition class could give us an opportunity to explore other issues of subjectivity and recipient design in academic writing.

The third component, taking more occasions to speak directly to one's readers, will not only go further towards reinforcing the idea of a more personal connection between author and audience, but could also be a logical move if we wanted to take academic conversations into a more conversational type of social media direction – some form of a cross between Academia.edu and Twitter. One thing that has really appealed to me about academic discourse (at least in the discipline of Rhetoric) has been the small-community feel and the perspective of research and publication as being like a conversation. For my own work, when presenting at a conference, I presented alongside friends who stood up at a lectern and spoke with authority. I, on the other hand, sat with my audience, saying things like, "This is what I do in my class; tell me if you do something different." The reaction I got from that seemed very positive – I got many more comments and questions at the end of the presentation than my authoritative friend. A social media platform dedicated to academic conversations could be a natural place

to practice a more human, personal voice in academic discourse, similar to what I did in my conference presentation. Many of my friends are reluctant to pose serious, academic questions on platforms like Facebook¹³, but a dedicated platform could not only encourage more academic discussion but could foster more-personal connections between scholars and audiences (I may have to suggest this to my programmer friends).

As my findings suggest that a one-unit increase in dialogic certainty (going from 0% dialogic certainty all the way to 100% in dialogic certainty) predicts a 0.00019 increase in engagement rate – roughly one additional like or retweet for every 5,000 followers an account has - dialogic language only seems to be preferred by a slight margin. I do not feel that any immediate paradigm shifts in how we conceptualize academic *ethos* are warranted in the teaching of first-year composition. It is certainly worth mentioning when talking about adapting writing to a specific audience, and I will discuss that more later in this chapter, but it doesn't necessitate a system-wide upheaval of expectations for voice in academic writing. These findings, however, could still be significant for other writing and communications classes – those more concerned with persuasion and interaction with the public outside of academic circles. Before committing to that, one way or the other, the findings do tell me that further research would be worthwhile.

¹³ I tried this once: At a time when we were discussing concepts of agency in a graduate class, I posed a few situations to my Facebook friends, including: "The baby sleeps peacefully when John is in the room. John doesn't need to do anything but be there. Jane notices the baby is awake and sends John into the room. The baby falls asleep. Who put the baby to sleep?" And, then, I asked them to imagine the same situation but John was actually a teddy bear. In response, I got no serious answers, only jokes and non-sequiturs.

Limitations of this Study and Possibilities for Further Research

In looking back over my results, there were a few pieces which did not fully satisfy me. My first issue was with the accuracy of the final machine learning evaluation we ran. I was very optimistic after we ran our test set – the machine learning seemed to capture the dialogic and authoritative distinction well. After we ran the full dataset, I made a quick check through the results but saw no reason to really question them. As I went further and further into my analysis, I began noticing some discrepancies between the results of the test run and the results of the full dataset run. The full dataset run still accurately captured all of the “I’m so sorry” and “thanks for sharing” tweets as dialogic, but it seemed to miss some of the more complicated tweets which I would have classified as dialogic. I would not say this was a large issue; in scanning through the first 1,000 tweets that had scores that were higher in favor of authoritative certainty by a margin of less than 0.1 (there were a total of 6,687), I only found around 20 I was concerned with.

One discrepancy that particularly bothered me was with @stopvaccinating’s tweet from August 30th, which I specifically called out in chapter 2. In that tweet, Larry Cook (@stopvaccinating) asks, “Has a pediatrician ever bullied, harassed, or demeaned you for asking about vaccine safety or for refusing to vaccinate?” In our test run, this tweet scored a dialogic strength of 0.909 and an authoritative strength of 0.260. As I felt this tweet invited participation and involvement by encouraging a response (regardless of how strongly the question favors a certain response) while also carrying within it a sense of empathy in expressing concern, I was happy that it was captured as high in dialogic certainty. This was one result that gave me confidence in the Bag of Words method. However, I later discovered that,

in our full dataset run, this tweet actually scored a 0.6878 in authoritative certainty and only a 0.1422 in dialogic certainty. So, the first thing that I would like to do, had I the time to rerun the data and redo the analysis, would be to find out what caused the discrepancies between our test run and our full dataset run and what we needed to do in order to replicate the results of our test run. In theory, we were training the machine in the same way each time. So, the system should have rated tweets in the full run the exact same way as it did in the test run, as the system does not do additional learning beyond the training set.

More generally speaking, I would also like to see the results of our Bag of Words learning refined. In looking at other tweets in the full run, I was very satisfied with how the machine learning captured tweets which it rated either high on the dialogic scale or high on the authoritative scale. I was, however, less satisfied with some of the determinations on tweets with mid-range certainty levels. As an example, on August 27, 2020, Toby Rogers (@uTobian) tweeted out to another user, “Apologies, I'm not the right person to answer that. Perhaps there is a naturopath or integrative/functional medicine doc on this thread that might want to weigh in?” The machine learning classified this with a slightly higher authoritative certainty rating than dialogic certainty – it gave this a 0.5231 authoritative certainty and a 0.5197 dialogic certainty. I would like to have seen this tweet rated much higher in dialogic certainty. Conversely, on September 8, 2020, Michelle Malkin tweeted, “When parental sovereignty is undermined, national sovereignty is undermined.” The machine learning rated this tweet in favor of being dialogic, giving it an authoritative certainty of 0.3781 and a dialogic certainty of 0.5135. I would have preferred to see this falling more on the authoritative side. These are just a couple examples, but I found others. I would like to see how tweets like these would have

been captured in the way we did our test run (neither were actually in the test set). Should it still be an issue, we could always pull more key tweets like these into our learning set to help refine the rating results. Doing so would, unfortunately, remove these key tweets from the data set (we cannot reuse any tweets used for the training within the actual dataset), so I might gather training tweets from other times or accounts.

Even if we had the opportunity to refine our Bag of Words results, the method itself does have some limitations. From the very beginning of our testing, there was a part of me that felt unsatisfied in reducing a style of writing just down to the issue of word choice. The Bag of Words method, obviously, only looks for the presence of certain words. It does not consider other features of language use or grammatical structures. I do know there are other, more complicated methods of language processing which can be used for machine learning. Word2Vec Embeddings is similar to BoW in that it looks at words, but it focuses more on word groupings rather than the mere presence of the words. There are Convolutional Neural Network (CNN) systems which look at hierarchical relationships within data. There's also Bidirectional Encoder Representations from Transformers (BERT) which, I believe, looks at contextual embedding through elements like word order. While these other methods may not necessarily be any more or less promising than Bag of Words on their own, I think we might get much more accurate ratings by combining multiple methods of machine learning and combining the scores from the different methods. Doing it that way, we would still use word choice as a method of evaluating tweets, but we would not have to limit ourselves to relying solely on word choice in evaluating style.

In thinking beyond computer learning, you may recall, from what I discussed in chapter 2, my original inclination in looking at how a dialogic style influences a reader's sense of an author's credibility was to set up more of an artificial writing environment and record reader responses through survey questions - I had thought about writing my own series of microblog statements, showing those statements to a group of survey participants, and asking those participants to rate each tweet based on various statements about the tweet's level of credibility. The statements would have been written to convey similar messages but be written in the different styles I was looking to examine. Ultimately, we decided that divorcing the messages from the greater context of the Twitter environment would have been too artificial and could produce misleading results. With the current study, one thing I found which I have not yet talked a lot about but could turn out to be one of the more significant findings was what was suggested about the coefficient of determination for the use of dialogic or authoritative style.

As mentioned in chapter 3, the data from this study shows that the coefficient of determination for the relationship between authoritative or dialogic certainty and estimated engagement rate is roughly 10% (11.7% for authoritative certainty and engagement, and 10.2% for dialogic certainty and engagement) for the tweets I collected. This means that, at the most, ten percent of any change in engagement rate can be attributed to changes in authoritative or dialogic styles within this type of science-denialist tweets. The other 90% is what comes from those other influences that are part of the Twitter environment. As discussed in chapter 2, studying this within the context of a live Twitter environment was really a way of forcing us to take any other influences of the Twitter environment into account. Now that we have that

coefficient of determination, and we know that these styles of building *ethos* can be responsible for, at the most, ten percent of potential change in engagement rate of vaccine-related tweets, can we now, more safely take the study into a more artificial environment and study reader impressions of authoritative and dialogic *ethea* without the influence of any other variables? Doing so, we can get reactions more clearly tied to the influence of this style while still understanding that we're only looking at a potential ten percent of the total engagement picture.

Beyond these refinements and other ways of looking at the influence of using dialogic language, there are other pieces of this puzzle that may be worth further study. In chapter 4, I mentioned a possible way to look at the evolution of these styles as an author builds followers, specifically to answer the question of if an authoritative author attracts more followers or if having more followers encourages the development of a more authoritative style. I suggested identifying newly created accounts with a low number of followers which seem to have potential for long-term involvement in the anti-vaccine conversation and then following those accounts from their inception and through their growth and development over the course of five to ten years. Assuming some of those accounts showed evidence of a progression in their language from less authoritative to more authoritative along with a growth in their number of followers, I would try to determine if their following grew after the evolution of their language or if their language evolved after their following grew.

While I find this to be an interesting question, I will admit that the full run of machine learning on my dataset does not seem to show the same evidence of a correlation between following size and the frequency of using a dialogic style that the data from my initial readers

suggested. If we recall from chapter 4, the findings that first inspired this questioning was that the Twitter users with fewer followers seemed to be responsible for more of the dialogic tweeting. The one-third of the accounts with the most followers on my list were responsible for 26% of the dialogic tweets my readers identified. The middle-third of the accounts that I was looking at were responsible for 31% of the dialogic tweets. And, the third of the accounts with the fewest followers were responsible for 42% of the dialogic tweets. If I break down the numbers from the full dataset in the same way that I approached the data from the initial 1000 tweets my human readers looked at, this is what I get:

Name	Followers	Total Dialogic	Total Authoritative	Ratio
Michelle Malkin	2,200,000	332	1874	1:5.6
Dr Joseph Mercola	290,500	293	2383	1:8.1
Frank Lipman MD	75,400	153	2927	1:19.1
Del Bigtree	47,600	63	316	1:5
Dr Sherri Tenpenny	43,800	399	2169	1:5.4
LotusOak	40,100	31	1506	1:48.6
The HighWire	37,800	199	2257	1:72.8
Childrens Health Defense	31,000	166	1921	1:11.6
Toby Rogers PhD, MPP	19,400	570	2253	1:4
Generation Rescue	18,500	997	1766	1:1.8
Barb Loe, NVIC	16,300	160	2780	1:17.4
Larry Cook	14,300	270	1978	1:7.3
sally	11,400	313	1420	1:4.5
Physicians for Info	11,400	87	1060	1:12.2
Inside Vaccines	10,900	586	2293	1:3.9
Vaxxed II: The People's Truth	10,700	111	826	1:7.4
Jefferey Jaxen	10,300	172	1907	1:11.1
Eileen Iorio	10,100	392	2174	1:5.5
Catie	7,500	465	1868	1:4
Noforcedvaccination	5,700	94	1321	1:14.1
Jennifer Margulis	4,900	370	2169	1:5.9
Vaxxed_Supporter	3,100	194	1457	1:7.5

Truth Lover	3,000	224	2192	1:9.8
Wayne Rohde	2,900	363	2165	1:6
One Pissed Off Mom	2,800	305	1245	1:4.1

Looking primarily at the ratios of the tweets, this does not seem to immediately suggest the same relationship of the most dialogic tweets necessarily coming primarily from the accounts with the fewest followers. However, I did not ask a statistician (my uncle) to look at these numbers, so there could very well be a pattern that I'm just not seeing.

One other thing that we could still do, even if just with the data I already have, would be to get the top features (in this case, words) that the machine learned to associate with authoritative and dialogic styles. Using surrogate models, such as SHAP or LIME, we could figure out which words are most strongly associated with each of the two styles. These surrogate models are used to make slight changes to the data and to measure the resulting changes in classification. We can then identify the most influential features to, in essence, deconstruct the authoritative and dialogic styles. As we've also coded emoji use in this, it also opens up ways of studying emoji use as part of these styles.

If we wanted to look beyond the authoritative and dialogic binary, it could be worthwhile to see what other linguistic style choices might have an influence on engagement rate within the data I've gathered. We could look at the highest performing tweets and break them down to look for other features and patterns we might see in the language and compositional elements. There may be other stylistic choices being made and attracting followers that would be worth looking into.

In retrospect, I feel like I'm seeing what I've done here more as a first step in a larger question that asks us to revisit some assumptions boyd, Gardner, and others have made about authenticity and credibility in digital spaces. I've established that there is a relationship (however weak) between the strengths of authoritative and dialogic language in a tweet and the performance of that tweet, as measured by engagement rate. I've also established that altering that language may account for up to ten percent of that tweet's performance. Based on these results, we can see the issue is more complicated than boyd suggested. However, as a composition teacher and a student of rhetoric, I still need to hang my hat on the idea that there are compositional elements at play beyond the at-a-glance features and confirmation bias as discussed earlier. Now, in order to really determine what this means for the teaching of persuasive writing and communication, we should still be looking further into how the authoritative/dialogic distinction works, how those styles may be better crafted, and what other stylistic elements may be contributing.

Without getting too far ahead of myself in directions to take this research, I would like to return to the question I opened this chapter with how this study might influence my approach to teaching first-year composition. Despite the slim margin by which I've found dialogic language to be preferred over authoritative language, the distinction between the two remains a concept worth discussing in a first-year composition course. As writers, we should still be conscious of the differences between dialogic and authoritative approaches and when each may be more appropriate. In the classroom, I will certainly want to include this in any discussions about tailoring language to specific audiences. Instead of automatically asking students to remove the epistemic hedging language and to write with more confidence in their

papers, we need get them to reflect on the value of authoritative or dialogic approaches with respect to their intended audiences. I already tend to focus heavily on getting students to think about purpose and how they want their readers to react to their writing. Asking students to examine how a particular expression (which might seem to be a hedge) could be used in different types of writing to achieve different types of effects. Talking about the authoritative and dialogic approaches would make for a good way to directly connect language choice with intentions of either lecturing to or partnering with their readers. It could also lead to conversations on how writing can build community and foster a sense of empathy rather than simply delivering information.

However, I think the findings suggest that we cannot rely on teaching dialogic language as any kind of sure-fire way for doctors and scientists to engage the public. Therefore, our responsibility, as teachers, still has to be to approach the issue of post-truth misinformation and disinformation from a news consumer's position. This means teaching our students to be responsible media consumers – a concept that fits well within the research-focused curriculum of first-year composition.

As I opened this dissertation with reference to danah boyd's essay, "Did Media Literacy Backfire," her perspective is also a logical entry point to the current discussion. As her title suggests, boyd was skeptical of the value in our approach to teaching media literacy. She does feel that teaching critical thinking skills that get students to question information production and dissemination are still important. However, boyd says, "We cannot fall back on standard educational approaches because the societal context has shifted." Boyd suggests that we, somehow, need to build the "social infrastructure" around our approach to information

gathering, but she does not discuss what this may look like. Recent research on the impact of teaching medial literacy has shown that boyd's skepticism is well founded.

Concepts in Media Literacy

Until recently, many scholars have taken a broad definition of media literacy, but it usually entailed a critical understanding of the production and dissemination of stories in the media. They have also used terms like media literacy, digital literacy, and even multiliteracy, if not interchangeably, then with much overlap. David Buckingham, author of *The Media Education Manifesto*, even admits that conceptualizing media literacy "often seems to be more of a rhetorical gesture than a concrete commitment" (29). S. Mo Jones-Jang compiled various definitions of and approaches to media-related literacies and broke them down into four main fields: media literacy, digital literacy, news literacy, and information literacy.

Jones-Jang identifies media literacy as the ability for people to access, analyze and produce informational stories through various forms of media. He further distinguishes media literacy through four precepts: first, media literate people will recognize that information in the media are shaped by perception of an event and will further shape others' perceptions of the event. Secondly, media literate people will understand that messages in the media are created in response to commercial, ideological, or political motivations. Third, they will appreciate the fact that each media medium will have its own unique conventions which are designed to meet unique expectations of the viewers/readers/listeners of that medium. Finally, they will understand that those audiences negotiate the meaning of the messages being disseminated.

Digital literacy, which Jones-Jang also cites as being referred to as online literacy and new media literacy, seems primarily focused on two elements: the first of those elements is simply the ability for people to adapt to new technologies. This is concerned with how easily a person can learn to navigate and participate in new platforms while also learning the language, terminology, and conventions of each. The second element in digital literacy is the understanding of the participatory nature of digital media. This further entails an appreciation of the significance and effects of user feedback, reactions, and other forms of user-generated content.

Similar to media literacy, news literacy begins with the recognition that news stories are produced with the intent of meeting commercial goals and respond to outside influences. Beyond that, news literacy involves developing the ability to find and recognize news stories. It also includes the ability to critically evaluate the message being conveyed within the news story.

Where media literacy, digital literacy, and news literacy are primarily knowledge-based skills, Jones-Jang distinguishes information literacy as being more of an active skill. When Jones-Jang discusses information literacy, he focuses primarily on a person's ability to find and retrieve information. This also entails the ability to compile a wide variety of different perspectives on a given message or story. So, where the other three literacies discussed emphasize the understanding and analysis of messages and the context within which they are produced, information literacy emphasizes the gathering of information.

Turning specifically to multiliteracies pedagogy, Jeff Share and Tatevik Mamikonyan design a course framework guided by six conceptual understandings with corresponding

questions to guide them which seems to cover all of the common concerns about the critical thinking piece of these literacies. Those concepts and questions are:

1. Social Constructivism – Who are all the possible people who made choices that helped create this text?
2. Languages/Semiotics – How was this text constructed and delivered/accessed?
3. Audience/Positionality – How could this text be understood differently?
4. Politics of Representation – What values, points of view, and ideologies are represented or missing from this text or influenced by the medium?
5. Production/Institutions – Why was this text created and/or shared?
6. Social and Environmental Justice – Whom does this text advantage and/or disadvantage? (41)

This framework addresses the common concerns of production and dissemination that many theorists and teachers have focused on up until recently.

Buckingham, however, cautions that this common approach to the teaching of media literacies tends to take on a “protectionist” approach (66). He observes that this approach often fails in its objectives primarily because students find it patronizing. As such, they will pay it lip-service in the classroom and then disregard anything they learned when in real-life situations. Buckingham suggests that media education should be guided by two key principals: first, that any class on media literacies needs to focus on students’ direct experience with media and begin with an open discussion on why they use different media and how they communicate through that media. What Buckingham seems to be getting at here is a way of ensuring an unbiased attempt to get students to interrogate their own media choices rather than falling into the trap of imposing a single, authoritative perspective. Secondly, he feels that

a class on media studies should teach through the creative production of the students' own messages. "Production," he says, "can offer a space to reflect on the personal and emotional dimensions of media use, and this can then feed back into critical analysis" (73).

I don't believe that the traditional approach suggested by Share and Mamikonyan and the approach suggested by Buckingham are mutually exclusive, nor even that they are necessarily conceptually divergent. Either way they both rely on inoculation theory – that exposing students to these issues in the classroom will carry over into the real world – and on strengthening the knowledge and critical thinking skills involved in media, digital, and news literacies. More recent research, however, suggests these approaches may not be as effective as we have been hoping.

Recent Studies on the Efficacy of Media-Related Literacies

In studying the long-term effectiveness of media literacies education, Michael Hameleers and Andrew Guess, *et al.* both did studies which measured student responses to fake news, and both seemed to come to the same conclusions. First, they both found that media literacies education resulted in a general distrust of all media sources by students. That distrust was somewhat more pronounced when dealing with fake news sources, but did still extend to all legitimate sources. Their main finding, however, again as in both studies, was that, even when students were able to identify false information, those students were still influenced by the message supported by that misinformation.

The findings show that exposure to a media literacy message significantly lowers the perceived accuracy of misinformation... However, exposure to a media literacy message

does not result in lower levels of agreement with the statements made in misinformation... Although news media literacy interventions can lower the perceived accuracy of misinformation, they do not decrease the overall levels of agreement with communicative untruthfulness (Hameleers 10).

So, it appears that the critical thinking skills fostered by teaching media literacy were not effective in combating the message conveyed through misinformation.

Another study performed by Jones-Jang, *et al.* provides a little more insight into those findings but also showed some positive effects under very specific conditions. Jones-Jang performed a similar experiment but did so by conducting separate trials for each of the four literacies described in the previous section: media literacy, digital literacy, news literacy, and information literacy. His findings match what Hameleers and Guess found, but only when looking at media, digital, and news literacy. Where his results greatly differed was in testing the group that practiced information literacy skills. The group that was taught information literacy showed a greatly improved ability to identify and reject messages based on misinformation. As previously mentioned, the difference here was that media, digital, and news literacies focused instruction on passive knowledge and understanding skills, whereas information literacy focused instruction on actively gathering multiple sources of information. To explain this difference, Jones-Jang cites the Dunning-Kreuger effect, suggesting that the students who were the most confident in their media literacy skills were the least likely to exercise those skills.

Based on his findings, it seems like logical reasoning to me that, if one is only exposed to one perspective, that that perspective would likely be adopted – even when the evidence supporting that perspective is rejected. Whereas, if one is able to gather multiple perspectives,

one can make better judgements on which perspectives to adopt – even when presented with misinformation. Jones-Jang concludes that, “media-related literacy concepts should place greater emphasis on the process of effective gathering of truthful and verified information, as well as the concept of evaluating any information regarding its authoritativeness and credibility” (382).

Conclusions

I began this dissertation with hopes of finding evidence that adopting a more dialogic *ethos* would be a way for experts to reach hesitant or hostile readers in order to open up a more constructive conversation than the flame-war we see so frequently today. While I did find some evidence to support this, that evidence only suggested a slight benefit and certainly not the “seismic shift” Howard Gardner alluded to. For the moment, it seems that ownership of this fight still has to lie with the way we are educating our students to consume their news. While many teachers have found frustration in addressing media literacy, there may be hope in a shift to the active skills of information literacy and providing students more practice in engaging with media and controversial or otherwise polarized issues. The benefits of dialogic language are still worth examining, but it will not be any kind of magic bullet with which to arm doctors and scientists as an updated approach to digital *ethos*.

References

- ADL.org. 5 4 2019. https://www.adl.org/blog/anti-vaccine-protesters-misappropriate-holocaust-era-symbol-to-promote-their-cause?fbclid=IwAR2jEbFw_3VU5nv_OTIJurjA3jAGwgtS7koa8hy804hzJJhjcQj827epy7s. 21 6 2021.
- ADL.org. 5 7 2020. <https://www.adl.org/blog/michelle-malkin-is-attempting-to-normalize-white-supremacy>. 27 7 2021.
- Androutsopoulos, Jannis. "Negotiating Authenticities in Mediatized Times." *Discourse, Context & Media* (2015): 74-77.
- Aristotle. *The Art of Rhetoric (translated by H.C.Lawson-Tancred)*. London: Penguin Classics, 1991.
- Bischoff, Laura A. *Dispatch.com*. 5 9 2021. <https://archive.ph/20210609201649/https://www.dispatch.com/story/news/2021/06/09/doctor-sherri-tenpenny-testimony-ohio-lawmakers-vaccines-magnetized-5-g/7616027002/>. 21 6 2021.
- Boudry, Maarten and Johan Braeckman. "Immunizing Strategies and Epistemic Defense Mechanisms." *Philosophia* (2011): 145-161.
- boyd, danah. "Did Media Literacy Backfire?" 5 Jan 2017. points.datasociety.net/did-media-literacy-backfire-7418c084d88d. Web. 16 Jan 2017.
- Brooke, Collin Gifford. "New Media Pedagogy." Brooke, Collin Gifford et al. *A Guide to Composition Pedagogies, 2nd Edition*. Oxford University Press, 2013. 177-193.

- Bucholtz, Mary. "Sociolinguistic Nostalgia and the Authentication of Identity." *Journal of Sociolinguistics* (2003): 398-416.
- Buck, Amber. "Examining Digital Literacy Practices on Social Networki Sites." *Research in the Teaching of English* (2012): 9-38.
- Buckingham, David. *The Media Education Manifesto*. Cambridge: Polity Press, 2019.
- Colbert, Stephen. "The Colbert Report." Comedy Central, 18 October 2005.
- Court, Andrew. *DailyMail.com*. 7 5 2020. <https://www.dailymail.co.uk/news/article-8604639/Dentists-says-mask-wearing-causing-tooth-decay-gum-disease.html>. 21 6 2021.
- Elliott-Maksymowicz, Katarzyna, Alexander Nikolaev and Douglas Porpora. "How Much Can You Say in a Tweet? An Approach to Political Argumentation on Twitter." *Humanities & Social Sciences Communications* (2021): 1-9.
- Erstad, Ola and Julian Sefton-Green Eds. *Idnetity, Community and Learning Lives in the Digital Age*. New York: Cambridge University Press, 2013.
- Flanagin, Andrew and Miriam Metzger. "The Perceived Credibility of Personal Web Page Information as Influenced by the Sex of the Source." *Computers in Human Behavior* (2003): 683-701.
- Frenkel, Sheera. *New York Times.com*. 24 7 2021.
<https://www.nytimes.com/2021/07/24/technology/joseph-mercola-coronavirus-misinformation-online.html>. 18 9 2021.
- Gardner, Howard. *Truth, Beauty, and Goodness Reframed: Educating for the Virtues of the Twenty-First Century*. New York: Basic Books, 2011.

- Gilbert, Richard, Jessica Foss and Nora Murphy. "Multiple Personality Order." Peachy, Anna and Mark Childs. *Reinventing Ourselves: Contemporary Concepts of Identity in Virtual Worlds*. New York: Springer, 2011. 213 - 234.
- Goffman, Erving. *The Presentation of the Self in Everyday Life*. New York: Anchor Books, 1959.
- Gray, Bethany and Douglas Biber. "Current Conceptions of Stance." Hyland, K. and C. Sancho eds. *Stance and Voice in Written Academic Genres*. Guinda: Palgrave Macmillan Limited, 2012. 15-33.
- Green Our Vaccines*. By Jennifer McCarthy. 4 5 2008.
- Gu, Zhongyi et al. "A Vaccine Crisis in the Era of Social Media." *National Science Review* (2018): 8-10.
- Guess, Andrew M. et al. "A Digital Media Literacy Intervention Increases Discernment Between Mainstream and False News in the United States and India." *PNAS* (2020): 15536-15545.
- Hameleers, Michael. "Separating Truth from Lies: Comparing the Effects of News Media Literacy Interventions and Fact-Checkers in Response to Political Misinformation in the US and Netherlands." *Information, Communication & Society* (2020).
- Han, Kyungsik. "How Do You Percieve This Author? Understanding and Modeling Authors' Communication Quality in Social Media." *PloS one* (2018): e0192061.
- Hussain, Azhar, Syed Ali, Madiha Ahmed and Sheharyar Hussain. "The ANti-Vaccination Movement: A Regression in Modern Medicine." *Cureus* (2018): e2919.
- Ickes, William. "Empathic Accuracy." *Journal of Personality* (1993): 587-610.
- Izadi, Ahmad. "The Epistemic Grounds of Face in Institutional Argumentative Talk-in-Interaction." *Journal of Pragmatics* (2018): 45-56.

Jackson, Jasper. *TheGuardian.com*. 8 2 2017.

<https://www.theguardian.com/technology/2017/feb/08/wikipedia-bans-daily-mail-as-unreliable-source-for-website>. 21 6 2021.

Johnson, Neil F. et al. "The Online Competition Between Pro- and Anti-Vaccination Views." *Nature* (2020): 1-5.

Jones, John. "Social Media and Persuasion: Crowdsourcing Arguments on Digital Networks." Noor Al-Deen, Hana and John Hendricks. *Social Media Usage and Impact*. Lanham: Lexington Books, 2012. 23-38.

Jones-Jang, S. Mo, Tara Mortensen and Jingjing Liu. "Does Media Literacy Help Identification of Fake News? Information Literacy Helps, but Other Literacies Don't." *American Behavioral Scientist* (2021): 371-388.

Kabani, Shama. *The Zen of Social Media Marketing*. Dallas: BenBella Books, Inc., 2013.

Kärkkäinen, Elise. "I Thought It Was Very Interesting. Conversational Formats for Taking a Stance." *Journal of Pragmatics* (2012): 2194-2210.

Kerpen, Dave. *Likeable Social Media*. New York: McGraw Hill Education, 2015.

Kreps, David. "My Social Networking Profile: Copy, Resemblance, or Simulacrum? A Poststructuralist Interpretation of Social Information Systems." *European Journal of Information Systems*, Vol. 19 (2010): 104-115.

Lachlan, Kenneth, Patric Spence and Autumn Edwards. "If You Are Quick Enough, I Will Think about It: Information Speed and Trust in Public Health Organizations." *Computers in Human Behavior* (2014): 377-380.

- Levenson, Robert W. and Anna M. Ruef. "Empathy: A Physiological Substrate." *Journal of Personality and Social Psychology* (1992): 234-246.
- Lipson, Peter. *ScienceBasedMedicine.org*. 28 9 2009.
- <https://sciencebasedmedicine.org/huffington-post-is-at-it-again/>. 27 7 2021.
- Massachusetts Department of Public Health. *Mass.gov*. 19 8 2020.
- <https://www.mass.gov/news/flu-vaccine-now-required-for-all-massachusetts-school-students-enrolled-in-child-care-pre-school-k-12-and-post-secondary-institutions>. 26 6 2021.
- McComiskey, Bruce. *Post-Truth Rhetoric and Composition*. Logan: Utah State Press, 2017.
- Mcintyre, Lee. *Post-Truth*. Cambridge: The MIT Press, 2018.
- Menczer, Filippo and Pik-Mai Hui. *TheConversation.com*. 2 4 2019.
- <https://theconversation.com/anti-vaxxers-appear-to-be-losing-ground-in-the-online-vaccine-debate-114406> . 21 6 2021.
- Mercola.com*. n.d. 27 7 2021.
- Metzger, Miriam and Andrew J. Flanagin. "Credibility and Trust of Information in Online Environments: The Use of Cognitive Heuristics." *Journal of Pragmatics* (2013): 210-220.
- Morris, Meredith, et al. "Tweeting is Believing?: Understanding Mircoblog Credibility Perceptions." *International Conference on Computer-Supported Cooperative Work and Social Computing*. ACM, 2012. 441-450.
- Muckelbauer, John. *The Future of Invention: Rhetoric, Postmodernism, and the Problem of Change*. Albany: State University of New York Press, 2008.

- Nambisan, Priya. "Information Seeking and Social Support in Online Health Communities: Impact on Patients' Perceived Empathy." *J Am Med Inform Assoc* (2011): 298-304.
- New York State Education Department. *NYSED.gov*. n.d.
<http://www.op.nysed.gov/prof/dent/dentcounts.htm>. 21 6 2021.
- Nichols, Tom. *The Death of Expertise*. New York: Oxford University Press, 2017.
- Osatuyi, Babajide. "Information Sharing on Social Media Sites." *Computers in Human Behavior* (2013): 2622-2631.
- Palfreman, Jon. "The Vaccine War." *Frontline*. 27 4 2010.
- Pérez-Paredes, Pascual and Camino Bueno-Alastuey. "A Corpus-Driven Analysis of Certainty Stance Adverbs: Obviously, Really and Actually in Spoken Narrative and Learner English." *Journal of Pragmatics* (2019): 22-32.
- Preece, Jenny and Kambiz Ghazati. "In Search of Empathy Online: A Review of 100 Online Communities." *AMCIS* (1998): 92-94.
- Preece, Jenny. "Empathic Communities: Balancing Emotional and Factual Communication." *Interacting with Computers* (1999): 63-77.
- . "Empathic Communities: Reaching Out Across the Web." *Interactions* (1998): 32-43.
- Sagolla, Dom. *140 Characters: A Style Guide for the Short Form*. Hoboken: John Wiley & Sons Inc., 2009.
- Searls, Doc and David Weinberger. "Markets are Conversations." Levine, Rick. *The Cluetrain Manifesto: The End of Business as Usual*. Cambridge: Perseus Books, 2000.

- Share, Jeff and Tatevik Mamikonyan. "Preparing English Teachers with Critical Media Literacy for the Digital Age." *Contemporary Issues in Technology and Teacher Education* (2020): 37-54.
- Shariff, Shafiza, Xiuzhen Zang and Mark Sanderson. "On the Credibility Perception of News on Twitter: Readers, Topics and Features." *Computers in Human Behavior* (2017): 785-796.
- Srikanth, Anagha. *TheHill.com*. 24 3 2021.
<https://archive.ph/20210325074028/https://thehill.com/changing-america/well-being/prevention-cures/544712-twelve-anti-vaxxers-are-responsible-for-two#selection-337.0-337.15>. 21 6 2020.
- Stern, Alexandra Minna and Howard Markel. "The History of Vaccines and Immunization: Familiar Patters, New Challenges." *Health Affairs* (2005): 611-621.
- Stevens, Jeffrey Edward and Melanie Fridl Ross. "Social Media: Helping Health Systems Build Empathy and Engagement." *AMWA Journal* (2016): 124-128.
- Stirling, Lesley and Lenore Manderson. "About You: Empathy, Objectivity and Authority." *Journal of Pragmatics* (2011): 1581-1602.
- Turkle, Sherry. *Alone Together*. New York: Basic Books, 2011.
- Vaxxed2.com. n.d. <https://www.vaxxed2.com/>. 21 6 2021.
- Vogel, Lauren. "Medicine Slow to Recognize Social Media as Window into the Patient Experience." *CMAJ* (2019): e87-e88.
- Westerman, David, Patric Spense and Brandon Van Der Heide. "A Social Network as Information: The Effect of System Generated Reports of Connectedness on Credibility on Twitter." *Computers in Human Behavior* (2011): 199-206.

Whitehead, Niel L. and Michael Wesch eds. *Human No More: Digital Subjectivities, Unhuman Subjects and the End of Anthropology*. Boulder: University Press of Colorado, 2012.

Wysocki, Anne Frances, et al. *Writing New Media*. Logan: Utah State University Press, 2004.

Yang, Jiang, et al. "Microblog Credibility Perceptions: COmparing the USA and CHina."

International Conference on Computer-Supported Cooperative Work and Social Computing. ACM, 2013. 575-586.

Young, Kirsty. "Managing Online Identity and Diverse Social Networks on Facebook." *Webology*, Vol 10, No. 2, Article 109 (2013): 1-18.

Zhao, Jing et al. "Trust, Empathy, Social Identity, and Contribution of Knowledge Within Patient Online Communities." *Behavior & Information Technology* (2013): 1041-1048.

Appendix: Machine Learning Code

```
In [17]: import numpy as np # linear algebra
import pandas as pd # data processing, CSV file I/O (e.g. pd.read_csv)

import re
import pickle
#from emot.emo_unicode import UNICODE_EMO, EMOTICONS
import emoji

from statistics import mean

import string

#Sklearn Libraries
from sklearn.model_selection import train_test_split
from sklearn.feature_extraction.text import TextVectorizer
from sklearn.feature_extraction.text import CountVectorizer
from sklearn.multioutput import MultiOutputClassifier
from sklearn.naive_bayes import MultinomialNB
from sklearn.linear_model import LogisticRegression
from sklearn.metrics import roc_auc_score

#Plotting Libraries
import matplotlib.pyplot as plt
import seaborn as sns

#To set precision to 3 decimals
pd.options.display.float_format = "{:,.2f}".format

In [18]: train_text = pd.read_csv('training_tweets.csv')
test_text = pd.read_csv("testing_tweets.csv")

train_text = train_text.dropna()
train_text = train_text.rename(columns={"Rating": "y", "Tweet(Title)": "comment_text"})
#train_text = train_text[~train_text.y.str.contains("Reject")]
train_text['y'] = train_text['y'].astype(str)
train_text = train_text.replace({'y': {'1': 'authoritative', '2': 'dialogic', '3': 'other', '4': 'other'}})

train_text = train_text[['y', 'comment_text']]
one_hot = pd.get_dummies(train_text['y'])
train_text = train_text.drop(columns='y')
train_text = train_text.join(one_hot)
```

```

test_text = test_text.rename(columns={"Tweet(Title)": "comment_text"})

print(train_text.shape, test_text.shape)
result = train_text.head(10)

#drop usernames
train_text = train_text.replace(to_replace='@[^\\s]+', value='', regex=True)

#replace emojis
train_text['comment_text'] = (train_text.apply(lambda row: emoji.demojize(row['comment_text']), delimiters="

result = train_text.head(10)
display(result)

```

(3478, 4) (53908, 9)

	comment_text	authoritative	dialogic	other
0	Fascinating report by veteran Pam Long on the ...	0	0	1
1	Treatment of acute childhood diarrhea w/ #Home...	1	0	0
2	Elon Musk says he won't take #coronavirus vacc...	0	0	1
3	Okay I've read it now. And it's terrifying. W...	0	1	0
4	#Diabetes induced by #vaccine should not be co...	1	0	0
5	Mercury: Highly Toxic, Cumulative and Still in...	1	0	0
6	United_States oncoming_fist_medium-dark_ski...	0	0	1
7	The Democratic Party at the state level (prob...	1	0	0
8	REPLY New entry in my tracker thread: Of cou...	1	0	0
9	THE GREATER GOOD: Another MUST WATCH Documenta...	1	0	0

In [19]: *#Looking at the distribution of target variables*
y_cols = ["authoritative", "dialogic", "other"]
train_text[y_cols].apply(pd.Series.value_counts, args = (True, False))

Out[19]:

	authoritative	dialogic	other
0	0.30	0.96	0.74

	authoritative	dialogic	other
1	0.70	0.04	0.26

```
In [20]: #Dependent Variable
X = train_text.comment_text
#Independent Variables
y = train_text[["authoritative","dialogic","other"]]
#Splitting for checking the performance of the models on a holdout dataset
X_train, X_val, y_train, y_val = train_test_split(X, y, shuffle = True, random_state = 123)
```

```
In [21]: X_train.shape, X_val.shape
```

```
Out[21]: ((2608,), (870,))
```

```
In [22]: #importing stop words like in, the, of so that these can be removed from texts
#as these words dont help in determining the classes(Whether a sentence is toxic or not)
stop_words = text.ENGLISH_STOP_WORDS
#Function for basic cleaning/preprocessing texts
def clean(doc):
    # Removal of punctuation marks (.,/\)[{} etc) and numbers
    doc = "".join([char for char in doc if char not in string.punctuation and not char.isdigit()])
    # Removal of stopwords
    doc = " ".join([token for token in doc.split() if token not in stop_words])
    return doc.lower()
```

```
In [23]: vect = CountVectorizer(max_features= 5000, preprocessor=clean)
X_train_dtm = vect.fit_transform(X_train)
X_val_dtm = vect.transform(X_val)

print(X_train_dtm.shape, X_val_dtm.shape)
```

```
(2608, 5000) (870, 5000)
```

```
In [24]: pd.DataFrame(X_train_dtm.A[:5], columns = vect.get_feature_names())
```

```
Out[24]: aaps ability able about absentee absolute absolutely abundant abuse accept ... zev zinc zippermouthface zoo zoor
```


	aaps	ability	able	about	absentee	absolute	absolutely	abundant	abuse	accept	...	zev	zinc	zippermouthface	zoo	zoor
0	0	0	0	1	0	0	0	0	0	0	0 ...	0	0		0	0
1	0	0	0	0	0	0	0	0	0	0	0 ...	0	0		0	0
2	0	0	0	0	0	0	0	0	0	0	0 ...	0	0		0	0
3	0	0	0	0	0	0	1	0	0	0	0 ...	0	0		0	0
4	0	0	0	0	0	0	0	0	0	0	0 ...	0	0		0	0

5 rows x 5000 columns

```
In [25]: #Initializing and fitting models on Training Data
#Naive Bayes Model
nb = MultiOutputClassifier(MultinomialNB()).fit(X_train_dtm, y_train)
#Logistic Regression Model (As we have unbalanced dataset, we use class_weight which will use inverse
#of counts of that class. It penalizes mistakes in samples of class[i] with class_weight[i] instead of 1)
lr = MultiOutputClassifier(LogisticRegression(class_weight='balanced', max_iter=3000)).fit(X_train_dtm, y_train)
```

```
In [26]: #Function for calculating roc auc with given actual binary values across target variables
#and the probability score made by the model
def calculate_roc_auc(y_test, y_pred):
    aucs = []
    #Calculate the ROC-AUC for each of the target column
    for col in range(y_test.shape[1]):
        aucs.append(roc_auc_score(y_test[:,col], y_pred[:,col]))
    return aucs
```

```
In [27]: #Creating an empty list of results
results = []
#Making predictions from all the trained models and measure performance for each
for model in [nb, lr]:
    #Extracting name of the model
    est = type(model.estimator).__name__
    #Actual output variables
    y_vals = y_val.to_numpy()
    #Model Probabilities for class 1 of each of the target variables
    y_preds = np.transpose(np.array(model.predict_proba(X_val_dtm))[:, :, 1])
    #Calculate Mean of the ROC-AUC
```

```

mean_auc = mean(calculate_roc_auc(y_vals,y_preds))
#Append the name of the model and the mean_roc_auc into the results list
results.append([est, mean_auc])

#Output the results as a table
pd.DataFrame(results, columns = ["Model", "Mean AUC"])

```

Out[27]:

	Model	Mean AUC
0	MultinomialNB	0.79
1	LogisticRegression	0.83

In [29]:

```

# Merging the test dataset with sample_submission to have all the columns:
# id, text_data and the target variables in one dataframe
# df_test = pd.merge(test_text, sample_submission, on = "id")

#display (test_text)
df_test = test_text
display(df_test)

# Transform the test dataset as well based on Bag of Words/ Count Vectorizer as the Logistic model would
# expect the same
X_test_dtm = vect.transform(df_test["comment_text"])
#Use the Logistic Regression model to output probabilities and take the probability for class 1
y_preds = np.transpose(np.array(lr.predict_proba(X_test_dtm))[:,1])
#Assign the predictions by the model in the final test dataset
df_test[y_cols] = y_preds
display(df_test)
#Drop Comment Text as the sample_submission doesnt have it and wouldnt be expected
#df_test.drop(["comment_text"], axis = 1, inplace = True)
#Save the dataset as a csv to submit it
df_test.to_csv("sample_submission.csv", index = False)

```

	Label	EntryID	Published	comment_text	FavouriteCount	Retweets	UserFollowersCount	Lil
0	lotusOak2	1,310,580,000,000,000.00	Mon Sep 28 14:09:22 +0000 2020	The Popularity of #Homeopathy by Jeremy Sherr ...	11	8	43284	

Label	EntryID	Published	comment_text	FavouriteCount	Retweets	UserFollowersCount	Lil
1	lotusOak2	1,310,580,000,000,000,000.00	Mon Sep 28 14:08:06 +0000 2020 Diane Harper, MD, #HPV Expert, on Risk of #Gar...	42	56	43284	
2	lotusOak2	1,310,580,000,000,000,000.00	Mon Sep 28 13:53:05 +0000 2020 #Vaccines Licensed for Use in the US https://...	36	44	43284	
3	lotusOak2	1,310,570,000,000,000,000.00	Mon Sep 28 13:31:00 +0000 2020 Infant Twins Die Simultaneously After #Vaccine...	114	182	43284	
4	lotusOak2	1,310,570,000,000,000,000.00	Mon Sep 28 13:18:21 +0000 2020 94% of All NYC #Coronavirus Patients Have Unde...	60	55	43284	
...
53903	delbigtree	715,726,000,000,000,000.00	Fri Apr 01 02:24:15 +0000 2016 REPLY @WakeTheFlockUp6 Thanks @Stutzy6	4	1	69039	
53904	delbigtree	715,725,000,000,000,000.00	Fri Apr 01 02:20:22 +0000 2016 Let's get this film heard! #Vaxxed #vaxxedthem...	16	25	69039	
53905	delbigtree	715,725,000,000,000,000.00	Fri Apr 01 02:18:57 +0000 2016 Thank YOU @LaPizzi63 @DrWakefield @AngelikaNew...	3	1	69039	
53906	delbigtree	435,194,000,000,000,000.00	Sun Feb 16 23:30:20 +0000 2014 http://t.co/0uVPg3TDFe	1	1	69039	

	Label	EntryID	Published	comment_text	FavouriteCount	Retweets	UserFollowersCount	Lil
53907	delbigtree	434,791,000,000,000,000.00	Sat Feb 15 20:45:19 +0000 2014	Hi!	3	2	69039	

53908 rows × 9 columns

	Label	EntryID	Published	comment_text	FavouriteCount	Retweets	UserFollowersCount	Lil
0	lotusOak2	1,310,580,000,000,000,000.00	Mon Sep 28 14:09:22 +0000 2020	The Popularity of #Homeopathy by Jeremy Sherr ...	11	8	43284	
1	lotusOak2	1,310,580,000,000,000,000.00	Mon Sep 28 14:08:06 +0000 2020	Diane Harper, MD, #HPV Expert, on Risk of #Gar...	42	56	43284	
2	lotusOak2	1,310,580,000,000,000,000.00	Mon Sep 28 13:53:05 +0000 2020	#Vaccines Licensed for Use in the US https://...	36	44	43284	
3	lotusOak2	1,310,570,000,000,000,000.00	Mon Sep 28 13:31:00 +0000 2020	Infant Twins Die Simultaneously After #Vaccine...	114	182	43284	
4	lotusOak2	1,310,570,000,000,000,000.00	Mon Sep 28 13:18:21 +0000 2020	94% of All NYC #Coronavirus Patients Have Unde...	60	55	43284	
...
53903	delbigtree	715,726,000,000,000,000.00	Fri Apr 01 02:24:15 +0000 2016	REPLY @WakeTheFlockUp6 Thanks @Stutzy6	4	1	69039	

	Label	EntryID	Published	comment_text	FavouriteCount	Retweets	UserFollowersCount	Lik
53904	delbigtree	715,725,000,000,000,000.00	Fri Apr 01 02:20:22 +0000 2016	Let's get this film heard! #Vaxxed #vaxxedthem...	16	25	69039	
53905	delbigtree	715,725,000,000,000,000.00	Fri Apr 01 02:18:57 +0000 2016	Thank YOU @LaPizzi63 @DrWakefield @AngelikaNew...	3	1	69039	
53906	delbigtree	435,194,000,000,000,000.00	Sun Feb 16 23:30:20 +0000 2014	http://t.co/0uVPg3TDFe	1	1	69039	
53907	delbigtree	434,791,000,000,000,000.00	Sat Feb 15 20:45:19 +0000 2014	Hi!	3	2	69039	

53908 rows × 12 columns

```
In [127...
#Assigning the feature names to an empty list
feat_impts = [vect.get_feature_names()]
#For all the models save the feature importances in the list. estimators_ would give the internal models used
for clf in lr.estimators_:
    feat_impts.append(clf.coef_.flatten())
#Saving the results in a dataframe
df_feats_impts = pd.DataFrame(np.transpose(np.array(feat_impts)), columns = ["word", "authoritative", "dialogic", "other"])
#Converting Feature Importance Columns from string to float
df_feats_impts[y_cols] = df_feats_impts[y_cols].astype("float32")
df_feats_impts.head()
```

```
Out[127...
   word authoritative dialogic other
0  aaps             0.14   -0.07 -0.13
1  ability          0.09   -0.02 -0.10
2   able         -0.28    0.32  0.09
3  about           0.83   -0.40 -0.71
```

	word	authoritative	dialogic	other
4	absentee	0.20	-0.05	-0.16

```
In [15]: #Creating Individual Feature Importance table by sorting on specific toxic-type column and selecting top 5 words
authoritative_fi = df_feats_impts[["word","authoritative"]].sort_values(by = "authoritative", ascending = False).head()
dialogic_fi = df_feats_impts[["word","dialogic"]].sort_values(by = "dialogic", ascending = False).head()

#toxic_fi = df_feats_impts[["word","toxic"]].sort_values(by = "toxic", ascending = False).head()
#severe_toxic_fi = df_feats_impts[["word","severe_toxic"]].sort_values(by = "severe_toxic", ascending = False).head()
#obscene_fi = df_feats_impts[["word","obscene"]].sort_values(by = "obscene", ascending = False).head()
#threat_fi = df_feats_impts[["word","threat"]].sort_values(by = "threat", ascending = False).head()
#insult_fi = df_feats_impts[["word","insult"]].sort_values(by = "insult", ascending = False).head()
#identity_hate_fi = df_feats_impts[["word","identity_hate"]].sort_values(by = "identity_hate", ascending = False).head()

#Plotting top 5 words based on coefficient values from the LR model
fig,(ax1, ax2) = plt.subplots(2,3,figsize=(10,5))
sns.barplot(x = "authoritative", y = "word", ax = ax1[0], data = authoritative_fi)
sns.barplot(x = "dialogic", y = "word", ax = ax1[1], data = dialogic_fi)

#sns.barplot(x = "severe_toxic", y = "word", ax = ax1[1], data = severe_toxic_fi)
#sns.barplot(x = "obscene", y = "word", ax = ax1[2], data = obscene_fi)
#sns.barplot(x = "threat", y = "word", ax = ax2[0], data = threat_fi)
#sns.barplot(x = "insult", y = "word", ax = ax2[1], data = insult_fi)
#sns.barplot(x = "identity_hate", y = "word", ax = ax2[2], data = identity_hate_fi)
plt.suptitle("Feature Importance")
fig.tight_layout(rect=[0, 0.03, 1, 0.95])
plt.show()
```

```
-----
NameError                                Traceback (most recent call last)
<ipython-input-15-cb1b1abf7aeb> in <module>
      1 #Creating Individual Feature Importance table by sorting on specific toxic-type column and selecting
    top 5 words
----> 2 authoritative_fi = df_feats_impts[["word","authoritative"]].sort_values(by = "authoritative", ascending = False).head()
      3 dialogic_fi = df_feats_impts[["word","dialogic"]].sort_values(by = "dialogic", ascending = False).head()
    d()
      4
      5

NameError: name 'df_feats_impts' is not defined
```

```
In [ ]:
```

Jeffery A Sternstein, A.B.D.

Education

University of Wisconsin-Milwaukee – (2014 - Present) – Milwaukee, WI

A.B.D. for Ph.D. in English with a focus on Rhetoric and Composition.

Dissertation Topic: A Posthuman Perspective on Digital Ethos – Defense scheduled for April 21st, 2022.

University of Iowa – (1996-1998, 1999-2000) – Iowa City, IA

M.A.T. in Secondary English Education

Comp Exam Topics: Creative Writing in Academic Settings; Popular Fiction; Myth, Legend and Fairy Tale.

University of Iowa – (1992-1996) – Iowa City, IA

B.A. with double major in Psychology and Literature, Science & the Arts.

College-Level Teaching Positions

University of Wisconsin-Milwaukee: College of General Studies – (2021 – Present) – Milwaukee, WI.

University of Wisconsin-Milwaukee: College of Letters and Science – (2015 – 2020, 2021) – Milwaukee, WI.

Oakton Community College – (2019 - 2021) – Des Plaines, IL

College of Lake County – (2020 - 2021) – Grayslake, IL

College-Level Courses Taught

Composition Courses

- First-Year Composition I - First part of the basic composition series, focusing on rhetorical analysis and introductory research skills.
 - Framed writing in consideration for audience, context and purpose.
 - Emphasized the rhetorical consideration of sources.
 - Structured assignments to build on one another to lead to a 5-7 page exploratory research paper.
 - Themes used for primary readings include: digital technology's effect on reading; systemic racism in the college composition class; social issues exacerbated by the COVID quarantine.
 - Schools: UWM:CGS, UWM:L&S, OCC, CLC.
 - Formats: 16-week face-to-face, 16-week asynchronous online, 12-week asynchronous online, 8-week face-to-face.

- First-Year Composition II - Second part of the basic composition series, focusing on analyzing scholarly texts and using those texts to critically interpret other ideas.
 - Framed writing in consideration for audience, context and purpose.
 - Emphasized engagement with scholarly sources.
 - Structured assignments to build on one another to lead to a 10 page argumentative research paper.
 - Themes used for primary readings include: systemic racism in the college composition class; media literacy.
 - Schools: UWM:CGS; UWM:L&S; OCC; CLC.
 - Formats: 16-week face-to-face, 16-week asynchronous online, 8-week face-to-face.

Business Writing Courses

- Introduction to Business Writing – Introductory course, focusing on effective workplace communication and genres.
 - Framed writing in consideration for audience, context and purpose.
 - Emphasized audience awareness.
 - Structured assignment sequence to incorporate multiple genres of office communication revolving around a single job description.
 - Schools: UWM:L&S.
 - Formats: face-to-face.
- Writing in the Professions: Writing and Social Media for Careers –200-level undergraduate course on designing social media campaigns and short-form writing.
 - Framed writing in consideration for audience needs and content value.
 - Emphasized the dialogic nature of writing for social media.
 - Structured assignments to build to a well-developed proposal for an organization's social media campaign.
 - Schools taught at: UWM:L&S.
 - Formats taught in: asynchronous online.
- Writing for the Health Sciences – 200-level undergraduate course, focusing on effective health field related workplace communication and genres.
 - Framed writing in consideration for audience, context and purpose.
 - Emphasized motivational approaches to eliciting behavior change.
 - Structured assignment sequence to incorporate multiple genres of office communication including doctor/patient communication, brochure design, and grant writing.
 - Schools: UWM:L&S.
 - Formats: face-to-face; asynchronous online.

Developmental Courses

- Reading and Academic Development – Developmental reading and study skills course for students two levels below standard college reading.
 - Framed reading in consideration for academic discourse.
 - Emphasized the active reading and inquiry.
 - Structured thematic assignments to encourage engagement and meaningful reflection.
 - Themes used for primary readings include: personal struggles and victimhood culture.
 - Schools: OCC.
 - Formats: face-to-face
- Academic Reading - Developmental reading and study skills course for at-risk students and students on academic probation.
 - Framed reading in consideration for academic discourse.
 - Emphasized the active reading and inquiry.
 - Structured thematic assignments to encourage engagement and meaningful reflection.
 - Themes used for primary readings include: digital technology's effect on reading.
 - Schools: UWM:CGS; OCC.
 - Formats: face-to-face
- Finding Your Pathway – Research-based course for at-risk students on academic goal setting, selecting a college major and preparing for a profession.
 - Emphasized critical thinking and awareness of the decision-making process.
 - Structured assignments to build on one another to lead to a 2 research papers on possible majors and professions.
 - Schools: UWM:CGS; OCC.
 - Formats: face-to-face.

High-School Teaching Positions:

Chicagoland Jewish High School – (2002-2003) – Morton Grove, IL

- Freshman Language Arts – Classical and Ancient Literature.
- Sophomore Language Arts – British Literature.
- Junior Language Arts – American Literature.
- Yearbook Advisor – Supervised the production of the school yearbook.

Cedar Rapids Community Schools – (2000-2001) – Cedar Rapids, IA.

- Freshman Language Arts – World Literature.

- Sophomore Language Arts – American Literature.
- After school Tae Kwon Do – Trained students in martial arts.

Adult-Education Positions:

NOVA of Japan – (2004-2005) – Maebashi-Shi, Gunma-Ken, Japan.

- Conversational English – Taught conversational English to adults and children.

Aon Hewitt – (2006–2014) – Lincolnshire, IL

- Trainer – Adapted and administered special training programs.
- On-Floor Supervisor – Assisted customer service representatives and provided feedback in one-on-one coaching and conferencing.

Additional Achievements:

“Definitions of Human Agency in Digital Composing Spaces.” Presented at *College Composition and Communication Conference (CCCC)*. March 2018.

“Posthuman Futures for Rhetoric, Composition, and Pedagogy.” Presented at *Thomas R. Watson Conference*. October 2018.

Experienced with academic delivery platforms such as Canvas, Desire2Learn (D2L) and Whiteboard.